Book of Abstracts

13th-17th June 2022

11th World Congress of the Bachelier Finance Society

Table of Contents

Louis Bachelier Lecture
June 13, 2022

Title: On the Emergence of Financial Engineering from Financial Science to Improve Sustainable Economic Growth and Stability—History and Future

Speaker: Robert C. Merton

Time: 9:45pm to 10:45pm (HK Time, GMT+8)
2:45pm to 3:45pm (London Time, GMT+1)
9:45am to 10:45am (Eastern Coast Time (US), GMT-4)

Abstract:
A well-functioning financial system plays a critical role in sustainable economic growth and stability. Financial innovation is the process of improving the financial system, and financial engineering is the process of applying financial science to solve real-world problems and implement financial innovation. Beginning in the 1950s, the field of finance underwent a remarkable transformation from a largely conceptual potpourri of anecdotes, rules of thumb, and manipulations of accounting data with wholly inadequate capital market data to test their validity to become a principles-based science, subjected to rigorous empirical examination using large data bases and employing some of the most sophisticated mathematical models of uncertainty and optimization. It was not however until the early 1970s with an explosion of financial innovation that financial engineering emerged to bring the models of finance science directly into the main stream of finance practice, and the two became inexorably intertwined.

In the subsequent half century, financial innovation became a central force generally improving the global financial system with considerable economic benefit having accrued from those changes. The scientific breakthroughs in finance in this period both materially shaped, and were shaped by, the extraordinary wave of innovations in finance practice. Today no major financial institution in the world, including central banks, can function without the computer-based mathematical models of modern financial science and the myriad of derivative contracts and markets used to execute risk-transfer transactions as well as extract price-and risk-discovery information.

In this lecture, we will explore the impact of financial science and engineering on the economy through a series of examples—from the past and forecasts for the future. These examples will illustrate how financial engineering can materially impact the real-economy, both in terms of efficiency gains and environmental sustainability. They will demonstrate how crisis can induce implementation of financial innovations, which not only address the immediate challenges created by the crisis but also create opportunity for continuing, long-term benefits to the economy. We will offer and demonstrate three future trends for financial engineering product and service solutions:

--- Solutions need not be simple; only simple for the users. All essential complexity should reside with the provider, and not with the users.

--- Solutions should be designed based on what the users already know

--- Goal-Based Investment Solutions: Deliver Superior Performance without a Higher Sharpe Ratio

Bio:
Robert C. Merton is the School of Management Distinguished Professor of Finance at the MIT Sloan School of Management and John and Natty McArthur University Professor Emeritus at Harvard University. He serves as Resident Scientist at Dimensional Fund Advisors Inc.

Merton received the Alfred Nobel Memorial Prize in Economic Sciences in 1997 for a new method to determine the value of derivatives. He is past president of the American Finance Association, a member of the National Academy of Sciences, and a Fellow of the American Academy of Arts and Sciences. Merton received the 2021-22 James R. Killian Jr. Faculty Achievement Award, MIT; the Robert A. Muh Award in the Humanities, Arts, and Social Sciences, MIT; Distinguished Alumni Award, California Institute of Technology; the Michael I. Pupin Medal for Service to the Nation, Columbia University; the Kolmogorov Medal, University of London; and the Hamilton Medal, Royal Irish Academy.

Merton has also been recognized for translating finance science into practice. He received the inaugural Financial Engineer of the Year Award from the International Association for Quantitative Finance (formerly International Association of Financial Engineers), which also elected him a Senior Fellow. He received the 2011 CME Group Melamed-Arditti Innovation Award, and the 2013 WFE Award for Excellence from World Federation of Exchanges. A Distinguished Fellow of the Institute for Quantitative Research in Finance ('Q Group') and a Fellow of the Financial Management Association, Merton received the Nicholas Molodovsky Award from the CFA Institute. He is a member of the Halls of Fame of the Fixed Income Analyst Society, Risk, and Derivative Strategy magazines. Merton received Risk’s Lifetime Achievement Award for contributions to the field of risk management and the 2014 Lifetime Achievement Award from the Financial Intermediation Research Society. He received the 2017 Finance Diamond Prize from Fundación de Investigación, IMEF, Mexico. Merton received the 2022 Plan Sponsor Council of America Lifetime Achievement Award for his contribution to retirement finance.

Merton received a BS in Engineering Mathematics from Columbia University, a MS in Applied Mathematics from California Institute of Technology, a PhD in Economics from Massachusetts Institute of Technology and numerous honorary degrees from US and foreign universities.

Plenary Talk II
June 13, 2022

Title: Learning to Trade: Data-Driven Quantitative Finance
Speaker: Hans Buehler
Time: 10:45pm to 11:45pm (HK Time, GMT+8)
3:45pm to 4:45pm (London Time, GMT+1)
10:45am to 11:45am (Eastern Coast Time (US), GMT-4)

Abstract:
We discuss our application of "AI" machine learning methods to optimal hedging of a portfolio of OTC derivatives in incomplete markets of liquid derivatives with friction under convex risk measures. The "Deep Hedging" approach differs from most of quantitative finance in that we focus on realistic market dynamics of the statistical measure, incorporate real life market frictions, and take into account the portfolio manager's risk aversion. Most importantly our framework is numerically accessible and used in a production environment.
As a side product, we present the first numerical implementation of a stochastic implied volatility model.

Bio:
Dr Hans Buehler was until recently global head of Equities "Analytics, Automation, and Optimization" as well as of the Equities, Sales, and Securities Services Quantitative Research teams at JP Morgan. He will soon take on a new role as Deputy CEO at XTX Markets.

Hans holds a PhD in Quantitative Finance from Technical University Berlin, and a Diploma in Stochastic Analysis from Humboldt University Berlin. He is a visiting professor at Technical University Munich. He was voted *Quant of the Year* by Risk Magazine in 2022 for his work on Machine Learning and AI for trading.

Plenary Talk I
June 14, 2022

Title: Adapted Wasserstein Distance for Robust Stochastic Optimization
Speaker: Beatrice Acciaio
Time: 9:45pm to 10:45pm (HK Time, GMT+8)
2:45pm to 3:45pm (London Time, GMT+1)
9:45am to 10:45am (Eastern Coast Time (US), GMT-4)

Abstract:
The causality constraint in optimal transport allows to capture the temporal structure of the transported objects, which is crucial when transporting measures on path spaces. In its symmetric formulation, this leads to the definition of adapted Wasserstein distances. I will present several applications of those concepts to illustrate their suitability for multiple problems in mathematical finance and stochastic analysis.

Bio:
Beatrice Acciaio is Professor of Mathematics at ETH Zurich since 2020. Before joining ETH, Beatrice was associate professor at the London School of Economics, and prior to that she has been part of several research groups, at the Technical University of Vienna, the University of Perugia, and the University of Vienna. Beatrice completed her PhD in 2006 under the supervision of Walter Schachermayer.
Beatrice's main areas of research are probability, mathematical finance, and optimal transport. Beatrice is member of the Council of the Bachelier Finance Society, she is Associate Editor for the SIAM Journal on Financial Mathematics, for Finance and Stochastics, for Mathematical Finance, and for the Bocconi & Springer Series on Mathematics, Statistics, Finance and Economics.

Plenary Talk II

June 13 Monday

Title: Time (In)Consistency of multivariate Problems
Speaker: Birgit Rudloff
Time: 10:45pm to 11:45pm (HK Time, GMT+8)
3:45pm to 4:45pm (London Time, GMT+1)
10:45am to 11:45am (Eastern Coast time (US), GMT-4)

Abstract:

In several time-inconsistent problems, the time-inconsistency is due to the fact that the underlying problem is multi-variate in some sense. Examples include the mean-risk problem, Nash equilibria in a dynamic game, or dynamic risk measure in a market with frictions. What unifies these examples is that one can formulate these problems with a set-valued value function.

In this talk, the Bellman's principle is extended to value functions that are set-valued. It is shown that the problems mentioned above do satisfy this Bellman's principle under reasonable assumptions and are thus actually time-consistent in a set-valued sense. That means, that problems that were previously thought to be time in-consistent, are in fact time-consistent in a sense that takes the multi-variateness of the underlying problem into account, and thus provides a more appropriate concept of time-consistency. Practical implications, economic interpretations, and the connection to the scalar time-(in)consistency are discussed. Numerical examples are given which lead in the discrete time setup to a sequence of vector optimization problems solved backwards in time.

Even a dynamic game fits into this structure, as we can show that the set of all Nash equilibria of a non-cooperative game coincides with the set of all Pareto solutions of a certain vector optimization problem. This is true for all shared constraint games, and in a similar form also for generalized games and vector-valued games. This is a join work with Zachary Feinstein and Gabriela Kovacova.

Bio:

Birgit Rudloff is a full professor of financial mathematics and optimization and the vice chair of the Institute for Statistics and Mathematics at Vienna University of Economics and Business. She is the speaker of the PhD label Mathematics in Economics and Business at WU Vienna, as well as co-PI of the Vienna Graduate School on Computational Optimization. Born in East-Germany, she did her Ph.D. in Financial Mathematics at Martin-Luther-University Halle-Wittenberg, spend one year at the Brazilian National Institute for Pure and Applied Mathematics (IMPA) in Rio de Janeiro, and worked as a PostDoc at TU Vienna. From 2006 until 2015 Birgit Rudloff was an Assistant Professor at Princeton University at the Department of Operations Research and Financial Engineering, and affiliated with the Bendheim Center for Finance, before coming to WU Vienna in 2015.

Her research centers around multivariate risks. She works on multivariate dynamic optimization problems, measuring and regulating systemic risk in banking networks, dynamic risk measures in markets with frictions, dynamic games, and time (in-)consistency, as well as on algorithms to solve vector optimization problems. Birgit Rudloff published articles in Financial Mathematics and Optimization in renowned journals including Operations Research, Finance and Stochastics, Bernoulli, SIAM Journal on Financial Mathematics, and Mathematical Programming.

Plenary Talk I
June 15, 2022

Title: Designing Stablecoins
Speaker: Steven Kou
Time: 9:45pm to 10:45pm (HK Time, GMT+8)
2:45pm to 3:45pm (London Time, GMT+1)
9:45am to 10:45am (Eastern Coast Time (US), GMT-4)

Abstract:
Existing cryptocurrencies are too volatile to be used as currencies for daily payments. Stablecoins, which are cryptocurrencies pegged to other stable financial assets such as the U.S. dollar, are desirable for payments within blockchain networks, often called the “Holy Grail of cryptocurrency.” By using the option pricing theory and the Ethereum platform that allows running smart contracts, we design several dual-class structures that are written on the ETH cryptocurrency and offer a fixed income crypto asset (class A coin), a stablecoin (class A′ coin) pegged to a traditional currency and leveraged investment instruments (class B and B′ coins). Our investigation of the values of stablecoins in the presence of jump risk and black-swan-type events shows the robustness of the design. The design has been implemented on the Ethereum platform.

Bio:
Steven Kou is a Questrom Professor in Management and Professor of Finance at Boston University. He teaches courses on FinTech and quantitative finance. Currently he is a co-area-editor for Operations Research and a co-editor for Digital Finance, and has served on editorial boards of many journals, such as Management Science, Mathematics of Operations Research, and Mathematical Finance. He is a fellow of the Institute of Mathematical Statistics and won the Erlang Prize from INFORMS in 2002. Some of his research results have been incorporated into standard MBA textbooks and have implemented in commercial software packages and terminals.

Plenary Talk II
June 15, 2022

Title: Overcoming the curse of dimensionality in computational finance: from nonlinear Monte Carlo to deep learning
Speaker: Arnulf Jentzen
Time: 10:45pm to 11:45pm (HK Time, GMT+8)
3:45pm to 4:45pm (London Time, GMT+1)
10:45am to 11:45am (Eastern Coast Time (US), GMT-4)

Abstract:
Partial differential equations (PDEs) are among the most universal tools used in modelling problems in nature and man-made com- plex systems. For example, stochastic PDEs are a fundamental ingredient in models for nonlinear filtering problems in chemical engineering and weather forecasting, deterministic Schroedinger PDEs describe the wave function in a quantum physical system, deterministic Hamiltonian-Jacobi-Bellman PDEs are employed in operations research to describe optimal control problems where companies aim to minimise their costs, and deterministic Black- Scholes-type PDEs are highly employed in portfolio optimization models as well as in state-of-the-art pricing and hedging models for financial derivatives. The PDEs appearing in such models are often high-dimensional as the number of dimensions, roughly speaking, corresponds to the number of all involved interacting substances, particles, resources, agents, or assets in the model. For instance, in the case of the above mentioned financial engineering models the dimensionality of the PDE often corresponds to the number of financial assets in the involved hedging portfolio. Such PDEs can typically not be solved explicitly and it is one of the most challenging tasks in applied mathematics to develop approximation algorithms which are able to approximatively compute solutions of high-dimensional PDEs. Nearly all approximation algorithms for PDEs in the literature suffer from the so-called ”curse of dimensionality” in the sense that the number of re- quired computational operations of the approximation algorithm to achieve a given approximation accuracy grows exponentially in the dimension of the considered PDE. With such algorithms it is impossible to approximatively compute solutions of high-dimensional PDEs even when the fastest currently available computers are used. In the case of linear parabolic PDEs and approximations at a fixed space-time point, the curse of dimensionality can be overcome by means of Monte Carlo approximation algorithms and the Feynman-Kac formula. In this talk we rigorously prove that deep neural networks do indeed overcome the curse of dimensionality in the case of a general class of semi-linear parabolic PDEs and we thereby prove, for the first time, that a general semi-linear parabolic PDE with a nonlinearity depending on the PDE solution can be solved approximatively without the curse of dimensionality.

Bio:
Arnulf Jentzen is appointed as a presidential chair professor at the School of Data Science and the Shenzhen Research Institute of Big Data at the Chinese University of

Hong Kong, Shenzhen (since 2021) and as a full professor at the Faculty of Mathematics and Computer Science at the University of Münster (since 2019). In 2004 he started his undergraduate studies in mathematics (minor field of study: computer science) at Goethe University Frankfurt in Germany, in 2007 he received his diploma degree at this university, and in 2009 he completed his PhD in mathematics at this university. The core research topics of his research group are machine learning approximation algorithms, computational stochastics, numerical analysis for high dimensional partial differential equations (PDEs), stochastic analysis, and computational finance. He is particularly interested in deep learning based algorithms for high dimensional approximation problems and different kinds of differential equations. At the moment he serves as an associate, division, or managing editor for the Annals of Applied Probability (AAP, since 2019), for Communications in Computational Phyiscs (CiCP, since 2021s), for Communications in Mathematical Sciences (CMS, since 2015), for Discrete and Continuous Dynamical Systems Series B (DCDS-B, since 2018), for the Journal of Applied Mathematics and Physics (ZAMP, since 2016), for the Journal of Complexity (JoC, since 2016), for the Journal of Machine Learning (JML, since 2021), for the Journal of Mathematical Analysis and Applications (JMAA, since 2014), for the SIAM / ASA Journal on Uncertainty Quantification (JUQ, since 2020), for the SIAM Journal on Scientific Computing (SISC, since 2020), for the SIAM Journal on Numerical Analysis (SINUM, since 2016), and for Partial Differential Equations and Applications (PDEA, since 2019). His research activities have been recognized by several scientific prizes. In particular, in 2020 he has been awarded the Felix Klein Prize from European Mathematical Society (EMS) and in 2022 he has been awarded the Joseph F. Traub Prize for Achievement in Information-Based Complexity. Further details on the activities of his research group can be found at the webpage http://www.ajentzen.de.

Plenary Talk I
June 16, 2022

Title: Differential Access to Dark Markets and Execution Outcomes

Speaker: Carole Comerton-Forde

Time: 9:45pm to 10:45pm (HK Time, GMT+8)
2:45pm to 3:45pm (London Time, GMT+1)
9:45am to 10:45am (Eastern Coast Time (US), GMT-4)

Abstract:
We compare dark pool trades across exchange-operated dark pools, where all trader types have equal access, and broker-operated dark pools, where brokers can restrict access to exclude certain traders, such as high frequency traders. Conditional on execution, trades on broker dark pools have less information leakage and adverse selection than trades on exchange-operated dark pools. Broker dark pools that do not allow high frequency traders have less information leakage than those that do. Differences in execution outcomes are concentrated in smaller trades. We conclude that the ability to segment order flow can improve execution outcomes for investors.

Bio:
Carole Comerton-Forde is Professor of Finance at University of Melbourne. Her research in market structure examines how the mechanics of the market, such as regulation and technology, impact prices, liquidity and trader behaviour.
Her research focuses on equities markets but spans many geographies including Europe, US, Canada, Hong Kong and Australia. Her current research interests include fragmentation, dark pools and the connection between the cost of raising capital and secondary market liquidity. Her research has been published in leading academic journals including the Journal of Finance, the Journal of Financial Economics and the Journal of Financial and Quantitative Analysis. Carole has previously held academic positions at UNSW, the Australian National University and University of Sydney, and visiting positions at New York University and the London School of Economics and Political Science. She was also Visiting Economist at the New York Stock Exchange. She has acted as a consultant for a number of stock exchanges and market regulators around the world. She is currently an economic consultant for the Australian Securities and Investments Commission, an Academic Advisor to the Plato Partnership and a Research Fellow at the Centre for Economic Policy Research.

Plenary Talk II
June 16, 2022

Title: Dynamic Trading: Price Inertia and Front-Running
Speaker: Yuliy Sannikov
Time: 10:45pm to 11:45pm (HK Time, GMT+8)
3:45pm to 4:45pm (London Time, GMT+1)
10:45am to 11:45am (Eastern Coast Time (US), GMT-4)

Abstract:
We build a linear-quadratic model of trading in a market with private information and heterogeneous agents. Agents receive private taste/inventory shocks and trade continuously. Agents have different costs of holding excessive inventory, which may stem from different absolute risk aversion. In equilibrium, trade is gradual. Trading speed depends on the number of participants and their size. Trade among large market participants is slower than that among small ones. Price has momentum due to the actions of large traders: it drifts down if sellers are more patient to trade than buyers. Traders infer total supply from trading flow of an individual. This inference affects the willingness to accept flow and the price impact of the individual trader. The model can also answer welfare questions, e.g. about the social costs and benefits of market consolidation. The model can accommodate individual inventory shocks as well as private information about common value.

Bio:
Yuliy Sannikov is a Ukrainian economist known for his contributions to mathematical economics, game theory, and corporate finance. He is an economics professor at the Stanford Graduate School of Business, and won both the 2015 Fischer Black Prize and 2016 John Bates Clark Medal.

Yuliy Sannikov develops new methods for analyzing continuous time dynamic games using stochastic calculus methods. His work paves elegant and powerful ways to resolve lots of problems. It has not only broken new ground in methodology, but also had a substantial influence on applied theory. He has significantly altered the toolbox available for studying dynamic games, and as a result of his contributions, new areas of economic inquiry have become tractable for rigorous theoretical analysis. The areas of application include the design of securities, contract theory, macroeconomics with financial frictions, market microstructure, and collusion. He is widely regarded as one of the few theorists in many years to have introduced a truly novel tool that changed the way theory is done.

He received his A.B. in mathematics from Princeton in 2000 and earned a Ph.D. in business administration from Stanford Graduate School of Business in 2004. Sannikov is also one of the few participants to win three gold medals at the International Mathematical Olympiad.

Plenary Talk I
June 17, 2022

Title: Mean field game of mutual holding and systemic risk

Speaker: Nizar Touzi

Time: 9:45pm to 10:45pm (HK Time, GMT+8)
2:45pm to 3:45pm (London Time, GMT+1)
9:45am to 10:45am (Eastern Coast Time (US), GMT-4)

Abstract:
We introduce a mean field model for optimal holding of a representative agent of her peers as a natural expected scaling limit from the corresponding $N-$agent model. The induced mean field dynamics appear naturally in a form which is not covered by standard McKean-Vlasov stochastic differential equations. We study the corresponding mean field game of mutual holding in the absence of common noise. Our first main result provides an explicit equilibrium of this mean field game, defined by a bang-bang control consisting in holding those competitors with positive drift coefficient of their dynamic value. As a second main result, we use this mean field game equilibrium to construct (approximate) Nash equilibria for the corresponding N-player game. All of these results extend to the defaultable agent setting. Our last main result is a characterization of the default probability.

Bio:
Nizar Touzi is Professor at Ecole Polytechnique in Paris. He is one of the world-leading experts in stochastic optimal control theory and mathematical finance, He was invited speaker at ICM (Hyderabad) in 2010, laureate of Best Young Researcher in Finance Award by the Europlace Institute of Finance in 2007, and laureate of Louis Bachelier Prize in 2012. He serves as co-editor and associate editor for several journals in the field of stochastic analysis and mathematical finance, including Finance and Stochastic, Stochastic Processes and their Applications, etc. More recently, his main research interest is in the principal-agent problem and mean field game/control theory.

Plenary Talk II
June 17, 2022

Title: Mean Field Games and Heterogeneous Interactions

Speaker: Daniel Lacker

Time: 10:45pm to 11:45pm (HK Time, GMT+8)
3:45pm to 4:45pm (London Time, GMT+1)
10:45am to 11:45am (Eastern Coast Time (US), GMT-4)

Abstract:
The now-widespread theory of mean field games provides a systematic framework for modeling stochastic dynamic games with many interacting players. It is fundamentally limited, however, to models in which players view each other as exchangeable. This talk discusses some recent efforts to push past this limitation by incorporating heterogeneous interactions, such as those governed by networks, which are relevant in applications to systemic risk. In the simplest case, each player is labeled by a vertex in a graph and interacts symmetrically with its nearest neighbors. Different methods and phenomena arise in sparse versus dense regimes, the latter being more tractable by taking advantage of well-developed limit theories for dense graphs, such as graphons. Constructions of approximate equilibria are possible, similar to those of mean field game theory, but only for sufficiently regular and dense graphs, or alternatively by weakening the notion of approximate equilibrium.

Bio:
Daniel Lacker is an assistant professor in the Department of Industrial Engineering and Operations Research at Columbia University. He was an NSF postdoctoral fellow in the Division of Applied Mathematics at Brown University from 2015 to 2017. Prior to that, he received his Ph.D. from Princeton University in 2015 and his B.S. from Carnegie Mellon University in 2010. He is a recipient of an NSF CAREER award and a SIAG-FME early career prize. Daniel's primary research areas are mean field games and interacting particle systems, which form the mathematical foundation for a wide range of models of large-scale interactions arising in physics, engineering, economics, and finance.

A Constrained Nonzero-Sum Stochastic Differential Game Application

E. Savku

University of Oslo, Department of Mathematics, Oslo, Norway

ARTICLE HISTORY
Compiled March 15, 2022

Extended Abstract

We develop an approach for two player constraint nonzero-sum stochastic differential game, which is modeled by Markov regime-switching jump-diffusion processes as follows:

$$\begin{aligned} Y(t) = {} & b(t,Y(t),\alpha(t),u_1(t),u_2(t))dt \\ & + \sigma(t,Y(t),\alpha(t),u_1(t),u_2(t))dW(t) \\ & + \int_{\mathbb{R}_0} \eta(t,Y(t-),\alpha(t-),u_1(t-),u_2(t-),z)\tilde{N}_\alpha(dt,dz) \\ & + \gamma(t,Y(t-),\alpha(t-),u_1(t-),u_2(t-))d\tilde{\Phi}(t), \qquad t \in [0,T]. \qquad (1) \\ Y(0) = {} & y_0 \in \mathbb{R}^N, \qquad (2) \end{aligned}$$

where $\alpha(\cdot)$ represents a time-inhomogenous, finite state Markov chain and its state space is $S = \{e_1, e_2, ..., e_D\}$, where $D \in \mathbb{N}$, $e_i \in \mathbb{R}^D$ and the jth component of e_i is the Kronecker delta δ_{ij} for each pair of $i, j = 1,2,...,D$.
We provide the relations between a usual stochastic optimal control setting and a Lagrangian method. In this context, we prove corresponding theorems for two different type of constraints, which lead us to find real valued and stochastic Lagrange multipliers, respectively.
Let $u_1 \in \Theta_1$ and $u_2 \in \Theta_2$ be two admissible control processes for Player 1 and Player 2, respectively. We define the performance criterion for each player as follows:

$$J_k(y,e_i,u_1,u_2) = E^{y,e_i}\left[\int_0^T f_k(s,Y(s),\alpha(s),u_1(s),u_2(s))ds + g_k(Y^{u_1,u_2}(T),\alpha(T))\right]$$

for each $e_i \in S$, $i = 1,2,\dots,D$, and both purpose to maximize their payoffs with respect to other player's best action as follows:

$$J_1(y,e_i,u_1^*,u_2^*) = \sup_{u_1\in\Theta_1} J_1(y,e_i,u_1,u_2^*), \qquad (3)$$

$$J_2(y,e_i,u_1^*,u_2^*) = \sup_{u_2\in\Theta_2} J_2(y,e_i,u_1^*,u_2), \qquad (4)$$

CONTACT E. Savku. Email: esavku@gmail.com

for each $e_i \in S$ and for all $y \in G$, where G is an open subset of $\mathbb{R}^N$ and corresponds to a solvency region for the state processes. Then, the pair of optimal control processes $(u_1^*, u_2^*) \in \Theta_1 \times \Theta_2$ is called a Nash equilibrium for the stochastic differential game of the system (1)-(2) and (3)-(4).
Our constrained nonzero-sum stochastic differential game is to find out (u_1^*, u_2^*) for the problems (3)-(4) subject to the system (1)-(2) and

$$\textbf{(i)}\ E^{y,e_i}[M_k(Y^{u_1,u_2}(T), \alpha(T))] = 0 \tag{5}$$

or

$$\textbf{(ii)}\ M_k(Y^{u_1,u_2}(T), \alpha(T)) = 0 \quad a.s., \tag{6}$$

where $M_k : \mathbb{R}^N \to \mathbb{R}$, are C^1 functions with respect to y and we assume that $E[M_k(Y^{u_1,u_2}(T), \alpha(T))] < \infty,\ k = 1,2$.
Finally, our unconstrained nonzero-sum stochastic differential game problem is described as follows:

$$\begin{aligned}\phi_k^{\lambda_k}(y,e_i) = J_k(y,e_i,u_1^{*,\lambda_1},u_2^{*,\lambda_2}) = \sup_{u_k \in \Theta_k} E^{y,e_i}\Bigg[\int_0^T f_k(t,Y(t),\alpha(t),u_1(t),u_2(t))dt \\ + g_k(Y^{u_1,u_2}(T),\alpha(T)) + \lambda_k M_k(Y^{u_1,u_2}(T),\alpha(T))\Bigg]\end{aligned} \tag{7}$$

for $k = 1,2$ and $e_i \in S,\ i = 1,2,\ldots,D$, subject to the system (1)-(2). Then, we illustrate our results for an example of cooperation between a bank and an insurance company, which is a popular, well-known business agreement type, called Bancassurance. By using stochastic maximum principle, we investigated optimal dividend strategy for the company as a best response according to the optimal mean rate of return choice of a bank for its own cash flow and vice versa. We found out a Nash equilibrium for this game and solved the adjoint equations explicitly for each state.
It is well known that the timing and the amount of dividend payments are strategic decisions for companies. The announcement of a dividend payment may reduce or increase the stock prices of a company. From the side of the bank, it is clear that creating a cash flow with high returns would be the main goal.
Hence, in our formulation, we provide an insight to both of the bank and the insurance company about their best moves in a bancassurance commitment under specified technical conditions.

Reference: E. Savku, A stochastic control approach for constrained stochastic differential games with jump and regimes. Submitted 2022.

A Data-Driven Deep Learning Approach for Options Market Making

Qianhui Lai[*] Xuefeng Gao[†] Lingfei Li[‡]

February 27, 2022

Abstract

We develop a data-driven approach for options market making. Using stock options data from CBOE, we find that both buy and sell orders exhibit strong self-excitation but insignificant cross-excitation. We show that a Hawkes process with a time-varying base intensity and a power law kernel provides good fit to the data of market orders for stock options. To solve the optimal market making problem for a single option, we approximate the market making strategy at each decision time by a neural network and train them to optimize the expected utility of market making profits. We study feature selection for the neural networks and compare the out-of-sample performance of the optimal neural network strategy trained from data generated by the Hawkes model and a simple model. We find that using the more realistic Hawkes model improves the out-of-sample performance substantially. Furthermore, utilizing the Hawkes process intensity or the expected number of market order arrivals computed under the Hawkes model as an additional input feature can improve the performance. We also show how to solve the market making problem for option portfolios with Greeks and inventory constraints using neural network approximation.

[*]Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: laiqianhui@sz.pku.edu.cn

[†]Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: xfgao@se.cuhk.edu.hk.

[‡]Corresponding author. Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: lfli@se.cuhk.edu.hk.

A Deep Learning Scheme for Multi-Asset Optimal Investment Problem with Transaction Cost

Maxim Bichuch* Ke Chen †

15th March 2022

In this paper, we present a deep learning algorithm for solving the problem of optimal investment with a proportional transaction cost. The optimal investment problem with a single risky asset was first solved by Merton (Merton, 1969) and the proportional transaction costs were incorporated by (Magill and Constantinides, 1976). However, there is still no known solution to the (correlated) multi-assets problem. Moreover, while the value function in the single risky and illiquid asset case is known to be (almost entirely) C^2, only C^1 regularity is known in the multi-assets case. These two issues together make a development of a robust numerical algorithm in the multi-asset case very difficult, and the proof of convergence of algorithms that exists in case of single illiquid asset cannot be readily generalized to prove convergence in the multi-assets case. Therefore, our goal is to develop a robust deep learning algorithm and to show its convergence in the correlated, risky and illiquid multi-assets market.

We consider the problem of maximizing the expected utility of the terminal wealth, over all admissible strategies. We assume the utility is a power utility. The market consists of one risk-free, liquid asset, growing at a constant risk-free rate, and of multiple correlated risky and illiquid assets. The latter are driven by a number of Brownian Motions, and each evolving as a geometric Brownian Motion. An agent is allowed to trade the assets, by selling (respectively buying) a risky asset, and buying (respectively selling) the risk-free asset, or any combination of the above. It is assumed that proportional transaction costs are charged for any such trade.

It is well known that the value function is the unique viscosity solution of the Hamilton-Jacobi-Bellman (HJB) equation. Therefore, we develop a numerical method to (approximately) find this solution. We construct the loss function as L^2 norms of each of the components of the HJB equation, together with the terminal condition. That is at interior points, we evaluate each component of the HJB equation and take the minimum value among the components, whereas at terminal time points, we compute the difference between the approximation function and terminal condition value. Finally, we set the sum of square these quantities at the mesh points as the loss function. We approximate the derivatives in the loss function using finite difference. For the convergence proof to work, we use the coordinate transformation and wide stencil to deal with the second-order derivatives. Using this loss function, we then deep learn the solution to the HJB equation and prove that it is a good approximation to the value function, and gives a good approximation to the optimal strategy.

Specifically, we prove that:

*Department of Applied Mathematics and Statistics, Whiting School of Engineering, Johns Hopkins University, 3400 North Charles St., Baltimore, MD 21218, USA mbichuch@jhu.edu.

†Department of Applied Mathematics and Statistics, Whiting School of Engineering, Johns Hopkins University, 3400 North Charles St., Baltimore, MD 21218, USA kchen78@jhu.edu.

1. For any small $\epsilon > 0$, with sufficiently large number of neurons on the hidden layer and hyperbolic tangent function as the activation function, there exists an output of the shadow network which makes the loss function smaller than the ϵ.

2. When the loss function is smaller than ϵ, the error between it and the solution to the finite difference approximation would be $O(\sqrt{\epsilon})$, which implies that the neural network solution converge to the viscosity solution of the HJB equation locally uniformly as the mesh size and ϵ go to 0.

Finally, we numerically illustrate the applicability and the robustness of this algorithm in the case of one and two risky and illiquid assets. We find the optimal strategies, given by the no-trade region. We investigate the effects of correlation, transaction costs, and changes to model parameters. Additionally, we verify our solution against the known asymptotic expansions in the case of one illiquid asset and in the case of two uncorrelated illiquid assets.

References

Michael JP Magill and George M Constantinides. Portfolio selection with transactions costs. *Journal of economic theory*, 13(2):245–263, 1976.

Robert C Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The review of Economics and Statistics*, pages 247–257, 1969.

A finite mixture modelling perspective for combining experts' opinions with an application to quantile-based risk measures

Despoina Makariou (d.makariou@lse.ac.uk),
Pauline Barrieu (p.m.barrieu@lse.ac.uk), and George Tzougas (g.tzougas@lse.ac.uk)
Department of Statistics, London School of Economics and Political Science

June 16, 2021

Extended Abstract

Due to the growing uncertainty in an abundance of contemporary societal settings, we often come across circumstances when an agent, who acts on behalf of another party, is called to make a decision by combining multiple and sometimes diverging sources of information that can be described as opinions. Moreover, the latter may take any form; from experts to forecasting methods or models (see Clemen and Winkler (2007)), and from now on, we may use these terms interchangeably when referring to an opinion. Opinions communicated to an agent can differ to varying degrees, and the level of confidence that an agent allocates to any given viewpoint is subjective. In quantitative risk management, the process of defining, measuring and managing operational risk is crucial since it formalizes the financial institutions' approaches to comply with the qualifying qualitative criteria of the Basel Capital Accord and Solvency Directive. This approach relies on the knowledge of experienced enterprise agents and risk management experts who are asked to provide opinions regarding plausible high-severity events. For instance, these opinions can be expressed as parameters of an assumed loss distribution. However, the company's risk profile, which could accord to a consensus of experts' individual judgements regarding the severity distribution, might often not be robustly estimated. The main reason for this is that when experts are presented with internal data and need to express probabilistic opinions about the same uncertain quantity of interest, there may be multiple sources of heterogeneity in their responses concerning the choice of models and their parameters and, in addition to these, the allocation of weights from the agent that are not considered as being embedded in the data-generative process of the uncertain quantity of interest based on which the agent needs to make a decision. In particular, each expert reports their opinion based on what their focus is, and if we assume that they report their opinions honestly, each believes that their opinion reflects best the true data-generative process. Therefore, since a major challenge in operational risk management is to evaluate the exposure of severe losses based on a weighted combination of a variety of opinions in the first place, it appears that it would make more sense to employ probabilistic models that reflect group structures.

In this paper, we present an alternative perspective for modelling of operational risk in an enterprise context by combining expert opinions based on finite mixture models. Finite mixtures models can provide a formal framework for clustering and classification that

can be effectively used within the opinions combination research setting. In particular, this versatile and easily extensible class of models can accommodate different sources of unobserved heterogeneity in the data-generative process of the uncertain quantity of interest by allowing for the mixture components to represent groups within which there is a concurrence of judgements. At this point, it is worth noting that finite mixtures models have not been applied in the area of opinion combinations, with the exception of Rufo et al. (2010), who employed Bayesian hierarchical models based on mixtures of conjugate prior distributions for merging expert opinions. Furthermore, it should be noted that Shevchenko and Wüthrich (2006) employed the Bayesian inference method for quantifying frequency and severity distributions in the context of operational risk. Their approach was based on specifying the prior distributions for the parameters of the frequency and severity distributions based on expert opinions or external data. Furthermore, Lambrigger et al. (2009) extended the framework of the previous paper by developing a Bayesian inference model that permits combining internal data, external data, and expert opinions simultaneously. The setup they proposed enlarged the Bayesian inference models of the exponential dispersion family (EDF) and their corresponding conjugate priors; see, for instance, Bühlmann and Gisler (2006), Chapter 2. However, to the best of our knowledge, the use of finite mixture models within the traditional frequentist approach for combining diverging opinions remains a largely uncharted research territory.

Our main contribution is that we consider that the component distributions can stem from different parametric families. The advantage of this formulation is that it allows the agent to obtain the aggregated opinion of a group of experts, based on a linear opinion pool, and account for the various sources of unobserved heterogeneity in the decision-making process in the following ways: (i) by assuming that the data are drawn from a finite mixture distribution with components representing different opinions about both the distribution family and its parameters regarding the uncertain quantity of interest, and (ii) via the mixing weights that reflect the quality of each opinion. Furthermore, when the proposed family of models is applied to internal data, it can enable the agent to utilize all the available information for accurately assessing the effectiveness of (i) the combination of the expert judgements and (ii) their own judgement about the weights that they intended to allocate to each expert—a concept not so dissimilar to the the main idea behind the long-established weights allocation approach of Cooke et al. (1991) and the scoring rules in general. Finally, the proposed family of models is used for numerically computing quantile-based risk measures, which are of interest in a variety of different types of insurance problems, such as setting premiums, insurance deductibles, and reinsurance cedance levels and determining reserves or capital and ruin probabilities. We compare our approach to the traditional weighted average one, and we find that they lead to different results.

References

Bühlmann, H. and Gisler, A. (2006). *A course in credibility theory and its applications*. Springer Science & Business Media.

Clemen, R. T. and Winkler, R. L. (2007). *Aggregating Probability Distributions*, page 154–176. Cambridge University Press.

Cooke, R. et al. (1991). *Experts in uncertainty: opinion and subjective probability in science*. Oxford University Press on Demand.

Lambrigger, D. D., Shevchenko, P. V., and Wüthrich, M. V. (2009). The quantification of operational risk using internal data, relevant external data and expert opinions. *arXiv preprint arXiv:0904.1361*.

Rufo, M., Pérez, C., and Martín, J. (2010). Merging experts' opinions: A bayesian hierarchical model with mixture of prior distributions. *European journal of operational research*, 207(1):284–289.

Shevchenko, P. V. and Wüthrich, M. V. (2006). The structural modelling of operational risk via bayesian inference: Combining loss data with expert opinions. *The Journal of Operational Risk*, 1(3):3–26.

A General Approach for Parisian Stopping Times under Markov Processes

Gongqiu Zhang* Lingfei Li†

December 28, 2021

Abstract

We propose a method based on continuous time Markov chain approximation to compute the distribution of Parisian stopping times and price Parisian options under general one-dimensional Markov processes. We prove the convergence of the method under a general setting and obtain sharp estimate of the convergence rate for diffusion models. Our theoretical analysis reveals how to design the grid of the CTMC to achieve faster convergence. Numerical experiments are conducted to demonstrate the accuracy and efficiency of our method for both diffusion and jump models. To show the versatility of our approach, we develop extensions for multi-sided Parisian stopping times, the joint distribution of Parisian stopping times and first passage times, Parisian bonds and for more sophisticated models like regime-switching and stochastic volatility models.

*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China. Email: zhanggongqiu@cuhk.edu.cn.

†Corresponding author. Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: lfli@se.cuhk.edu.hk.

A general approach for solving behavioral optimal investment problems with non-concave utility functions and probability weighting

Xiuchun Bi * Lvning Yuan † Zhenyu Cui ‡ Jiacheng Fan §
Shuguang Zhang ¶

Abstract In this paper, we consider the optimal investment problem with both probability distortion/weighting and general non-concave utility functions with possibly finite number of inflection points. Our model contains the model under cumulative prospect theory (CPT) as a special case, which has inverse S-shaped probability weighting and S-shaped utility function (i.e. one inflection point). Existing literature has shown the equivalent relationships (strong duality) between the concavified problem and the original one by either assuming the presence of probability weighting or the non-concavity of utility functions, but not both. In this paper, we combine both features and propose a step-wise relaxation method to handle general non-concave utility functions and probability distortion functions. The necessary and sufficient conditions on eliminating the duality gap for the Lagrange method have been provided under this circumstance. We have applied this solution method to solve several representative examples in mathematical behavioral finance including the CPT model, Value-at-Risk based risk management (VAR-RM) model with probability distortions, Yarri's dual model and the goal reaching model. We obtain a closed-form optimal trading strategy for a special example of the CPT model, where a "distorted" Merton line has been shown exactly. The slope of the "distorted" Merton line is given by an inflation factor multiplied by the standard Merton ratio, and an interesting finding is that the inflation factor is solely dependent on the probability distortion rather than the non-concavity of the utility function.

Keywords: Non-concave Utility, Probability Distortion, Concavification, Lagrange Duality, Relaxation Method.

MSC(2020): Primary: 49J55, 91B16; Secondary: 49K45, 91G80.

*School of Mathematics and Statistics, Guizhou University of Finance and Economics, Guiyang Guizhou, 550025, PR China. Email: xcbi@ustc.edu.cn

†PhD candidate at Department of Statistics and Finance, School of Management, University of Science and Technology of China, China. Email: yln1994@mail.ustc.edu.cn

‡Corresponding author. School of Business, Stevens Institute of Technology, 1 Castle Point on Hudson, United States. Email: zcui6@stevens.edu

§Postdoctoral Research Fellow, Chinese University of Hong Kong, Hong Kong. Email: jcfan@cuhk.edu.hk

¶Co-Corresponding author. Department of Statistics and Finance, School of Management, University of Science and Technology of China, China. Email: sgzhang@ustc.edu.cn

A General Method for Analysis and Valuation of Drawdown Risk under Markov Models

Gongqiu Zhang* Lingfei Li†

December 31, 2021

Abstract

Drawdown risk is a major concern in financial markets. We develop a novel algorithm to solve the first passage problem of the drawdown process for general one-dimensional Markov processes (including time-inhomogeneous ones) as well as regime-switching and stochastic volatility models. We compute its Laplace transform based on continuous time Markov chain (CTMC) approximation and numerically invert the Laplace transform to obtain the first passage probabilities and the distribution of the maximum drawdown. We prove convergence of our algorithm for general Markovian models and provide sharp estimate of the convergence rate for a general class of jump-diffusion models. We apply the algorithm to three financial applications: (1) price maximum drawdown options and hedge the risk of selling such options with a highly volatile asset as the underlying; (2) calculate the Calmar ratio for investment analysis; (3) quantify the contribution of each asset to the portfolio's drawdown risk when the assets follow multivariate exponential Lévy models. Various numerical experiments document the computational efficiency of our method.

*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China. Email: zhanggongqiu@cuhk.edu.cn.

†Corresponding author. Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: lfli@se.cuhk.edu.hk.

A linear-rational multi-curve term structure model with stochastic spread

José Da Fonseca* Komi Edem Dawui† Yannick Malevergne‡

February 24, 2022

Abstract

We develop a linear-rational multi-curve term structure model based on the Wishart affine process. Within this framework, we compute the zero-coupon bond price, the swap rate and develop a semi-closed form expression for the price of a swaption that can be computed efficiently. We also show how more complex interest rate derivatives such as the constant maturity swap and the constant maturity swap spread option can be priced. We provide a swaption price approximation that is fast to evaluate and accurate. We illustrate how well the model performs on real data by rolling a calibration using a 3-month long sample of at-the-money swaption data. We find that the estimated parameters are remarkably stable and the calibration procedure is robust. In particular, we show that thanks to the specific Wishart properties the model can handle the stochastic correlation between the OIS term structure and the Euribor-OIS spread term structure.

JEL Classification: G12; G13; C61
Keywords: Interest rate model, Multi-Curve, Wishart process, Stochastic spread, Swaption market.

*Auckland University of Technology, Business School, Department of Finance, Private Bag 92006, 1142 Auckland, New Zealand. Phone: +64 9 9219999 extn 5063. Email: jose.dafonseca@aut.ac.nz and PRISM Sorbonne EA 4101, Université Paris 1 Panthéon - Sorbonne, 17 rue de la Sorbonne, 75005 Paris, France.

†The World Bank, 1818 H Street NW, Washington, D.C., 20433, United States of America. Email: edawui@worldbank.org and PRISM Sorbonne EA 4101, Université Paris 1 Panthéon - Sorbonne, 17 rue de la Sorbonne, 75005 Paris, France. The content of this article represents the views of the authors, and does not represent the opinion of the World Bank Group.

‡PRISM Sorbonne EA 4101, Université Paris 1 Panthéon - Sorbonne, 17 rue de la Sorbonne, 75005 Paris, France. Email:yannick.malevergne@univ-paris1.fr

1 Introduction

Following the global financial crisis (GFC), interest rate models were revisited to take into account the widening of the spread between the overnight interest swap (OIS) term structure given by Eonia swaps in the European interest rate market and the Euribor term structure. It led to what is commonly named nowadays the multi-curve models. These models are more challenging as they need to specify the dynamic of each curve but also the interaction between these curves making the problem a multidimensional problem that is naturally more complicated. Interest rate derivatives such as caps/floors and swaptions become more difficult to price and manage due to this additional complexity while exotic derivatives are even more challenging.

Notice that during the pre-GFC period (*i.e.*, single-curve models), swaptions were already difficult to price as it is a product that is intrinsically multidimensional. Even if there are some approximation formulas for the swaption price, see for example Collin-Dufresne and Goldstein (2002), Singleton and Umantsev (2002) or Schrager and Pelsser (2006), the numerical difficulty is such that only very few empirical studies on the swaption market are available in the literature. It is in sharp contrast with the equity derivatives literature where comparisons between different model specifications were extensively performed. As a consequence, it should not come as a surprise that in the multi-curve case, which is more challenging numerically, the swaption market is barely analyzed.

In Rogers (1997), the author shows how standard interest rate models fit into the framework of the potential approach.[1] Roughly speaking, it amounts to conveniently choose a stochastic process and a function. The author also shows how to generate new interest rate models. Among those new models, the linear-rational model, which owes its name to the fact that the zero-coupon bond is a linear-rational function of the state variable, is of particular interest as the pricing of swaptions is extremely simple and at par, in terms of computational difficulty, with caps/floors. Several works investigated the linear-rational interest rate framework, Nakamura and Yu (2000) and Macrina (2014) in the single-curve case and for the multi-curve case Nguyen and Seifried (2015), Macrina and Mahomed (2018) and Filipović et al. (2017) to name a few. However, among multi-curve works, those performing an empirical analysis of the swaption market are not so many. To the best of our knowledge, such kind of results can be found only in Nguyen and Seifried (2015) who calibrate the model using 1-day of ATM swaption quotes, Crépey et al. (2015b) who calibrate the model using 4 days of swaption quotes (with different strikes, so not only ATM swaptions) and Filipović et al. (2017) who calibrate simultaneously 866 weekly ATM swaption quotes.[2]

The swaption market is an important interest rate derivatives market. According to Skantzos and Garston (2019), as of June 2018, the monthly trading volume of the interest rate options market is approximately 1.5 trillion USD, two thirds of which comes from swaption trades and a further 125 billion USD from the cap/floor market. As such, the swaption market is a major component of the interest rate derivatives market. Therefore, building a model that can be calibrated easily on swaption data so that its performance can be analyzed is extremely important.

Following the works of Rogers (1997) and Filipović et al. (2017), we propose a linear-rational multi-curve term structure model based on the Wishart process. Within this framework, the zero-coupon bond price, whose value depends on the risk-free overnight interest swap (OIS) curve, and the spread between the Euribor and OIS curves are linear-rational functions of the state variable given by a Wishart process. The model enables a stochastic correlation between these two curves, as the Wishart process allows a non trivial correlation between its components, and thus the model captures such a kind of dependency that is prevalent in the EUR interest rate derivatives market. We derive a pricing formula for swaptions whose numerical cost is at par with caps and floors. We also show how more exotic interest rate derivatives can be priced in that framework. We develop an approximation for the swaption price, which enables a fast pricing, by adjusting Collin-Dufresne and Goldstein

[1] By standard interest rate models we are referring to the exponential affine framework that builds upon Duffie and Kan (1996) and constitutes the dominant part of the interest rate literature. We refer to Da Fonseca et al. (2013), Moreni and Pallavicini (2014), Morino and Runggaldier (2014), Crépey et al. (2015a), Grbac et al. (2015), Grasselli and Miglietta (2016), Cuchiero et al. (2016), Cuchiero et al. (2019) or Alfeus et al. (2020) just to name a few.

[2] The data used in Crépey et al. (2015b) are also used in Crépey et al. (2015b) but they need to rely on Singleton and Umantsev (2002) to price swaptions at the model is standard exponential affine.

(2002)'s methodology. We fully exploit the computational simplicity of the model by rolling a calibration over a 3-month sample of daily ATM swaption prices. The calibrated parameters are extremely stable thereby showing the ability of the model to handle the daily fluctuation of the data. We show how information regarding the correlation between the two curves can be extracted from swaptions but also the advantages of the Wishart process to manage this dependency. The results convincingly demonstrate the performance of the Wishart process as modeling tool for interest rate derivatives.

The structure of the paper is as follows. Section 2 reminds the main analytical properties of the Wishart process. In section 3, the interest rate model is specified and we make explicit the pricing formulas for different interest rate products. Section 4 presents the data and illustrates how well the model performs in practice. Section 5 concludes the paper while proofs and tables are gathered in the appendix.

2 The Wishart process

Given a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{P})$ we denote by $\mathbb{E}[\,\cdot\,]$ (resp. $\mathbb{E}_t[\,\cdot\,] := \mathbb{E}[\,\cdot\,|\mathcal{F}_t]$) the expectation (resp. conditional expectation) under the probability measure $\mathbb{P}$. The Wishart process, proposed in Bru (1991) and introduced in finance in Gouriéroux and Sufana (2010), satisfies the matrix stochastic differential equation

$$dx_t = (\omega + mx_t + x_t m^\top)dt + \sqrt{x_t}dw_t\sigma + \sigma^\top dw_t^\top \sqrt{x_t}\,, \tag{1}$$

where x_t is an $n \times n$ matrix that belongs to the set of positive definite matrices denoted $\mathbb{S}_n^{++}$, m, σ belong to the set of $n \times n$ real matrices denoted $\mathsf{M}(n)$, $\{w_t; t \geq 0\}$ is a matrix Brownian motion of dimension $n \times n$ (*i.e.*, a matrix of n^2 independent scalar Brownian motions) under the probability measure $\mathbb{P}$ and $\cdot^\top$ stands for the matrix transposition.[3] The matrix $\omega \in \mathbb{S}_n^{++}$ satisfies certain constraints involving $\sigma^\top\sigma$ to ensure the positiveness of the matrix process x_t. Note that the transpositions in Eq. (1) are necessary to preserve the symmetry of the solution. The quantity $\sqrt{x_t}$ is well defined since $x_t \in \mathbb{S}_n^{++}$. The matrix m is such that $\{\Re(\lambda_i^m) < 0; i = 1, \ldots, n\}$ where $\lambda_i^m \in Spec(m)$ for $i = 1, \ldots, n$ and $Spec(m)$ is the spectrum of the matrix m while $\Re(\,\cdot\,)$ stands for the real part. The matrix σ belongs to $\mathsf{GL}_n(\mathbb{R})$ the general linear group over $\mathbb{R}$ (*i.e.*, the set of real invertible matrices). Thanks to the invariance of the law of the Brownian motion to rotations and the polar decomposition of σ, we can assume that $\sigma \in \mathbb{S}_n^{++}$.

The infinitesimal generator of the Wishart process is given by (Bru, 1991):

$$\mathcal{G} = \mathrm{tr}[(\omega + mx + xm^\top)D + 2xD\sigma^2 D]\,, \tag{2}$$

where D is the $(n \times n)$ matrix operator $D_{ij} := \partial_{x_{ij}}$.

Using Eq. (1), one can establish the following relations for the quadratic covariations of the components of the Wishart process:

$$\begin{aligned}
d\langle x_{11,.}, x_{11,.}\rangle_t &= 4x_{11,t}(\sigma_{11}^2 + \sigma_{12}^2)dt\,, &(3)\\
d\langle x_{22,.}, x_{22,.}\rangle_t &= 4x_{22,t}(\sigma_{12}^2 + \sigma_{22}^2)dt\,, &(4)\\
d\langle x_{12,.}, x_{12,.}\rangle_t &= x_{11,t}(\sigma_{12}^2 + \sigma_{22}^2)dt + 2x_{12,t}(\sigma_{11}\sigma_{12} + \sigma_{12}\sigma_{22})dt + x_{22,t}(\sigma_{11}^2 + \sigma_{12}^2)dt\,, &(5)\\
d\langle x_{11,.}, x_{12,.}\rangle_t &= 2x_{11,t}(\sigma_{11}\sigma_{12} + \sigma_{12}\sigma_{22})dt + 2x_{12,t}(\sigma_{11}^2 + \sigma_{12}^2)dt\,, &(6)\\
d\langle x_{12,.}, x_{22,.}\rangle_t &= 2x_{12,t}(\sigma_{12}^2 + \sigma_{22}^2)dt + 2x_{22,t}(\sigma_{11} + \sigma_{22})\sigma_{12}dt\,, &(7)\\
d\langle x_{11,.}, x_{22,.}\rangle_t &= 4x_{12,t}(\sigma_{11}\sigma_{12} + \sigma_{12}\sigma_{22})dt\,. &(8)
\end{aligned}$$

Bru (1991) showed that the Wishart process is affine, that is the moment generating function is exponentially affine in the state variable. More precisely, the moment generating function is given by

$$\Phi(t, \theta_1, \theta_2, x) := \mathbb{E}\left[\exp\left(\mathrm{tr}[\theta_1 x_t] + \int_0^t \mathrm{tr}[\theta_2 x_u]du\right)\right]\,, \tag{9}$$

[3] By definition, w_t is an $(n \times n)$ matrix Brownian motion if and only if $\forall u, v \in \mathbb{R}^n$, $(w_t u, w_t v)$ is a vector Brownian motion with covariance structure $\mathrm{cov}_t\,[dw_t u, dw_t v] = u^\top v I_n dt$ with I_n the $n \times n$ identity matrix.

where θ_1, θ_2 belong to $\mathbb{S}_n$ the set of real $n \times n$ symmetric matrices, $\mathrm{tr}[\,\cdot\,]$ stands for the trace of a matrix and $\mathbb{E}\,[\,\cdot\,] := \mathbb{E}\,[\,\cdot\,|x_0 = x]$.

Following Grasselli and Tebaldi (2008), it is possible to prove that

$$\Phi(t, \theta_1, \theta_2, x_0) = \exp\left(\mathrm{tr}[a(t, \theta_1, \theta_2)x_0] + b(t, \theta_1, \theta_2)\right) , \tag{10}$$

with the deterministic functions $(a(t, \theta_1, \theta_2), b(t, \theta_1, \theta_2))$, where $a(t, \theta_1, \theta_2)$ is an $n \times n$ matrix function and $b(t, \theta_1, \theta_2)$ a scalar function, satisfying the system

$$a' = am + m^\top a + 2a\sigma^2 a + \theta_2 \,, \tag{11}$$

$$b' = \mathrm{tr}[\omega a] \,, \tag{12}$$

with initial conditions $a(0, \theta_1, \theta_2) = \theta_1$ and $b(0, \theta_1, \theta_2) = 0$. As usual $\cdot'$ denotes the time derivative.

Eq. (11) is a Matrix Riccati ordinary differential equation (ODE) whose solution is given by

$$a(t, \theta_1, \theta_2) = (\theta_1 A_{12}(t) + A_{22}(t))^{-1}(\theta_1 A_{11}(t) + A_{21}(t)) \,, \tag{13}$$

where

$$\begin{pmatrix} A_{11}(t) & A_{12}(t) \\ A_{21}(t) & A_{22}(t) \end{pmatrix} := \exp\left\{ t \begin{pmatrix} m & -2\sigma^2 \\ \theta_2 & -m^\top \end{pmatrix} \right\} , \tag{14}$$

Eq. (12), along with the corresponding initial condition, leads to $b(t)$ after integration.

We denote by e_{ij} the basis of $\mathsf{M}(n)$, it is the $n \times n$ matrix with 1 in the (i, j) place and zero elsewhere, then $x_{ij,t} = \mathrm{tr}[e_{ij} x_t]$. Then,

$$d\mathbb{E}[x_t] = (\omega + m\mathbb{E}[x_t] + \mathbb{E}[x_t]m^\top)dt \,, \tag{15}$$

that leads if $n = 2$ and m is diagonal to the ODEs

$$d\mathbb{E}[x_{11,t}] = (\omega_{11} + 2m_{11}\mathbb{E}[x_{11,t}])dt \,, \tag{16}$$

$$d\mathbb{E}[x_{22,t}] = (\omega_{22} + 2m_{22}\mathbb{E}[x_{22,t}])dt \,, \tag{17}$$

and we conclude that $\mathbb{E}[x_{11,t}]$ only depends on ω_{11}, m_{11} and $x_{11,0}$ and not on $x_{12,0}$. Similarly, $\mathbb{E}[x_{22,t}]$ only depends on ω_{22}, m_{22} and $x_{22,0}$ and not on $x_{12,0}$. For $x_{12,t}$, we get

$$d\mathbb{E}[x_{12,t}] = (\omega_{12} + (m_{11} + m_{22})\mathbb{E}[x_{12,t}])dt \,. \tag{18}$$

It implies that even if $x_{12,0} = 0$, we can have $\mathbb{E}[x_{12,t}] \neq 0$ for $t > 0$ if $\omega_{12} \neq 0$.

If the process $(x_t)_{t\geq 0}$ is stationary then $\bar{x}_\infty = \lim_{t\to+\infty} \mathbb{E}[x_t]$ satisfies the matrix equation

$$m\bar{x}_\infty + \bar{x}_\infty m^\top = -\omega \,. \tag{19}$$

The Wishart process was initially defined and analyzed in Bru (1991) under the assumption that $\omega = \beta\sigma^2$ with $\beta \in \mathbb{R}_+$ such that $\beta \geq n+1$ to ensure that $x_t \in \mathbb{S}_n^{++}$. Hereafter, this specification will be referred to as the Bru case. It was later extended in Mayerhofer et al. (2011) (see also Cuchiero et al. 2011) to the case $\omega \in \mathbb{S}_n^{++}$ and proved that if

$$\omega \succeq \beta\sigma^2 \,, \tag{20}$$

with $\beta \geq n + 1$ (where Eq. (20) means that $\omega - \beta\sigma^2 \in \mathbb{S}_n^{++}$) then $x_t \in \mathbb{S}_n^{++}$. From a financial modeling point of view, the advantage of having ω not so tightly related to the volatility matrix σ is that they are naturally estimated using different financial products, which gives the model a flexibility that is often necessary in the applications.

The moment generating function Eq. (10) gives the Laplace transform of the process x_t as the following proposition shows.

Proposition 2.1. *Define*

$$\Xi_t := -\frac{1}{2}\int_0^t e^{(t-s)m}(-2\sigma^2)e^{(t-s)m^\top}ds\,, \tag{21}$$

$$\Lambda_t := \Xi_t^{-1}e^{mt}x_0e^{m^\top t}\,, \tag{22}$$

then the Laplace transform of x_t in the Bru case (i.e., $\omega = \beta\sigma^2$) rewrites as

$$\mathbb{E}_{x_0}\left[\operatorname{etr}(-\theta_1 x_t)\right] = \det\left(I + 2\Xi_t\theta_1)\right)^{-\beta/2}\operatorname{etr}\left(-\frac{\Lambda_t^\top}{2}2\Xi_t\theta_1(I+2\Xi_t\theta_1)^{-1}\right)\,, \tag{23}$$

for $\theta_1 \in \mathbb{S}_n^{++}$.

The Laplace transform Eq. (23) is known to be associated with the density of a non-central Wishart distribution. Indeed, if X is a random variable with non-central Wishart distribution, it takes values in $\mathbb{S}_n^{++}$ and its law is denoted by $W_n(\beta,\Xi,\Lambda)$, with $\beta \geq n$, $\Xi \in \mathbb{S}_n^{++}$ and $\Lambda \in \mathsf{M}(n)$. The density of X, reported in Gupta and Nagar (2000, Eq. 3.5.1 p. 114) for example, is given by

$$f(X) = \frac{2^{-\frac{n\beta}{2}}}{\Gamma_n(\beta/2)}\det(\Xi)^{-\frac{\beta}{2}}\operatorname{etr}\left(-\frac{\Lambda}{2} - \frac{\Xi^{-1}X}{2}\right)\det(X)^{\frac{\beta-n-1}{2}}{}_0F_1\left(\frac{\beta}{2};\frac{1}{4}\Lambda\Xi^{-1}X\right)\,, \tag{24}$$

with $X \in \mathbb{S}_n^{++}$, $\Gamma_n(z)$ with $z \in \mathbb{C}$ the multivariate gamma function defined in Gupta and Nagar (2000, Eq. 1.4.5 p. 18) and ${}_0F_1(a;Z)$ with $a \in \mathbb{C}$ and $Z \in \mathsf{M}(n)$ is the hypergeometric function of matrix argument (see Gupta and Nagar, 2000, p. 34 for a definition). According to Gupta and Nagar (2000, Theorem 1.4.1 p. 19) the following relation between the multivariate gamma and the standard gamma function (of scalar argument) holds $\Gamma_n(z) = \pi^{\frac{1}{4}n(n-1)}\prod_{i=1}^n\Gamma\left(z-(i-1)/2\right)$. In Gupta and Nagar (2000), $\beta \in \mathbb{N}$ while one consequence of Bru (1991) is to extend to the case $\beta \in \mathbb{R}$ with $\beta \geq n+1$ (see Mayerhofer 2019 and reference therein). An efficient numerical algorithm to compute hypergeometric function of matrix argument appears in Koev and Edelman (2006) and its first use in quantitative finance can be found in Kang et al. (2017).

3 A multicurve model

3.1 The OIS and Euribor-OIS term structure curves

We follow Filipović et al. (2017), who build upon the potential approach proposed by Rogers (1997), in order to develop a two-curve model based on a 2×2 Wishart process (*i.e.*, $n=2$). First, we define a pricing kernel as:[4]

$$\zeta_t := e^{-\alpha t}(1 + x_{11,t})\,, \tag{25}$$

with $\alpha \in \mathbb{R}_+$ and $(x_{11,t})_{t\geq 0}$ is the $(1,1)^{th}$ element of a Wishart process $(x_t)_{t\geq 0}$. According to Rogers (1997), the pricing kernel can be rewritten as follows. Define the positive function $f : \mathbb{S}_2^{++} \to \mathbb{R}^+$ such that $f(x) := 1+\operatorname{tr}[e_{11}x]$. Define $g(x) := (\alpha - \mathcal{G})f(x)$, it is a positive function for α large enough (*i.e.*, $\alpha > \operatorname{tr}[\omega]$).

The pricing kernel allows us to compute the time t value of a collateralized zero-coupon bond with maturity T, denoted $P(t,T)$, that is given by

$$P(t,T) := \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_s ds}\right] = \mathbb{E}_t\left[\frac{\zeta_T}{\zeta_t}\right]\,, \tag{26}$$

$$= e^{-\alpha(T-t)}\frac{1+\mathbb{E}_t[x_{11,T}]}{1+x_{11,t}}\,, \tag{27}$$

[4]Note that Filipović et al. (2017) suggests to consider a function of the form $e^{-\alpha t}(a_0 + a_1 x_{11,t})$ with $a_0 > 0$ and $a_1 > 0$ but for identification reasons, clearly explained in Filipović et al. (2017, Theorem 5), one needs to impose $a_0 = 1$ and $a_1 = 1$.

with $\mathbb{E}_t^{\mathbb{Q}}[\cdot]$ the (conditional) expectation under the risk neutral probability $\mathbb{Q}$ equivalent to $\mathbb{P}$ under which zero-coupon bond prices are martingale.

The expectation in Eq. (27) can be explicitly computed thanks to the affine property of the Wishart process and for the particular specification adopted here is very simple as the following proposition shows.

Proposition 3.1. *Under the assumption that the matrix m in Eq.* (1) *is diagonal, Eq.* (16) *holds and the zero-coupon bond is given by*

$$P(t,T) = e^{-\alpha(T-t)} \frac{b_1(T-t) + a_1(T-t)x_{11,t}}{1 + x_{11,t}}\,, \tag{28}$$

with $b_1(t) := 1 + \frac{\omega_{11}}{2m_{11}}(e^{2m_{11}t} - 1)$ and $a_1(t) := e^{2m_{11}t}$.

Note that the zero-coupon bond price only depends on $x_{11,t}$, which is a consequence of the diagonal form chosen for m. This assumption is for convenience of the exposition only and can be relaxed at the expense of more cumbersome formulas. One striking property of linear-rational models, whether they are built upon the standard affine process of Duffie and Kan (1996) or the Wishart process, is that the bond price does not depend on the volatility structure of the process. This has a strong consequence in terms of model implementation as it enables the calibration of the parameters ω_{11}, m_{11} and $x_{11,0}$, from the bond yield curve only.

According to Rogers (1997, Eq. 2.4), the short rate is given by

$$\begin{aligned} r_t &= \frac{(\alpha - \mathcal{G})f}{f}\,, && (29)\\ &= \alpha - \frac{\omega_{11} + 2m_{11}x_{11,t}}{1 + x_{11,t}}\,, && (30) \end{aligned}$$

and is positive by construction as α is such that $g(x)$ is positive. Also, as m has negative eigenvalues then $x_{11,t}$ is stationary and it is straightforward to check from Eq. (28) the following result

$$\lim_{T\to+\infty} -\frac{1}{T-t}\ln P(t,T) = \alpha\,, \tag{31}$$

so that α is the infinite-maturity zero-coupon bond yield as in Filipović et al. (2017). It gives a very simple way to estimate the parameter α from the zero-coupon bond price.

The discount factor $P(T, T+\delta)$ is related to the time T overnight indexed swap (OIS) rate with maturity $T+\delta$ by the formula

$$\mathrm{OIS}(T, T+\delta) = \frac{1}{\delta}\frac{1 - P(T,T+\delta)}{P(T,T+\delta)}. \tag{32}$$

The above formula holds for an OIS with maturity less than one year.

Additionally we consider the Euribor rate $L(T, T+\delta)$, which is the rate at time T for the period $[T, T+\delta]$. Let us denote by $\mathrm{Spread}(T, T+\delta)$, the spread between the Euribor and OIS rates, this is the difference between $L(T,T+\delta)$ and $\mathrm{OIS}(T,T+\delta)$ and is often called the Euribor-OIS spread. Before the global financial crisis, the spread was negligible but after the crisis it widened significantly and a multi-curve interest rate model aims at taking into account that spread and its stochastic evolution. The Euribor-OIS spread is defined by:

$$\begin{aligned} \mathrm{Spread}(T,T+\delta) &:= L(T,T+\delta) - \mathrm{OIS}(T,T+\delta)\,, && (33)\\ &= L(T,T+\delta) - \frac{1}{\delta}\left(\frac{1}{P(T,T+\delta)} - 1\right). && (34) \end{aligned}$$

Similar to the approach in the appendix of Filipović et al. (2017), we specify for the time T deflated value of the Euribor-OIS spread payment at time $T+\delta$ a linear functional of the stochastic process $(x_{22,t})$, the $(2,2)^{th}$

element of the Wishart process. More precisely, the deflated time-T value of the Euribor-OIS spread time-$T+\delta$ payment is defined as[5]

$$\zeta_T P(T, T+\delta)\delta \text{Spread}(T, T+\delta) = e^{-\alpha T} x_{22,T}. \tag{35}$$

Once the deflated value of the Euribor-OIS spread payment at a future date is specified, its expectation gives the value of the spread as a (linear-rational) function of the process as shown in the next proposition.

Proposition 3.2. *The time-t value of the Euribor-OIS spread payment set at time T and made at time $T+\delta$, simply called the (time-t value) Euribor-OIS spread, is given by:*

$$A(t, T, T+\delta) = \frac{1}{\zeta_t}\mathbb{E}_t\left[\zeta_T P(T, T+\delta)\delta \text{Spread}(T, T+\delta)\right], \tag{36}$$

$$= \frac{1}{\zeta_t}\mathbb{E}_t\left[e^{-\alpha T} x_{22,T}\right], \tag{37}$$

$$= e^{-\alpha(T-t)}\frac{b_2(T-t) + a_2(T-t)x_{22,t}}{1 + x_{11,t}}, \tag{38}$$

with $b_2(t) := \frac{\omega_{22}}{2m_{22}}(e^{2m_{22}t} - 1)$ and $a_2(t) := e^{2m_{22}t}$ if we assume that m is diagonal.

Notice that according to Eq. (27) the OIS term structure depends on $x_{11,t}$ whilst the Euribor-OIS spread depends on $x_{22,t}$ and as these two processes are stochastically correlated, the model proposed here is an interest rate multi-curve model with stochastic spread, as Filipović et al. (2017)'s model is, but with the additional important property that the Euribor-OIS spread is correlated with the OIS term structure.[6] What is more, the correlation can take any sign thanks to the property of the Wishart process. Let us now have a closer look at the rich correlation structure generated by the model.

3.2 Correlation structure

In Filipović et al. (2017), the time T deflated value of the Euribor-OIS spread payment at time $T+\delta$ given by Eq. (35) is an affine function of a standard vector affine process that is *independent* of the standard vector affine process that drives the OIS term structure given by Eq. (25). This independence is motivated by the empirical finding in Filipović and Trolle (2013) which states that the OIS term structure and the Libor-OIS spread term structure are not correlated. For the Eonia-Euribor market the empirical correlation between the Eonia and Euribor-Eonia curves is not null (as we shall see later in the empirical section), it is the main motivation to introduce the Wishart process to capture such a dependency.

Let us denote $F(T_1 - t, x_{11,t}) := P(t, T_1)$ the bond price with maturity T_1 given by Eq. (28). One can check that

$$\partial_{x_{11}} F = e^{-\alpha(T_1-t)}\left(1 - \frac{\omega_{11}}{2m_{11}}\right)\frac{(e^{2m_{11}(T_1-t)} - 1)}{(1+x_{11})^2} \leq 0, \tag{39}$$

since $m_{11} < 0$, whilst the Euribor-OIS spread $G(T_2 - t, x_{11,t}, x_{22,t}) := A(t, T_2, T_2+\delta)$ satisfies

$$\partial_{x_{11}} G = -\frac{G}{1+x_{11}} \leq 0, \tag{40}$$

$$\partial_{x_{22}} G = e^{-\alpha(T_2-t)}\frac{a_2(T_2-t)}{1+x_{11}} \geq 0, \tag{41}$$

[5]Filipović et al. (2017, online appendix) suggests to specify the right hand side of Eq. (35) as $e^{-\alpha T}(1 + x_{22,T})$ but we found that specification rather inconvenient as the left hand side of Eq. (35) can be arbitrarily small, if for example the spread is small, and as $x_{22,t}$ is a positive process it can lead to calibration problems. In fact, the specification Eq. (35) matches the one of Rogers (1997, Example 3.7).

[6]Let us stress the fact that even if the factor driving the discounted spread in Eq. (35) is independent of the factor driving the pricing kernel in Eq. (25), the expected spread given by Eq. (38) does depend on the factor driving the pricing kernel as the presence of $x_{11,t}$ in Eq. (38) clearly shows.

therefore the instantaneous covariance between the OIS zero-coupon bond and the Euribor-OIS spread is given by

$$d\langle P(\cdot,T_1),A(\cdot,T_2,T_2+\delta)\rangle_t=\partial_{x_{11}}G\partial_{x_{11}}Fd\langle x_{11,\cdot},x_{11,\cdot}\rangle_t+\partial_{x_{22}}G\partial_{x_{11}}Fd\langle x_{11,\cdot},x_{22,\cdot}\rangle_t\,. \tag{42}$$

Suppose that $\sigma_{12}=0$, then Eq. (8) implies that the right hand side of Eq. (42) comprises only the leftmost term that is positive thanks to Eq. (3), Eq. (39) and Eq. (40). We conclude that the OIS zero-coupon bond and the Euribor-OIS spread are positively correlated in that particular case. Notice that even if x_{22} is independent of x_{11} the Euribor-OIS spread depends on x_{11} as Eq. (38) clearly shows. Thanks to the second term in Eq. (42), the correlation between the OIS zero-coupon bond and the Euribor-OIS spread of the Wishart multi-curve model proposed here can display any sign. Indeed, Eq. (8), Eq. (39) and Eq. (41) imply that the sign of the second term is $-\mathrm{sign}(x_{12}\sigma_{12})$. So if x_{12} and σ_{12} have the same signs, the second term in Eq. (42) can lead, if it is large enough in absolute terms, to a negative correlation between the OIS bond price and the Euribor-OIS spread. As such, the Wishart multi-curve model possesses a stochastic basis whose correlation with the OIS term structure is stochastic and can take any sign.

Also of interest is the instantaneous covariance of the Euribor-OIS spread term structure. Let $\tau_1=T_1-t$ and $\tau_2=T_2-t$ two maturities and $A(t,T_1,T_1+\delta)$ and $A(t,T_2,T_2+\delta)$ the Euribor-OIS spreads with time to maturity τ_1 and τ_2, respectively. The instantaneous covariance between those two Euribor-OIS spreads is given by

$$\begin{aligned}\mathrm{COV}(\tau_1,\tau_2)=&\,\partial_{x_{11}}G(\tau_1)\partial_{x_{11}}G(\tau_2)4x_{11,t}(\sigma_{11}^2+\sigma_{12}^2)+\partial_{x_{22}}G(\tau_1)\partial_{x_{22}}G(\tau_2)4x_{22,t}(\sigma_{12}^2+\sigma_{22}^2)\\&+\left(\partial_{x_{11}}G(\tau_1)\partial_{x_{22}}G(\tau_2)+\partial_{x_{22}}G(\tau_1)\partial_{x_{11}}G(\tau_2)\right)4x_{12,t}\sigma_{12}(\sigma_{11}+\sigma_{22})\,.\end{aligned} \tag{43}$$

From Eq. (39) and Eq. (41) we deduce that the first two terms of the right hand side of Eq. (43) are positive whilst the last term's sign is $-\mathrm{sign}(x_{12,t}\sigma_{12})$. If $\sigma_{12}\neq 0$, the covariance between the Euribor-OIS spreads depends on a factor that does not impact the OIS term structure nor the Euribor-OIS term structure. It is an unspanned stochastic volatility factor (USV). Further to this, the Wishart multicurve model's additional factor x_{12} can take any sign so the last term of Eq. (43) can mitigate the first two terms that are always positive.

3.3 Swaption pricing

The pricing of nonlinear derivatives is important as they are used to calibrate the model on liquid products such as caps/floors and swaptions, often called vanilla products, so that the calibrated model can then be used to price exotic derivatives. It is commonly said that exotic products are priced "consistently" with vanilla products. With exponential affine models, the pricing of caps/floors is often simple from a numerical point of view but, in contrast, the pricing of swaptions is often excessively difficult.

In order to derive the value of a swaption in the Wishart model, let us firt compute the time-t value, denoted $C(t,T,T+\delta)$, of a floating coupon fixed at time T and paying $\delta L(T,T+\delta)$ at time $T+\delta$ as

$$\begin{aligned}C(t,T,T+\delta)&=\frac{1}{\zeta_t}\mathbb{E}_t\left[\zeta_{T+\delta}\delta L(T,T+\delta)\right]\,, &(44)\\&=\frac{1}{\zeta_t}\mathbb{E}_t\left[\zeta_T P(T,T+\delta)\delta L(T,T+\delta)\right]\,, &(45)\\&=P(t,T)-P(t,T+\delta)+A(t,T,T+\delta)\,. &(46)\end{aligned}$$

Then, let us consider an interest rate swap starting at T_0 and maturing at T_{n_1} where the Euribor based floating leg payment dates are $T_1,\cdots,T_{n_1}$, with $T_j-T_{j-1}=\delta$ for $j=1,\ldots,n_1$, the fixed leg payment rate K and the fixed leg payment dates are $t_1,\cdots,t_{m_1}=T_{n_1}$, $t_i-t_{i-1}=\Delta$ for $i=1,\ldots,m_1$ and $t_0=T_0$. The time $t<T_0$ value of the floating leg of the swap is $\sum_{j=1}^{n_1}C(t,T_{j-1},T_j)=P(t,T_0)-P(t,T_{n_1})+\sum_{j=1}^{n_1}A(t,T_{j-1},T_j)$ while the fixed leg value is $\Delta K\sum_{i=1}^{m_1}P(t,t_i)$. So the fixed-rate payer swap value at time t is

$$\Pi_t^{\mathrm{swap}}=P(t,T_0)-P(t,T_{n_1})+\sum_{j=1}^{n_1}A(t,T_{j-1},T_j)-\Delta K\sum_{i=1}^{m_1}P(t,t_i)\,. \tag{47}$$

The time-t forward swap rate, denoted $S_t^{T_0,T_{n_1}}$, is

$$S_t^{T_0,T_{n_1}} = \frac{P(t,T_0) - P(t,T_{n_1}) + \sum_{j=1}^{n_1} A(t,T_{j-1},T_j)}{\Delta \sum_{i=1}^{m_1} P(t,t_i)} \,. \tag{48}$$

Remark 3.3. *The spot swap rate can be obtained from Eq.* (48) *by taking* $t = T_0$ *and, combined with the zero-coupon bonds extracted from the OIS curve, allows the computation of the current time value of the Euribor-OIS spread, that is the terms* $\{A(T_0,T_{i-1},T_i+\delta); i = 1,\dots,n_1\}$.[7] *These terms can then be used in Eq.* (38) *to estimate the parameters* ω_{22}, m_{22} *and* x_{22,T_0}. *This calibration strategy is consistent with the structure of the model that suggests to stage the estimation procedure.*

Given Eq. (47) for the fixed-rate payer swap value at time t, we can derive the value of the corresponding swaption. A striking property of the linear-rational model based on the affine process (whether it be vector or matrix) is the relative simplicity of the swaption pricing formula as the following proposition shows.

Proposition 3.4. *The value at time* $t < T_0$ *of the European payer swaption with maturity* T_0 *and swap tenor* $T_{n_1} - T_0$ *is given by*

$$\begin{aligned}\Pi_t^{swaption} &= \mathbb{E}_t\left[\frac{\zeta_{T_0}}{\zeta_t}(\Pi_{T_0}^{swap})_+\right], \\ &= \frac{e^{-\alpha(T_0-t)}}{1+x_{11,t}}\mathbb{E}_t\left[(B(T_0,T_{n_1}) + A_1(T_0,T_{n_1})x_{11,T_0} + A_2(T_0,T_{n_1})x_{22,T_0})_+\right],\end{aligned} \tag{49}$$

with

$$\begin{aligned}B(T_0,T_{n_1}) :=& b_1(T_0-T_0) - e^{-\alpha(T_{n_1}-T_0)}b_1(T_{n_1}-T_0) + \sum_{j=1}^{n_1} e^{-\alpha(T_{j-1}-T_0)}b_2(T_{j-1}-T_0) \\ &- K\Delta\sum_{i=1}^{m_1} e^{-\alpha(t_i-T_0)}b_1(t_i-T_0)\,,\end{aligned} \tag{50}$$

$$A_1(T_0,T_{n_1}) := a_1(T_0-T_0) - e^{-\alpha(T_{n_1}-T_0)}a_1(T_{n_1}-T_0) - K\Delta\sum_{i=1}^{m_1} e^{-\alpha(t_i-T_0)}a_1(t_i-T_0)\,, \tag{51}$$

$$A_2(T_0,T_{n_1}) := \sum_{j=1}^{n_1} e^{-\alpha(T_{j-1}-T_0)}a_2(T_{j-1}-T_0)\,. \tag{52}$$

Remark 3.5. *As already noticed, the OIS term structure and the Euribor-OIS spread term structure can be used to estimate* α *and the diagonal terms of* x_0, ω *and* m. *The swaption prices then allows to calibrate* σ *as well as the off diagonal terms of* x_0 *and* ω *(the off-diagonal terms of* m *are still assumed to be zero).*

As aforementioned, the pricing of the swaption involves a linear functional of the state variable. It sharply contrasts with the classical approach based on the exponential affine framework where the computation of a sum of exponential functions of the state variable is involved for which no simple procedure is available. There are approximation algorithms such as those presented in Singleton and Umantsev (2002) and Schrager and Pelsser (2006) that freeze certain coefficients or the approximation of the density through the cumulant expansion of Collin-Dufresne and Goldstein (2002). In the linear-rational approach, the pricing of a swaption only requires the density of an affine function of the state variables which is known in closed form as we shall see below.

As usual, the caplet pricing formula is obtained by considering a swaption with one fixed payment. More precisely, a caplet with maturity T_0 on the Euribor rate $L(T_0,T_0+\delta)$, pays at time $T_1 = T_0+\delta$ the difference

[7] We remind the reader that T_0 is the current time in that particular case.

$L(T_0, T_0+\delta) - K$, if it's positive, where K is the strike of the caplet. Indeed, standard computations show

$$\Pi_t^{\text{caplet}} = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{T_0+\delta} r_u du}\delta(L(T_0,T_0+\delta)-K)_+\right], \tag{53}$$

$$= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{T_0} r_u du}P(T_0,T_0+\delta)\delta(L(T_0,T_0+\delta)-K)_+\right], \tag{54}$$

$$= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{T_0} r_u du}P(T_0,T_0+\delta)\delta\left(\text{Spread}(T_0,T_0+\delta)+\frac{1}{\delta}\left(\frac{1}{P(T_0,T_0+\delta)}-1\right)-K\right)_+\right], \tag{55}$$

$$= \frac{1}{\zeta_t}\mathbb{E}_t\left[\zeta_{T_0}(1-P(T_0,T_0+\delta)+A(T_0,T_0,T_0+\delta)-K\delta P(T_0,T_0+\delta))_+\right], \tag{56}$$

that is the expression of an option on a swap with one payment, which is a payer swaption.

One striking property of the linear-rational model is that the computational cost of a swaption is at par with the one of the caplet. Indeed, Eq. (49) clearly shows that only the terminal law of an affine function of the marginal of the process is needed, which can be carried out very easily using a Fourier transform. Define $\Lambda(T_0, T_{n_1})$ the 2×2 matrix with $(A_1(T_0,T_{n_1}), A_2(T_0,T_{n_1}))^\top$ on its diagonal (and 0 elsewhere) and denote the scalar variable $Y = B(T_0,T_{n_1}) + \text{tr}[\Lambda(T_0,T_{n_1})x_{T_0}]$ then the expectation in Eq. (49) rewrites $\mathbb{E}_t[(Y)_+]$. The characteristic function of Y is given by $\Phi_Y(u) = \mathbb{E}_t\left[e^{\mathrm{i}uY}\right] = e^{\mathrm{i}uB(T_0,T_{n_1})}\Phi(T_0-t, \mathrm{i}u\Lambda(T_0,T_{n_1}), 0, x_t)$ with Φ defined by Eq. (9). We have

$$\mathbb{E}_t[(Y)_+] = \frac{1}{\pi}\int_0^{+\infty}\Re\left(\frac{\Phi_Y(u+\mathrm{i}u_i)}{(\mathrm{i}(u+\mathrm{i}u_i))^2}\right)du, \tag{57}$$

with $u_i < 0$. That latter constraint on the integration axis corresponds to a similar constraint in Filipović et al. (2017, Theorem 4).

At that level, the choice of the stochastic process for the state variables is essential. In Crépey et al. (2015b), the authors use exponential martingales based on the Brownian motion and, therefore, need the density of their sum that is not known in closed form and have to rely on a multidimensional integration. In Nguyen and Seifried (2015), a two-factor model is proposed, there it is called the multi-curve rational lognormal model, and leads to a two-dimensional integration of the bivariate Gaussian distribution. As a result, these n-dimensional models imply integrating the n-dimensional Gaussian distribution with the numerical difficulties that come with it when n is larger than two. In the linear-rational based on the Wishart process for the Bru case (*i.e.*, $\omega = \beta\sigma^2$) one could compute the expectation by integrating the distribution Eq. (24) but it will remain numerically tedious. Instead, the formula above shows that in the linear-rational model based on the affine process as presented in Filipović et al. (2017) or the Wishart model as presented here, the pricing of a swaption leads to a one-dimensional integration, irrespective of the size of the model.

3.4 CMS option pricing

The vanilla swaption proved to be surprisingly simple to value in the linear-rational Wishart model and a natural question is whether other exotic products can be also easily priced in that framework. Looking at the interest rates derivatives actively traded on the market, the constant maturity swap is certainly the most obvious choice to consider. Following the academic literature (*e.g.*, Brigo and Mercurio 2006, Chapter 13.7) we now recall the characteristics of that product.

Consider a constant maturity swap (CMS) with tenor dates $T_0, \cdots, T_{n_1}$, with $T_j - T_{j-1} = \delta$. The two legs of the CMS have the same payment dates $T_1, \cdots, T_{n_1}$. At a payment date T_{j+1}, with $j = 0, \ldots, n_1 - 1$, one leg pays the Euribor rate resetting at time T_j plus a fixed spread K, while the other leg pays the swap rate $S_{T_j}^{T_{j,0},T_{j,n_s}}$, which is the swap rate with tenor structure and payment dates $T_{j,l} = T_j + l\delta_s$ with $l = 0, \ldots, n_s$ for the floating leg and $t_{j,k} = T_j + k\Delta_s$ for $k = 0, \ldots, m_s$ for the fixed leg and $T_{j,n_s} = t_{j,m_s}$. We suppose that δ and δ_s are equal so that there is no need to introduce another factor (or several factors) to handle the two tenor structures. In practice δ, δ_s and Δ_s are different but we stress the fact that all the computations below can be performed for that more general case without additional significant difficulty.

Proposition 3.6. *The time-t value of the constant maturity swap receiving the Euribor (plus a fixed rate K) leg and paying the swap leg is therefore given by*

$$\Pi_t^{\text{CMS}} = P(t,T_0) - P(t,T_{n_1}) + \sum_{j=1}^{n_1} A(t,T_{j-1},T_j) + \delta K \sum_{j=1}^{n_1} P(t,T_j)$$
$$- \sum_{j=0}^{n_1-1} \delta \mathbb{E}_t \left[\frac{\zeta_{T_{j+1}}}{\zeta_t} S_{T_j}^{T_j,T_{j,n_s}} \right], \tag{58}$$

with

$$\mathbb{E}_t \left[\frac{\zeta_{T_{j+1}}}{\zeta_t} S_{T_j}^{T_j,T_{j,n_s}} \right] = \frac{e^{-\alpha(T_{j+1}-t)}}{1+x_{11,t}} \mathbb{E}_t \left[\frac{c_0 + c_1 x_{11,T_j} + c_2 x_{22,T_j} + c_{12} x_{11,T_j} x_{22,T_j} + c_{11} x_{11,T_j}^2}{\mu_0 + \mu_1 x_{11,T_j}} \right], \tag{59}$$

and c_0, c_1, c_2, c_{12}, c_{11}, μ_0 and μ_1 defined in Eqs. (116-122) in the appendix.

To compute the value of the CMS, the expectation Eq. (59) needs to be evaluated but its simple structure, a rational function, combined with the affine property of the Wishart process enable an explicit computation thanks to the following well known remark.

Remark 3.7. *Suppose that we know the moment generating function of the vector (X,Y), that is $G(z_1,z_2) = \mathbb{E}\left[e^{z_1 X + z_2 Y}\right]$. To compute $\mathbb{E}\left[\frac{X}{Y}\right]$, the relation $1/y = \int_0^{+\infty} e^{-sy} ds$ leads to $\mathbb{E}\left[\frac{X}{Y}\right] = \int_0^{+\infty} \mathbb{E}\left[X e^{-sY}\right] ds$, and using the propriety of the moment generating function, we get $\mathbb{E}\left[\frac{X}{Y}\right] = \int_0^{+\infty} \partial_{z_1} \mathbb{E}\left[e^{z_1 X - sY}\right] ds|_{z_1=0}$. As $\mathbb{E}\left[e^{z_1 X - sY}\right]$ is known and can possibly be derived explicitly with respect to z_1, we obtain a quasi closed form expectation of the ratio of the two random variables.*

Proposition 3.8. *The integral representation of the ratio of two random variables combined with the moment generating function of the Wishart process Eq.* (10) *give explicit expressions for the expectations:*

$$\mathbb{E}_t \left[\frac{x_{11,T_j}}{\mu_0 + \mu_1 x_{11,T_j}} \right], \quad \mathbb{E}_t \left[\frac{x_{22,T_j}}{\mu_0 + \mu_1 x_{11,T_j}} \right], \quad \mathbb{E}_t \left[\frac{x_{11,T_j} x_{22,T_j}}{\mu_0 + \mu_1 x_{11,T_j}} \right], \quad \mathbb{E}_t \left[\frac{x_{11,T_j}^2}{\mu_0 + \mu_1 x_{11,T_j}} \right]. \tag{60}$$

The CMS naturally serves as an underlying for interest rates derivatives but instead of the standard call/put on an CMS rate what is frequently found is the CMS spread single-option which involves two CMS rates as its name suggests. Let us denote by $\Pi_t^{\text{CmsSpSO}}(T_1, n_{s_1}, n_{s_2}, K)$ the t-value of a CMS spread call option with single expiration date T_1 and strike K. It is an option whose value at time T_1 is based on the difference between the swap rate $S_{T_1}^{T_{1,0},T_{1,n_{s_1}}}$, starting at time T_1 and ending at time $T_{1,n_{s_1}} > T_1$, and the spot swap rate $S_{T_1}^{T_1,T_{1,n_{s_2}}}$ starting at time T_1 and ending at time $T_{1,n_{s_2}} > T_1$.

The underlying swap $S_{T_1}^{T_{1,0},T_{1,n_{s_1}}}$ floating leg's tenor and payment dates are $T_{1,l} = T_1 + l\delta_s$ with $l = 0, \ldots, n_{s_1}$ whilst the fixed leg's tenor and payment dates are $t_{1,k} = t_1 + k\Delta_s$ with $k = 0, \ldots, m_{s_1}$ and we further have that $T_{1,0} = T_1$, $t_{1,0} = t_1 = T_1$ and $T_{1,n_{s_1}} = t_{1,m_{s_1}}$ which imply that both legs start and end at the same time. The swap $S_{T_1}^{T_{1,0},T_{1,n_{s_2}}}$ is defined similarly.

Those two CMS rates are the underlyings of the CMS spread single-option whose pricing formula is in the next proposition.

Proposition 3.9. *The option time-t value, denoted Π_t^{CmsSpSO} for simplicity, is given by:*

$$\Pi_t^{\text{CmsSpSO}} = \mathbb{E}_t \left[\frac{\zeta_{T_1}}{\zeta_t} \left(S_{T_1}^{T_{1,0},T_{1,n_{s_1}}} - S_{T_1}^{T_{1,0},T_{1,n_{s_2}}} - K \right)_+ \right], \tag{61}$$
$$= \frac{e^{-\alpha(T_1-t)}}{1+x_{11,t}} \mathbb{E}_t \left[\left(g^1(x_{11,T_1}, x_{22,T_1}) - g^2(x_{11,T_1}, x_{22,T_1}) - K(1+x_{11,T_1}) \right)_+ \right], \tag{62}$$

with

$$g^i(x_{11,T_1}, x_{22,T_1}) = \frac{c_0^i + c_1^i x_{11,T_1} + c_2^i x_{22,T_1} + c_{12}^i x_{11,T_1} x_{22,T_1} + c_{11}^i x_{11,T_1}^2}{\mu_0^i + \mu_1^i x_{11,T_1}}, \tag{63}$$

and for $i \in \{1,2\}$ with c_0^i, c_1^i, c_2^i, c_{12}^i, c_{11}^i μ_0^i and μ_1^i for $i \in \{1,2\}$ given in the appendix.

Similar to caplets (or floorlets) that are not traded individually but as a component of a cap (floor), the CMS spread single-option is traded through a CMS spread multi-option which is just a portfolio of CMS spread single-options and is defined as follows. Let $\Pi_t^{\text{CmsSpMO}}(T_1, T_{n_1}, n_{s_1}, n_{s_2}, K)$ be the t-value of the multi CMS spread call option with exercise dates $T_1, \cdots, T_{n_1}$, with $T_j - T_{j-1} = \delta$ and strike K. It is a sum of CMS spread single call option with maturity dates $T_1, \ldots, T_{n_1}$. All the options' two underlying swaps have the same tenor structures. Using the previous definition, the option time t-value denoted Π_t^{CmsSpMO}, for simplicity and when no confusion is possible, is given by:

$$\Pi_t^{\text{CmsSpMO}} = \sum_{j=1}^{n_1} \Pi_t^{\text{CmsSpSO}}(T_j, n_{s_1}, n_{s_2}, K) \tag{64}$$

Remark 3.10. *Unfortunately, the pricing formula of the CMS spread single-option Eq.* (62) *is not as simple as the swaption pricing formula. Notice, however, that it only involves $(x_{11,T}, x_{22,T})$, the marginal distribution of the process at time T and not the process path from t to T. In the Bru case, the marginal distribution of the process can be expressed, when the parameter β is an integer, as the square of a matrix Gaussian distribution as is therefore computable by Monte Carlo very efficiently. When β is not an integer, Ahdida and Alfonsi (2013) derived an exact and fast simulation algorithm.*

3.5 Swaption price approximation

The pricing of a swaption in the standard exponential affine framework is known to be notoriously tedious as it involves the density of a sum of exponentials of random variables. Only in some very specific cases, typically when the state variable is one dimensional, the swaption price can be computed easily. In Collin-Dufresne and Goldstein (2002), the authors propose an approximation of the swaption price by approximating the density of a coupon bearing bond. Their result crucially relies on the affine property of the process driving the interest rates. In the approach adopted here, Proposition 3.4 shows that the pricing of a swaption is simple as it only requires a one-dimensional integration, the affine property of the Wishart process enables us to derive an approximation in the spirit of Collin-Dufresne and Goldstein (2002) and derive an even faster option pricing formula. Notice that the affine property is used in two different ways. In Collin-Dufresne and Goldstein (2002), the authors use the fact that the expected value of the exponential of an affine variable is exponential affine whereas here we use the fact that the expected value of a polynomial function of a given order of an affine process can be expressed as a polynomial function of the process, in order words the set of polynomials is stable for the infinitesimal generator of the Wishart process.

To establish the approximation, we need the two following lemmas.

Lemma 3.11. *Let $y(t)$ a function solution of the ordinary differential equation*

$$\frac{dy(t)}{dt} = \kappa y(t) + \sum_{i=1}^{l} \nu_i (\bar{a}_i + \bar{b}_i e^{\kappa_i t}), \tag{65}$$

with $\kappa \neq \kappa_i$ $\forall i \in \{1, \ldots, l\}$ and $\nu_i, \bar{a}_i, \bar{b}_i$ $\forall i \in \{1, \ldots, l\}$ some constants then it can be integrated to

$$y(t) = \sum_{i=1}^{l+1} \bar{c}_i + \bar{d}_i e^{\kappa_i t}, \tag{66}$$

with $\kappa_{n+1} = \kappa$ *and*

$$\bar{c}_i = 0 \quad i = 1, \dots, l\,, \tag{67}$$

$$\bar{d}_i = \frac{\nu_i \bar{b}_i}{\kappa_i - \kappa} \quad i = 1, \dots, l\,, \tag{68}$$

$$\bar{c}_{n+1} = -\sum_{i=1}^{l} \frac{\nu_i \bar{a}_i}{\kappa}\,, \tag{69}$$

$$\bar{d}_{n+1} = y(0) + \sum_{i=1}^{l} \frac{\nu_i \bar{a}_i}{\kappa} - \sum_{i=1}^{l} \frac{\nu_i \bar{b}_i}{\kappa_i - \kappa}\,. \tag{70}$$

For notational convenience, it is useful to introduce $\bar{\sigma} = \sigma^2$ and notice that in Eqs. (3–8), $(\sigma_{11}^2 + \sigma_{12}^2) = (\sigma^2)_{11} = \bar{\sigma}_{11}$, $(\sigma_{12}^2 + \sigma_{22}^2) = (\sigma^2)_{22} = \bar{\sigma}_{22}$ and $(\sigma_{11}\sigma_{12} + \sigma_{12}\sigma_{22}) = (\sigma^2)_{12} = \bar{\sigma}_{12}$. It can be seen that the infinitesimal generator $\mathcal{G}$ given by Eq. (2) only involves σ^2. The following lemma shows that the affine property of the Wishart process implies a simple expression for the expected value of a polynomial function of the process.

Lemma 3.12. *Let us denote* $g(t,i,k,j) = \mathbb{E}[x_{11,t}^i x_{12,t}^k x_{22,t}^j]$ *where* $(x_{11,t}, x_{12,t}, x_{22,t})$ *are the components of a* 2×2 *Wishart process then using the brackets Eqs. (3-8) and Itô's Lemma we get*

$$\frac{dg(t,i,k,j)}{dt} = (i2m_{11} + k(m_{11} + m_{22}) + 2jm_{22})\, g(t,i,k,j) \tag{71}$$

$$+ (i\omega_{11} + 2ik\bar{\sigma}_{11} + 2i(i-1)\bar{\sigma}_{11})\, g(t,i-1,k,j) \tag{72}$$

$$+ (k\omega_{12} + k(k-1)\bar{\sigma}_{12} + 2ik\bar{\sigma}_{12} + 2jk\bar{\sigma}_{12})\, g(t,i,k-1,j) \tag{73}$$

$$+ (j\omega_{22} + 2j(j-1)\bar{\sigma}_{22} + 2jk\bar{\sigma}_{22})\, g(t,i,k,j-1) \tag{74}$$

$$+ \frac{k(k-1)}{2}\bar{\sigma}_{22} g(t,i+1,k-2,j) \tag{75}$$

$$+ \frac{k(k-1)}{2}\bar{\sigma}_{11} g(t,i,k-2,j+1) \tag{76}$$

$$+ 4ij\bar{\sigma}_{12} g(t,i-1,k+1,j-1)\,. \tag{77}$$

Notice that Eqs. (72–77) involve polynomials with degree lower or equal to $i+k+j-1$ whilst Eq. (71) involves a polynomial of degree $i+k+j$, it is a consequence of the affine property of the Wishart process. As $g(t,1,0,0)$, $g(t,0,1,0)$ and $g(t,0,0,1)$ can be written in the form $\bar{a}_0 + \bar{b}_0 e^{\kappa t}$ then we deduce by recurrence that $g(t,i,k,j)$ solves an ODE of the form Eq. (65) and therefore Lemma 3.11 applies.

Starting from Eq. (49), define $Y_{T_0} = B(T_0, T_{n_1}) + A_1(T_0, T_{n_1}) x_{11,T_0} + A_2(T_0, T_{n_1}) x_{22,T_0}$, to apply Collin-Dufresne and Goldstein (2002)'s swaption price approximation one needs to compute the q^{th} moment of Y_{T_0} that is simply given by

$$\mathbb{E}[Y_{T_0}^q] = \sum_{l_0+l_1+l_2=q} \binom{q}{l_0, l_1, l_2} B(T_0, T_{n_1})^{l_0} A_1(T_0, T_{n_1})^{l_1} A_2(T_0, T_{n_1})^{l_2} \mathbb{E}[x_{11,T_0}^{l_1} x_{22,T_0}^{l_2}]\,. \tag{78}$$

As $\mathbb{E}[x_{11,T_0}^{l_1} x_{22,T_0}^{l_2}] = g(T_0, l_1, 0, l_2)$ and is known thanks to Lemma 3.12, the moments are also known.

Remark 3.13. *Notice that although only terms of the form* $\mathbb{E}[x_{11,T_0}^{l_1} x_{22,T_0}^{l_2}] = g(T_0, l_1, 0, l_2)$ *are needed to determine the moment* $\mathbb{E}[Y_{T_0}^q]$*, Lemma 3.12 shows that these terms depend on moments involving* x_{12,T_0} *(through Eq. 77).*

Starting from Collin-Dufresne and Goldstein (2002, Eq. 17), which represents an approximation of the density of Y_{T_0}, given by

$$\sum_{j \geq 0} \gamma_j \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi c_2}} (y - c_1)^j e^{-\frac{(y-c_1)^2}{2c_2}} dy\,, \tag{79}$$

13

the expectation in Eq. (49) can be approximated by

$$\mathbb{E}_t\left[(Y_{T_0})_+\right] \sim \sum_{j\geq 0}\gamma_j\int_0^{+\infty}\frac{1}{\sqrt{2\pi c_2}}y(y-c_1)^j e^{-\frac{(y-c_1)^2}{2c_2}}\,dy\,, \tag{80}$$

$$= \sum_{j\geq 0}\gamma_j\sqrt{c_2}c_2^{j/2}\int_{\frac{-c_1}{\sqrt{c_2}}}^{+\infty} z^{j+1}\frac{1}{\sqrt{2\pi}}e^{-\frac{z^2}{2}}\,dz\,, \tag{81}$$

$$+ \sum_{j\geq 0}\gamma_j c_1 c_2^{j/2}\int_{\frac{-c_1}{\sqrt{c_2}}}^{+\infty} z^{j}\frac{1}{\sqrt{2\pi}}e^{-\frac{z^2}{2}}\,dz\,, \tag{82}$$

$$= \sum_{j\geq 0}\gamma_j\lambda_{j+1} + c_1\sum_{j\geq 0}\gamma_j\lambda_j\,, \tag{83}$$

where $\{\gamma_j; j\in\mathbb{N}\}$ are related to the cumulants through Eqs. (B.18–B.25) in Collin-Dufresne and Goldstein (2002), $\{c_j; j\in\mathbb{N}\}$ are the cumulants of the variable Y_{T_0} that are related to the moments of that variable through Eqs. (A.1–A.7) in Collin-Dufresne and Goldstein (2002) whilst $\{\lambda_j; j\in\mathbb{N}\}$ are related to the normal density/cumulative distribution (see Collin-Dufresne and Goldstein, 2002, Eqs. B.10–B.17).

4 Model implementation

4.1 The data

This study considers the Euro market and the data comprise the OIS term structure, the Euribor term structure and the ATM swaption prices for the period 04/10/2011 to 12/03/2012. For the term structures, either OIS or Euribor, we restrict to a maturity smaller than 15 years.[8] For the OIS, we use Eonia rates that have floating and fixed legs that pay annually (when the swap's maturity is larger than 1 year). For the Euribor rates, the floating leg pays semi-annually while the fixed leg pays annually. Table I reports the basic descriptive statistics, mean and standard deviation, for each term structure of interest rate. Both term structures are increasing and as expected the Euribor curve is above the OIS curve reflecting its credit risk component. For both curves, long term rates display lower standard deviations.

[Insert Table I here]

The swaption data is usually quoted in terms of normal or log-normal volatility. In our data set, the normal volatility quotes are converted into price using the Bachelier formula for at-the-money call option, it is the market practice and the approach used in Filipović et al. (2017, online appendix). Section A.1 of the appendix presents the basic formulas. By definition the at-the-money swaption is the option with a strike equals to the forward swap rate that can be synthesized using two spot swap rates which are quoted (see A.2 of the appendix for the details). Furthermore, for the swaption strike we follow Filipović et al. (2017) and set it to the model-implied forward swap rate. We consider the swaption maturities 1Y, 2Y, 3Y, 4Y and 5Y whilst for the swap tenor we restrict to 1Y, 2Y, 3Y, 4Y and 5Y. Table II reports mean and standard deviation of the normal implied volatility for each option. For a given swap tenor, the implied volatility is increasing with the swaption maturity while for a given swaption maturity, the implied volatility is increasing with the swap tenor for swaption maturities smaller or equal to two years and decreasing for swaption maturities larger or equal to three years. Regarding the standard deviations, for a given swaption maturity, the standard deviation decreases as the swap tenor increases while for a given swap tenor the standard deviation decreases as the swaption maturity increases.

[Insert Table II here]

[8] The current framework is designed to generate positive interest rates, indeed Rogers (1997) generalizes the work of Constantinides (1992) that focuses on nominal interest rates, whilst they have been close to zero and even negative during a recent period. This forces us to consider the period beginning of 2012.

4.2 Calibration results and analysis

For the implementation, we follow the common market practice of performing a daily calibration and rolling it but we take into account the specifics of the model by staging the estimation procedure. More precisely, we proceed as follows. First, relying on Eq. (31), α is estimated as the long-term zero-coupon bond yield. Then the parameters $x_{11,t}$, ω_{11} and m_{11} are estimated by solving the optimization problem

$$\min \frac{1}{N}\sum_{i=1}^{N}\left(P_{\text{model}}(t,T_i)-P_{\text{market}}(t,T_i)\right)^2 , \tag{84}$$

where $P_{\text{market}}(t,T_i)$ is the market price at time t of a zero-coupon with maturity T_i, obtained by bootstrapping the OIS term structure, whilst $P_{\text{model}}(t,T_i)$ stands for the corresponding model price given by Eq. (28) and N is the number of zero-coupon prices available for that day. Using the Euribor swap rates along with the OIS zero-coupon bond market prices, we extract the market spreads given by Eq. (38) and then calibrate for each day the parameters $x_{22,t}$, ω_{22} and m_{22} by solving the optimization problem

$$\min \frac{1}{N}\sum_{j=1}^{N}\left(A_{\text{model}}(t,T_{j-1},T_j)-A_{\text{market}}(t,T_{j-1},T_j)\right)^2 . \tag{85}$$

Lastly, using the swaptions we calibrate $\sigma_{11},\sigma_{12},\sigma_{22}$ and $x_{12,t}$, ω_{12} by solving

$$\min \frac{1}{N}\sum_{i=1}^{N}\frac{1}{M_i}\sum_{j=1}^{M_i}\left(\sigma_{\text{model}}(t,T_i,T_{i,j})-\sigma_{\text{market}}(t,T_i,T_{i,j})\right)^2 , \tag{86}$$

with $\sigma_{\text{model}}(t,T_i,T_{i,j})$ the swaption model (normal) implied volatility for day t, swaption maturity T_i and swap tenor $T_{i,j}-T_i$ given by Eq. (49), $\sigma_{\text{market}}(t,T_i,T_{i,j})$ stands for the corresponding market (normal) implied volatility while N is the number of swaption maturities and M_i is the number of tenors for the i^{th} maturity available for that day.

The mean value as well as the standard deviation of the estimated parameters are reported in Table III while Table IV reports the eigenvalues of x, ω and σ in order to provide a sanity check of the estimates. Table V reports the correlations associated with x, ω and σ as well as the long term mean value $\bar{x}_\infty$ given by Eq. (19). Table VI contains the average as well as the standard deviation of the root mean square errors of the calibrations Eqs. (84-86).

[Insert Table III here]

[Insert Table IV here]

[Insert Table V here]

[Insert Table VI here]

As per Eq. (31), the value of α corresponds to the long term yield and the mean value is equal to 2.4% with a small standard deviation. Combined with Eq. (30) we find that the model short term rate is not positive. Let us point out that Filipović et al. (2017, Table IA.III, online appendix) proceed the other way around; α is such that the short term rate is positive and therefore the natural question is whether the long term yield is accurately fitted. The mean values of x_{11} and x_{22} are positive with small standard deviations. The value of x_{12} is on average negative and when combined x_{11} and x_{22} leads to matrix that has positive eigenvalues according to Table IV whilst the correlation associated with the matrix x is on average equal to -0.423 as shown in Table V. The mean values of ω_{11} and ω_{22} are positive with a small standard deviation for the ω_{11} but a rather large (compared to the mean) one for ω_{22}. The value of ω_{12} is on average positive and leads to a matrix ω which has positive eigenvalues according to Table IV, the correlation associated with the matrix ω is on average equal to 0.223 as shown in Table V. We find that for each day m_{11} and m_{22} are negative with the mean estimated values reported in Table III along with the standard deviations that are small. All the elements of σ are positive with small standard deviations, the eigenvalues of σ reported in Table IV are positive and the correlation associated with the matrix σ is 0.497, that is rather strong. As expected all the matrices belong to $\mathbb{S}_2^{++}$ and the correlations associated with these matrices give an indication of the dependency between the factors and therefore the curves in the model. Notice that $\bar{x}_\infty$ given by Eq. (19) is a positive definite matrix whose correlation associated with the off-diagonal term is 0.208 according to Table V, there is a change in the correlation sign between the long term mean value of the Wishart process and its initial value.

Tables III, IV and V also report min and max values for the estimates. They confirm that the model properties deduced from the mean and the standard deviation values are not only valid on average over the sample but also punctually in time for each day.

The calibration errors are reported in Table VI, they are overall very reasonable if we take into account the parsimony of the model. Regarding the standard deviation of the errors, compared to the mean it is small for the OIS curve but rather large for the spread and translates the large standard deviation observed for ω_{22}. The rather large (compared to the mean) standard deviation of the spread calibration error is due to the calibration procedure that is sequential. Any variation in the OIS calibration error will impact the spread calibration error that builds upon it.[9] For the swaptions, the error is 36.24 with a small standard deviation showing the ability of the model to capture the daily variation of the data.

Although the calibration errors reported in Table VI are quite reasonable, they could be improved in two directions. First, the fit of the initial yield curve at the first stage (see Eq. 84) can be improved following Filipović et al. (2017) who suggests to add a time dependent function as follows

$$\tilde{P}(t,T) = e^{-\int_t^T \nu_1(u)du} P(t,T)\,, \tag{87}$$

and it implies to redefine the pricing kernel as

$$\tilde{\zeta}_T = e^{-\int_t^T \nu_1(u)du} \zeta_T\,. \tag{88}$$

Second, a similar approach applies to improve the fit of the Euribor-OIS spread term structure. Indeed, suppose we observe the market spreads $\{\tilde{A}(t,T_i,T_i+\delta); i=1,\ldots,n_1\}$ (extracted from the Euribor-OIS spread curve) and that the model implied spread A is given by the formula Eq. (38) above with α already estimated (from the OIS curve) then the fit can be improved by adding a function $(\nu_2(u); u \in [t,\,T])$ such that

$$\tilde{\zeta}_T \tilde{P}(T,T+\delta)\delta \mathrm{Spread}(T,T+\delta) = e^{-\int_0^T \nu_2(u)du} e^{-\alpha T} x_{22,T}\,, \tag{89}$$

which leads to

$$\tilde{A}(t,T,T+\delta) = e^{-\int_t^T \nu_2(u)du} A(t,T,T+\delta)\,. \tag{90}$$

Improving the fit by making some of the parameters time dependent is common in the interest rate literature and is central to the Heath-Jarrow-Morton approach, see Brigo and Mercurio (2006) for more details. Notice that the improvements mentioned above barely change the swaption calibration error. This latter can be reduced by making the volatility σ time dependent, such a time dependent parameter strategy is used in Nguyen and Seifried (2015), but this has far reaching numerical consequences that are beyond the scope of that work.

Once the model is calibrated, it is relevant to analyze the distribution of the variable $Y_{T_0} = B(T_0,T_{n_1}) + A_1(T_0,T_{n_1})x_{11,T_0} + A_2(T_0,T_{n_1})x_{22,T_0}$ that is involved in the swaption pricing in Eq. (49) and whose moments are known and given by Eq. (78). Using the characteristic function of that variable, we report in Figures 1-2 its density for two pairs of maturity/tenor: (1 year, 1 year) and (5 years, 5 years). These pairs are the extremes of the swaption data reported in Table II. All the other pairs look similar to those reported. The figures shows two distributions that are uni-modal and slightly asymmetric. It suggests that the first three moments of the distribution could be used in the swaption price approximation developed in Collin-Dufresne and Goldstein (2002) and produce reasonably accurate results. The approximate density given by Eq. (79) is reported for the

[9] Notice that it seems to affect more ω_{22} than x_{22}.

maturity/tenor pair (1 year, 1 year) in Figure 1 and in Figure 2 for maturity/tenor pair (5 years, 5 years). For the first pair, the approximate density is very close to the true one, suggesting that a fairly accurate swaption price can be obtained using the first three moments, while for the second pair the difference will translate into a swaption price approximation that can deviate substantially from the exact price. To assess that latter point, we follow the details of section 3.5 and price the options using the calibrated parameters and the approximation formula Eq. (83). We restrict to the first three cumulants, as we found that taking higher cumulants deteriorates the results, then the average root mean square swaption pricing error between the normal volatility computed using the exact swaption pricing formula and the normal volatility using the approximate swaption pricing formula is 9.410.[10] That error reduces to 7.301 when restricted to swaptions with maturity equal to 1 year and increases to 9.724 when restricted to swaptions with a maturity larger (or equal) to 4 years.

[Insert Figure 1 here]

[Insert Figure 2 here]

To the matrices x, ω and σ correspond certain correlation matrices with off-diagonal terms reported in Table V confirming that the model does not have a diagonal structure. In particular, $\sigma_{12} \neq 0$ implies that the last term in Eq. (42) does not vanish (*i.e.*, $\langle x_{11,.}, x_{22,.}\rangle_t$ depends linearly on σ_{12} according to Eq. (8)) while the last term in Eq. (43) also remains. Furthermore, x_{12} is on average negative, as the calibrated value or the correlation associated with x shows, combined with σ_{12} that is positive, we deduce that the last term in Eq. (42) is positive and contributes to increase the covariance between the two curves. Whether that covariance is mainly driven by the first term or the second term of Eq. (42) determines the importance of the off-diagonal terms x_{12} and σ_{12} of the Wishart process and therefore the degree of dependency that exists between the two diagonal terms x_{11} and x_{22} of the Wishart process or factors and, by extension, the OIS and Euribor-OIS curves.

The importance of off-diagonal terms illustrates the result of Benabid et al. (2009) according to which the law of the diagonal terms of the Wishart process for a given time t (*i.e.*, $(x_{t,11}, x_{t,22})$), which are the only terms involved in the argument of the characteristic function $\Phi_Y(.)$ in Eq. (57) as $\Lambda(T_0, T_{n_1})$ is diagonal, is not the product of two noncentral Chi-squared distributions. As a consequence, the off-diagonal terms x_{12}, ω_{12} and σ_{12} of the Wishart process are essential for the model to capture the dependency between the OIS and the Euribor-OIS curves and Table V, which reports the correlations associated with these matrices, clearly show that they are significant.

To further illustrate the importance of the correlation between the two curves, and the relevance of the Wishart process to capture that dependency, we compare the market correlation with the model correlation. Following Eq. (48), let us denote $\bar{A}(t, T_{n_1}) = \sum_{j=1}^{n_1} A(t, T_{j-1}, T_j)$ the sum of spreads up to T_{n_1} involved in a swap contract with maturity T_{n_1} and $P(t, T_{n_1})$ the OIS zero-coupon bond with maturity T_{n_1}. We are interested in

$$\mathrm{Corr}(dP(t, T_{n_1}), d\bar{A}(t, T_{n_1}))\,, \tag{91}$$

the correlation between $P(t, T_{n_1})$ increments and $\bar{A}(t, T_{n_1})$ increments. Thanks to Eq. (42), it is known that it can take any sign but also that it is driven by two terms, the first one depending only on the diagonal terms of the model while the second one depends on off-diagonal terms. To assess the quality of the linear-rational Wishart model, we compare the market correlation with the model correlation (*i.e.*, the correlation given by the calibrated model) and report the results in Table VII. The market correlation, reported in the line "Market", is positive and declines with the zero-coupon bond/spread maturity as the table shows. As the zero-coupon bond and the spread are calibrated using only the drift of the process (see Eqs. 84–85), the correlation Eq. (91) can be computed using the zero-coupon bond and the spread given by the calibrated model. It is named "Model drift" in Table VII and it shows the ability of the model to capture the correlation between the two curves. So far the volatility σ was not used and therefore the right hand side of Eq. (42) was not exploited. It allows us to check the consistency of the model. Indeed, calibrating σ as well as the off-diagonal terms of x and ω, by solving

[10]The fact that higher order cumulants do not improve the results as in Collin-Dufresne and Goldstein (2002) is likely to be related to distribution of the variable that is close to a noncentral chi-squared distribution while in Collin-Dufresne and Goldstein (2002) the variable is a sum of exponentials of normal variables/noncentral chi-squared variables as they use an exponential affine model.

Eq. (86), we can compute the model correlation given by the right hand side of Eq. (42), the results are reported in "Model vol" in Table VII, and check whether the model correlation is close to the market correlation.[11] The values reported in line "Model vol" are consistent with both the values reported in the line "Model drift" and the values reported in the line "Market". The values confirm that the model is coherent and can handle the non trivial dependency that exists between the two curves.

[Insert Table VII here]

To clarify further the analysis of the model and, in particular, the importance of the non-diagonal terms, we redo the calibration Eq. (86) only for σ_{11} and σ_{22} while the off-diagonal parameters (*i.e.*, x_{12}, ω_{12} and σ_{12}) are set to zero. Notice that as the first two calibrations Eqs. (84-85) do not depend on off-diagonal terms, their results remain valid for the diagonal model so we do not need to redo these calibrations. A diagonal structure for σ makes the model similar to a vector affine model and therefore belongs to the category considered in Filipović et al. (2017). The calibrated parameters are reported in Table VIII while the calibration error is in Table IX. The estimated parameters are different but in line with those of Table III while the calibration error is 36.51 bps, marginally larger than the one in Table VI. Even thought one might consider that difference as negligible and opt for the diagonal model, a look at the correlation Eq. (91) computed using the right hand side of Eq. (42) and reported in Table X clearly shows the problem. Indeed, the values are around 0.038 for all maturities, far from the market values. They show that between the two components that comprise Eq. (42), the main one is by far the second term. Using a diagonal model, one might consider not making an important loss in terms of fit of swaptions data but it will definitively affect the pricing of products that depend on the correlation between the OIS curve and Euribor-OIS curve.[12] In conclusion, the non trivial dependency between the two curves can be handled by the linear-rational Wishart model that provides, compared to the standard vector affine process, an additional factor that is crucial.

[Insert Table VIII here]

[Insert Table IX here]

[Insert Table X here]

5 Conclusion

We propose a linear-rational multi-curve term structure model based on the Wishart process. Following Filipović et al. (2017)'s modeling strategy that is based on the potential approach presented in Rogers (1997), we use the Wishart process to build a multi-curve model that allows for a stochastic correlation between the curves and develop the pricing formulas for interest rate products commonly traded on the market such as interest swaptions. One striking property of the model is that the swaptions have the same computational cost as caps/floors. Pricing formulas for more complex interest rate derivatives such as the constant maturity swap and the option on a constant maturity swap spread are also derived. Thanks to the affine property of Wishart process, a swaption price approximation in the spirit of Collin-Dufresne and Goldstein (2002) is also developed. To analyze the model empirical properties, we perform a daily calibration of the model using a three-month sample of OIS term structure, Euribor-OIS term structure and ATM swaption prices. The calibration errors are stable and show the model's ability to handle the data fluctuations. The estimated parameters lead to a model that possesses the right statistical properties. The estimated parameters have small standard deviations, the model is therefore robust. What is more, the estimated parameters illustrate the ability of the model to capture the non null relationship that exists between the OIS curve and the Euribor-OIS spread curve that critically relies on the Wishart process properties. Overall, the results clearly underline the quality of the linear-rational model based on the Wishart process.

[11] When analyzing the right hand side of Eq. (42) ω_{12} is not needed.

[12] A byproduct of that result is that one needs to go beyond calibration error. That point of view is not new and is in the spirit of the discussion developed Brigo and Mercurio (2006, Section 6.5).

References

A. Ahdida and A. Alfonsi. Exact and high-order discretization schemes for Wishart processes and their affine extensions. *The Annals of Applied Probability*, 23(3):1025–1073, 2013. doi: 10.1214/12-AAP863.

M. Alfeus, M. Grasselli, and E. Schlögl. A consistent stochastic model of the term structure of interest rates for multiple tenors. *Journal of Economic Dynamics and Control*, 114:103861, 2020. doi: 10.1016/j.jedc.2020.103861.

A. Benabid, H. Bensusan, and N. El Karoui. Wishart stochastic volatility : Asymptotic smile and numerical framework. 2009. URL `http://hal.archives-ouvertes.fr/hal-00458014/fr`.

D. Brigo and F. Mercurio. *Interest Rate Models - Theory and Practice*. Springer-Verlag, 2 edition, 2006. doi: 10.1007/978-3-540-34604-3.

M. F. Bru. Wishart processes. *Journal of Theoretical Probability*, 4(4):725–751, 1991. doi: 10.1007/BF01259552.

P. Collin-Dufresne and R. Goldstein. Pricing swaptions within an affine framework. *Journal of Derivatives*, 3(1-2):167–179, 2002.

G. M. Constantinides. A theory of the nominal term structure of interest rates. *Review of Financial Studies*, 5(4):531–552, 1992. doi: 10.1093/rfs/5.4.531.

S. Crépey, Z. Grbac, N. Ngor, and D. Skovmand. A lévy HJM multiple-curve model with application to cva computation. *Quantitative Finance*, 15(3):401–419, 2015a. doi: 10.1080/14697688.2014.942232.

S. Crépey, A. Macrina, T. M. Nguyen, and D. Skovmand. Rational multi-curve models with counterparty-risk valuation adjustments. *Quantitative Finance*, 16(6):847–866, 2015b. doi: 10.1080/14697688.2015.1095348.

C. Cuchiero, D. Filipović, E. Mayerhofer, and J. Teichmann. Affine processes on positive semidefinite matrices. *The Annals of Applied Probability*, 21(2):397–463, 2011. doi: 10.1214/10-AAP710.

C. Cuchiero, C. Fontana, and A. Gnoatto. A general HJM framework for multiple yield curve modeling. *Finance and Stochastics*, 20(2):267–320, 2016. doi: 10.1007/s00780-016-0291-5.

C. Cuchiero, C. Fontana, and A. Gnoatto. Affine multiple yield curve models. *Mathematical Finance*, 29(2):568–611, 2019. doi: 10.1111/mafi.12183.

J. Da Fonseca, A. Gnoatto, and M. Grasselli. A flexible matrix Libor model with smiles. *Journal of Economic Dynamics and Control*, 37(4):774–793, 2013. doi: 10.1016/j.jedc.2012.11.006.

D. Duffie and R. Kan. A yield-factor model of interest rates. *Mathematical Finance*, 6(4):379–406, 1996. doi: 10.1111/j.1467-9965.1996.tb00123.x.

D. Filipović and A. B. Trolle. The term structure of interbank risk. *Journal of Financial Economics*, 109(3):707–733, 2013. doi: 10.1016/j.jfineco.2013.03.014.

D. Filipović, M. Larsson, and A. B. Trolle. Linear-rational term structure models. *The Journal of Finance*, 72(2):655–704, 2017. doi: 10.1111/jofi.12488.

C. Gouriéroux and R. Sufana. Derivative pricing with multivariate stochastic volatility. *Journal of Business and Economic Statistics*, 28(3):438–451, 2010. doi: 10.1198/jbes.2009.08105.

M. Grasselli and G. Miglietta. A flexible spot multiple-curve model. *Quantitative Finance*, 16(10):1465–1477, 2016. doi: 10.1080/14697688.2015.1108521.

M. Grasselli and C. Tebaldi. Solvable affine term structure models. *Mathematical Finance*, 18(1):135–153, 2008. doi: 10.1111/j.1467-9965.2007.00325.x.

Z. Grbac, A. Papapantoleon, J. Schoenmakers, and D. Skovmand. Affine LIBOR models with multiple curves: Theory, examples and calibration. *SIAM Journal on Financial Mathematics*, 6(1):984–1025, 2015. doi: 10.1137/15M1011731.

A. Gupta and D. Nagar. *Matrix Variate Distributions*. Chapman & Hall/CRC, 2000.

C. Kang, W. Kang, and J. M. Lee. Exact simulation of the Wishart multidimensional stochastic volatility model. *Operations Research*, 65(5):1190–1206, 2017. doi: 10.1287/opre.2017.1636.

P. Koev and A. Edelman. The efficient evaluation of the hypergeometric function of a matrix argument. *Mathematics of computation*, 75(254), 2006.

A. Macrina. Heat kernel models for asset pricing. *International Journal of Theoretical and Applied Finance*, 17(07):1450048, 2014. doi: 10.1142/S0219024914500484.

A. Macrina and O. Mahomed. Consistent valuation across curves using pricing kernels. *Risk*, 6(18), 2018. doi: https://doi.org/10.3390/risks6010018.

E. Mayerhofer. On wishart and noncentral Wishart distributions on symmetric cones. *Transactions of the AMS*, 371(10):7093–7109, 2019. doi: 10.1090/tran/7754.

E. Mayerhofer, O. Pfaffel, and R. Stelzer. On strong solutions of positive definite jump-diffusions. *Stochastic processes and their applications*, 121(9):2072–2086, 2011. doi: 10.1016/j.spa.2011.05.006.

N. Moreni and A. Pallavicini. Parsimonious HJM modelling for multiple yield curve dynamics. *Quantitative Finance*, 14(2):199–210, 2014. doi: 10.1080/14697688.2013.829242.

L. Morino and W. Runggaldier. On multicurve models for the term structure. *Nonlinear Economic Dynamics and Financial Modelling: Essays in Honour of Carl Chiarella*, pages 275–290, 01 2014. doi: 10.1007/978-3-319-07470-2_16.

N. Nakamura and F. Yu. Interest rate, currency and equity derivatives valuation using the potential approach. *International Review of Finance*, 1(4):269–294, 2000. doi: 10.1111/1468-2443.00019.

T. A. Nguyen and F. T. Seifried. The multi-curve potential model. *International Journal of Theoretical and Applied Finance*, 18 (7):1550049, 2015. doi: https://doi.org/10.1142/S0219024915500491.

L. C. G. Rogers. The potential approach to the term structure of interest rates and foreign exchange rates. *Mathematical Finance*, 7(2):157–176, 1997. doi: 10.1111/1467-9965.00029.

D. F. Schrager and A. A. J. Pelsser. Pricing swaptions and coupon bond options in affine term structure models. *Mathematical Finance*, 16(4):673–694, 2006. doi: 10.1111/j.1467-9965.2006.00289.x.

K. J. Singleton and L. Umantsev. Pricing coupon-bond options and swaptions in affine term structure models. *Mathematical Finance*, 12(4):427–446, 2002. doi: 10.1111/j.1467-9965.2002.tb00132.x.

N. Skantzos and G. Garston. The perfect smile. Technical report, Deloitte, 2019.

C. Van Loan. Computing integrals involving the matrix exponential. *IEEE Transactions on Automatic Control*, 23(3):395–404, 1978. doi: 10.1109/TAC.1978.1101743.

A Appendix

A.1 Black formula for swaption pricing

Let us consider a swap starting at T_0 and ending at T_{n_1}, with floating leg payment dates $(T_j)_{j=1,\cdots,n_1}$ (and reset dates $(T_j)_{j=0,\cdots,n_1-1}$) and fixed rate leg payment dates given by $(t_i)_{i=1,\cdots,m_1}$ with $T_{j+1}-T_j=\delta$, $t_{i+1}-t_i=\Delta$, $t_{m_1}=T_{n_1}$ and $t_0=T_0$. At time t, the floating leg value is given by $P(t,T_0)-P(t,T_{n_1})+\sum_{j=1}^{n_1}A(t,T_{j-1},T_j)$ and the fixed leg value is $K\sum_{i=1}^{m_1}\Delta P(t,t_i)$, where K is the fixed rate.

We can therefore derive the time-t forward swap rate, $S_t^{T_0,T_{n_1}}$ as:

$$S_t^{T_0,T_{n_1}}=\frac{P(t,T_0)-P(t,T_{n_1})+\sum_{j=1}^{n_1}A(t,T_{j-1},T_j)}{\Delta\sum_{i=1}^{m_1}P(t,t_i)}\,. \tag{92}$$

Let us denote $An_t^{T_0,T_{n_1}}=\sum_{i=1}^{m_1}\Delta P(t,t_i)$ the annuity. The swap rate is a martingale under the swap numeraire (also called annuity numeraire) and if we assume that the swap rate follows a normal process, its dynamic under the swap numeraire can then be written as:

$$dS_t^{T_0,T_{n_1}}=\sigma dW_t. \tag{93}$$

This leads to $S_T^{T_0,T_{n_1}}\sim\mathcal{N}\left(S_t^{T_0,T_{n_1}},\,\sigma\sqrt{T-t}\right)$. The value V_t^{swaption} at time t of a swaption associated with the swap described above, with option expiry date T_0, is given by (below we simplify the notation by replacing $An_t^{T_0,T_{n_1}}$ with An_t):

$$\begin{aligned}
V_t^{\text{swaption}}&=An_t\mathbb{E}_t\left[\frac{1}{An_{T_0}}\left(P(T_0,T_0)-P(T_0,T_{n_1})+\sum_{j=1}^{n_1}A(T_0,T_{j-1},T_j)-K\sum_{i=1}^{m_1}\Delta P(T_0,t_i)\right)_+\right], && (94)\\
&=An_t\mathbb{E}_t\left[\left(S_{T_0}^{T_0,T_{n_1}}-K\right)_+\right], && (95)\\
&=\sum_{i=1}^{m_1}\Delta P(t,t_i)\left(\left(S_t^{T_0,T_{n_1}}-K\right)\Phi\left(\frac{S_t^{T_0,T_{n_1}}-K}{\sigma\sqrt{T_0-t}}\right)+\sigma\sqrt{T_0-t}\,\Phi'\left(\frac{S_t^{T_0,T_{n_1}}-K}{\sigma\sqrt{T_0-t}}\right)\right), && (96)
\end{aligned}$$

where Φ is the cumulative distribution function of the standard normal variable and Φ' its derivative. When the swaption is at the money then its price simplifies to

$$V_t^{\text{swaption}}=\sum_{i=1}^{m_1}\Delta P(t,t_i)\sqrt{\frac{T_0-t}{2\pi}}\sigma\,. \tag{97}$$

It is a market practice to quote the swaption price through its normal volatility σ.

A.2 Synthesizing a forward swap with two spot swaps

Following the notation of the swap above, let us consider the swap $S_0^{T_0=0,T_{n_1}}$ (the swap starting at time $T_0=0$ and ending at T_{n_1}, $T_0=0<T_{n_1}$) with the floating leg reset and payment dates $T_0,T_1,\cdots,T_{n_1}$, with $T_j-T_{j-1}=\delta$, and the fixed leg payment dates $t_1,\cdots,t_{m_1}=T_{n_1}$ and $t_i-t_{i-1}=\Delta$, $(T_0=t_0=0)$.

Similarly, we consider a second swap $S_0^{T_0=0,T_{n_2}}$ with floating leg reset and payment dates $T_0,T_1,\cdots,T_{n_2}$, with $T_j-T_{j-1}=\delta$, and fixed leg payment dates $t_1,\cdots,t_{m_2}=T_{n_2}$ and $t_i-t_{i-1}=\Delta$, (with $t_{m_2}=T_{n_2}$, and $T_0=0$).

The par swap rates are given by:

$$\begin{aligned}
S_0^{0,T_{n_1}}&=\frac{1-P(0,T_{n_1})+\sum_{j=1}^{n_1}A(0,T_{j-1},T_j)}{\Delta\sum_{i=1}^{m_1}P(0,t_i)}, && (98)\\
S_0^{0,T_{n_2}}&=\frac{1-P(0,T_{n_2})+\sum_{j=1}^{n_2}A(0,T_{j-1},T_j)}{\Delta\sum_{i=1}^{m_2}P(0,t_i)}. && (99)
\end{aligned}$$

Suppose that $T_{n_1}<T_{n_2}$, then the forward starting swap rate $S_0^{T_{n_1},T_{n_2}}$ with floating leg reset and payment dates $T_{n_1},T_{n_1+1},\cdots,T_{n_2}$ and fixed leg payment dates $t_{m_1+1},\cdots,t_{m_2}=T_{n_2}$ can be expressed as function of $S_0^{T_0=0,T_{n_1}}$ and $S_0^{T_0=0,T_{n_2}}$ as follows

$$\begin{aligned}
S_0^{T_{n_1},T_{n_2}}&=\frac{P(0,T_{n_1})-P(0,T_{n_2})+\sum_{j=n_1+1}^{n_2}A(0,T_{j-1},T_j)}{\Delta\sum_{i=m_1+1}^{m_2}P(0,t_i)}, && (100)\\
&=\frac{S_0^{T_0,T_{n_2}}\left(\Delta\sum_{i=1}^{m_2}P(0,t_i)\right)-S_0^{T_0,T_{n_1}}\left(\Delta\sum_{i=1}^{m_1}P(0,t_i)\right)}{\Delta\sum_{i=m_1+1}^{m_2}P(0,t_i)}. && (101)
\end{aligned}$$

For the model calibration purpose, we will consider spot swap rates ($t=T_0=0$). Further, in order to apply the formula (101) above, the start date T_{n_1} of the underlying swap of the swaption $S_{t=0}^{T_{n_1},T_{n_2}}$ should be one of the payment date of the spot swap $S_{t=0}^{T_0=0,T_{n_2}}$.

A.3 Tables

Table I: Descriptive statistics

Maturity	0.5	1	3	5	7	10	12	15
					OIS			
Mean	0.466	0.460	0.713	1.189	1.598	1.995	2.187	2.365
Std. dev.	0.150	0.138	0.170	0.187	0.174	0.149	0.142	0.140
					Euribor			
Mean	1.575	1.423	1.423	1.806	2.143	2.464	2.618	2.756
Std. dev.	0.183	0.176	0.189	0.200	0.182	0.159	0.152	0.151

Note: Mean value and standard deviation of the OIS and Euribor term structures (with the maturity expressed in years). Rates are expressed in percentage and the data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table II: Swaption volatilities

Swap tenor	1	2	3	4	5
			1Y		
Mean	70.72	75.06	79.24	84.12	88.06
Std. dev.	16.34	15.49	13.79	11.31	9.80
			2Y		
Mean	85.51	85.83	87.87	90.34	91.49
Std. dev.	12.39	10.43	8.71	7.12	6.36
			3Y		
Mean	94.74	92.02	91.29	91.41	92.07
Std. dev.	8.27	6.52	5.55	5.07	4.41
			4Y		
Mean	96.58	92.52	91.44	90.91	90.63
Std. dev.	5.74	4.71	4.11	3.49	3.14
			5Y		
Mean	95.04	91.28	89.70	88.97	88.49
Std. dev.	3.71	3.33	2.98	2.72	2.60

Note: Mean value and standard deviation of the normal implied swaption volatilities (expressed in basis points) for the swaption maturities 1Y, 2Y, 3Y, 4Y and 5Y (in years) and swap tenors (in years). The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table III: Calibrated parameters

Param.	Mean	Std. dev.	Min	Max
α	0.024	1.440×10^{-3}	0.021	0.027
x_{11}	0.125	0.037	0.045	0.218
x_{12}	-1.121×10^{-2}	4.038×10^{-3}	-2.713×10^{-2}	-6.675×10^{-3}
x_{22}	5.745×10^{-3}	6.590×10^{-4}	0.004	0.006
ω_{11}	0.130	0.029	0.068	0.203
ω_{12}	1.797×10^{-3}	2.011×10^{-3}	1.950×10^{-4}	0.011
ω_{22}	4.660×10^{-4}	2.320×10^{-4}	5.026×10^{-5}	9.610×10^{-4}
m_{11}	-0.375	0.016	-0.416	-0.352
m_{22}	-0.181	0.053	-0.284	-0.069
σ_{11}	0.050	0.013	0.019	0.121
σ_{12}	0.024	6.904×10^{-3}	0.009	0.038
σ_{22}	0.047	0.085	0.020	0.067

Note: Mean value, standard deviation, min value and max value of the calibrated parameters obtained by rolling the daily calibration. The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table IV: Eigenvalues of the matrices

	First				Second			
Param.	Mean	Std. dev.	Min	Max	Mean	Std. dev.	Min	Max
x	0.126	0.037	0.046	0.219	4.601×10^{-3}	9.890×10^{-4}	0.002	0.006
ω	0.130	0.029	0.068	0.203	4.130×10^{-4}	2.150×10^{-4}	5.000×10^{-5}	9.050×10^{-4}
σ	0.056	0.028	0.014	0.124	0.041	0.024	0.010	0.092

Note: Mean value, standard deviation, min value and max value of the eigenvalues of the estimated parameters. The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table V: Correlation associated with parameters

Param.	Mean	Std. dev.	Min	Max
x	-0.423	0.124	-0.837	-0.258
ω	0.223	0.190	0.063	0.885
σ	0.497	0.118	0.292	0.754
$\bar{x}_\infty$	0.208	0.183	0.058	0.867

Note: Mean value, standard deviation, min value and max value of the correlation associated with the estimated parameters with $\bar{x}_\infty$ defined in Eq. (19). The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table VI: Calibration errors

	OIS error	Spread error	Swaption error
Mean	115.04	3.58	36.24
Std. dev.	8.63	2.11	4.27

Note: Mean value and standard deviation of the daily root mean square errors of the calibrations. OIS error stands for the square root of the error Eq. (84) expressed in basis points, Spread error stands for the square root of the error Eq. (85) expressed in basis points and Swaption error stands for the square root of the error Eq. (86) expressed in basis points. The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

23

Table VII: Market vs. model correlations

Maturity	1	3	5	7	10	12	15
Market	0.514	0.270	0.241	0.174	0.184	0.181	0.128
Model drift	0.304	0.239	0.175	0.159	0.157	0.157	0.158
Model vol	0.353	0.342	0.322	0.301	0.272	0.256	0.235

Note: Market and model correlations, given by Eq. (91), between the OIS zero-coupon bond with maturity T and the sum of spreads up to maturity T for different values for T (in years). "Model drift" stands for the computation of Eq. (91) using model OIS zero-coupon bonds and model spreads that solve Eq. (84) and Eq. (85), respectively. "Model vol" stands for the computation of Eq. (91) using the right hand side of Eq. (42) and the calibrated parameters solving Eq. (86) (along with x_{11} and x_{22} obtained from Eq. (84) and Eq. (85), respectively). The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table VIII: Calibrated parameters for a diagonal model

Param.	Mean	Std. dev.	Min	Max
σ_{11}	0.058	0.010	0.035	0.079
σ_{22}	0.035	0.005	0.023	0.046

Note: Mean value, standard deviation, min value and max value of the calibrated parameters obtained by rolling the daily calibration for a model with diagonal parameters. The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table IX: Calibration errors for a diagonal model

	Swaption error
Mean	36.51
Std. dev.	2.92

Note: Mean value and standard deviation of the daily root mean square errors of the calibration for a diagonal model (*.i.e.*, σ_{11} and σ_{22}). Swaption error stands for the square root of the error Eq. (86) expressed in basis points. The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

Table X: Model correlation for a diagonal model

Maturity	1	3	5	7	10	12	15
Model vol	0.039	0.039	0.039	0.038	0.038	0.038	0.037

Note: Model correlation, named "Model vol" in the table, for a diagonal model given by Eq. (91) between the OIS zero-coupon bond with maturity T and the sum of spreads up to maturity T for different values for T (in years). The correlation is computed using Eq. (42) and the calibrated parameters. The calibrated parameter x_{11} is obtained by solving Eq. (84), x_{22} is obtained by solving Eq. (85) while σ_{11} and σ_{22} are obtained by solving Eq. (86). The data sample period is 4/10/2011 to 12/03/2012 at daily frequency.

A.4 Figures

Figure 1: Density of the variable Y with maturity 1Y and tenor 1Y

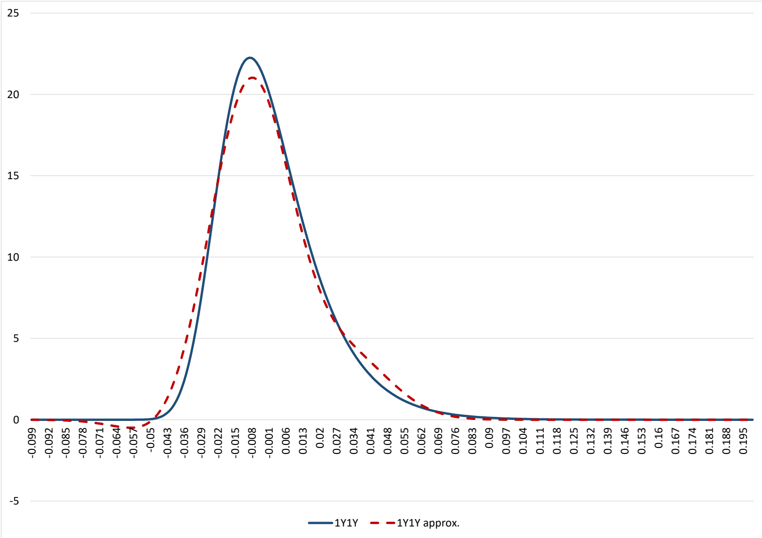

Note: Density of the variable Y defined as $Y_{T_0} = B(T_0, T_{n_1}) + A_1(T_0, T_{n_1})x_{11,T_0} + A_2(T_0, T_{n_1})x_{22,T_0}$ involved in the pricing formula Eq. (49) for a swaption with maturity $1Y$ and tenor $1Y$ given by the solid blue line (1Y1Y) and in red dash line (1Y1Y approx.) its approximation using the first three cumulants and Eq. (79). The parameters used to compute the density are those of 4/10/2011.

Figure 2: Density of the variable Y with maturity 5Y and tenor 5Y

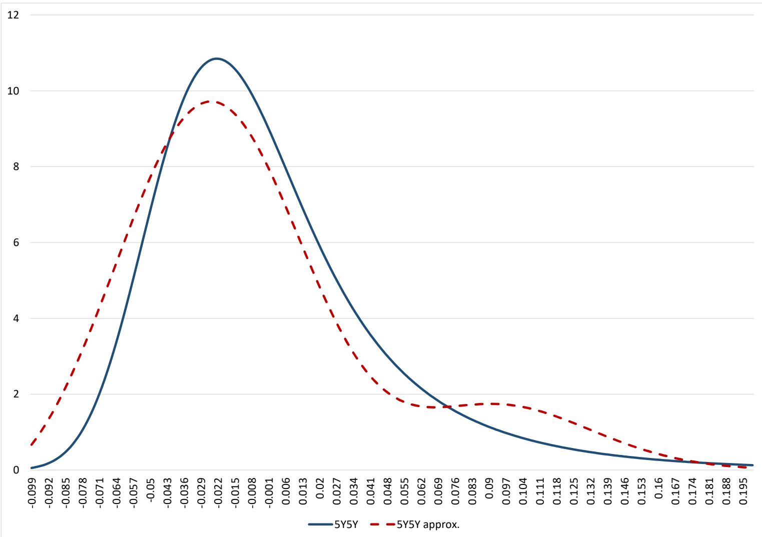

Note: Density of the variable Y defined as $Y_{T_0} = B(T_0, T_{n_1}) + A_1(T_0, T_{n_1})x_{11,T_0} + A_2(T_0, T_{n_1})x_{22,T_0}$ involved in the pricing formula Eq. (49) for a swaption with maturity $5Y$ and tenor $5Y$ given by the solid blue line (5Y5Y) and in red dash line (5Y5Y approx.) its approximation using the first three cumulants and Eq. (79). The parameters used to compute the density are those of 4/10/2011.

A.5 Proofs

Proof of Proposition 2.1. Thanks to the property of the exponential function, when $\theta_2 = 0_n$ (with 0_n the $n \times n$ null matrix) Eq. (14) leads to the system of matrix ODEs (*e.g.*, see Van Loan 1978)

$$A'_{11} = mA_{11} - 2\sigma^2 A_{21}\,, \tag{102}$$
$$A'_{12} = mA_{12} - 2\sigma^2 A_{22}\,, \tag{103}$$
$$A'_{21} = -m^\top A_{21}\,, \tag{104}$$
$$A'_{22} = -m^\top A_{22}\,, \tag{105}$$

and the initial conditions $A_{11}(0) = I_n$, $A_{12}(0) = 0_n$, $A_{21}(0) = 0_n$ and $A_{22}(0) = I_n$ (with I_n the $n \times n$ identity matrix). Solving these ODEs lead to: $A_{21}(t) = 0$, $A_{11}(t) = e^{mt}$, $A_{22}(t) = e^{-m^\top t}$ and

$$A_{12}(t) = \int_0^t e^{(t-s)m}(-2\sigma^2)e^{-sm^\top}\,ds\,. \tag{106}$$

As a result, $e^{\mathrm{tr}(a(t)x_0)}$ in Eq. (10) after some transformations is given by

$$\mathrm{etr}\left(e^{m^\top t}\left(\theta_1 \int_0^t e^{(t-s)m}(-2\sigma^2)e^{(t-s)m^\top}ds + I\right)^{-1}\theta_1 e^{mt}x_0\right)\,, \tag{107}$$

where $\mathrm{etr(A)} := e^{\mathrm{tr}(A)}$.

In the Bru case the term $e^{b(t)}$ in Eq. (10) rewrites as

$$\mathrm{etr}\left(-\frac{\beta}{2}m^\top t\right)\left(\mathrm{etr}\left(\log(\theta_1 A_{12} + A_{22})\right)\right)^{-\beta/2}\,, \tag{108}$$

and thanks to the relation $\det(e^A) = e^{\mathrm{tr}(A)}$ we get

$$e^{b(t)} = \det\left(I + \theta_1 \int_0^t e^{(t-s)m}(-2\sigma^2)e^{(t-s)m^\top}ds\right)^{-\beta/2}\,. \tag{109}$$

Combining Eq. (107) and Eq. (109) gives the moment generating function of x_t.

Consider Eq. (9) with $\theta_2 = 0_n$ and θ_1 replaced with $-\theta_1$ with $\theta_1 \in \mathbb{S}_n^{++}$, it is the Laplace transform of x_t that is given by

$$\begin{aligned}\mathbb{E}_x\left[\mathrm{etr}(-\theta_1 x_t)\right] &= \det\left(I - \theta_1 \int_0^t e^{(t-s)m}(-2\sigma^2)e^{(t-s)m^\top}ds\right)^{-\beta/2} \\ &\times \mathrm{etr}\left(-e^{mt}x_0 e^{m^\top t}\theta_1\left(I - \int_0^t e^{(t-s)m}(-2\sigma^2)e^{(t-s)m^\top}ds\theta_1\right)^{-1}\right)\,, \end{aligned}\tag{110}$$

and defining Ξ_t as in Eq. (21) and Λ_t as in Eq. (22) lead to the result after reorganizing the terms. □

Proof of Proposition 3.6. It is known from Eq. (48) that the swap rate is given by

$$S_{T_j}^{T_{j,0},T_{j,n_s}} = \frac{P(T_j,T_j) - P(T_j,T_{j,n_s}) + \sum_{l=1}^{n_s} A(T_j,T_{j,l-1},T_{j,l})}{\Delta_s \sum_{k=1}^{m_s} P(T_j,t_{j,k})}\,. \tag{111}$$

The time t-value (with $t \le T_0$) of the leg that pays the Euribor rate plus a fixed rate K is given by

$$P(t,T_0) - P(t,T_{n_1}) + \sum_{j=1}^{n_1} A(t,T_{j-1},T_j) + \delta K \sum_{j=1}^{n_1} P(t,T_j)\,, \tag{112}$$

while the time t-value (with $t \le T_0$) of the leg that pays the swap rate is given by

$$\mathbb{E}_t\left[\sum_{j=0}^{n_1-1} \frac{\zeta_{T_{j+1}}}{\zeta_t}\delta S_{T_j}^{T_j,T_{j,n_s}}\right]\,, \tag{113}$$

and taking into account Eq. (111), it leads to evaluate

$$\mathbb{E}_t\left[\frac{\zeta_{T_{j+1}}}{\zeta_t}\frac{P(T_j,T_j) - P(T_j,T_{j,n_s}) + \sum_{l=1}^{n_s} A(T_j,T_{j,l-1},T_{j,l})}{\Delta_s \sum_{k=1}^{m_s} P(t_i,t_{i,k})}\right]\,. \tag{114}$$

Taking into account Eq. (28), Eq. (38) and $\mathbb{E}_{T_j}\left[e^{-\alpha T_{j+1}}\zeta_{T_{j+1}}\right] = e^{-\alpha T_{j+1}}(b_1(\delta_s) + a_1(\delta_s)x_{11,T_j})$, the above expectation is equal to

$$\frac{e^{-\alpha(T_{j+1}-t)}}{1 + x_{11,t}}\mathbb{E}_t\left[\frac{c_0 + c_1 x_{11,T_j} + c_2 x_{22,T_j} + c_{12}x_{11,T_j}x_{22,T_j} + c_{11}x_{11,T_j}^2}{\mu_0 + \mu_1 x_{11,T_j}}\right]\,, \tag{115}$$

with

$$c_0 = b_1(\delta_s)\left(b_1(T_j - T_j) - e^{-\alpha(T_{j,n_s}-T_j)}b_1(T_{j,n_s} - T_j) + \sum_{l=1}^{n_s} e^{-\alpha(T_{j,l-1}-T_j)}b_2(T_{j,l-1} - T_j)\right), \tag{116}$$

$$c_1 = b_1(\delta_s)\left(a_1(T_j - T_j) - e^{-\alpha(T_{j,n_s}-T_j)}a_1(T_{j,n_s} - T_j)\right) + c_0\frac{a_1(\delta_s)}{b_1(\delta_s)}, \tag{117}$$

$$c_2 = b_1(\delta_s)\sum_{l=1}^{n_s} e^{-\alpha(T_{j,l-1}-T_j)}a_2(T_{j,l-1} - T_j), \tag{118}$$

$$c_{12} = a_1(\delta_s)\sum_{l=1}^{n_s} e^{-\alpha(T_{j,l-1}-T_j)}a_2(T_{j,l-1} - T_j), \tag{119}$$

$$c_{11} = a_1(\delta_s)\left(a_1(T_j - T_j) - e^{-\alpha(T_{j,n_s}-T_j)}a_1(T_{j,n_s} - T_j)\right), \tag{120}$$

$$\mu_0 = \Delta_s\sum_{k=1}^{m_s} e^{-\alpha(t_{j,k}-t_i)}b_1(t_{i,k} - t_i), \tag{121}$$

$$\mu_1 = \Delta_s\sum_{k=1}^{m_s} e^{-\alpha(t_{j,k}-t_i)}a_1(t_{i,k} - t_i), \tag{122}$$

which is the announced result. □

Proof of Proposition 3.8. Remark 3.7 allows the computation of the expectations

$$\mathbb{E}_t\left[\frac{x_{11,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0}\partial_z\Phi(\tau,\theta_1,0,x_t)ds|_{z=0}, \tag{123}$$

$$\mathbb{E}_t\left[\frac{x_{22,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0}\partial_z\Phi(\tau,\theta_2,0,x_t)ds|_{z=0}, \tag{124}$$

$$\mathbb{E}_t\left[\frac{x_{11,T_j}x_{22,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0}\partial^2_{z_1 z_2}\Phi(\tau,\theta_3,0,x_t)ds|_{z_1=0,z_2=0}, \tag{125}$$

$$\mathbb{E}_t\left[\frac{x^2_{11,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0}\partial^2_{zz}\Phi(\tau,\theta_4,0,x_t)ds|_{z=0}, \tag{126}$$

with $\tau = T_j - t$ and

$$\theta_1 = (z - \mu_1 s)e_{11}, \tag{127}$$
$$\theta_2 = ze_{22} - \mu_1 se_{11}, \tag{128}$$
$$\theta_3 = (z_1 - \mu_1 s)e_{11} + z_2 e_{22}, \tag{129}$$
$$\theta_4 = (z - \mu_1 s)e_{11}. \tag{130}$$

The next step is to differentiate the moment generating function of the Wishart process which can be explicitly carried out. From Eq. (13) (if we drop the dependency of the matrices A_{ij} on t), if θ_1 is given by Eq. (127) then the computation of Eq. (123) leads to

$$\frac{d}{dz}a(t,\theta_1,0) = -(\theta_1 A_{12} + A_{22})^{-1}e_{11}A_{12}a(t) + (\theta_1 A_{12} + A_{22})^{-1}e_{11}A_{11}. \tag{131}$$

We might denote the above derivative as $d_z a(t,\theta_1,0)$. Under the hypothesis that $\omega = \beta\sigma^2$, consider $z \to \mathrm{tr}\left[\log(\theta_1 A_{12} + A_{22})\right]$, denote $c = \theta_1 A_{12} + A_{22}$ and $l = c - I_n$ then using the Taylor expansions for $\ln(I_n + X)$ and $(I_n + X)^{-1}$ we get $\frac{d}{dz}\mathrm{tr}[\ln c] = \frac{d}{dz}\mathrm{tr}[\ln(I_n + l)] = \mathrm{tr}[\frac{d}{dz}\{l - l^2/2 + \ldots\}] = \mathrm{tr}[\frac{dl}{dz}\{I_n - l + l^2 - \ldots\}] = \mathrm{tr}[\frac{dc}{dz}c^{-1}]$ (thanks to $\mathrm{tr}[\frac{dl}{dz}l] = \mathrm{tr}[l\frac{dl}{dz}]$) and as a result

$$\frac{d}{dz}b(t,\theta_1,0) = -\frac{\beta}{2}\mathrm{tr}\left[e_{11}A_{12}(\theta_1 A_{12} + A_{22})^{-1}\right]. \tag{132}$$

We might denote the above derivative as $d_z b(t,\theta_1,0)$. Combining these two derivatives, we get the expression for $\partial_z\Phi(T - t,\theta_1,0,x_t) = (\mathrm{tr}[d_z a(t,\theta_1,0)x_t] + d_z b(t,\theta_1,0))\,\Phi(T - t,\theta_1,0,x_t)$ that is involved in Eq. (123) and when evaluated at $z = 0$ then $\theta_1 = -\mu_1 se_{11}$ and it leads to:

$$\mathbb{E}_t\left[\frac{x_{11,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0}\left(\mathrm{tr}[d_z a(\tau,\theta_1,0)x_t] + d_z b(\tau,\theta_1,0)\right)\Phi(\tau,\theta_1,0,x_t)ds|_{z=0}, \tag{133}$$

with $\tau = T_j - t$.

To compute Eq. (124), we just need to replace in Eq. (131) and Eq. (132) e_{11} with e_{22} whilst when the derivative is evaluated at $z = 0$ then $\theta_2 = -\mu_1 se_{11}$, it leads to:

$$\mathbb{E}_t\left[\frac{x_{22,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0}\left(\mathrm{tr}[d_z a(\tau,\theta_2,0)x_t] + d_z b(\tau,\theta_2,0)\right)\Phi(\tau,\theta_2,0,x_t)ds|_{z=0}, \tag{134}$$

with $\tau = T_j - t$.

To compute the second derivative involved in Eq. (125) for θ_3 defined in Eq. (129), we derive Eq. (131) and Eq. (132) with respect to z_2, it leads to:

$$\begin{aligned}\frac{d^2}{dz_1 z_2} a(t,\theta_3,0) &= (\theta_3 A_{12} + A_{22})^{-1} e_{22} A_{12} (\theta_3 A_{12} + A_{22})^{-1} e_{11} A_{12} a(t,\theta_3,0) \\ &+ (\theta_3 A_{12} + A_{22})^{-1} e_{11} A_{12} (\theta_3 A_{12} + A_{22})^{-1} e_{22} A_{12} a(t,\theta_3,0) \\ &- (\theta_3 A_{12} + A_{22})^{-1} e_{11} A_{12} (\theta_3 A_{12} + A_{22})^{-1} e_{22} A_{11} \\ &- (\theta_3 A_{12} + A_{22})^{-1} e_{22} A_{12} (\theta_3 A_{12} + A_{22})^{-1} e_{11} A_{11}\,, \end{aligned} \tag{135}$$

and

$$\frac{d^2}{dz_1 z_2} b(t,\theta_3,0) = \frac{\beta}{2} \mathrm{tr}\left[e_{11} A_{12} (\theta_3 A_{12} + A_{22})^{-1} e_{22} A_{12} (\theta_3 A_{12} + A_{22})^{-1} \right]\,, \tag{136}$$

and when the derivative is evaluated at $z = 0$ then $\theta_3 = -\mu_1 s e_{11}$ and it gives:

$$\begin{aligned}\mathbb{E}_t\left[\frac{x_{11,T_j} x_{22,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] &= \int_0^{+\infty} e^{-s\mu_0} \Big(\big(\mathrm{tr}[d_{z_1} a(\tau,\theta_3,0) x_t] + d_{z_1} b(\tau,\theta_3,0)\big)\big(\mathrm{tr}[d_{z_2} a(\tau,\theta_3,0) x_t] + d_{z_2} b(\tau,\theta_3,0)\big) \\ &+ \big(\mathrm{tr}[d^2_{z_1 z_2} a(\tau,\theta_3,0) x_t] + d^2_{z_1 z_2} b(\tau,\theta_3,0)\big)\Big) \Phi(\tau,\theta_3,0,x_t) ds|_{z_1=0,z_2=0}\,, \end{aligned} \tag{137}$$

with $\tau = T_j - t$.

We might denote the above derivatives as $d^2_{z_1 z_2} a(t,\theta_3,0)$ and $d^2_{z_1 z_2} b(t,\theta_3,0)$, respectively. For the last derivative involved in Eq. (126) for θ_4 defined in Eq. (130), we derive Eq. (131) and Eq. (132) with respect to z. For the former, it amounts to consider Eq. (135) with e_{22} replaced with e_{11} and θ_3 with θ_4 whilst for the latter it amounts to consider Eq. (136) with e_{22} replaced with e_{11} and θ_3 with θ_4, and when these derivatives are evaluated at $z = 0$ then $\theta_4 = -\mu_1 s e_{11}$ and it gives:

$$\mathbb{E}_t\left[\frac{x^2_{11,T_j}}{\mu_0 + \mu_1 x_{11,T_j}}\right] = \int_0^{+\infty} e^{-s\mu_0} \big(\mathrm{tr}[d_z a(\tau,\theta_4,0) x_t] + d_z b(\tau,\theta_4,0)\big)^2 \Phi(\tau,\theta_4,0,x_t) ds|_{z=0}\,, \tag{138}$$

$$+ \int_0^{+\infty} e^{-s\mu_0} \big(\mathrm{tr}[d^2_{zz} a(\tau,\theta_4,0) x_t] + d^2_{zz} b(\tau,\theta_4,0)\big) \Phi(\tau,\theta_4,0,x_t) ds|_{z=0}\,, \tag{139}$$

with $\tau = T_j - t$. As a results, the expectations Eq. (60) are known up to a one dimensional integration. □

Proof of Proposition 3.9. Starting from Eq. (61) and replacing $S^{T_{1,0},T_{1,n_{s_1}}}_{T_1}$ and $S^{T_{1,0},T_{1,n_{s_2}}}_{T_1}$ with Eq. (111) leads to the announced results after performing computations similar to those of Proposition 3.6 but with

$$c^i_0 = \left(b_1(T_1 - T_1) - e^{-\alpha(T_{1,n_{s_i}} - T_1)} b_1(T_{1,n_{s_i}} - T_1) + \sum_{l=1}^{n_{s_1}} e^{-\alpha(T_{1,l-1} - T_1)} b_2(T_{1,l-1} - T_1) \right)\,, \tag{140}$$

$$c^i_1 = \Big(a_1(T_1 - T_1) - e^{-\alpha(T_{1,n_{s_i}} - T_1)} a_1(T_{1,n_{s_i}} - T_1) \Big) + c^i_0\,, \tag{141}$$

$$c^i_2 = \sum_{l=1}^{n_{s_i}} e^{-\alpha(T_{1,l-1} - T_1)} a_2(T_{1,l-1} - T_1)\,, \tag{142}$$

$$c^i_{12} = \sum_{l=1}^{n_{s_i}} e^{-\alpha(T_{1,l-1} - T_1)} a_2(T_{1,l-1} - T_1)\,, \tag{143}$$

$$c^i_{11} = \Big(a_1(T_1 - T_1) - e^{-\alpha(T_{1,n_{s_i}} - T_1)} a_1(T_{1,n_{s_i}} - T_1) \Big)\,, \tag{144}$$

$$\mu^i_0 = \Delta_s \sum_{k=1}^{m_{s_i}} e^{-\alpha(t_{1,k} - t_1)} b_1(t_{1,k} - t_1)\,, \tag{145}$$

$$\mu^i_1 = \Delta_s \sum_{k=1}^{m_{s_i}} e^{-\alpha(t_{1,k} - t_1)} a_1(t_{1,k} - t_1)\,. \tag{146}$$

□

28

A model free approach to continuous-time finance

Henry Chiu, Rama Cont

May 2021

Abstract

We present a non-probabilistic, pathwise approach to continuous-time finance based on functional calculus. In the presence of complete model uncertainty (including the continuous evolution of price paths), we first obtain the analytical analogues of the (probabilistic) classical notions in mathematical finance and show that generic domains of functional calculus is inherently arbitrage free. We then demonstrate how one may use this fundamental property to obtain the optimal hedging strategy by a (fully non-linear) path dependent equation and solved the case for Asian option explicitly.

A NEW APPROACH TO ESTIMATING LOSS-GIVEN-DEFAULT DISTRIBUTION BY MASAHIKO EGAMI AND RUSUDAN KEVKHISHVILI

ABSTRACT. We propose a new method for estimating loss-given-default distribution implied in the current credit market based on a new idea of using a last passage time. Since the market standard is the predetermined constant loss rate (60%) for all firms, it is hard to estimate loss-given-default distribution by just observing the market information such as default time distribution, default probability, and CDS spreads. To overcome this difficulty, we construct a *hybrid model* with the last passage time of the *leverage ratio* (defined as the ratio of a firm's assets over its debt) to a certain level. The last passage time, which is *not* a stopping time, allows us to appropriately model the timing of severe firm-value deterioration which is not apparent to credit market participants. Under minimal and standard assumptions, our model captures asset value dynamics close to default, while treating default as an unexpected event. We explicitly obtain the distribution of the leverage ratio at default time. As we do not introduce additional recovery risk factors, both the model and the procedure for calibrating model parameters to the credit market are simple. We illustrate this procedure in detail by calculating the loss-given-default distribution implied in the quoted CDS spreads.

0.1. **Last Passage time.** Given the fixed probability space $(\Omega, \mathscr{G}, \mathbb{P})$, let the amount of debt be given by a positive deterministic process $(B_t)_{t\geq 0}$ satisfying $\mathrm{d}B_t = rB_t\mathrm{d}t$ with constant growth rate r. As in the firm-value approach, the nonnegative regular diffusion V (with random switching) represents the asset process and we define the *leverage process* as $\left(\frac{V}{B}\right)_{t\geq 0}$. We assume that the dynamics of V is given by

$$\mathrm{d}V_t = \mu(V_t, I(t))\mathrm{d}t + \sigma(V_t, I(t))\mathrm{d}W_t, \tag{1}$$

where W is a standard Brownian motion. The process I is defined as $I(t) = 1_{\{(V/B)_t \leq \alpha\}}$ and the process V is measurable with respect to $\mathscr{G}$. We define

$$L_\alpha := \sup\left\{0 \leq t < \infty : \frac{V_t}{B_t} = \alpha\right\}.$$

The level $\alpha > 0$ represents "distress level" and is to be calibrated to the market quoted 5 year default probability. The modeling in (1) allows us to incorporate the fact that the firm's business operations are altered when it is under financial distress. In fact, after the last passage time L_α, business environment surrounding the company severely alters and the firm is not able to recover to normal operations. As is well known, the last passage time is not a stopping time and this is, in fact, the reason why L_α is appropriate for modeling loss-given-default distribution: market observers (such as regular shareholders) do not know whether L_α has occurred, nor see when this switch in the dynamics of the asset value takes place. Therefore, the last passage time is relevant in the sense that it accounts for the reality that market observers are facing.

0.2. **Unexpected Default time.** After the time L^α, as in the intensity-based model, default occurs abruptly. We let $\lambda : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be a nonnegative piecewise continuous function and define the moment inverse of integral functional by

$$\tau := \min\left(s : \int_0^s \lambda\left(\frac{V_u}{B_u}\right)\mathrm{d}u = J\right). \tag{2}$$

Fig. 1. (a) The density function of τ in $\xi = L_\alpha + \tau \circ \theta_{L_\alpha}$ and (b) the density of the loss given default distribution

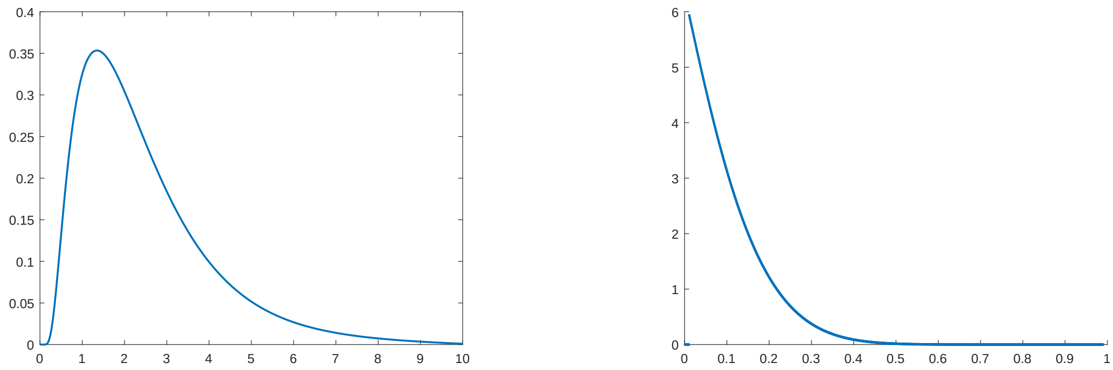

(a) (b)

that satisfies $\int_0^\infty \lambda\left(\frac{V_s}{B_s}\right) \mathrm{d}s = \infty \quad \mathbb{P}-a.s$, where J is a unit exponential random variable independent of V. Then the default time ξ is written as

$$\xi := L_\alpha + \tau \circ \theta_{L_\alpha},$$

where θ denotes the shift operator. In our paper, we set $\lambda(x) := 1_{x \le d}$ with $d \le \alpha$, so that the integral in (2) represents the occupation time of the process $(\frac{V}{B})_{t \ge 0}$ under a certain level. Occupation time of a region is widely used in the credit-risk-related literature such as Albrecher et al. (2011). The innovative part of our usage is that we consider occupation time of a region after the last passage time L_α. In this way, we look at the occupation time of a 'dangerous zone' after the firm's credit conditions have deteriorated.

0.3. **Loss-given-default distribution.** In our implementation of the model, we set a geometric Brownian motion in (1):

$$\mu_0 v = \mu(v,0) \neq \mu(v,1) = \mu_1 v, \quad \text{and} \quad \sigma(v,0) = \sigma(v,1) = \sigma v, \quad v \in (0,+\infty)$$

with $\mu_0, \mu_1 \in \mathbb{R}$ and $\sigma > 0$. Subsequent to L^α, the diffusion $X_t := \frac{1}{\sigma} \ln\left(\frac{V_t}{B_t}\right)$ has the following form that represents the fact that it does not return to level α:

$$\mathrm{d}X_t = m \coth\left(m\left(X_t - \alpha^*\right)\right) \mathrm{d}t + \mathrm{d}W_t \quad \text{where } \alpha^* = \frac{1}{\sigma}\ln(\alpha), \quad m := \frac{\mu_1 - \frac{1}{2}\sigma^2 - r}{\sigma}.$$

In particular, α becomes the entrance boundary for X. It is this process X that is evaluated in (2). Under this setup, we obtain in *explicit forms* the distribution of τ and $\left(\frac{V_\xi}{B_\xi}\right)$, from which we compute the loss given default distribution. We implement our model by taking an example of Ford Motor Company, using past 1000 daily observations from October 1, 2021, together with $r = 0.0093$. The estimated parameters are $\alpha = 1.7, \mu_0 = \mu_1 = 0.01022, \sigma = 0.11821$, and the threshold level for the occupation time computation is 1.4. See the density functions in Figure 1.

References

Albrecher, H., Gerber, H. U., Shiu, E. S. W., 2011. The optimal dividend barrier in the Gamma-Omega model. European Actuarial Journal 1, 43–55.

A note on product-convolution for generalized subexponential distributions

Dimitrios Konstantinides[a], Remigijus Leipus[b], Jonas Šiaulys[c,]

[a] *Department of Statistics and Actuarial - Financial Mathematics, University of the Aegean, Karlovassi, GR-83 200 Samos, Greece*
[b] *Institute of Applied Mathematics, Vilnius University, Naugarduko 24, Vilnius LT-03225, Lithuania*
[c] *Institute of Mathematics, Vilnius University, Naugarduko 24, Vilnius LT-03225, Lithuania*

Abstract

In this paper we consider the stability property of the class of generalized subexponential distributions with respect to product-convolution. Assuming that the primary distribution is in the class of generalized subexponential distributions, we find conditions for the second distribution in order that that their product-convolution belongs to the class of generalized subexponential distributions as well. The similar problem for the class of generalized subexponential positively decreasing-tailed distributions is considered.

Keywords: product-convolution, subexponential distribution, generalized subexponential distribution, positively decreasing-tailed distribution

2000 MSC: 60E05, 60G57, 62P05

References

[1] Albin, J.M.P., Sundén, M., 2009. On the asymptotic behavior of Lévy processes, Part I: Subexponential and exponential processes. Stoch. Proc. Appl. 119, 281—304.
[2] Assmussen, S., 2003. Applied Probability and Queues. 2nd ed. Springer, New York.
[3] Assmussen, S., Albrecher, H., 2010. Ruin Probabilities. World Scientific, Singapure.
[4] Baltrūnas, A., Omey, E., Van Gulck, S., 2006. Hazard rates and subexponential distributions. Publ. de l'Institute Math. 80, 29–46.
[5] Bardoutsos, A.G., Konstantinides, D.G., 2011. Characterization of tails through hazard rate and convolution closure properties. J. Appl. Probab. 48A, 123–132.
[6] Berkes, I., Györfi, L., Kevei P., 2017. Tail probabilities of St. Petersburg sums, trimmed sums, and their limit. J. Theor. Probab. 30, 1104–1129.
[7] Bingham, N.H., Goldie, C.M., Teugels, J.L., 1987. Regular Variation. Cambridge University Press, Cambridge.
[8] Chen Y., Ng K.W., Tang Q., 2005. Weighted sums of subexponential random variables and their maxima. Adv. Appl. Probab. 37, 510–522.
[9] Cheng, D., Wang, Y., 2012. Asymptotic behavior of the ratio of tail probabilities of sum and maximum of independent random variables. Lith. Math. J. 52, 29–39.
[10] Chistyakov, V.P., 1964. A theorem on sums of independent, positive random variables and its applications to branching processes. Theory Probab. Appl. 9, 640–648.
[11] Cline, D.B.H., 1986. Convolution tails, product tails and domain of attraction. Probab. Theory Relat. Fields 72, 529–557.
[12] Cline, D.B.H., Samorodnitsky, G., 1994. Subexponentiality of the product of independent random variables. Stoch. Proc. Appl. 49, 75–98.
[13] Cui, Z., Wang, Y., 2020. On the long tail property of product convolution. Lith. Math. J. 60, 315–329.
[14] Dirma, M., Paukštys, S., Šiaulys, J., 2021. Tails of the moments for sums with dominatedly varying random summands, Mathematics 9, 824.
[15] Embrechts, P., Klüppelberg, C., Mikosch, T., 1997. Modelling Extremal Events: for Insurance and Finance. Springer. New York.
[16] Feller, W., 1969. One-sided analogues of Karamata's regular variation. L'Enseignement Mathematique 15, 107–121.
[17] Foss, S., Korshunov, D., Zachary, S., 2013. An Introduction to Heavy-Tailed and Subexponential Distributions. 2nd edition, Springer-Verlag, New York.
[18] Galambos, J., Simonelli, I., 2004. Product of Random Variables: Applications to Problems of Physics and to Arithmetical Functions. CRC Press, Boca Raton, London, New York.
[19] Gerber, H.U., 1979. An Introduction to Mathematical Risk Theory. Huebner Foundation Monographs, Philadelphia.
[20] Goldie, C.M., 1978. Subexponential distributions and dominated variation tails. J. Appl. Probab. 15, 440–442.

Email addresses: `konstant@aegean.gr` (Dimitrios Konstantinides), `remigijus.leipus@mif.vu.lt` (Remigijus Leipus), `jonas.siaulys@mif.vu.lt` (Jonas Šiaulys)

Preprint submitted to Elsevier *February 12, 2022*

[21] Hao, X., Tang, Q., 2008. A uniform asymptotic estimate for discounted aggregate claims with subexponential tails. Insur. Math. Econ. 43, 116–120.
[22] Klüppelberg, C. 1989. Subexponential distributions and characterization of related classes. Probab. Theor. Rel. Fields 82, 259–269.
[23] Klüppelberg, C. 1990. Asymptotic ordering of distribution functions and convolution semigroups. Semigroup Forum 40, 77–92.
[24] Konstantinides, D.G., 2008. A class of heavy tailed distributions. J. Numer. Appl. Math. 96, 127–138.
[25] Konstantinides, D.G., Tang, Q., Tsitsiashvili, G., 2002. Estimates for the ruin probability in the classical risk model with constant interest force in the presence of heavy tails. Insur. Math. Econ. 31, 447–460.
[26] Leipus, R., Paukštys, S., Šiaulys, J.. 2021. Tails of higher-order moments of sums with heavy-tailed increments and application to the Haezendonck-Goovaerts risk measure. Stat. Probab. Lett. 170, 108998.
[27] Lin, J., Wang, Y., 2012. New examples of heavy-tailed O-subexponential distributions and related closure properties. Stat. Probab. Lett. 82, 427–432.
[28] Markovich N., 2007. Nonparametric Analysis of Univariate Heavy-Tailed Data. Wiley. New York.
[29] Matuszewska, W., 1964. On generalization of regularly increasing functions. Studia Mathematica 24, 271–279.
[30] Rolski, T., Schmidli, H., Schmidt, V., Teugels, J.F., 1999. Stochastic Processes for Insurance and Finance. Wiley. Chichester.
[31] Schmidli, H., 2017. Risk Theory. Springer, New York.
[32] Shimura, T., Watanabe, T., 2005. Infinite divisibility and generalized subexponentiality. Bernoulli 11, 445–469.
[33] Tang, Q., 2004. Asymptotics for the finite time ruin probability in the renewal model with consistent variation. Stoch. Models 20, 281–297.
[34] Tang, Q., 2006. On convolution equivalence with applications. Bernoulli 12, 535–549.
[35] Tang, Q., 2006. The subexponentiality of products revisited. Extremes 9, 231–241.
[36] Tang Q., Tsitsiashvili G., 2003. Precise estimates for the ruin probability in finite horizon in a discrete-time model with heavy-tailed insurance and financial risks, Stoch. Process. Appl. 108, 299–325.
[37] Tang, Q., Tsitsiashvili, G., 2003. Randomly weighted sums of subexponential random variables with application to ruin theory, Extremes 6, 171–188.
[38] Tang, Q., Tsitsiashvili, G., 2004. Finite- and infinite-time ruin probabilities in the presence of stochastic returns on investments, Adv. Appl. Probab. 36, 1278–1299.
[39] Watanabe, T., 2008. Convolution equivalence and distributions of random sums, Probab. Theory Relat. Fields. 142, 367–397.
[40] Watanabe, T., Yamamuro, K., 2010. Ratio of the tail of an infinitely divisible distribution on the line to that of its Lévy measure. Electron. J. Probab. 15, 44–74.
[41] Xu, H., Foss, S., Wang, Y., 2015. Convolution and convolution-root properties of long-tailed distributions, Extremes 18, 605–628.
[42] Zwart, A.P., 2001. Queueing systems with heavy tails. Technische Universiteit Eindhoven. https://doi.org/10.6100/IR547196

A positivity preserving scheme for the alpha-CEV process.

Libo Li, Guanting Liu

March 15, 2022

Abstract

We propose a positivity preserving implicit numerical scheme for the jump-extended CEV (and CIR) process, whose jumps are governed by a compensated spectrally positive alpha-stable process, with alpha taking values in $(1, 2)$.

Key words and phrases. Euler-Maruyama scheme, positivity preserving implicit scheme, CBI process, alpha-CIR process and alpha-CEV process.

AMS 2000 subject classification.

Introduction

The CIR and the CEV processes are non-negative processes widely used in the modelling of interest rates, default rates and volatility. Extending with a compensated spectrally positive α-stable process, The affine α-CIR process was introduced by Jiao et al. to capture the persistency of low interest rate, self-exciting and large jump behaviours exhibited by sovereign interest rates and power markets [1, 2]. Li and Taguchi studied a positivity preserving scheme for the α-CIR process and proved strong convergence [3], though bounded $(\alpha-)$-moments of the scheme hasn't been obtained, and truncation was necessary in the method of proof. Our positivity preserving scheme for the α-CEV process employs a similar design as that in Li and Taguchi [3], though we successfully bounded $(\alpha-)$-moments of the scheme, lifted the truncation, and improved the method of proof significantly.

Consider the following stochastic differential equation:

$$X_t = x_0 + \int_0^t (a - kX_s)ds + \sigma_1 \int_0^t (X_s)^\gamma dW_s + \sigma_2 \int_0^t (X_s)^{\frac{1}{\alpha}} dZ_s, \tag{1}$$

where $x_0 > 0$, a, σ_1, σ_2 are non-negative, $k > 0$, $\gamma \in (1/2, 1)$, and $\alpha \in (1, 2)$. Process $\{W_t, t \geq 0\}$ is a standard Brownian motion. Process $\{Z_t, t \geq 0\}$, independent of W, is a compensated spectrally positive α-stable process, so its Lévy measure ν has the form

$$\nu(x) = \frac{1}{x^{\alpha+1}} \mathbb{1}_{(0,\infty)}(x)dx,$$

and its characteristic triple is $(\gamma_0, 0, \nu)$, with drift $\gamma_0 = -\int_1^\infty x\nu(dx) = 1/(1-\alpha)$. Using the corresponding compensated Poisson random measure $\tilde{N}$, we can express process Z in the form

$$Z_t = \int_0^t \int_0^\infty z\tilde{N}(dz, ds).$$

We also assume $a - \sigma_1^2/2 > 0$ and $2\gamma < \alpha$. Then using the work of Li and Mytnik (Theorem 2.3) [5], we know that the SDE (1) has a pathwise unique non-negative strong solution, which we simply denote as $\{X_t, t \geq 0\}$.

2

1 The Numerical Scheme

Consider the uniform time grid $0 = t_0 < t_1 < ... < t_n = T$, and on the grid we propose the positivity-preserving numerical scheme X^n, given by the following recursion:

$$
\begin{aligned}
X^n_{t_0} &= x_0, \\
X^n_{t_{i+1}} &= \left[\frac{\sigma_1 (X^n_{t_i})^{q-\frac{1}{2}}\Delta W_{t_i} + \sqrt{\sigma_1^2 (X^n_{t_i})^{2q-1}(\Delta W_{t_i})^2 + 4(1+k\Delta t)|D_{t_{i+1}}|}}{2(1+k\Delta t)}\right]^2, \\
D_{t_{i+1}} &= X^n_{t_i} + \left(a - \frac{\sigma_1^2}{2}(X^n_{t_i})^{2q-1}\right)\Delta t + \sigma_2 (X^n_{t_i})^{\frac{1}{\alpha}}\Delta Z_{t_i}.
\end{aligned} \tag{2}
$$

2 Strong convergence

Lemma 2.1. *For $\beta \in (1,\alpha)$, the β-moment of the discretization scheme X^n is finite. That is*

$$\mathbb{E}\left[\sup_{s\le T}(X^n_s)^\beta\right] < \infty.$$

Theorem 2.1. *There exists some constant $C_T > 0$ such that*

$$\sup_{t\le T}\mathbb{E}|X_t - X^n_t| \le C_T n^{-\frac{1}{2}q(\alpha)},$$

where $q(\alpha)$ is given by

$$q(\alpha) = \begin{cases} \frac{\alpha_-}{2} & \alpha \in (1,\sqrt{2}] \\ \frac{1}{\alpha} & \alpha \in (\sqrt{2}, 2) \end{cases}$$

and the constant $\alpha_- \in (1,\alpha)$ can be chosen arbitrarily close to α.

References

[1] Jiao, Y., Ma, C. and Scotti, S.: Alpha-CIR model with branching processes in sovereign interest rate modelling. Finance and Stochastics **21** (3), 789-813 (2017).

[2] Jiao, Y., Ma, C., Scotti, S. and Sgarra, C.: A branching process approach to power markets. Energy Economics **79**, 144-156 (2019).

[3] Li, L. and Taguchi, D.: On a positivity preserving numerical scheme for jump-extended CIR process: the alpha-stable case. BIT Numerical Mathematics **59** (3), 747–774 (2019).

[4] Li, L. and Taguchi, D.: On the Euler-Maruyama scheme for spectrally one-sided Lévy driven SDEs with Hölder continuous coefficients. Statistics and Probability Letters **146**, 15-26 (2019).

[5] Li, Z. and Mytnik, L.: Strong solutions for stochastic differential equations with jumps. Annales de l'Institut Henri Poincaré - Probabilités et Statistiques **47** (4), 1055-1067 (2011).

A Reinforcement Learning Framework for Behavioral Portfolio Selection under Cumulative Prospect Theory

Nixie S. Lesmana[1], Pun Chi Seng[1], and Tang Jingxiang[1]

[1] *Nanyang Technological University, Singapore*

December 31, 2021

Abstract

We develop an RL framework for multi-period behavioral portfolio selection model under S-shaped utility maximization with inverse-S-shaped probability weighting. We define a corresponding value function

$$V^{\boldsymbol{\pi}}(t, x_t) \doteq \int_B^{+\infty} (x-B)^\alpha \omega_+(\mathbb{P}_{t,x_t}(W_T^{\boldsymbol{\pi}} \in dx)) + \int_{-\infty}^{B} -\lambda(B-x)^\alpha \omega_-(\mathbb{P}_{t,x_t}(W_T^{\boldsymbol{\pi}} \in dx)) \tag{1}$$

where $\{W_t^{\boldsymbol{\pi}} : t = 0, 1, \ldots, T\}$ captures the wealth evolution under policy $\boldsymbol{\pi}$. Due to the existence of probability weighting i.e. the non-linear, state-dependent transformation functions (ω_+, ω_-) of the terminal wealth distribution, we face time-inconsistency (TIC) issue. We apply the policy improvement technique of SPERL [LP21] that has been shown to terminate to subgame perfect equilibrium (SPE) policy for any given TIC Q-function $Q^{\boldsymbol{\pi}}(t, x_t, u_t)$, defined as usual from a value function (as in (1)) by fixing $\pi_t(x_t) = u_t$. We note however that SPERL's policy evaluation draws from extended Bellman equation in [BMZ14] which does not cover probability weighting. Thus, we borrow the technique from distributional RL [BDM17] to establish a recursive relationship between terminal wealth distributions in $\{\mathbb{P}_{t,x_t}(W_T^{\boldsymbol{\pi}} \in dx) : t = 0, 1, \ldots, T\}$. We propose a training algorithm based on our developed RL framework and empirically test its performance under the environment setup in [SCL15] when semi-analytical solution of piece-wise linear feedback policy is available.

References

[BDM17] Marc G. Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning, 2017.

[BMZ14] Tomas Björk, Agatha Murgoci, and Xun Yu Zhou. Mean-variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance*, 24(1):1–24, January 2014.

[LP21] Nixie S. Lesmana and Chi Seng Pun. A subgame perfect equilibrium reinforcement learning approach to time-inconsistent problems. *Working paper at NTU Singapore*, 2021.

[SCL15] Yun Shi, Xiangyu Cui, and Duan Li. Discrete-time behavioral portfolio selection under cumulative prospect theory. *Journal of Economic Dynamics and Control*, 61:283–302, 2015.

A sensitivity analysis of the long-term expected utility of optimal portfolios

Hyungbin Park* and Stephan Sturm†

Abstract

In this study, we discuss the sensitivity of the long-term expected utility of optimal portfolios for an investor with constant relative risk aversion. Under an incomplete market given by a factor model, we consider the utility maximization problem with long-time horizon. The main purpose is to find the long-term sensitivity, that is, the long-term behavior of partial derivatives of the optimal expected utility with respect to underlying model parameters. We investigate several types of sensitivities: the partial derivatives with respect to the initial factor, the drift function and the volatility function.

To achieve this, we combine several techniques: the duality approach, the dynamic programming principle, the ergodic Hamilton–Jacobi–Bellman equation and the Hansen–Scheinkman decomposition. The factor model induces a specific eigenpair of an operator, and this eigenpair does not only characterize the long-term behavior of the optimal expected utility but also provides an explicit representation of the expected utility on a finite time horizon. We conclude that this eigenpair therefore determines the long-term sensitivity. As examples, explicit results for several market models such as the Kim–Omberg model for stochastic excess returns and the Heston stochastic volatility model are presented.

*Department of Mathematical Sciences and RIMS, Seoul National University, 1, Gwanak-ro, Gwanak-gu, Seoul, Republic of Korea. hyungbin@snu.ac.kr, hyungbin2015@gmail.com

†Department of Mathematical Sciences, Worcester Polytechnic Institute, 100 Institute Rd, Worcester, MA, USA. ssturm@wpi.edu

A Two-Step Framework for Arbitrage-Free Prediction of the Implied Volatility Surface

Wenyong Zhang[*], Lingfei Li[†], Gongqiu Zhang[‡]

December 31, 2021

Abstract

We propose a two-step framework for predicting the implied volatility surface over time without static arbitrage. In the first step, we select features to represent the surface and predict them over time. In the second step, we use the predicted features to construct the implied volatility surface using a deep neural network (DNN) model by incorporating constraints that prevent static arbitrage. We consider three methods to extract features from the implied volatility data: principal component analysis, variational autoencoder and sampling the surface, and we predict these features using LSTM. Using a long time series of implied volatility data for S&P500 index options to train our models, we find two feature construction methods, sampling the surface and variational autoencoders combined with DNN for surface construction, are the best performers in out-of-sample prediction. In particular, they outperform a classical method substantially. Furthermore, the DNN model for surface construction not only removes static arbitrage, but also significantly reduces the prediction error compared with a standard interpolation method. Our framework can also be used to simulate the dynamics of the implied volatility surface without static arbitrage.

[*]Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: wyzhang@se.cuhk.edu.hk.

[†]Corresponding author. Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: lfli@se.cuhk.edu.hk.

[‡]School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China. Email: zhanggongqiu@cuhk.edu.cn.

A Unified Framework for Regime-Switching Models

Ning Cai (HKUST) Steven Kou (Boston University) Yingda Song (SJTU)

I. Introduction

Despite the prevalence of regime-switching, only relatively simple models have been used to study regime switching, because adding regime switching to a model inevitably increases the dimension of the problem leading to more computational complexity. For example, many papers focus on relatively simple regime-switching models such as regime-switching geometric Brownian motions (GBMs) and their applications to the valuation of plain vanilla European options (see, e.g., [1], [2], [3], [4], [5], [6], [7], and [8]); of course, it would be more realistic to consider more sophisticated regime-switching models as well as the pricing of a wider range of financial derivatives (e.g., barrier, lookback and Asian options).

Although in general it is difficult to reduce a high dimensional model to a low dimensional one, we show that regime-switching finite state Markov chains, which are widely-used two-dimensional Markov processes, can be represented as a functional of one-dimensional Markov chains. Intuitively, this is possible because both Markov chains have a finite number of states and the functional is a bijection. The functional can lead to very simple analytical computations (as in the case of first passage times of a stochastic process), or more complicated computations (as in the case of the time averages of a stochastic process) yet still with analytical solutions.

More precisely, the contribution of this paper is threefold.

(i) Methodologically, we establish a unified analytical approximation framework for computions under *general regime-switching Markov models*.

(ii) We derive analytical approximations to for pricing of path-dependent options. Numerical results suggest that our computational methods are accurate, fast and simple to implement.

(iii) Our results indicate that regime changes can have significant modeling implications.

Our unified computational method can be viewed as a significant generalization of the elegant CTMC approximation method proposed by [9] (which is for barrier option pricing under general one-dimensional Markov models) to more general quantities (including average time integrals) under *two-dimensional*, regime-switching Markov models.

II. A Unified Analytical Approximation Framework

Consider a two-dimensional Markov process $\{(S_t, Z_t) : t \geq 0\}$, where S_t denotes the value of the state variable of interest at time t for any $t \geq 0$, and $\{Z_t : t \geq 0\}$ follows a continuous-time Markov chain with a finite state space $\mathcal{S}_Z = \{1, 2, \cdots, R\}$ and the transition rate matrix $\mathbf{\Lambda} = (\lambda_{ij})_{R\times R}$. Here $\{Z_t\}$ accounts for the regime changes of the dynamics of $\{S_t\}$ and hence is referred to as the *regime process*. In each regime, $\{S_t\}$ follows a general one-dimensional time homogeneous Markov process with the state space $\mathcal{S} \subset [0, \infty)$.

In this paper, we propose the following *unified three-step analytical approximation framework* for computations under general regime-switching Markov models.

A Unified Three-Step Analytical Approximation Framework for Computations under General Regime-Switching Markov Models

1) Construct a regime-switching CTMC that is sufficiently close to the original general regime-switching Markov model.
2) Convert the computational problems under the approximate regime-switching CTMC obtained in Step 1 into other computational problems under a newly constructed one-dimensional CTMC.
3) Derive analytical solutions to the resulting computational problems under the one-dimensional CTMC constructed in Step 2.

The approximation of Step 1 builds on the recent elegant work of [9]. In the remainder of this section, we shall focus on the implementation of Step 2, i.e., how to convert a computational problem under a regime-switching CTMC into another computational problem under a one-dimensional CTMC. Step 3 will be discussed in the next section.

For ease of exposition, we define a regime-switching CTMC (S, Z) as a special, regime-switching Markov process, where the regime process Z is a CTMC with the state space $\mathcal{S}_Z = \{1, 2, \cdots, R\}$ and the transition rate matrix $\mathbf{\Lambda} = (\lambda_{ij})$, and moreover, for each $i = 1, 2, \cdots, R$, given that Z stays in regime i, S is a CTMC with the state space $\mathcal{S} = \{x_1, x_2, \cdots, x_N\}$ with $0 \leq x_1 < \cdots < x_N$ and the transition rate matrix $\mathbf{G}_i$.

The idea is to convert the computation of the expectation of a (possibly path-dependent) functional contingent on a regime-switching CTMC (S, Z) into the computation of the expectation of another (possibly path-dependent) functional contingent on a newly constructed one-dimensional CTMC $\{Y_t : t \geq 0\}$ with the state space $\mathcal{S}_Y = \{1, 2, \cdots, N \cdot R\}$

and the transition rate matrix

$$\mathbf{G} = \begin{pmatrix} \lambda_{11}\mathbf{I}_N + \mathbf{G}_1 & \lambda_{12}\mathbf{I}_N & \cdots & \lambda_{1R}\mathbf{I}_N \\ \lambda_{21}\mathbf{I}_N & \lambda_{22}\mathbf{I}_N + \mathbf{G}_2 & \cdots & \lambda_{2R}\mathbf{I}_N \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{R1}\mathbf{I}_N & \lambda_{R2}\mathbf{I}_N & \cdots & \lambda_{RR}\mathbf{I}_N + \mathbf{G}_R \end{pmatrix}, \tag{1}$$

where $\mathbf{I}_N$ is the $N \times N$ identity matrix. Note that the block matrix $\mathbf{G}$ has a special structure that the off-diagonal blocks are scalar matrices.

Theorem 1: Define two functions: $\chi : \mathcal{S}_Y \to \mathcal{S}$ by $\chi(n) = x_k$ and $\pi : \mathcal{S}_Y \to \mathcal{S}_Z$ by $\pi(n) = i$, where $k \in \{1, \cdots, N\}$ and $i \in \mathcal{S}_Z$ are the unique integers such that $n = (i-1)N + k$. Then the expectation of a (possibly path-dependent) functional ψ contingent on the regime-switching CTMC (S, Z) is equal to the expectation of a (possibly path-dependent) functional $\psi_{\chi,\pi}$ contingent on the one-dimensional CTMC Y, i.e.,

$$\mathbb{E}[\psi(S,Z)|S_0 = x_k, Z_0 = i] = \mathbb{E}[\psi_{\chi,\pi}(Y)|Y_0 = (i-1)N+k], \tag{2}$$

where the functional $\psi_{\chi,\pi}(\cdot)$ is defined as $\psi_{\chi,\pi}(Y) := \psi(\chi(Y), \pi(Y))$ and the stochastic processes $\chi(Y)$ and $\pi(Y)$ are defined as $\chi(Y) = \{(\chi(Y))_t := \chi(Y_t) : t \geq 0\}$ and $\pi(Y) = \{(\pi(Y))_t := \pi(Y_t) : t \geq 0\}$.

III. Pricing of Asian Options

In this section we exemplify Step 3 by considering the pricing of Asian call options, which depends on the distributions of time integrals of the underlying asset price process over certain time intervals. The price of an Asian call option with the maturity T and the strike price K at time 0 under the regime-switching CTMC (S, Z) is given by

$$V_a(T,K) = e^{-rT}\mathbb{E}\left[\left(\frac{1}{T}\int_0^T S_t dt - K\right)^+\right].$$

Theorem 2: (Laplace Transforms of Asian Option Prices under the Regime-Switching CTMC) Suppose that $S_0 = x_k$ and $Z_0 = i$. For any complex θ such that $\Re(\theta) > 0$, the Laplace transform of the Asian option price $V_a(T, K)$ with respect to K is given by

$$\int_0^{+\infty} e^{-\theta K} V_a(t,K) dK = e^{-rT}\left(\frac{x_k(e^{rT}-1)}{\theta r T} - \frac{1}{\theta^2} + \frac{\mathbf{e}'_{i,x_k} \cdot \mathrm{Exp}(T\mathbf{G} - \theta\mathbf{D}) \cdot \mathbf{1}}{\theta^2}\right),$$

where $\mathbf{D} = (d_{ij})_{N \cdot R \times N \cdot R}$ is a diagonal matrix with $d_{jj} = \chi(j)$ for $j = 1, \cdots, N \cdot R$.

The Asian option prices can be computed by inverting the closed-form Laplace transforms in Proposition 2 via Laplace inversion algorithms. Some numerical results and the comparison with those obtained via the Monte Carlo simulation method are given in Table I. It turns out that our analytical approximation method is accurate and fast.

TABLE I
THE PRICES OF ASIAN OPTIONS UNDER A REGIME-SWITCHING CEV MODEL WITH TWO REGIMES

Strikes	Simulation	Std.Err.	RS-CTMC	Abs.Err.	CPU time (secs)
90	12.367	0.009	12.380	0.013	1.5
100	5.198	0.007	5.210	0.012	1.4
110	1.482	0.004	1.491	0.009	1.4

Notes. The prices of Asian call options under a regime-switching CEV model with two regimes obtained via our analytical approximation method with $N = 100$ and the Monte Carlo simulation method with the sample size 1,000,000 and the step size for time discretization 10^{-4}. The parameters are: $r = 0.05$, $\sigma_1 = 0.15$, $\beta_1 = -0.1$, $\sigma_2 = 0.25$, $\beta_2 = -0.2$, $S_0 = 100$, $Z_0 = 1$ and $T = 1$.

IV. Conclusion

This paper provides a unified analytical approximation framework for accurate and efficient computations under general regime-switching Markov models. This unified computational framework makes it possible to study many issues in different areas when regime changes are taken into account and this is demonstrated by the example of Asian option pricing. Details about the theories and more applications can be found in the complete version of this paper.

References

[1] N. P. Bollen, "Valuing options in regime-switching models," *The Journal of Derivatives*, vol. 6, no. 1, pp. 38–49, 1998.

[2] X. Guo, "An explicit solution to an optimal stopping problem with regime switching," *Journal of Applied Probability*, vol. 38, no. 2, pp. 464–481, 2001.

[3] R. J. Elliott, L. Chan, and T. K. Siu, "Option pricing and Esscher transform under regime switching," *Annals of Finance*, vol. 1, pp. 423–432, oct 2005.

[4] D. D. Yao, Q. Zhang, and X. Zhou, "A regime-switching model for European options," in *Stochastic Processes, Optimization, and Control Theory: Applications in Financial Engineering, Queueing Networks, and Manufacturing Systems*, ch. 14, pp. 281–300, Springer US, 2006.

[5] A. Jobert and L. C. G. Rogers, "Option pricing with markov-modulated dynamics," *SIAM Journal on Control and Optimization*, vol. 44, no. 6, pp. 2063–2078, 2006.

[6] P. Boyle and T. Draviam, "Pricing exotic options under regime switching," *Insurance: Mathematics and Economics*, vol. 40, no. 2, pp. 267–282, 2007.

[7] M. Konikov and D. B. Madan, "Option pricing using variance gamma Markov chains," *Review of Derivatives Research*, vol. 5, pp. 81–115, 2002.

[8] R. J. Elliott and C.-J. U. Osakwe, "Option pricing for pure jump processes with Markov switching compensators," *Finance and Stochastics*, vol. 10, no. 2, pp. 250–275, 2006.

[9] A. Mijatović and M. Pistorius, "Continuously monitored barrier options under Markov processes," *Mathematical Finance*, vol. 23, no. 1, pp. 1–38, 2013.

Special Sessions in memory of Prof. Mark H.A. Davis

A weak law of large numbers for dependent random variables

Walter Schachermayer

Every sequence $f_1, f_2, \cdots$ of random variables with $\lim_{M\to\infty}\big(M\ \sup_{k\in}\mathbb{P}(|f_k|>M)\big)=0$ contains a subsequence $f_{k_1}, f_{k_2}, \cdots$ that satisfies, along with all its subsequences, the weak law of large numbers: $\lim_{N\to\infty}\big((1/N)\sum_{n=1}^{N} f_{k_n}-D_N\big)=0$, in probability. Here D_N is a "corrector" random variable with values in $[-N,N]$, for each $N\in$; these correctors are all equal to zero if, in addition, $\liminf_{k\to\infty}\mathbb{E}\big(f_k^2\,\mathbf{1}_{\{|f_k|\le M\}}\big)=0$ holds for every $M\in(0,\infty)$. Joint work with Ioannis Karatzas.

Achieving Mean-Variance Efficiency by Continuous-Time Reinforcement Learning

We conduct extensive empirical analysis to evaluate the performance of the recently developed continuous-time actor-critic reinforcement learning algorithms by Jia and Zhou (2021) in asset allocation. We propose an efficient implementation of the algorithms in a dynamic mean-variance portfolio selection setting, and compare it with other widely used strategies, including the 1/N portfolio strategy, buy-and-hold the market portfolio, and many variants to the static mean-variance models. To make the algorithm more stable and practical, we adopt an offline pre-training stage, use off-policy learning for online decision making, and incorporate constraints in terms of leverage and rebalancing frequency. Using S&P 500 data from Jan 2000 to Dec 2019, we demonstrate that our online algorithm with the aforementioned modifications outperforms the others, in terms of annualized return, Sharpe ratio, Sortino ratio, Calmar ratio and recovery time from maximum drawdown, for various periods under consideration, including the period of the 2008 financial crisis. In addition, the gross returns of the unconstrained portfolios can reach the pre-specified target returns on average and remain competitive under all criteria.

Adversarial Reinforcement Learning: A Duality-Based Approach for Optimal Control

Nan Chen Mengzhou Liu * H. Mete Soner †

The deep neural network approximation, combined with Monte Carlo simulation, has demonstrated a great potential in solving high-dimensional stochastic control problems (e.g., Han and E (2016) and Han et al. (2018)). However, Reppen and Soner (2021) find that this deep learning approximation (DLA) based method faces the classical bias-variance trade-off. Especially, when the training data is not sufficient relative to the complexity of the networks, deep neural networks networks tend to circumvent the essential restriction of the adaptedness in learning decisions.

In this paper, we propose an adversarial deep reinforcement learning algorithm (ADRL) to address this overlearning issue. It is based on the literature of information-relaxation duality (see, e.g., Rogers (2007), Brown et al. (2010), Brown and Haugh (2017), and Chen et al. (2020)). By relaxing the adapted requirement on control policies and incorporating a Lagrangian martingale penalty into the objective function, we reformulate the stochastic control problem as a min-max game between the decision maker and the adversary. ADRL aims to learn the optimal policy and penalty iteratively.

To fix the idea, let us consider the following reinforcement learning (RL) problem

$$V^*(\boldsymbol{s}) = \max_{\pi\in\Pi} V^\pi(\boldsymbol{s}) = \max_{\pi\in\Pi} \mathbb{E}\left[\sum_{t=0}^{T-1} r_t(\boldsymbol{s}_t, \pi_t) + R(\boldsymbol{s}_T)\Big|\boldsymbol{s}_0 = \boldsymbol{s}\right], \tag{1}$$

where Π be the collection of all non-anticipative policies and the evolution of the state $\{\boldsymbol{s}_t : t = 0, \cdots, T\}$ follows the dynamic

$$\boldsymbol{s}_{t+1} = f_t(\boldsymbol{s}_t, \boldsymbol{a}_t, \xi_{t+1}).$$

The DLA approach initiated by Han and E (2016) specifies a multilayer neural network $\varpi_t(\cdot, \theta_t)$ to approximate the optimal policy in episode t. Once we have the training dataset $\mathcal{L}^n := \{\xi^{(1)}, \cdots, \xi^{(n)}\}$ with $\xi^{(l)} = (\xi_1^{(l)}, \cdots, \xi_T^{(l)})$ for $1 \le l \le n$, Han and E (2016) propose to use stochastic gradient descent (SGD) method with backpropagation to solve the following optimization problem for training the optimal θ:

$$v^{\mathrm{DLA},k}(\mathcal{L}^n) := \max_\theta \frac{1}{n}\sum_{l=1}^{n}\left[\sum_{t=0}^{T-1} r_t(\boldsymbol{s}_t^{(l)}, \varpi_t(\boldsymbol{s}_t^{(l)}, \theta_t)) + R(\boldsymbol{s}_T^{(l)})\right], \tag{2}$$

where $\{\boldsymbol{s}_t^{(l)}, t = 0, \cdots, T\}$ is the state sequence obtained by applying the approximate policy $\varpi_t(\cdot, \theta_t)$ to the sample $\xi^{(l)}$. Note that the solution to the right-hand side of (2) depends on the entire training dataset ξ, violating the non-anticipative requirement for the policy. Let k represent the number of neurons in the hidden layers of the neural network. Reppen and Soner (2021) show that

$$\lim_{k\to+\infty} v^{\mathrm{DLA},k}(\mathcal{L}^n) = \frac{1}{n}\sum_{i=1}^{n}\max_{\boldsymbol{a}_t}\left[\sum_{t=0}^{T-1} r(\boldsymbol{s}_t^{(l)}(\boldsymbol{a}), \boldsymbol{a}_t) + R(\boldsymbol{s}_T^{(l)}(\boldsymbol{a}))\right]. \tag{3}$$

Thus, sufficiently deep or wide neural networks overlearn and are thus able to overperform on the training data.

Our idea is to apply information relaxation (IR) based duality theory to deal with the above anticipativeness issue. The theory starts with the construction of Lagrangian martingale penalties. In particular, given any sequence of functions $W = (W_0, W_1, \cdots, W_T)$, construct

$$z_t^W(\boldsymbol{a}, \xi) := W_t(\boldsymbol{s}_t) - \mathbb{E}[W_t(f_{t-1}(\boldsymbol{s}_{t-1}, \boldsymbol{a}_{t-1}, \xi_t))|\boldsymbol{s}_{t-1}] \tag{4}$$

Let

$$z^W(\boldsymbol{a}, \xi) := \sum_{t=0}^{T} z_t^W(\boldsymbol{a}, \xi).$$

*Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N. T., Hong Kong. Email: nchen@se.cuhk.edu.hk, mzliu@se.cuhk.edu.hk

†Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ, 08450, USA. Email: soner@princeton.edu

It is easy to see that z^W is the sum of martingale differences. In addition, we have the following strong duality relation (see Rogers (2007) and Brown et al. (2010)):

$$V^*(s) = \min_W \mathbb{E}\left[\max_{\boldsymbol{a}=\{\boldsymbol{a}_t\}_{t=0}^T \in \mathbb{A}^T} \left(\sum_{t=0}^{T-1} r(\boldsymbol{s}_t, \boldsymbol{a}_t) + R(\boldsymbol{s}_T) - z^W(\boldsymbol{a}, \xi)\right) \Big| \boldsymbol{s}_0 = \boldsymbol{s}\right]. \tag{5}$$

Note that the right-hand side of (3) is a special case of the above dual formulation with $W = 0$.

The min-max formulation of (5) enables a game theoretic interpretation for the strong duality. Imagine that there exists an adversary whose objective is to minimize the expected (penalized) reward of the agent by selecting a proper W, while the agent tries to maximize the rewards using anticipative policies. Inspired by this intuition, we propose the following adversarial deep reinforcement learning algorithm to solve (5) to determine V^*.

Use a set of networks $\{\varrho_t(\cdot, \phi_t) : t \in \mathbb{T}\}$, parametrized by $\phi_t \in \Phi_t$, to approximate the adversary's optimal choice on W. Divide $\mathcal{L}^n$ into m disjoint batches: $\mathcal{L}^n = \mathcal{K}_1 \cup \cdots \cup \mathcal{K}_m$. We use one batch in each iteration. Suppose that the current value of the parameters used in the adversary network is estimated to be $\{\phi_t^{b-1}\}$ after $b-1$ iterations. A generic iteration of ADRL algorithm consists of two stages:

- **Action Stage**. Construct the penalty $z^{b-1}(\boldsymbol{a}, \xi)$ from $\varrho_t(\cdot, \phi_t^{b-1})$. Solve the optimization

$$\boldsymbol{a}^b(\xi) = (\boldsymbol{a}_0^b(\xi), \cdots, \boldsymbol{a}_{T-1}^b(\xi)) := \arg\max_{\boldsymbol{a}=\{\boldsymbol{a}_t\}_{t=0}^{T-1}} \left[\sum_{t=0}^{T-1} r(\boldsymbol{s}_t, \boldsymbol{a}_t) + R(\boldsymbol{s}_T) - z^{b-1}(\boldsymbol{a}, \xi)\right]. \tag{6}$$

- **Adversarial Stage**. Identify an unbiased estimator Y^b to the following gradient such that

$$\mathbb{E}[Y^b] = \nabla_\phi \mathbb{E}\left[\max_{\boldsymbol{a}=\{\boldsymbol{a}_t\}_{t=0}^{T-1}} \left(\sum_{t=0}^{T-1} r(\boldsymbol{s}_t, \boldsymbol{a}_t) + R(\boldsymbol{s}_T) - z^{b-1}(\boldsymbol{a}, \xi)\right) \Big| \boldsymbol{s}_0 = \boldsymbol{s}\right].$$

  Update $\phi^b = \phi^{b-1} - \alpha_b Y^b$, where α_b is the learning rate.

After the training process, we can obtain $v^{\text{ADRL},k}(\mathcal{L}^n)$ as the in-sample estimate of V^*.

We show in the paper

Theorem 1

$$\mathbb{E}\left[\lim_{k\to\infty} v^{DLA,k}(\mathcal{L}^n)\right] \geq \mathbb{E}\left[\lim_{k\to\infty} v^{ADRL,k}(\mathcal{L}^n)\right] \geq V^*.$$

Theorem 1 clearly shows the advantage of our ADRL algorithm in alleviating the overlearning issue encountered by DLA: it will generate an estimate closer to the true value V^* when we apply overcomplicated networks to $\mathcal{L}^n$ to learn. In addition,

Theorem 2

$$\lim_{k\to\infty} \lim_{n\to\infty} v^{ADRL,k}(\mathcal{L}^n) = V^*(\boldsymbol{s}).$$

Several numerical experiments in the paper illustrate the efficiency of ADRL.

Algorithmic Collusion in Financial Markets

Álvaro Cartea[a,b], Patrick Chang[a], José Penalva[c,a]

[a] *Oxford-Man Institute of Quantitative Finance*
[b] *Mathematical Institute, University of Oxford*
[c] *Universidad Carlos III*

The rise in popularity and adaptation of reinforcement learning brings upon the need to study multi-agent systems with care. Particularly when artificial intelligence is used in systems that can affect social welfare such as the use of AI to price goods. Chen et al. (2016) find that more than a third of the vendors have automated their pricing for the 1,600 best-selling items on Amazon, while Assad et al. (2020) find that the use of algorithmic pricing in duopoly markets lead to an increase in prices in Germany's retail gasoline market. Therefore, it is important for policy makers and regulators to study and prevent tacit collusion between algorithms to maintain social welfare.

At first glance, it is not obvious how to cast the interaction between independent stochastic learning algorithms as a system of deterministic interactions. However, with the stochastic approximation techniques in Norman (1972); Benaïm (1999); Kushner and Yin (2003) one can model the dynamics of the learning algorithms as a system of deterministic ordinary differential equations (ODEs). Börgers and Sarin (1997) formalises the connection between Cross' learning model and the replicator dynamics from evolutionary game theory, while Beggs (2005) and Hopkins and Posch (2005) prove that Erev and Roth's model converges to the adjusted or Maynard Smith version of the replicator dynamics.

Since then, many authors have made connections between different learning algorithms and deterministic ODEs. Tuyls et al. (2003) and Sato and Crutchfield (2003) derive the dynamics for Q-learning with a softmax activation for mapping the Q-values to the policy, Gomes and Kowalczyk (2009) derive the dynamics for Q-learning with the ϵ-greedy policy, Kaisers and Tuyls (2010) introduce the frequency adjusted Q-learning to address the discrepancy between Q-learning and the replicator-mutation dynamics, and Kasbekar and Proutiere (2010) derive the dynamics for the EXP3 algorithm. However, these latter connections are derived in an ad-hoc manner, which is the source of discrepancies between Q-learning and the replicator-mutation dynamics.

In this paper, we apply a simplified framework based on the work from Benaïm (1999) to formalise the aforementioned connections and we show how to compute the correct expectations to obtain the appropriate ODEs. Specifically, we demonstrate the stochastic approximation framework on different stochastic processes that arise from the algorithms such as the policies or the Q-values from Q-learning. In the process, we formally complete the proofs for Q-learning, frequency adjusted Q-learning and the EXP3 algorithm. In addition, we prove that synchronous Q-learning also converges to the replicator-mutation dynamics and demonstrate an alternative approach compared to Börgers and Sarin (1997) and Beggs (2005) who apply results from Norman (1972).

The ODEs bring the additional benefit of interpretability, along with the rich literature surrounding evolutionary game dynamics which we use to understand the long-term behaviour of the ODEs in pricing games. Our approach complements the existing literature on tacit collusion between independent learning algorithms in pricing games. Specifically, Calvano et al. (2020, 2021) finds that Q-learning leads to supracompetitive prices in both a Bertrand and Cournot oligopoly. We argue that this is a result of Q-learning getting "stuck" as we show that Q-learning can converge to outcomes that are not a Nash equilibrium. Specifically, we show that in a symmetric pricing game, Q-learning can converge to asymmetric actions.

We also explain the mechanics behind the results from Asker et al. (2022) where they find that synchronous Q-learning leads to competitive prices while standard Q-learning leads to supracompetitive prices. This is because synchronous Q-learning converges to the replicator-mutation dynamics, which we show converges to a Nash equilibrium in a pricing game with sufficiently large exploitation.

Equipped with a deterministic approach to model the interactions between independent learning algorithms, we use the ODEs to study tacit collusion in pricing games through a stylised market-making game. We study the impact of tick size, which relates to the number of actions available to the algorithm, and its role in preventing tacit collusion. We use the size of the basins of attraction to find a probability estimate for the likelihood of a competitive outcome, and we use the convergent policies to estimate the expected excess profits. We find that the replicator dynamics all converge to a pure Nash equilibrium, but most do not converge to the competitive Nash equilibrium. However, for the replicator dynamics, a smaller tick size leads to lower excess profits because the possible set of Nash equilibria becomes closer to the Bertrand Nash equilibria. Finally, we find that Q-learning does not necessarily converge to Nash equilibria and leads to higher excess profits compared to algorithms that behave like replicator dynamics.

References

Asker, J., Fershtman, C., Pakes, A., 2022. Artificial Intelligence and Pricing: The Impact of Algorithm Design. AEA Papers and Proceedings Forthcoming.

Assad, S., Clark, R., Ershov, D., Xu, L., 2020. Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market. Available at SSRN 3682021 CESifo Working Paper No. 8521.

Beggs, A., 2005. On the convergence of reinforcement learning. Journal of Economic Theory 122, 1–36. doi:https://doi.org/10.1016/j.jet.2004.03.008.

Benaïm, M., 1999. Dynamics of stochastic approximation algorithms, in: Azéma, J., Émery, M., Ledoux, M., Yor, M. (Eds.), Séminaire de Probabilités XXXIII, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 1–68.

Börgers, T., Sarin, R., 1997. Learning Through Reinforcement and Replicator Dynamics. Journal of Economic Theory 77, 1–14.

Calvano, E., Calzolari, G., Denicolo, V., Pastorello, S., 2020. Artificial intelligence, algorithmic pricing, and collusion. American Economic Review 110, 3267–97.

Calvano, E., Calzolari, G., Denicoló, V., Pastorello, S., 2021. Algorithmic collusion with imperfect monitoring. International Journal of Industrial Organization 79, 102712.

Chen, L., Mislove, A., Wilson, C., 2016. An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace, in: Proceedings of the 25th International Conference on World Wide Web, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE. p. 1339–1349.

Gomes, E.R., Kowalczyk, R., 2009. Dynamic Analysis of Multiagent Q-Learning with ϵ-Greedy Exploration, in: Proceedings of the 26th Annual International Conference on Machine Learning, Association for Computing Machinery, New York, NY, USA. p. 369–376.

Hopkins, E., Posch, M., 2005. Attainability of boundary points under reinforcement learning. Games and Economic Behavior 53, 110–125. doi:https://doi.org/10.1016/j.geb.2004.08.002.

Kaisers, M., Tuyls, K., 2010. Frequency Adjusted Multi-Agent Q-Learning, in: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC. p. 309–316.

Kasbekar, G., Proutiere, A., 2010. Opportunistic medium access in multi-channel wireless systems: A learning approach, in: 2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 1288–1294.

Kushner, H.J., Yin, G., 2003. Stochastic approximation and recursive algorithms and applications. 2 ed., Springer.

Norman, M.F., 1972. Markov Processes and Learning Models. Mathematics in Science and Engineering : a series of monographs and textbooks, Academic Press.

Sato, Y., Crutchfield, J.P., 2003. Coupled replicator equations for the dynamics of learning in multiagent systems. Phys. Rev. E 67, 015206.

Tuyls, K., Verbeeck, K., Lenaerts, T., 2003. A Selection-Mutation Model for Q-Learning in Multi-Agent Systems, in: Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems, Association for Computing Machinery, New York, NY, USA. p. 693–700.

Algorithmic market making in foreign exchange cash markets with hedging and market impact

Alexander Barzykin* Philippe Bergault† Olivier Guéant‡

Key words: Market making, Algorithmic trading, Stochastic optimal control, Viscosity solutions

In financial markets, liquidity has traditionally been provided by a specific category of agents who, on a continuous and regular basis, set prices at which they agree to buy or sell assets and securities. These agents, called market makers or dealers, play a key role in the price formation process in all markets but their exact role and behavior depend on the considered asset class.

In most order-driven markets, such as stock markets, traditional exchanges have converted from open outcry communications between human traders to electronic platforms organized around all-to-all limit order books, and computers now handle almost all market activity. Official market makers and traditional market making companies still make money by providing liquidity to the market but they are now, somehow, in competition with all market participants who can post liquidity-providing orders. In quote-driven markets, electronification has also been one of the major upheavals of the last decade, with important consequences for market makers / dealers. In foreign exchange (FX) cash markets for instance, dealers set up their own private electronic platforms enabling clients to directly send them requests for quotes (RFQ) and to be connected to their stream (RFS). Many dealer-to-dealer and all-to-all plaforms have also emerged, therefore blurring the frontier between OTC and organized markets.

Alongside the multifaceted electronification of trading means, most human market makers have been replaced by market making algorithms. This evolution has naturally gone along with the development of many market making models in the academic literature. In 2008, largely inspired by a paper from the 1980s by Ho and Stoll [7], Avellaneda and Stoikov [3] proposed a stochastic optimal control model to determine the optimal bid and ask quotes that a single-asset risk-averse market maker should set. The authors paved the way to a new literature on market making that complements the contributions of the economic literature on the topic. The resulting new models can be divided into two groups: those adapted to the problem of a market maker in a limit order book and those adapted to OTC markets.

In most academic models adapted to OTC markets, market makers are pure internalizers: they buy assets at the bid price they quote and sell them at the ask price they quote – ideally earning the difference between these two prices. Of course, market makers seldom buy and sell simultaneously: they carry inventory and bear price risk. The problem faced by market makers in these models is already a subtle dynamic optimization problem in which market makers must mitigate the risk associated with price changes by skewing their quotes as a function of their inventory. In practice, however, market makers in FX cash markets have an additional way to manage their inventory risk since they can partially or completely hedge it by trading in liquidity pools on the Dealer-to-Dealer (D2D) segment of the market and in a variety of all-to-all platforms.

The co-existence of requests for quotes / requests for stream and liquidity pools has seldom been studied in the academic literature on optimal market making in OTC markets (the only instance we found beyond our paper is the very recent paper [4] that proposes a reinforcement learning approach). The trade-off faced by

*HSBC, 8 Canada Square, Canary Wharf, London E14 5HQ, United Kingdom, alexander.barzykin@hsbc.com

†Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, philippe.bergault@etu.univ-paris1.fr

‡Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, olivier.gueant@univ-paris1.fr

dealers in FX markets between internalization and externalization is nevertheless discussed in the literature. It is discussed by Butz and Oomen in [5] on the basis of queuing theory, though not with optimized quotes. It is also abundantly discussed on empirical grounds in the recent BIS Triennial Survey that concludes on the growing prevalence of internalization (see [8]). A wide spectrum of behaviors is documented, from pure externalization to large ratios of internalization. It is noteworthy that even though internalization ratios for top trading centers exceed 80%, hedging through externalization remains an essential component of risk management for any dealer. In particular, this features should be included in FX optimal market making models.

The main goal of our paper is to build an optimal market making strategy that includes the possibility for the market maker to hedge by buying and selling (in continuous time) in a liquidity pool, in order to better mitigate inventory risk. By trading in a liquidity pool, the market maker adds certainty to inventory risk management but that comes with execution costs and market impact, in part due to revealing of trading intentions to a wider audience. Our setup is inspired by Almgren-Chriss-like models of optimal execution (see Almgren et al. [2, 1], and Guéant [6] for a general presentation). More precisely, compared to existing optimal OTC market making models, ours includes a new form of control – in addition to the bid and ask quotes – that represents the trading rate of the market maker in a liquidity pool and features (i) execution costs to proxy transaction costs and nonlinear liquidity costs, and (ii) permanent market impact (assumed to be linear in the trading rate).

We present our market making model involving a currency pair for which the market maker has a classical market making quoting activity together with the possibility to hedge risk by trading in a liquidity pool. We then introduce the stochastic optimal control problem of the market maker. We characterize the associated value function as the unique continuous viscosity solution of a Partial Integro-Differential Equation of the Hamilton-Jacobi type. We illustrate our model numerically and discuss the results. In particular, we highlight the existence of a threshold of inventory under which it is not optimal for the market maker to trade in the liquidity pool.

References

[1] Robert Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied mathematical finance*, 10(1):1–18, 2003.

[2] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.

[3] Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.

[4] Alexey Bakshaev. Market-making with reinforcement-learning (sac). *arXiv preprint arXiv:2008.12275*, 2020.

[5] Maximilian Butz and Roel Oomen. Internalisation by electronic fx spot dealers. *Quantitative Finance*, 19(1):35–56, 2019.

[6] Olivier Guéant. *The Financial Mathematics of Market Liquidity: From optimal execution to market making*, volume 33. CRC Press, 2016.

[7] Thomas Ho and Hans Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial economics*, 9(1):47–73, 1981.

[8] Andreas Schrimpf and Vladyslav Sushko. Fx trade execution: complex and highly fragmented. *BIS Quarterly Review, December*, 2019.

An Enhanced Indexation Investment Strategy based on graph reinforcement learning

Zhao ziran[1], Cao hongduo[1]

(1. Department of Finance and Investment, Sun Yat-Sen Business School, Sun Yat-Sen University, Guangzhou 510275, China)

Abstract: Portfolio theory was first put forward by Markowitz in 1968. In recent years, there have been literatures that apply machine learning to financial market transactions and even portfolio management. Essentially, portfolio management is a decision-making process that allocates funds to different financial assets effectively and reasonably. The goal is to reduce the risk as much as possible while achieving the expected return at the same time. In this paper, a novel enhanced indexation model which aims to trade stocks that make up the index is proposed by combining graph neural network with deep reinforcement learning. Specifically, following Jiang, Xu, Liang (2017), we set the output of the deep neural network model as the weight of each asset in the portfolio. Taking both the characteristics of the stock itself and interaction between stocks into account , the action in our reinforcement learning framework is to determine the optimal investment weight of each component stock in the index according to a policy function that can be approximated by deep neural network model. Furthermore, we introduce a time series method called variational mode decomposition(VMD) to extract stock's own characteristics. Experiments on Chinese A stock market compare return of the adjusted portfolio with benchmark index and that of portfolio formed by other deep learning model, showing that our method outperforms the index and other similar trading strategies.

Key words: Portfolio; enhanced indexation model; graph; deep reinforcement learning

In the capital asset pricing model and many of its descendants, agents earn *risk premia* for holding assets whose payoffs are uncertain. A number of influential empirical studies (Amihud and Mendelson, 1986; Brennan and Subrahmanyam, 1996; Pástor and Stambaugh, 2003) suggest that – in reality – agents are also compensated for holding securities that are difficult to trade. To wit, if one sorts assets based on various measures of liquidity, then the returns earned by portfolios composed of the more liquid assets are systematically lower than for portfolios of less liquid ones.

The theoretical underpinnings of these *liqiuidity premia* have been studied in an active literature going back to the seminal work of Constantinides (1986). His work (and many more recent studies) takes a partial equilibrium approach, where the asset price dynamics are specified exogenously. Liquidity premia then refer to the amount by which the risky assets' expected returns have to be increased compared to a hypothetical frictionless version of the asset, in order to offset the utility losses caused by trading costs.

Another strand of research derives equilibrium asset prices with transaction costs endogenously by matching supply and demand (Vayanos, 1998; Vayanos and Vila, 1999; Lo et al., 2004; Sannikov and Skrzypacz, 2016; Weston, 2018; Isaenko, 2020; Herdegen et al., 2020). This allows to study how changes in liquidity feed back into asset prices, e.g., how liquidity premia are affected by the reduction of the fees charged by an exchange or the introduction of a financial transaction tax. Our paper also follows this equilibrium approach.

Yet, equilibrium models with a single illiquid risky asset still cannot say anything about the *cross section* of liquidity premia across a spectrum of different assets – that is, the subject of the empirical work of Amihud and Mendelson (1986); Brennan and Subrahmanyam (1996); Pástor and Stambaugh (2003). Equilibrium models with several illiquid assets lead to formidable computational challenges. These difficulties are of course only exacerbated if one moves beyond two (representative) agents that are typically assumed for tractability. To wit, even the most tractable models with linear state dynamics and quadratic transaction costs (Garleanu and Pedersen, 2016; Isaenko, 2020; Sannikov and Skrzypacz, 2016) then lead to coupled systems of matrix Riccati equations. Whereas general well-posedness results are available for partial equilibrium models (Ankirchner and Kruse, 2015; Bank et al., 2017; Garleanu and Pedersen, 2016; Kohlmann and Tang, 2002) or for models with exogenously given constant volatility (Bouchard et al., 2018), the only known results concerning the existence of equilibrium prices require the restrictive assumption that the agents' preferences are sufficiently similar (Herdegen et al., 2020), even in the case of only a single illiquid asset and just two agents.

In the present study, we establish the existence of equilibrium prices for an arbitrary number of illiquid risky assets that are traded by an arbitrary number of agents. These agents have mean-variance preferences as in Garleanu and Pedersen (2013, 2016) and trade to share the risk inherent in the fluctuations of their endowment streams, subject to a deadweight quadratic transaction cost as in Almgren and Chriss (2001); Garleanu and Pedersen (2013, 2016). For assets that pay exogenous liquidating dividends at a finite terminal time, the "Radner equilibrium" where the agents act as price takers then can be characterized by a fully-coupled system of forward-backward stochastic differential equations (FBSDEs). If the terminal dividends and the volatilities of the agents' endowment streams are linear in the driving Brownian motions, then this FBSDE system can be reduced to a fully-coupled system of matrix-valued ordinary differential equations of Riccati form.

For the simplest case of a single risky asset traded by two agents, existence for this system has been established using Picard iteration by Herdegen et al. (2020). However, even in this low-dimensional setting, establishing the convergence of the iteration scheme requires the restrictive assumption that the agents' risk aversions are sufficiently similar. In this paper, we show that this assumption is superfluous, in that the matrix Riccati system has a unique global solution even for an arbitrary number of agents and risky assets.

In order to facilitate the calibration of the model to time-series data, we complement this main result with rigorous asymptotic expansions. In the practically relevant limiting regime of small transaction costs, this leads to explicit formulas for the impact of illiquidity on price levels, volatilities, and the cross section of liquidity premia that are earned by assets with different trading costs.

To bring these theoretical results to life, we test them using an empirical case study following Acharya and Pedersen (2005). To wit, we sort the large-cap stocks in the S&P index by Amihud's "ILLIQ" measure for liquidity (Amihud, 2002), leading to three risky portfolios with high, medium, and low liquidity. In the frictionless version of our model, equilibrium returns solely compensate for risk and turn out to be very similar for all three portfolios. Using our asymptotic expansions, the calibration of the frictional version of the model to time series of prices *and* trading volumes is still feasible. When trading costs are taken into account, the equilibrium returns of the high-liquidity portfolio are indeed decreased in line with the data, whereas the returns of the low liquidity portfolio are increased. However, to match the magnitude of the liquidity premia observed empirically in our model, the risk aversion coefficients of the agents need to be rather heterogenous. In line with the partial equilibrium literature, this suggests that additional features such as market closure (Dai et al., 2016), unobservable regime shifts (Chen et al., 2020), or state-dependent trading costs (Acharya and Pedersen, 2005; Lynch and Tan, 2011) also play an important role in this context.

From a methodological perspective, this paper studies a competitive equilibrium model with an arbitrary number of agents and risky assets whose expected returns and volatilities are determined endogenously in equilibrium. This model necessarily leads to a nonstandard matrix-valued Riccati system. As the latter jointly determines optimal strategies and market-clearing prices, it is *not* associated with any concave optimization problem, so that wellposedness results from the literature do not apply. We therefore rigorously establish the existence and uniqueness of the system via a taylor-made comparison argument applied to the operator norm of the solution. As a by-product, this existence proof also allows us to establish explicit leading-order approximations of the equilibrium in the practically relevant limiting regime of small transaction costs.

An Equilibrium Model of Imperfect Hedging with Transaction Costs

January 24, 2020

Abstract

There has been a long standing debate on how transactions taxes, or more generally how transaction costs, affect financial markets. To address this question theoretically requires an equilibrium model where markets clear, however, this is technically challenging and, thus, the impact of transaction costs on asset pricing in equilibrium is rarely studied. In a framework with proportional transaction costs, we propose a simple equilibrium model which is analytically tractable.

In our model, two agents have linear exposures to non-traded risks and trade a zero net supply security. The zero net supply security has exponential payoffs in a Brownian motion which is correlated with the agents' endowments. This can be thought of as a derivative security with convex payoffs. Equilibrium is described by the risk exposure (volatility) and the risk premium (drift) of the derivative security. In equilibrium these adjust so that the sum of buy boundaries and sell boundaries for the different agents sum to zero. One important feature of our assumption about the payoffs is that while buy and hold has finite utility, sell and hold leads to negative infinite utility.

In equilibrium, high risk exposure leads to smaller costs of setting up and unwinding the initial and terminal hedge positions, whereas the highest total rebalancing costs are attained at an intermediate level of risk exposure when trading is the most frequent. For any positive transaction cost there always exist no trade equilibria where risk exposure is so small that the position required to hedge would be large and the transaction costs outweigh any benefit of hedging or earning a risk premium. There are also equilibria

with trade. When agents have different risk aversion, the less risk averse agent shorts the security and earns a risk premium. Curiously, when agents have the same risk aversion, the equilibrium is a no trade equilibrium. The reason is the asymmetry in the payoffs to the long and short positions. In equilibrium both agents cannot be long and there is no trade.

To assess the liquidity premium, we examine the equilibrium Sharpe ratio of the traded security and compare it to that in an equilibrium with no transaction cost. When the agents have very different (similar) risk aversion, equilibrium liquidity premia and equilibrium risk exposure tend to be low (high), and initial and terminal transaction costs tend to be high (low). In particular, with similar risk aversion, the liquidity premium can be quite large, on the order of 20 basis points for a 1 percent transaction cost. The intuition is that when risk aversions are similar, the investors have less incentive to trade, so higher liquidity premium is endogenized in equilibrium as compensations to induce trade from the less risk averse investor. Notably, the total rebalancing costs are bell-shaped with respect to the difference in risk aversion. It is because when the difference in risk aversion is very low, the portfolio holding in risky asset is close the zero, hence rebalanced number of shares tends to be low. One the other hand, when the difference in risk aversion is high, the no trade zone is very wide, leading to no need of rebalancing.

We also study the expected trading profits for a market maker as a function of the level of transaction costs. The interesting trade off is a small transaction cost leads to a lot of trade but each trade is not very profitable while a large transaction cost leads to fewer trades but each trade is more profitable. Our analysis suggests an interesting equilibrium effect. In equilibrium with lower transaction costs the risk exposure tends to be lower, positions tend to be larger, and rebalancing costs tend to be lower. With larger transaction costs, the risk exposure tends to be larger, positions tend to be smaller, and rebalancing costs tend to be higher. Expected profits for the market maker tend to be higher when rebalancing costs are relatively high; the profit maximizing transaction cost tends to be surprisingly high.

An exit contract optimization problem

Xihao He * Xiaolu Tan † Jun Zou ‡

January 31, 2022

Abstract

We study an exit contract design problem, where one provides a universal exit contract to multiple heterogeneous agents, with which each agent chooses an optimal (exit) stopping time. The problem consists in optimizing the universal exit contract w.r.t. some criterion depending on the contract as well as the agents' exit times. Under a technical monotonicity condition, and by using Bank-El Karoui's representation of stochastic processes, we are able to transform the initial contract optimization problem into an optimal control problem. The latter is also equivalent to an optimal multiple stopping problem and the existence of the optimal contract is proved. We next show that the problem in the continuous-time setting can be approximated by a sequence of discrete-time ones, which would induce a natural numerical approximation method. We finally discuss the optimaization problem over the class of all Markovian and/or continuous exit contracts.

Key words: Optimal stopping, stochastic optimal control, Bank-El Karoui's representation, contract theory.

MSC2010 subject classification: 60G40, 60G07, 49M25

*Department of Mathematics, The Chinese University of Hong Kong. xihaohe@math.cuhk.edu.hk.
†Department of Mathematics, The Chinese University of Hong Kong. xiaolu.tan@cuhk.edu.hk.
‡Department of Mathematics, The Chinese University of Hong Kong. zou@math.cuhk.edu.hk.

An Investment Theory with Lags and Adjustment Costs *

Wei Jiang[†] Shuaijie Qian[‡] Jialu Shen[§]

Abstract

We propose a stochastic control model to study corporate investment with generalized investment frictions, including investment lags and various of adjustment costs. We find that the dominance of the "good news principle" or "bad news principle" is determined by the joint effect of investment lags and adjustment costs, reconciling the results in Bernanke (1983) and Bar-Ilan and Strange (1996). Meanwhile, we resolve disputes between the net present value rule and the real option method of making investment decisions, and we find that the accuracy of the NPV rule depends on both investment lags and the opportunity cost of adjustment. Moreover, we calibrate our model with aggregated firm data and show that the co-existence of investment lags and the opportunity cost of adjustment is the key to explaining the correlation between investment and lagged profit.

JEL Classification: D21, E22, E32.
Key words: Bad News Principle, Investment Lags, Adjustment Costs, NPV Method, Real Option.

*We are grateful to Shan Huang, who provided excellent research assistance.
[†]Hong Kong University of Science and Technology. E-mail: weijiang@ust.hk
[‡]National University of Singapore. E-mail: e0046854@u.nus.edu
[§]University of Missouri. E-mail: jshen@missouri.edu

1 Introduction

How shall firms make the optimal investment decisions? Existing literature provides contradictory suggestions. Bernanke (1983) proposes "bad news principle" and finds that a firm's investment decision is only affected by bad news anticipated in the coming periods. He suggests that firms facing various adjustment costs shall be cautious about investing given the volatile return of investment and shall optimally delay investment decisions to avoid possible negative shocks in the future. However, considering investment lags, Bar-Ilan and Strange (1996a) show that uncertainty could have a positive effect on investment and come up with a "good news principle". They recommend firms to hasten investment under increasing uncertainty. Regardless of the discrepancy in results, they both consider optimal investment decisions as a real option (RO) problem and capture the uncertainty in capital returns.

Alternatively, a simple net present value (NPV) rule has been widely adopted by managers, although it assumes that the investment is either reversible or irreversible. Firms make investment decisions purely depending on the difference between the present value of cash inflows and the present value of cash outflows over a period of time. Majd and Pindyck (1987) show that the NPV rule can cause gross error compared with the RO approach. On the contrary, Milne and Whally (2000) advocate the NPV rule being an appropriate guide to investment timing decisions when long investment lags exist and suggest that the NPV rule and the RO approach deliver consistent decisions.

One possible reason why these papers deliver conflicting results is that they only focus on certain investment frictions. Under a more general model setting, which principle will be dominant and what is the true relationship between the NPV rule and the RO approach? To answer these two questions, this paper develops a dynamic stochastic control model in a general setting and makes four specific contributions. First, our model considers various investment frictions, including investment lags, the opportunity cost of adjustment, transaction costs, and costly reversibility. These frictions affect the optimal investment decisions from various perspectives: Investment lags can weaken or reverse the effect of price uncertainty. Specifically, when there are lags, investing inertia can be overridden, and an increase in uncertainty may hasten the decision to invest. Ignoring investment lags is inconsistent with an important fact that investments take time to implement; Opportunity cost of adjustment captures the need for plant restructuring, worker retraining, and organizational restructuring during periods of intensive investment; Transaction cost makes sure that the equilibrium capital stock can not be reached instantaneously; Costly reversibility indicates that firms face higher costs in cutting than in expanding capital stocks. Irreversible investment and costlessly reversible investment are considered as two extreme cases of costly reversibility. Because of this feature, the optimal investment decision is a two-trigger policy and inaction is preferred when the marginal profit of capital is strictly between the two trigger values. These particular features allow the investigation not only on the optimal investment timing but also on the optimal capital sale. We solve the model with an explicit characterization and provide an asymptotic analysis for the optimal control policies.

Second, the model is able to reconcile the "bad news principle" and "good news principle" for the optimal investment timing. We find that investment lags and the opportunity cost of adjustment jointly determine which principle is controlling. For optimal investment timing decisions, without opportunity cost of adjustment, investment lags advance investment, i.e., "Good news principle" is dominant. With the presence of opportunity cost of adjustment, investment lags delay investment, which favours the "bad news principle". The rationale behind this is that the effects of uncertainty on the opportunity cost of waiting for investment depends on both investment lag and opportunity cost of adjustment. The opportunity cost of waiting equals to the foregone profit from newly formed capital, netting the forgone cost from the existing capital stock driven by the opportunity cost of adjustment. Without opportunity cost of adjustment, opportunity cost of waiting consists solely of the forgone profit, and increases in uncertainty, leading to the "good news principle". With opportunity cost of adjustment, both the forgone cost exists and may outweigh the forgone profit. Thus, increase in uncertainty can lead to a lower opportunity cost of waiting, which may suggests the "bad news principle".

Third, we resolve disputes between the NPV rule and RO method of making investment decisions, complementing the analysis of Milne and Whally (2000) in a more general setting. We provide the sufficient conditions under which firms prefer investing immediately instead of waiting. Under this circumstance, the NPV rule is asymptotically the same as the RO method regarding investment timing decisions. For other circumastances, the NPV rule provides consistent conclusions with the RO method with a simple adjustment in the presence of large investment lag. Additionally, the opportunity cost of adjustment has asymmetric effects on investment boundaries for NPV and RO methods when investment lag is large (Proposition 6).

Fourth, we calibrate the parameters of our model using the firm data from Compustat with the Simulated Method of Moments (SMM), and investigate the impact of investment lag and the opportunity cost of adjustment. We observe that the contemporaneous correlation between investment ratio and profit is 0.202, and the largest serial correlation between investment ratio and lagged profit, 0.239, occurs at the lag-1 serial correlation. This suggests that firms' investment depends on the past firm fundamentals, which implies investment lag exists. The co-existence of investment lags and the opportunity cost of adjustment can explain the correlation between investment and the firms' lagged profit, while ignoring opportunity costs of adjustment or investment lag mismatches these lagged correlations.

Lastly, we extend the model to incorporate the abandonment option. When there are investment projects under development, firms need to re-examine their optimal exit choice in terms of both capital sale of the existing capital and the abandonment of ongoing projects. We find the abandonment option advances investment decisions since the firm is able to take more risk in investment by holding an abandonment option. Moreover, the abandon barrier is increasing with respect to the investment lag.

The paper draws on several strands of the literature. It builds on the classical real option studies of corporate investment with different kinds of investment frictions. Specifically, among all frictions, adjustment costs, irreversibility and investment lags have attracted a large amount of interest because of their importance on firms' investment strategy. Bernanke (1983) first discusses irreversible decisions under uncertainty and proposes the "bad news principle" for optimal investment timing. McDonald and Siegel (1986) further examine the value of waiting if an investment decision is delayed. Since then, considerable literature has attempted to study optimal investment timing with various adjustment costs (e.g., Abel and Eberly (1994), Abel and Eberly (1996), Cooper and Haltiwanger (2006)).[1] Meanwhile, Bar-Ilan and Strange (1996a) study the effect of investment lag on optimal investment timing and proposed the "good news principle" to show that uncertainty may advance a firm's investment in the presence of lags, contrary to the previous findings of the "bad news principle". Bar-Ilan and Strange (1998) consider a two-stage sequential investment decision problem. Recently, Kalouptsidi (2014) extends the analysis to the bulk shipping industry with a rich dataset and investigates the optimal investment timing and amount under implementation lags.

This paper contributes to this strand of real option literature by developing a model with several new features. (1) We consider a firm similar to that in Abel and Eberly (1994) and study its optimal investment strategy, but with additional investment frictions, in-

[1]Abel and Eberly (1994) studied optimal investment timing with fixed adjustment costs and irreversibility; Abel and Eberly (1996) further extended this analysis to investment with costly reversibility; Cooper and Haltiwanger (2006) summarized this line of research by studying investment under most kinds of adjustment costs.

cluding non-convex adjustment cost (see, Abel and Eberly, 1994; Cooper and Haltiwanger (2006)), costly reversibility (see Abel and Eberly, 1996), and investment lags (see Bar-Ilan and Strange, 1996a; Bar-Ilan and Strange, 1998). (2) We not only examine how existing capital stock affects investment choices by using capital stock as a state variable (i.e., Bernanke, 1983), but also investigate how various investment choices alter the productivity of existing capital during investment lags, i.e., opportunity cost of adjustment (e.g., Cooper and Haltiwanger, 2006; Sakellaris, 2004). Existing studies usually neglect either the latter effect (see Abel and Eberly, 1994; Abel and Eberly, 1996) or both effects (see Bernanke, 1983; Bar-Ilan and Strange, 1996a; Bar-Ilan and Strange, 1996b; Bar-Ilan and Strange, 1998; Kalouptsidi, 2014). (3) Besides investment lags, our model also assumes that the optimal investment amount is decided only at the end of the construction period. This assumption helps us to capture the effect of cost overrun, which is usually observed during project investment (e.g., John and Saunders, 1983; Ganuza, 2007). In contrary to our model, most existing models only assume that the investment amount is decided at the beginning of an investment (i.e., Bar-Ilan and Strange, 1996a, Bar-Ilan and Strange, 1998, and Kalouptsidi, 2014).

In addition, this paper connects to the literature that focuses on the differences between the RO method and the NPV method (e.g., Majd and Pindyck, 1987 and Milne and Whally, 2000). Majd and Pindyck (1987) set a maximum rate of investment which leads to a lag in completion of the investment. They show that using the NPV rule to determine investment timing has gross errors compared with the RO approach. With the same framework, Milne and Whally (2000) improve the numerical method and show that actually, the NPV method can be a good approximate guide to investment. However, Milne and Whally (2000) only consider a few types of investment frictions. Therefore, the true relationship between the real option method and the NPV method in a more general investment environment remains unclear. Our paper fills this gap. We show under which circumstance the NPV method can be an approximation of the RO method and provide a simple way to adjust the NPV rule when two approaches do not provide consistent conclusions.

This paper also relates to the literature regarding the decision-making problems with implementation lags or delay. Subramanian and Jarrow (2001) study the problem of a trader who wants to liquidate a position with execution lags in an illiquid market. Dai, Huang and Keppo (2019) consider a bank model with delayed recapitalizations of its opaque bank assets. Bruder and Pham (2009) study dividend payments with both decision lags and execution delays over a finite horizon. Bayraktar and Egami (2010) provide a new mathematical characterization for impulse control problems for dividend payments with implementation lag. Yao et al. (2002), He et al. (2015), and Shi et al. (2015) apply the general stochastic control methods to the financial engineering area. Our paper differs from these papers because of our application to firm-level investment decisions.

The paper is organized as follows. The baseline model is given in Section 2, followed by the model solution in Section 3. The model is calibrated with data in Section 4, based on which a quantitative analysis is conducted in Section 5. In addition, Section 6 compares the RO method and the NPV rule, and Section 7 discusses the effects of abandonment option on investment decisions. Section 8 concludes the paper.

2 Baseline Model

We consider a firm that produces products at time t with capital stock K_t, and variable factors of productions. The firm sells all of its products, and the price of its output is determined by a production function that depends on the random variable X_t. Assume that the operating profit of the firm, i.e., revenue minus the cost of the variable factors of production, is given by

$$\Pi(K_t, X_t) = \frac{h}{1-\gamma} X_t^{\gamma} K_t^{1-\gamma}, \tag{1}$$

where $h > 0$ and $0 < \gamma < 1$. The specification in equation (1) can be derived for a firm with a Cobb-Douglas production function facing an iso-elastic demand curve. Such setting has been used, for example, by Bertola (1988) and Dixit (1991) in their analyses of irreversible investment, and by Abel and Eberly, 1996 in analyses of investment with costly reversibility .

The random variable X_t measures the productivity shock [2] and evolves according to a geometric Brownian motion:

$$\frac{dX_t}{X_t} = \mu dt + \sigma d\mathcal{Z}_t;$$

where $\mathcal{Z}_t$ is a standard Brownian motion. The parameters μ and $\sigma > 0$ are the mean and volatility of the productivity shock.

The firm faces two types of adjustments in its capital stocks: capital investment and capital sale. There are two types of frictions in capital adjustments: (1) implementation

[2] In general, random variable X_t can be viewed as a combination of both demand shock and productivity shock; see, e.g., Abel and Eberly (1997).

lags (i.e., investment takes time to implement) and (2) capital adjustment costs (i.e., both lumpy investment and capital sale suffer from adjustment costs).

We assume the lag in investment is Δ_I, implying that for any decision of investing at time t, there is a construction lag Δ_I such that the capital will be successfully installed at time $t + \Delta_I$. The optimal investment amount will also be determined at the time $t+\Delta_I$, which is consistent with budget overrun commonly observed in project investment literature (see, John and Saunders, 1983; Ganuza, 2007).

The firm's capital adjustment policy ω is a collection

$$\omega = \{L_t, (\tau_j^I, I_j); j = 1, \cdots, \}. \tag{2}$$

where L_t, τ_j^I and I_j are the cumulative capital sales, the starting time of the $j-$th investment process, and the $j-$th investment amount, respectively.

Following the strategy ω, the evolution of capital stock K_t at time t satisfies:

$$K_t = K_0 - \int_0^t \delta K_s ds + \sum_j I_j \mathbf{1}_{\{\tau_j^I + \Delta_I \leqslant t\}} - L_t, \tag{3}$$

where K_0 is the capital stock at time 0 and δ is the depreciation rate of capital stock. The second term on the right side is the cumulative capital depreciation. The third term shows the cumulative capital investment at time t. The last term denotes the cumulative capital sell at time t.

Furthermore, we assume two types of capital adjustment costs in our model: transaction cost and opportunity cost of adjustment. On the one hand, investment faces both proportional transaction cost, which is linearly proportional to the amount of investment, and fixed transaction cost.[3] For instance, the firm can purchase capital at a constant price $1 + b_I$ where $b_I \geqslant 0$ and it pays a fixed cost $c_I K$ at the end of the investment lag. The fixed transaction cost allows investment to be lumpy form. To capture capital irreversibility, we also assume capital sale incurs transaction costs (b_L), that is, a firm sells capital at a constant price $1 - b_L$ where $1 \geqslant b_L \geqslant 0$. Moreover, during the investment period, the firm's operating profit also suffers from an opportunity cost of adjustment, which is $\lambda\Pi(K_t, X_t), \lambda < 1$. This type of cost implies that firm's productivity falls by a factor of λ during the capital investment period. Studies by Power (1998) and Sakellaris

[3]Generally, as discussed in Cooper and Haltiwanger, 2006, adjustment cost could be a general function of I_t, K_t, X_t, which may consists of different types of adjustment costs, such as convex and non-convex cost of adjustment. For simplicity, we only consider the non-convex cost transaction cost, including both proportional and fixed transaction costs.

6

(2004) provide empirical evidence that firm's productivity is lower during periods of large investment. For simplicity, we assume the external financial market is frictionless and therefore capital adjustment costs can be deducted directly from shareholder's present value.

Let $V(K, X)$ denote the shareholder's value. Shareholder chooses investment timing and amount (τ^I, I), capital sale strategy L to maximize the expected present value of its cash flow discounted by a constant rate r over an infinite horizon:

$$\begin{aligned} \sup_{\omega=\{L_t,(\tau_j^I,I_j);j=1,\cdots,\}} \mathbb{E}\Big[& \textstyle\int_0^\infty e^{-rt}(1-\lambda 1_{\{t\in[\tau_j^I,\tau_j^I+\Delta_I]\}})\Pi(K_t,X_t)dt \\ & -\textstyle\sum_j e^{-r(\tau_j^I+\Delta_I)}((1+b_I)I_{\tau_j^I+\Delta_I}+c_I K_{\tau_j^I+\Delta_I}) + \int_0^\infty e^{-r\tau_j^L}(1-b_L)dL_t\Big]. \end{aligned} \tag{4}$$

The expectation takes risk into account (i.e., under the risk-neutral measure $\mathbb{Q}$). The cash flow at time t in the objective function (4) equals the operating profit, $\Pi(K_t, X_t)$, minus the following three terms: (1) opportunity cost of adjustment $\lambda\Pi(K_t, X_t)$ during the investment period, (2) the cost of purchasing capital at a constant unit price $(1 + b_I)$ and (3) a fixed transaction cost $c_I K_t$, plus the proceeds received from selling capital at a constant unit price $(1 - b_L)$. To ensure well-posedness, we assume the discount rate r and productivity growth rate μ satisfies

$$\mu < r < 1. \tag{5}$$

This condition ensures that the firm' value function under optimal strategy is finite.[4]

Our model encompasses several existing models as special cases. If we assume zero transaction cost in capital sale, i.e., $b_L = 0$, then we have the case of costless reversible investment. If we assume the sale price of capital to be zero, i.e., $b_L = 1$, then we have the case of irreversible investment. If we assume no investment lag, i.e., $\Delta_I = 0$, then our model degenerates to corporate investment with adjustment costs studied by Abel and Eberly (1996). If both capital adjustment cost and investment lag are zero, our model pins down to the neoclassical capital investment without adjustment costs as in Jorgenson (1963).

[4] See, e.g., arguments in Abel and Eberly (1996)

3 Theoretical Analysis

To begin, we show that the firm value decreases in the investment lag. Intuitively, the presence of investment lag restrains firm's flexibility of investment and therefore lowers the firm's value.

Proposition 1. *The firm value $V(K,X)$ is non-increasing in the construction lag Δ_I.*

In what follows we assume that (5) holds. Problem (4) is a mixture of singular control and impulse control problem (see, e.g., Bruder and Pham (2009)). The corresponding value function $V(K,X)$ is associated with the following Hamilton-Jacobi-Bellman (HJB) equaiton:

$$\max\{\mathcal{L}V, \mathcal{M}V - V, \mathcal{P}V - (1-b_L)\} = 0, \tag{6}$$

in the region where $K > 0, X > 0$, and the operators $\mathcal{L}$, $\mathcal{M}$, and $\mathcal{P}$ are defined as

$$\begin{aligned}
\mathcal{L}V(K,X) &= \frac{1}{2}\sigma^2 X^2 V_{XX} + \mu X V_X - \delta K V_K - rV + \frac{h}{1-\gamma} X^\gamma K^{1-\gamma}, \\
\mathcal{M}V(K,X) &= E\Big[\int_0^{\Delta_I} e^{-rt}(1-\lambda) \times \Pi(K_t, X_t)dt \\
&\quad + \sup_{I \geqslant 0} e^{-r\Delta_I}(V(K_{\Delta_I} + I, X_{\Delta_I}) - (1+b_I)I - c_I K_{\Delta_I})\Big], \\
\mathcal{P}V(K,X) &= V_K(K,X).
\end{aligned}$$

The operator $\mathcal{L}V(K,X)$ corresponds to the case without dis/investment decision at time t. $\mathcal{M}V(K,X)$ corresponds to the expected firm value if the firm chooses to invest. The investment amount is lumpy and determined at the end of investment lag. $\mathcal{P}V(K,X)$ corresponds to the expected firm value if the firm chooses to disinvest. Different from the investment amount, the capital sale is determined instantaneously such that marginal value of capital equals the cost of capital sale.

Using the homogeneity property of the value function, we can reduce the dimension of the problem by the following transformation:

$$V(K_t, X_t) = K_t v(y_t), \tag{7}$$

where $y_t = X_t/K_t$ is the ratio of productivity to capital stock and $v(y_t)$ is the firm's market-to-book value.[5] Substituting into equation (6), we obtain

$$\max\{\mathcal{L}_1 v, \mathcal{M}_1 v - v, \mathcal{P}_1 v - (1 - b_L)\} = 0, \tag{8}$$

in the region where $y > 0$, and the operators $\mathcal{L}_1$, $\mathcal{M}_1$, and $\mathcal{P}_1$ are defined as

$$\begin{aligned}
\mathcal{L}_1 v(y) &= \frac{1}{2}\sigma^2 y^2 v''(y) + (\mu+\delta) y v'(y) - (r+\delta) v(y) + \frac{h}{1-\gamma} y^\gamma, \\
\mathcal{M}_1 v(y) &= E\Big[\int_0^{\Delta_I} e^{-(r+\delta)t}(1-\lambda)\frac{h}{1-\gamma} y_t^\gamma dt \\
&\quad + \sup_i \left\{ e^{-(r+\delta)\Delta_I} \left((1+i)\, v\left(\frac{y_{\Delta_I}}{1+i}\right) - (1+b_I)i - c_I \right) \right\} \Big], \\
\mathcal{P}_1 v(x) &= v(y) - y v'(y),
\end{aligned}$$

where $i = \frac{I}{K_{\Delta_I}}$ is the investment to capital ratio chosen at the end of the investment period.

In the absence of investment lag, i.e., $\Delta_I = 0$, there exists an investment barrier y_I and a capital sale barrier y_L. The firm is optimal to sell capital if the productivity to capital ratio, $\frac{X}{K}$, falls below y_L, and invests if its productivity to capital ratio exceeds y_I. In addition, there exists a target level s_I such that, whenever firm decides to invest, it invests a lump-sum immediately in raising its productivity-to-capital ratio to this target level s_I. Then the firm's value and optimal policies associated with the endogenous boundaries y_I, y_L and target level s_I can be characterized explicitly; see Proposition 2. Unlike Abel and Eberly (1996) where investment amount is infinitesemial, our model without investment lag allows lumpy investment and can be viewed as a an extension of Abel and Eberly (1996) with a fixed investment cost.

Proposition 2. *When there is no investment lag, i.e.,* $\Delta_I = 0$*, the firm value is* $V(K,X) = Kv(y)$ *with* $v(y) = Ay^\gamma + C_1 y^{1-\alpha_+} + C_2 y^{1-\alpha_-}$ *for* $y \in [y_L, y_I]$*, where the parameters* $A, \alpha_\pm, C_1, C_2$ *are given by* (A.1)-(A.4) *in Appendix A. The the boundaries* y_I *and* y_L *and the investment target* s_I *are determined by* (A.5)- (A.8)*. Moreover, the optimal strategy* $\omega^* = \{L_t^*, (\tau_j^{I*}, I_j^*)\}$ *is given by*

$$L_t^* = \int_0^t 1_{\{y_\xi = y_L\}} dL_\xi,$$

[5]A firm's market-to-book ratio is related to Tobin's q. In our model, if we neglect the fixed investment cost, i.e., $c = 0$, then Tobin's q equals to the market-to-book ratio divided by the capital purchasing price $1 + b_I$.

$$\begin{aligned}\tau_j^{I*} &= \inf\{t > \tau_{j-1}^{I*} : X_t/K_t \geqslant y_I\}, \ \forall j = 1, 2, \cdots, \\ I_j^* &= \max\{X_{\tau_j^{I*}}/s_I - K_{\tau_j^{I*}-}, 0\}, \ \forall j = 1, 2, \cdots,\end{aligned}$$

with $\tau_0^{I*} := -\Delta_I$. *The investment-to-capital ratio is given by* $i_j^* = I_j^*/K_{\tau_j^{I*}} = \max\{1 - s_I/y_I, 0\}$.

In the presence of investment lag, i.e., $\Delta_I > 0$, we have an semi-explicit solution to (8). Similarly, there exists endogenous boundaries y_I, y_L and target level s_I that determine the firm's optimal policies. Unlike the model without investment lags, when decides to investment at time t, the firm with investment lag will invest a stocastic lump-sum at time $t + \Delta_I$ such that the firm will raise its productivity-to-capital ratio to an optimal level at the end of the investment period. Our model captures the effect of cost overrun commonly observed in project investment (see, i.e., John and Saunders (1983) and Ganuza (2007)) by allowing the optimal investment amount to be decided at the end of construction period.

Proposition 3. *In the presence of investment lag, i.e.,* $\Delta_I > 0$, *the firm value is* $V(K, X) = Kv(y)$ *where*

$$v(y) = \begin{cases} f_1(y) & y > y_I \\ H(y) & y_L \leqslant y \leqslant y_I \\ f_2(y) & y < y_L, \end{cases}$$

where $H(\cdot), f_1(\cdot), f_2(\cdot)$ *are given by* (A.9)-(A.11) *in Appendix A, and the boundaries* y_I *and* y_L *and target level* s_I *are determined by* (A.12)-(A.15) *in Appendix A. Moreover, the optimal strategy* $\omega^* = \{L_t^*, (\tau_j^{I*}, I_j^*)\}$ *is given by*

$$\begin{aligned}L_t^* &= \int_0^t 1_{\{y_\xi = y_L\}} dL_\xi, \\ \tau_j^{I*} &= \inf\{t > \tau_{j-1}^{I*} + \Delta_I : X_t/K_t \leqslant y_I\}, \ \forall j = 1, 2, \cdots, \\ I_j^* &= \max\{X_{\tau_j^{I*}+\Delta_I-}/s_I - K_{\tau_j^{I*}+\Delta_I-}, 0\}, \ \forall j = 1, 2, \cdots,\end{aligned}$$

with $\tau_0^{I*} := -\Delta_I$. *The investment-to-capital ratio is given by*

$$i_j^* = I_j^*/K_{\tau_j^{I*}+\Delta_I} = 1 - s_I/\max\{y_{\tau_j^{I*}+\Delta_I-}, s_I\}.$$

As showed by Proposition 3, the optimal investment amount determined at the end of the investment period has two implications. If the firm's fundamental at the end of the j-th investment period, i.e., at time $\tau_j^{I*} + \Delta_I-$, exceeds the target level s_I, i.e., $y_{\tau_j^{I*}+\Delta_I-} > s_I$, then the firm will invest a lump-sum such that $y_{\tau_j^{I*}+\Delta_I} = s_I$. If the firm's fundamental falls below s_I, i.e., $y_{\tau_j^{I*}+\Delta_I-} \leqslant s_I$ then the firm's optimal capital adjustment amount is zero.

To better understand the firm's investment decision and how it is related to the marginal value of capital, we provide the following characterization.

Proposition 4. *Given the investment boundary y_I, capital sale boundary y_L, and investment target s_I defined in Proposition 3, the firm's marginal value of capital satisfies*

$$V_K = v(y) - yv'(y) = A(1-\gamma)y^\gamma + C_1|\alpha_+|y^{1-\alpha_+} - C_2|\alpha_-|y^{1-\alpha_-}, y \in [y_L, y_I]. \qquad (9)$$

where A, C_1, C_2 and $\alpha_\pm$ are given by (A.1)-(A.4).

The marginal value of capital equals the present value of expected marginal profit of capital under optimal behavior. Equation (9) suggests this expected present value is the sum of three components: 1) $A(1-\gamma)y^\gamma$ is the present value of marginal output of capital if the firm were prevented from purchasing or selling capital; 2) If the firm has the option of capital sale, the opportunity to reduce the future capital stock increases the expeted future marginal profit of capital and hence increases the present value of expected future marginal profit of capital by $C_1|\alpha_+|y^{1-\alpha_+}$; and 3) If the firm has the option to invest, this opportunity to increase the capital stock in future decreases expected future marginal profit of capital and decreases the present value of expected future maginal profit of capital by $C_2|\alpha_-|y^{1-\alpha_-}$. Note that the effects of the investment option on the marginal value of capital are negative, and the effects of the capital sale option are positive.

We then characterize the conditional mean of the optimal investment amount explicitly as follows:

Proposition 5. *The mean of the optimal investment amount at the beginning of the investment period is given as*

$$\mathbb{E}_{t=\tau_j^{I*}}\left[i_j^*\right] \quad = \quad \Phi(d_1) - s_I/y_I e^{-(\mu+\delta-\sigma^2)\Delta_I}\Phi(d_2) \qquad (10)$$

where $d_1 = \frac{\ln y/s_I + (\mu+\delta-1/2\sigma^2)\Delta_I}{\sigma\sqrt{\Delta_I}}, d_2 = \frac{\ln y/s_I + (\mu+\delta-3/2\sigma^2)\Delta_I}{\sigma\sqrt{\Delta_I}}$.

4 Calibration

In this section, we describe our dataset, introduce the method to construct aggregate data, and calibrate the parameters of our model with the simulated method of moments

11

(SMM) proposed by Lee and Ingram (1991).

We use the annual Compustat data from 1980 to 2020. We first include all firms whose financial fundamentals are reported in US dollar and can be fully observed. If end-of-period capital is missing, we drop the firms. Then we exclude firms with observations of negative profit, capital or investment. We focus our analysis on large firms, as large firms are usually considered as stable firms, which have less or no financing frictions. So we also drop those firms with sales no more than 1 million. Finally we end up with 277 firms.

We use the firm data related to the key variables in our model: investment, capital stock and profit. We define investment as capital expenditure, and estimate capital stock as total assets minus cash and short-term investments. Profit is the operating income before depreciation. All items are deflated by annual CPI, removing the effects of inflation.

4.1 Aggregate Investment

Note our model is built at the micro level. Because time-series investment data at the micro level is unstable, we want to use macro data that is aggregated from micro data for our calibration. The same approach has been employed by other studies, such as Kydland and Prescott (1982) and Zhou (2000).

Assume there are n firms in our data. We denote by $\widetilde{I}_t^{(k)}$, $\widetilde{K}_t^{(k)}$ and $\widetilde{\Pi}_t^{(k)}$ for firm k's investment, capital stock, and profit, respectivly, where $k = 1, 2, \cdots, n$. Based on the firm-level data, we can estimate aggregate investment as $\widetilde{I}_t = \sum_{k=1}^{n} \widetilde{I}_t^{(k)}$; aggregate profit as $\widetilde{\Pi}_t = \sum_{k=1}^{n} \widetilde{\Pi}_t^{(k)}$; aggregate capital stock $\widetilde{K}_t = \sum_{k=1}^{n} \widetilde{K}_t^{(k)}$. With these aggregate variables, we are able to calculate capital scaled investment and profit: $\widetilde{i}_t = \widetilde{I}_t / \widetilde{K}_t, \widetilde{\pi}_t = \widetilde{\Pi}_t / \widetilde{K}_t$.

In practice, the accounting principles (see, e.g., GAAP ARB 45) usually requires firms that engage in investment projects with long-term construction to comply with a percentage-to-completion method. That is, in the firm's finanncial statement, investment during the investment period is amortized as a percentage of an estimated total investment. Meanwhile, firm's capital stocks and profit are reported based on the amortized investment.

However, our model assumes that firms recognize the investment when the construction is completed and investment amount is decided at the end of each investment period. To fill the gap between the investment in our model and the amortized investment in financial statement, we introduce the following method to adjust the simulated variables from our model.

We denote by I_t, K_t, Π_t for investment, capital stock, profit at time t from our model.

Then we have $I_t = \sum_{j=1} I_j^* 1_{t=\tau_j^I+\Delta_I}$, where τ_j is the $j-$th optimal investment timing and I_j^* is the associated $j-$th optimal investment amount (see, i.e., Proposition 3). Then, the amortized investment and the associated capital stocks and profit functions are given by

$$\widehat{I}_t = g(\mathbb{E}_t i_j^* 1_{\{\tau_j^I \leqslant t \leqslant \tau_j^I + \Delta_I\}}); \tag{11}$$

$$\widehat{\Pi}_t = \left(1 - \lambda 1_{\{\tau_j \leqslant t \leqslant \tau_j + \Delta_I\}}\right) \Pi_t, \tag{12}$$

$$\widehat{K}_t = K_t + \widehat{I}_t dt. \tag{13}$$

The equation (11) provides the accounting technology $g(\cdot)$[6] for the amortized investment amount $\widehat{I}_t$, and implies that the firm's investment at time t depends on the firm's estimated total investment amount $\mathbb{E}_t i_j^* 1_{\{\tau_j^I \leqslant t \leqslant \tau_j^I + \Delta_I\}}$ if the firm is in the j-th investment period. Here we assume a linear amortization strategy, i.e., $g(\cdot) = \frac{1}{\Delta_I} \mathbb{E}_t i_j^* 1_{\{\tau_j^I \leqslant t \leqslant \tau_j^I + \Delta_I\}}$.[7] The equation (12) defines the adjusted profit $\widehat{\Pi}$ which captures the effect of opportunity cost of adjustment on the firm's profit. The equation (13) defines the adjusted capital stock to be the summation of the model simulated capital plus the amortized investment. Based on the adjusted investment, capital stock and profit from our model, we can calculate the scaled adjusted investment and profit from model as $\widehat{i}_t = \widehat{I}_t / \widehat{K}_t, \widehat{\pi}_t = \widehat{\Pi}_t / \widehat{K}_t$.

We estimate parameters by using the SMM approach. This approach requires selecting moments and then match these moments from aggregated data $\widetilde{i}_t, \widetilde{\pi}_t$ to those from adjusted model variables $\widehat{i}_t, \widehat{\pi}_t$. Theoretical and practical considerations suggest that the selected moments should be relevant in identifying the parameters and are also precisely estimated. We focus on the mean and serial correlation of scaled investment $\widetilde{i}_t$ and scaled profit $\widetilde{\pi}_t$, and lagged correlation between them.

More specifically, the first group of moments are $E(\widetilde{i}_t)$ and $E(\widetilde{\pi}_t)$, which are direct measurements of scaled investment and profit. The second group of moments we are interested in are the serial correlations in investment and profit, i.e., $corr(\widetilde{i}_t, \widetilde{i}_{t-1})$ and $corr(\widetilde{\pi}, \widetilde{\pi}_{t-1})$. It is well established in the literature that the serial correlations of investment and profit are sensitive to the structure of opportunity costs of adjustment. Given

[6] The accounting technology we use here should satisfy $\int_{\tau_j^I}^{\tau_j^I+\Delta_I} g(\mathbb{E}_t i_j^* 1_{\{\tau_j^I \leqslant t \leqslant \tau_j^I + \Delta_I\}}) = i_j^*$, implying that the summation of all reported investment during the investment period equals to the actual investment amount which is realized in the end of construction period. Since the investment lag is short as we estimated (less than 2 years), we neglect the time value in the accounting technology.

[7] For the ease of simulation, we further assume that the accountant can, approximately, have a perfect estimation of the total investment amount, i.e. $g(\cdot) = \frac{1}{\Delta_I} \mathbb{E}_t i_j^* 1_{\{\tau_j^I \leqslant t \leqslant \tau_j^I + \Delta_I\}} \approx \frac{1}{\Delta_I} i_j^* 1_{\{\tau_j^I \leqslant t \leqslant \tau_j^I + \Delta_I\}}$. This assumption can be relaxed and our results still hold. For instance, when we introduce some random pertubation in the accounting technology such that the accountant can not perfectly estimate the total investment amount, our calibration results do not change quite much.

the current investment and profit, a high opportunity cost lowers the likelihood to investment in the near future and increases the likelihood of a higher future profit, which hence changes the value of serial correlations. The third group of moments include the correlation between investment and profit, $corr(\widetilde{i}_t, \widetilde{\pi}_t)$, and the correlation between investment and lagged profit, $corr(\widetilde{i}_t, \widetilde{\pi}_{t-\ell}), \ell = 1, \cdots, n_0$, as this group reflects the covariance structure between investment and profits. This group of moment is sensitive to investment lag and in this sense satisfies the relevance criterion.

4.2 Calibrated Parameters

The model calibration relies on the model solution provided in Proposition 3. We first use our aggregate data $\widetilde{i}_t, \widetilde{\pi}_t$ to estimate the vector of moments $\widetilde{\Psi}$. Then we find the parameter vector $\widehat{\Theta}$ that minimizes the distance between the empirical moments $\widetilde{\Psi}$ and the simulated moments, which we denoted as $\widehat{\Psi}(\widehat{\Theta})$. Formally, we solve

$$L(\Theta) = \min_{\Theta}(\widetilde{\Psi} - \widehat{\Psi}(\Theta))'W(\widetilde{\Psi} - \widehat{\Psi}(\Theta)), \tag{14}$$

where the weighting matrix W is a diagonal matrix with diagonal element equal to the inverse of the variance of the targeted moments.

Based on the model solution given in Proposition 3, the function $\widehat{\Psi}(\widehat{\Theta})$ can be obtained analytically. It relies on the analytical model solution and the adjusted model variables $\widehat{i}_t, \widehat{\pi}_t$. We solve the minimization problem (14) using an annealing algorithm to reduce the risk of convergence to a local minimum.

To estimate the model, we include 8 moments in both $\widetilde{\Psi}$ and $\widehat{\Psi}$ vectors (see Table 2). We estimate the structural parameters $\Theta = (r, \gamma, \delta, \mu, \sigma, b_I, b_L, \Delta_I, c_I, \lambda)$ by solving the minimization problem (14). For the ease of calibration, we set $c_I = b_I = b_L$. To investigate the importance of investment lags and opportunitiy cost of adjustment, we consider three specifications. In the baseline model, we assume both investment lag Δ_I and opportunity cost of adjustment λ exist. In model 1, we assume $\Delta_I > 0, \lambda = 0$. In model 2, we assume $\Delta_I = \lambda = 0$. These are taken as leading specifications in the literature, and thus our estimation provides insights into which is more capable of capturing relevant features of the data.[8]

[8] Cooper and Haltiwanger (2006) develop a discrete-time model which is equivalent to our continuous-time model with $\Delta_I = 1$. They consider both $\lambda = 0$ and $\lambda > 0$ cases.

Table 1: **Summary of Parameters** The table summarized the symbols for the key parameters and the values calibrated from data for different models. The values in the "Baseline" column are calibrated from the baseline model. The values in the "Model 1" and "Model 2" are calibrated using model 1 ($\lambda = 0$) and model 2 ($\Delta_I = \lambda = 0$), respectively. The parameter values are annualized.

Variable	Symbol	Baseline	Model 1 ($\lambda = 0$)	Model 2 ($\Delta_I = \lambda = 0$)
Risk free rate	r	0.08	0.08	0.07
Mean of productivity shock	μ	0.07	0.07	0.06
Volatility of productivity shock	σ	0.31	0.25	0.35
Rate of depreciation	δ	0.01	0.02	0.01
Output elasticity of capital	γ	0.38	0.50	0.70
Investment lag	Δ_I	1.70	1.90	–
Opportunity cost of adjustment	λ	0.10	–	–
Proportional investment cost	b_I	0.08	0.10	0.25
Proportional capital sale cost	b_L	0.08	0.10	0.25
Fixed investment cost	c_I	0.08	0.10	0.25

Table 2: **Data and simulated moments** This table reports the data moments from the Compustat dataset and the simulated moments from baseline model, model 1, and model 2. The moments are annualized.

	Compustat	Baseline	Model 1($\lambda = 0$)	Model 2 ($\Delta_I = \lambda = 0$)
$E[i]$	0.068	0.053	0.063	0.058
Corr(i, i_{-1})	0.647	0.450	0.473	0.077
E[π]	0.156	0.123	0.159	0.244
Corr(π, π_{-1})	0.742	0.723	0.697	0.699
Corr(i, π)	0.202	0.142	0.261	0.546
Corr(i, π_{-1})	0.239	0.336	0.426	0.254
Corr(i, π_{-2})	0.202	0.238	0.257	0.178
Corr(i, π_{-3})	0.172	0.134	0.111	0.129
$L(\Theta)$		5.565	5.920	44.666

We report the estimated parameters in Table 1. Column 1 is for our baseline model, which includes both investment lag and opportunity costs of adjustment. The investment lag Δ_I is estimated to be 1.7, which suggests on average, firms takes 1.7 years to complete their investment projects. Opportunity cost λ is estimated to be 0.10, which indicates that once firms decide to invest, the profit will go down by 10%. When we consider no opportunity cost but still allow investment lag, as shown in model 1, we find investment lag increases from 1.7 to 1.9. In order to match the observed moments, the model 1 requires more time for the firm to finish the same projects.

Table 2 compares the data moments with the simulated moments from the baseline model and two alternatives. We find that overall the baseline model fits the moments the best among the three model specifications. We are particularly interested in the following moments: contemporaneous and lagged correlation between investment and profit.

From the Compustat data, we observe a key feature in the correlation between investment and profit or lagged profit: the largest value, 0.239, occurs at the lag-1 correlation (i.e., $corr(i_t, \pi_{t-1})$), followed by a slightly lower lag-2 correlation, 0.202. As we will show in Table 3, this feature suggests that, on average, firms take around one to two years to complete their investment projects. Our baseline model simulation captures this key feature and generates the largest value occuring at the lag-1 correlation, 0.336. This simulation result is also consistent with the estimated investment lag 1.7.

On the other hand, Model 1 estimates a longer investment lag, 1.9, and mismathes the contemporaneous and lagged correlation between investment and profit. Model 2 generates moments inconsistent with data and has the highest loss value among the three models. It shows that contemporaneous correlation is the highest among all reported correlations between investment and profit. The intuition is that, in the absence of both investment lag and opportunity cost of adjustment, the firm is willing to invest as soon as they observe higher productivity (profit), which suggests a high contemporaneous correlation.

To better understand the effect of different investment frictions on the firm's stationary moments, we report simulated moments for different values of Δ_I in Table 3 and those for different λ in Table 4. Table 3 shows that increase in Δ_I can change the largest value among all correlations between investment and profit. When investment lag is small $\Delta_I = 0.4$, the largest value is obtained by the contemporaneou correlation. As Δ_I increases to 1.2, the largest value is obtained by the lag-1 correlation. When Δ_I further increases to 2.4, the lag-2 correlation has the highest value. This property suggests that which lagged correlation achieves the highest value could corresponds to a certain level of investment lag. On the other hand, Table 4 shows that increase in λ (from 0 to 0.5) decreases the contemporaneous and lag-1 correlation between investment and profit (i.e., $corr(i, \pi)$ and $corr(i, \pi_{-1})$), and increases the lag-3 correlation (i.e., $corr(i, \pi_{-3})$). Overall, these simulation exercises show that the joint existence of investment lag and opportunity cost of adjusmtent can capture the key feature in the correlations between investment and profit or lagged profit.

Table 3: **The effect of investment lag Δ_I on simulated moments in the baseline model.** We employ the parameters in the baseline model and simulate moments for different values of Δ_I.

Δ_I	$E[i]$	$\mathrm{Corr}(i, i_{-1})$	$\mathrm{E}[\pi]$	$\mathrm{Corr}(\pi, \pi_{-1})$	$\mathrm{Corr}(i, \pi)$	$\mathrm{Corr}(i, \pi_{-1})$	$\mathrm{Corr}(i, \pi_{-2})$	$\mathrm{Corr}(i, \pi_{-3})$
0.4	0.069	0.096	0.127	0.662	0.556	0.186	0.105	0.057
0.8	0.064	0.186	0.125	0.811	0.413	0.262	0.151	0.081
1.2	0.058	0.294	0.124	0.817	0.254	0.321	0.193	0.105
1.6	0.053	0.419	0.123	0.737	0.159	0.336	0.229	0.129
2	0.051	0.525	0.122	0.697	0.099	0.321	0.264	0.152
2.4	0.048	0.591	0.121	0.699	0.050	0.284	0.288	0.176
2.8	0.047	0.635	0.120	0.715	0.010	0.244	0.295	0.204

Table 4: **The effect of opportunit cost of adjustment λ on simulated moments in the baseline model.** We employ the parameters in the baseline model and simulate moments for different values of λ.

λ	$E[i]$	$\mathrm{Corr}(i, i_{-1})$	$\mathrm{E}[\pi]$	$\mathrm{Corr}(\pi, \pi_{-1})$	$\mathrm{Corr}(i, \pi)$	$\mathrm{Corr}(i, \pi_{-1})$	$\mathrm{Corr}(i, \pi_{-2})$	$\mathrm{Corr}(i, \pi_{-3})$
0	0.050	0.454	0.122	0.724	0.246	0.367	0.215	0.103
0.1	0.053	0.450	0.123	0.723	0.142	0.336	0.238	0.134
0.2	0.055	0.447	0.123	0.721	0.053	0.304	0.251	0.155
0.3	0.058	0.444	0.123	0.718	-0.022	0.273	0.254	0.166
0.4	0.060	0.440	0.124	0.713	-0.086	0.247	0.257	0.175
0.5	0.062	0.438	0.124	0.707	-0.140	0.221	0.254	0.179

5 Quantitative Analysis

Section 3 provides semi-explicit solution for the optimization problem 4 and the associated HJB equation of (6). Since there is no full-explicit solution for the PDE systems, in this section, we employ numerical methods to investigate the effects of different types of frictions on the firm's investment and capital sale deisions. We use the calibrated parameters for the baseline model in Table 1 in the following quantitative analysis.

5.1 Market-to-book ratio and marginal value of capital

In Figure 1, we plot the firm's market-to-book ratio, $v(y) = V(K, X)/K$, and its marginal value of capital, $V_K = v(y) - yv'(y)$, against the firm's productivity-to-capital ratio $y = X/K$.

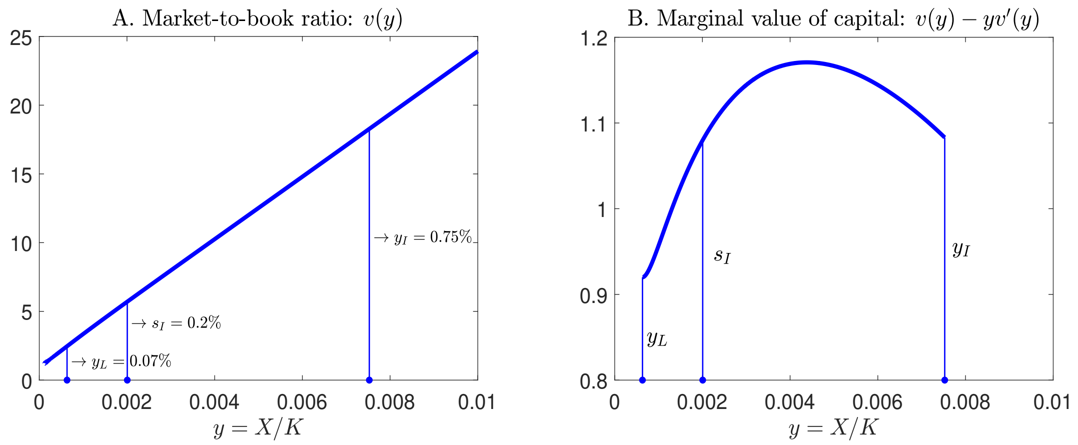

Figure 1: **Firm's market-to-book ratio and marginal value of capital.** . The left panel shows the market-to-book value defined as $V(K, X)/K = v(y)$, and the right panel shows the marginal value of capital defined as $V_K = v(y) - yv'(y)$. In both panels, the three vertical lines from left to right mark the capital sale boundary y_L, investment target s_I, and investment boundary y_I, respectively.

In panel A, we plot the firm's market-to-book ratio (solid line) and its decision boundaries (vertical lines). The vertical lines from left to right mark the capital sale boundary ($y_L = 0.07\%$), investment target ($s_I = 0.2\%$), and investment boundary ($y_I = 0.75\%$), respectively. When the firm's productivity-to-capital ratio is high (reaching $y_I = 0.75\%$), the firm will optimally enter into an investment period with an investment lag $\Delta_I = 1.7$ years. The high productivity-to-capital ratio might be driven by the high level of productivity X or the low level of capital stock K. Once investment decision is made, the firm will involve in buying equipments, building plants, hiring and training workers, etc, aim-ing to install some level of capital in the end of investment period. During the investment period, the firm's profit suffers a loss, $\lambda\Pi(K, X)$, its capital in place keeps depreciating, and the productivity-to-capital ratio y evolves stochastically with productivity shocks. In the end of the investment period, the investment amount will be finally determined and the newly invested capital will be in place (after paying both proportional and fixed investment costs) such that the firm's productivity-to-capital ratio y reaches $s_I = 0.2\%$. For instance, if the firm's y at the end of investment period is 0.5%, then the firm will invest a lumpy amount $i = 1 - 0.2\%/0.5\% = 0.6$. Standing at the beginning of the investment period, the investment amount is usually stochastic as it depends on the stochastic y at the end of the investment period and the endogenous investment target s_I.

The lumpy investment in our model is consistent with empirical observations (e.g., Cooper and Haltiwanger (2006)). The investment amount could be large since investment lag is long (1.7 years from our calibration) and investment frequency could be low. We will provide detailed discussion on investment amount in Section 5.3. In addition, our model is consistent with the commonly observed budget overrun in project investment literature (e.g., John and Saunders (1983) and Ganuza (2007)) where the capital will only be in place at the end of the investment period.

When the firm's productivity-to-capital ratio is low (reaching $y_L = 0.07\%$), the firm will optimally sell its capital such that the y will not drop below y_L. This is consistent with Abel and Eberly (1996) finding that firm sells capital when capital stock is too high or productivity is too low. When the firm's fundametal y stays in the region (y_L, y_I), the firm is in an inaction region where it takes no investment or capital sale actions, and solely relies on the capital in place to generate profits for shareholders.

Panel B plots the marginal value of capital $V_K = v(y) - yv'(y)$. As is shown, the marginal value of capital exhibits a hump shape. When the firm's productivity-to-capital ratio is low, the marginal value of capital increases against the productivity and is always higher than the marginal cost of capital sale, i.e., $1 - b_L$. When the firm's productivity is high, the marginal value of capital decreases against the productivity. This is closely related to Proposition 4 where we show the marginal value of capital consists of three parts: present value of marginal profit of capital without capital adjustment options plus the effects of investment option and the effects of capital sale option. When y is close to the capital sale boundary y_L, the effects of the capital sale option dominate and the marginal value of capital increases. When y is close to the investment boundary y_I, the effects of investment option dominate and the marginal value of capital decreases. Moreover, the optimal investment target level s_I is at such a level where the marginal value of capital equals to the marginal cost of capital, i.e., $1 + b_I = 1.08$. The existence of investment target level is due to the non-zero investment lag, and it suggests a lumpy investment amount determined at the end of an investment period.

In Figure 2, we compare the firm that follows our baseline model ($\Delta_I = 1.7, \lambda = 0.1$) with other two firms either without opportunity cost of capital adjustment ($\Delta_I = 1.7, \lambda = 0$) or without investment lag ($\Delta_I = \lambda = 0$). Except the specified parameters, the other two firms use the same parameters as the baseline in Table 1.

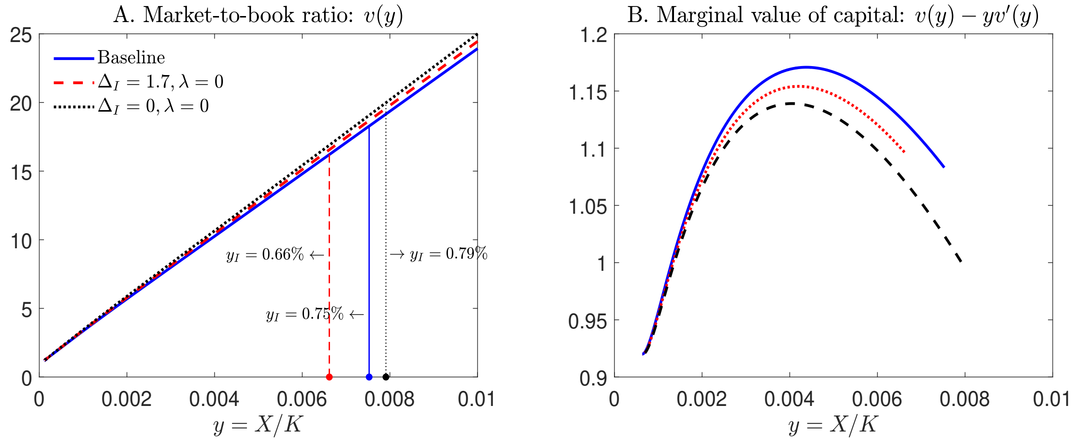

Figure 2: **Comparison of the baseline with the firm without opportunity cost of capital adjustment ($\Delta_I = 1.7, \lambda = 0$) and the firm without investment lag ($\Delta_I = 0, \lambda = 0$).** The left panel shows the market-to-book and the right panel shows the marginal value of capital. In the left panel, the vertical lines mark the investment boundary y_I for different firms.

In panel A, we plot three firms' market-to-book ratio and their investment boundaries. The firm with $\Delta_I = \lambda = 0$ has the highest market-to-book ratio since it has the least types of frictions among the three firms. It is closely related to the standard theory where investment is affected by uncertainty (e.g., Pindyck (1993) and Dixit (1991). This firm has an investment boundary $y_I = 0.79\%$ when only productivity shock plays a role. Once the investment lag is introduced in the firm with $\Delta_I = 1.7, \lambda = 0$, the firm's market-to-book ratio is lower than the previous one which has no investment lag. The firm's investment decision is hastened and the investment boundary decreases to be $y_I = 0.66\%$. This is because investment lag increases the firm's opportunity cost of waiting which depends on the growing productivity in future. The high opportunity cost of waiting makes the firm rush to invest in order to avoid facing high productivity while its capital stock is not in place. When we further introduce opportunity cost of adjustment $\lambda = 0.1$ as in the baseline model, it shows the lowest market-to-book ratio among these three models. The baseline firm's investment will be delayed and the investment boundary increases to be $y_I = 0.75\%$ since its opportunity cost of waiting, although increases by investment lag, is offset by the opportunity cost of adjustment.

Panel B shows the marginal value of capital for the three models. For a given level of y, the baseline model has the highest marginal value of capital. This is because the baseline has more investment frictions than the other two firms. The more investment frictions are, the more difficult capital formulation is and the higher the marginal value of capital is. This is in line with Proposition 4 and suggests that the effects of the option to invest on the marginal value of capital (which is negative) will be lower if there are more investment frictions.

5.2 Good / Bad news principles for investment

Without investment lag, the conventional bad news principle suggests that the increase in productivity volatility (uncertainty) delays investment decisions (Bernanke, 1983; Dixit (1991)). In the presence of investment lag, Bar-Ilan and Strange (1996a) argue that the conventional result is reversed to be a good news principle which suggests an increase in uncertainty may hasten the decision to invest. In this part, we examine numerically the effect of undertainty on investment decisions in our framework. We show the existence of opportunity cost of adjustment may reverse the good news principle and lead to the bad news principle.

In Figure 3, we plot the effect of uncertainty on investment boundaries for different levels of investment frictions. Panel A plots the cases with large investment lag ($\Delta_I = 4$). The black dashed line, which represents investment boundaries for a firm without investment lag and opportunity cost of adjustment ($\Delta_I = \lambda = 0$), increases in productivity volatility σ. It implies that more uncertainty delays investment, which is consistent with the bad news principle. When only the investment lag is introduced (red dotted line with $\lambda = 0$), it shows that the increase in undertainty not necessary delays investment. For the low levels of volatility, increasing undertainty increases the boundary as usual, but the effect is much smaller than the case without lag (black dased line). For higher levels of volatility, the effect is reversed where increasing uncertainty decreases the investment boundary, suggesting that more uncertainty can hasten investment. The implication from the red dotted line is consistent with the good news principle suggested by Bar-Ilan and Strange (1996a).

Our main result is that the opportunity cost of investment can dramatically change the effect of productivity uncertainty, in the presence of investment lag, on investment.

As shown by the blue solid line in panel A, when a 10% opportunity cost of adjustment is introduced ($\lambda = 0.1$), increasing uncertainty increases the boundary with a higher effect than without λ. For higher levels of volatility (from 30% to 40%), the increasing uncertainty still increases the boundary. Therefore, the presence of λ reverses the good news principle to be the bad news principle for volatilities varying from 10% to 40%. With a even higher λ (blue dash-doted line with $\lambda = 0.2$), the investment boundaries are even higher than the black dashed line for different levels of volatility. This implies that higher opportunity cost not only delays investment decisions but also enhances the effect that higher uncertainty delays investment decisions.

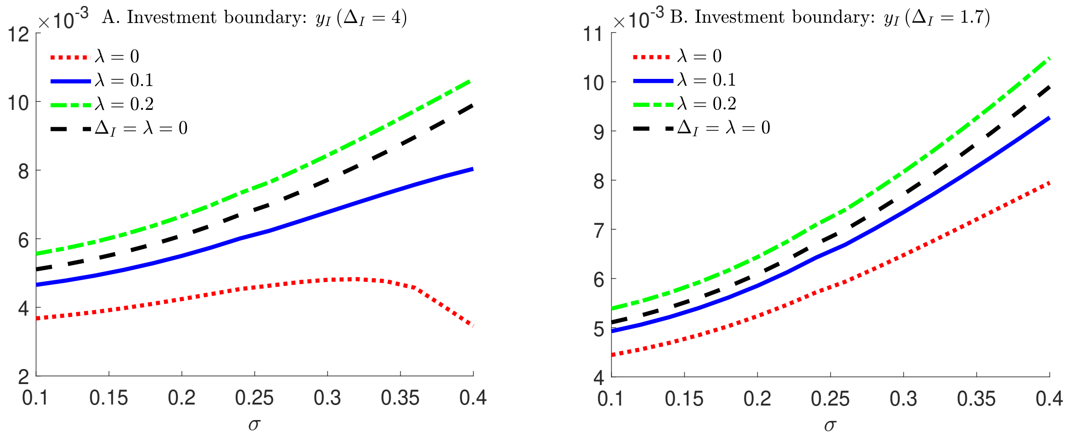

Figure 3: **The effect of uncertainty on investment boundaries.** The left panel show the results with a large investment lag ($\Delta = 4$) and the right panel is for a small investment lag ($\Delta_I = 1.7$). All other unspecified parameter values are from the baseline in Table 1.

Panel B plots the investment boundaries for a small investment lag calibrated from data ($\Delta_I = 1.7$). Quantitatively, the findings in panel A for large investment lag still hold. First, small investment lag hastens a firm's invest decision and hastens relatively more when uncertainty is larger (by comparing the red dotted line with the black dashed line). However, a small investment lag is not able to reverse the effect of uncertainty on investment boundaries as panel A does. In addition, opportunity cost of adjustment delays firm's investment and delay more for higher uncertainty (by comparing the blue solid line to the red dotted line). Furthermore, a even higher opportunity cost of adjustment further delays the firm's investment to be even later than the firm without investment lag (by comparing blue dash-dotted line to the black dashed line).

Our findings bridge the bad news principle and the good news principle, and we show

22

that whether a firm's investment follows good or bad news principle depends entirely on the interactions between investment lag and opportunity cost of adjustment. The investment decision is made by weighing the opportunity cost of waiting, which is the forgone profit during the period of inaction, and the benefit of waiting. When there is no investment lag, increase in uncertainty raises the benefit of waiting but doesn't affect the opportunity cost of waiting. Therefore, higher uncertainty delays investment and suggests a bad news principle. When there is investment lag, the opportunity cost of waiting also increases with uncertainty. Hence increase in uncertainty is possible to lead to ealier investment, which suggests a good news principle. In our framework, the opportunity cost of waiting equals to the foregone profit from newly formed capital netting the forgone cost from the existing capital stock due to the opportunity cost of adjustment. Although the increase in uncertainty increases both the foregone profit and forgone cost, the forgone cost from existing capital stock may outweigh the forgone profit generated by newly formed capital. Thus, it is possible that opportunity cost of waiting decreases with uncertainty, and increase in uncertainty can lead to later investment.

In Figure 4, we plot the capital sale boundaries for different parameters. Panel A plots the results for a relatively large investment lag ($\Delta_I = 4$). Without investment lag (the black dashed line), increase in uncertainty raises the benefit of waiting and hence delays capital sale decisions. In the presence of investment lag (red dotted line), capital formulation is more difficult than without lag, and firm's capital sale boundaries are getting lower. In the presence of both investment lag and opportunity cost of adjustment (blue solid line), capital formulation is getting even more difficult, which pushes the capital sale boundaries even lower. For all three cases, uncertainty delays capital selling decisions, which suggests bad news principle in capital sale. The findings also hold in panel B for a smaller investment lag ($\Delta_I = 1.7$).

5.3 Investment moments

In Figure 5, we plot the mean of investment amount conditioning on that investment decision is made.

Panel A plots the conditional mean of investment amount for two firms, the one with opportunity cost of adjustment (blue solid line) and the one without (red dashed line), for different levels of volatility.[9] For both firms, the mean of investment amount

[9]Note that the level of conditional mean for the baseline mode (varies from 0.7 to 0.75) is much higher than the simulated mean of investment in Table 2. This is because the moments reported in Table 2 are unconditional moments.

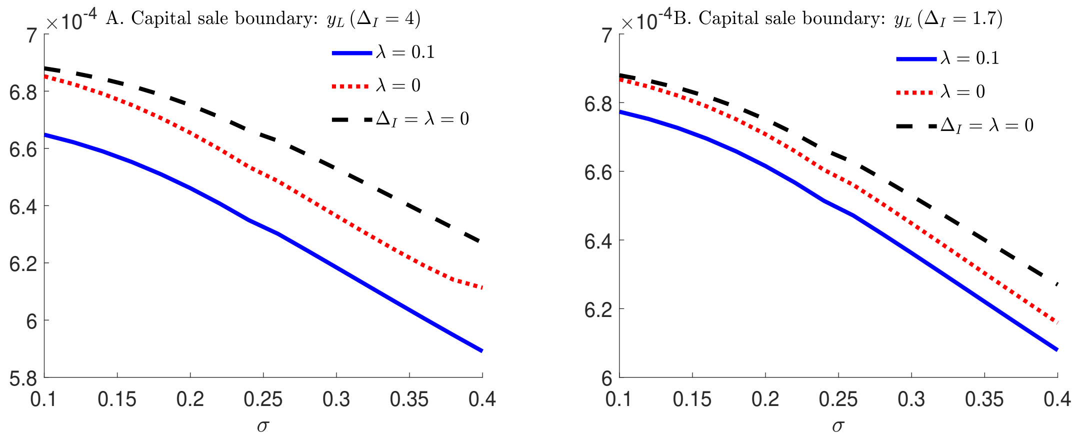

Figure 4: **The effect of uncertainty on capital sale boundaries.** The left panel show the results with a large investment lag ($\Delta = 4$) and the right panel is for a small investment lag ($\Delta_I = 1.7$). All other unspecified parameter values are from the baseline in Table 1.

decreases in uncertainty. This is because higher uncertainly lowers the likelihood that the scaled productivity-to-capital ratio (which follows geometric Brownian motion during the investment period) in the end of an investment period is higher than any given level, and hence lowers the mean of the investment amount. With opportunity cost of adjustment, the firm's investment boundaries are higher than that without. Meanwhile, the firm's investment target is lower with opportunity cost of adjustment than that without, because that the marginal value of capital is higher for the former (see, e.g., Figure 2 panel B). A higher investment boundary together with a lower investment target imply a higher conditional mean of investment amount. Therefore, the baseline model has a higher mean investment amount than the other firm.

Panel B plots the conditional mean of investment amount for the baseline model for different values of opportunit cost of adjustment λ. It shows that the conditional mean of investment amount increases in λ. The intuition is that increase in λ increases the investment boundary and meanwhile decreases the investment target. Thus the investment amount which depends on the distance between investment boundary and investment target increases in λ.

6 Net Present Value Method

Investment decision-making literature is interested in whether the net present value (NPV) method is a good approximation of the real option (RO) method in the presence of

24

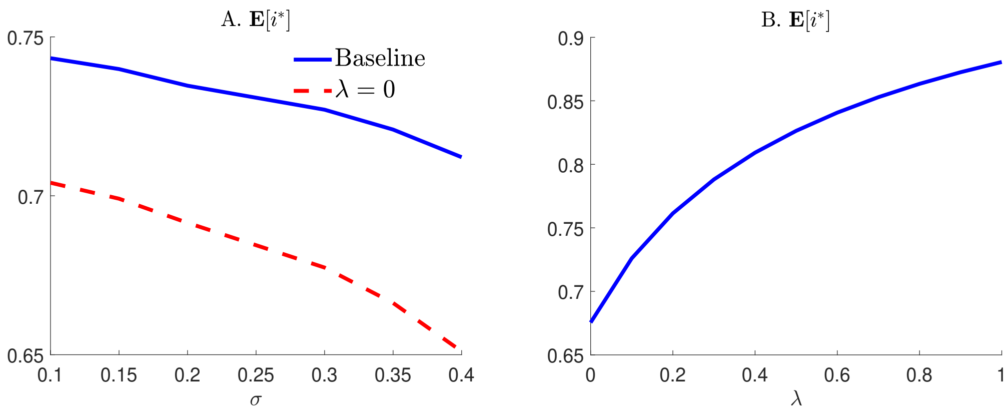

Figure 5: **Conditional mean of investment amount.** The conditional mean is characterized in Proposition 5. All unspecified parameters value use those of baseline in Table 1.

investment lag (see, i.e., Majd and Pindyck, 1987; Milne and Whally, 2000). In this section, we first introduce the net present value method into our framekwork. We then compare the optimal investment strategies from the NPV method and real option (RO) method, and examine asymptotically the effect of opportunity cost of adjustment on investment boundaries.

First, we define the NPV method according to the production function (1). In the NPV setting, the firm faces an one-off investment opportunity, and decision should be made immediately. It makes an investment decision by comparing the present value of future profit from investing with that from not investing, ignoring the uncertainty in future profit. In contrast, real option method refers to our baseline model which assumes that firm has the right to decide the optimal timing of investment. If a firm is uncertain about the future profit and investment is partially irreversible, it would prefer waiting rather than investing immediately.

Then, we define by V_1, the value function without investing, V_2 the value function when the firm chooses to invest immediately with an amount I, as follows:

$$\begin{aligned} V_1(K_0, X_0) &= \mathbb{E}\left[\int_0^\infty e^{-rt}\Pi(K_t, X_t)dt\right] \\ V_2(K_0, X_0) &= \mathbb{E}\left[\max_{I\geqslant 0}\int_0^{\Delta_I} e^{-rt}(1-\lambda)\Pi(K_t, X_t)dt\right. \\ &\qquad \left. + \int_{\Delta_I}^\infty e^{-rt}\Pi(K_t + I, X_t)dt - (1+b_I)e^{-r\Delta_I}I - e^{-r\Delta_I}c_I K_{\Delta_I}\right] \end{aligned}$$

25

Note that both V_1, V_2 are homogeneous in X and K with degree one. We assume $V_i(K,X) = Kv_i(y), y = X/K,\ i = 1,2$. The net present value method requires that the optimal investment boundary, denoted by y_{NPV}, equates the value of investing and that of not investing, i.e., $v_1(y_{NPV}) = v_2(y_{NPV})$. For the ease of notation, we denote by y_{RO} the optimal investment boundary from the baseline model (i.e., y_I given in Proposition 3) through this section.

We have the following asymptotic results when investment lag is large. The asymptotic analysis reveals the intrinsic impact of the existence of opportunity cost of adjustment λ on firm's investment decisions.

Proposition 6. *When investment lag is large, we have the following results*

(i) The optimal investment boundary from real option method satisfies

$$y_{RO} \approx \begin{cases} \dfrac{c_I\left(1+\frac{1}{|\alpha_-|}\right)}{\gamma A B^{1-\gamma}} e^{-(\delta+\mu)\Delta_I} & \text{if } \lambda = 0 \\ \left(\dfrac{\lambda\left(1+\frac{1-\gamma}{|\alpha_-|}\right)}{\gamma B^{1-\gamma}}\right)^{\frac{1}{1-\gamma}} e^{\frac{r-\mu}{1-\gamma}\Delta_I} & \text{if } \lambda > 0,\ , \end{cases}$$

where $B = \left(\frac{(1+b_I)}{(1-\gamma)A}\right)^{-1/\gamma}$*, and* A, α_- *are given in Appendix A.2. In particular, when* $\lambda = 0$, $\lim_{\Delta_I\to\infty} y_{RO} = 0$ *and when* $\lambda > 0$, $\lim_{\Delta_I\to\infty} y_{RO} = +\infty$.

(ii) The optimal investment boundary by net present value method satisfies

$$y_{NPV} \approx \begin{cases} \dfrac{c_I}{\gamma A B^{1-\gamma}} e^{-(\mu+\delta)\Delta_I} & \text{if } \lambda = 0 \\ \left(\dfrac{\lambda}{\gamma B^{1-\gamma}}\right)^{\frac{1}{1-\gamma}} e^{\frac{r-\mu}{1-\gamma}\Delta_I} & \text{if } \lambda > 0. \end{cases}$$

In particular, when $\lambda = 0$, $\lim_{\Delta_I\to\infty} y_{NPV} = 0$ *and when* $\lambda > 0$, $\lim_{\Delta_I\to\infty} y_{NPV} = +\infty$.

(iii) The ratio of the optimal investment boundary derived from the real option method and the net present value method satisfies

$$\lim_{\Delta_I\to\infty} \frac{y_{RO}}{y_{NPV}} = \begin{cases} 1 + \dfrac{1}{|\alpha_-|} & \text{if } \lambda = 0 \\ \left(1 + \dfrac{1-\gamma}{|\alpha_-|}\right)^{\frac{1}{1-\gamma}} & \text{if } \lambda > 0. \end{cases}$$

(iv) Given the parameter constaint (5), the investment boundaries from RO method

26

and NPV method are asymptotically equal, i.e., $\lim_{\Delta_I \to \infty} \frac{y_{RO}}{y_{NPV}} = 1$*, when*

$$r/\mu \to +\infty, \sigma \to 0, \delta \to 0.$$

The proof is relegated to Appendix B. Proposition 6 (i)-(ii) shows that the nature of the asymptotic behavior of optimal investment boundary (whether $\to +\infty$ or $\to 0$) relies on the existence of opportunity cost of adjustment λ. When there exists positive λ, both y_{NPV} and y_{RO} approach infinity as investment lag goes to infinity, implying the presence of λ will significantly delay the firm's investment for large investment lag, no matter which decision-making method is used. This is because the opportunity cost of adjustment lowers the opportunity cost of waiting to be close to zero when investment lag goes to infinity. In contrast, when there is no λ, both y_{NPV} and y_{RO} approach zero, implying the absence of λ makes the firm to invest as soon as possible when faces large investment lag. The reason is that extermely large investment lag will increase the opportunity cost of waiting to be infinity.

In addition, Proposition 6 (iii) compares the y_{RO} and y_{NPV} asymptotically and shows that the limit of the ratio between y_{RO} and y_{NPV} also depends on the existence of λ. We show that, asymptotically, the investment boundary from RO method is higher than that from NPV method, i.e., $y_{RO} > y_{NPV}$, since $|\alpha_-| > 0$. In general, it shows the ratio of y_{RO} to y_{NPV} depends on the value α_-. In practice, when investment lag is large, our asymptotic analysis further suggests that one may simply approximate the RO method by multiplying the NPV boundary with a coefficient given in Proposition 6.

Proposition 6 (iv) further provides sufficient conditions, i.e., $r/\mu \to +\infty, \delta \to 0$ and $\sigma \to 0$, such that the RO investment boundary y_{RO} equals the NPV boundary y_{NPV} asymptotically, i.e., $\lim_{\Delta_I \to +\infty} \frac{y_{RO}}{y_{NPV}} = 1$. Our findings have interesting implications to the disputation about whether NPV method is a good approximation of the RO method in the presence of investment lag (see, e.g., Majd and Pindyck, 1987; Milne and Whally, 2000). On the one hand, we find asymptotically whether the NPV method is a good approximation to the RO method depends on the parameter values. On the other hand, our findings concur with Milne and Whally (2000) that NPV method may be a good approximation of RO method for some conditions. We show when the risk free rate r is high, a firm grows slowly ($\mu \to 0$) with small volatility ($\sigma \to 0$), depreciates slowly ($\delta \to 0$), and faces large investment lag, the NPV method tends to be a good approximation of the RO method.

In Figure 6, we provide numerical analysis for the investment boundaries from both

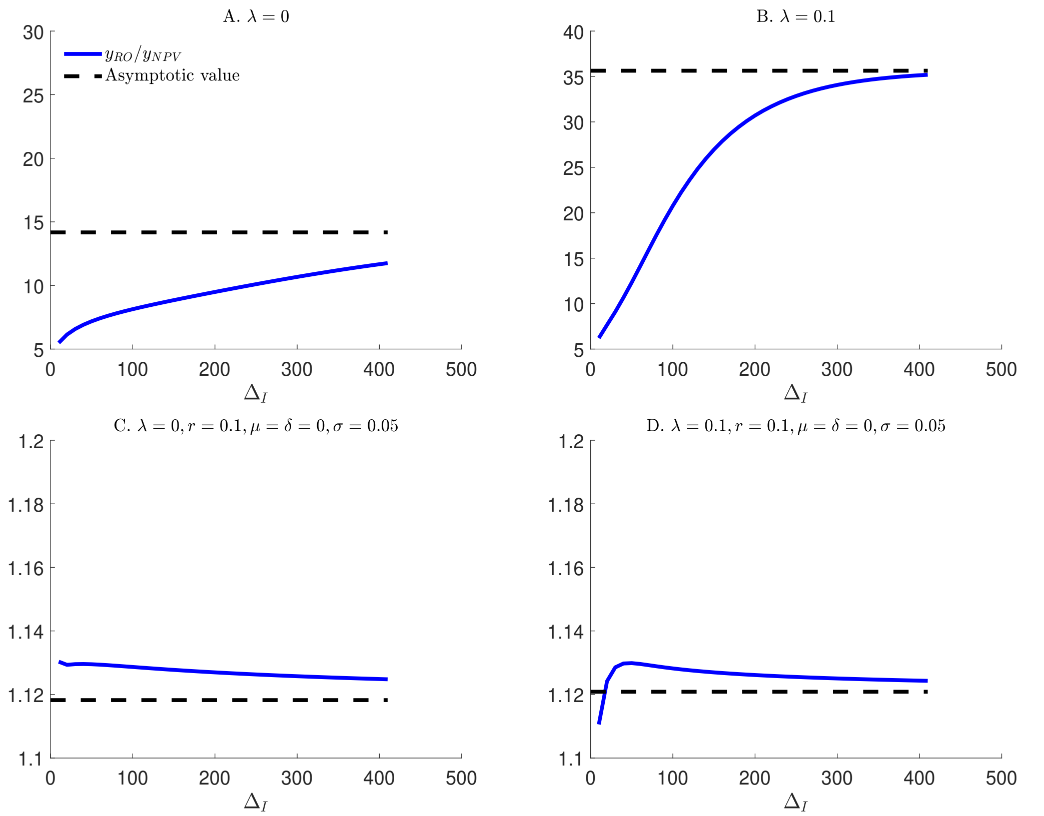

Figure 6: **Ratio of optimal investment boundaries of NPV method and RO method**. In each panel, the blue solid line is the exact ratio based on model solution in Proposition 3, and the dark dashed line is the asymptotic value given in Proposition 6 (iii). For each panel, except the specified parameters, all other parameter values are from the baseline in Table 1.

model solution and asymptotical analysis. All four panels show that for different parameter values, the exact ratio obtained from Proposition 3 approximates to the asymptotic value from Propostion 6 when investment lag is large. Panel A and B show the cases that NPV investment boundaries are quite different from RO boundaries (i.e., the ratio y_{RO}/y_{NPV} is large). Panel C and D show the cases that, when the conditions provided in Proposition 6 (iv) are satisfied, the NPV investment boundaries are close to the RO boundaries (i.e., the ratio y_{RO}/y_{NPV} approaches 1) in the presence of large investment lag.

7 Effect of Abandonment Option

In this section, we investigate the influence of abandonment option on the firm's investment strategies. In our baseline model, once investment decision is made, the firm enters into an investment period during which no abandonment decision can be made. However, when the investment lag is very large or the firm's productivity shock is very volatile, the productivity as well as the profit may change dramatically in an unfavorable direction in the investment period. It then generates needs for the firm to reconsider its investment decision and to consider abaondoning the investment or selling capital in the investment period. To accomodate such needs, we extend our model to the case that the firm is endowed with an abandonment option when it decides to invest. The abandonment option is associated with a sink cost $p > 0$, which can be treated as a down payment once investment decision is made. In our baseline model, we assume there is no down payment at all.

Firm's abandonment option is a series of impulse control $\omega_A = \{\tau_k^A, k = 1, \cdots\}$ where τ_j^A is the timing of abandonment when firm enters into investment at time τ_j^I for some j with $\tau_k^A \in [\tau_j^I, \tau_j^I + \Delta_I]$. Then the firm's capital stock evolves as follows:

$$K_t = K_0 - \int_0^t \delta K_s ds + \sum_j 1_{\{\tau_k^A \notin [\tau_j^I, \tau_j^I + \Delta_I]\}} I_j \mathbf{1}_{\{\tau_j^I + \Delta_I \leqslant t\}} - L_t \tag{15}$$

where τ_k^I, I_k and L_t are defined in (3). The firm's optimality satisfies:

$$\begin{aligned} \max_{\omega, \omega_A} E\Bigg[& \int_0^\infty e^{-rt}(1 - \lambda 1_{\{t \in [\tau_j^I, \tau_j^I + \Delta_I]\}})\Pi(K_t, X_t)dt \\ & - \sum_j e^{-r(\tau_j^I + \Delta_I)}((1 + b_I)I_j + c_I K_{\tau_j^I + \Delta_I}) 1_{\{\tau_k^A \notin [\tau_j^I, \tau_j^I + \Delta_I]\}} \\ & - \int_{\tau_j^I}^{\tau_k^A} p e^{-rt} dt 1_{\{\tau_k^A \in [\tau_j^I, \tau_j^I + \Delta_I]\}} + \sum_j e^{-r\tau_j^L}(1 - b_L) dL_t \Bigg]. \end{aligned} \tag{16}$$

The abandonment option can be view as an American style option, and it is activated once the firm decides to investment which is subject to an investment lag. This American style option has maturity Δ_I and is exercised with a cost w. The operator $\mathcal{M}V$ given in

HJB equation (6) now becomes

$$\mathcal{M}V(K,X) = \max\Bigg\{\mathbb{E}\left[\sup_{I>0} e^{-r\Delta_I}\Big(V(K_{\Delta_I}+I, X_{\Delta_I}) - ((1+b_I)I_{\tau_i^I+\Delta_I} + c_I K_{\tau_i^I+\Delta_I})\Big)\right],$$
$$\mathbb{E}\left[\sup_{\tau_k^A\in[0,\Delta_I]} e^{-r\tau_k^A} V(K_{\tau_k^A}, X_{\tau_k^A}) - \int_0^{\tau_k^A} pe^{-rt}dt\right]\Bigg\}. \quad (17)$$

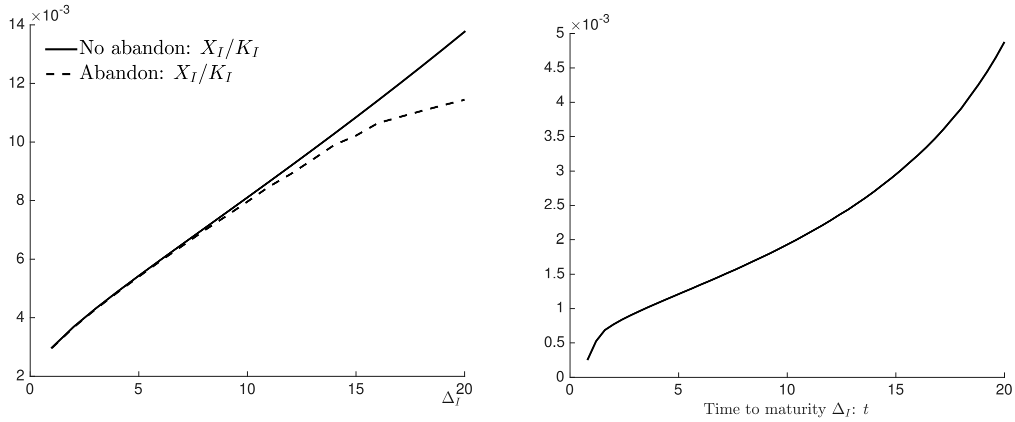

Figure 7: **Investment boundaries and abandonment boundaries with abandonment option** In the left panel, we plot the investment boundaries y_I for the baseline model (solid line) and the one with abandonment option (dashed line). The right panel plots the abandonment boundaries against the investment lag Δ_I. Here we choose $c_I = \lambda = 0$.

In Figure 7, we provide numerical results of the optimal investment boundaries and the optimal excises boundaries for the abandonment opion. In the left panel, we find abandon option advances investment decision since the firm is able to take more risk in investment by holding this option. In the right panel, we find that the abandon barrier is increasing with respect to the investment lag Δ_I, the reason is, when Δ_I is larger, the corporation will face more uncertainty and it is easier to abandon.

8 Conclusion

In this paper, we provide a unified investment theory that demonstrates how the presence of various investment frictions influence a firm's optimal investment decisions. Our model reconciles the "good news principle" and "bad news principle" in exsiting litreature, and show that which principle dominates depends on the joint effect of investment lag and opportunity cost of adjustment. Moreover, we show that whether the NPV rule provides consistent conclusions with the RO method depends on the parameter values. We provide sufficient conditions under which NPV method approximates to the RO method. Our calibration study shows that our unified model can well explain the patterns in the correlation between investment and the lagged profit.

References

Abel, Andrew B. and Eberly, Janice C. 1994. A Unified Model of Investment under Uncertainty. *American Economic Review* **84** 1369-1384.

Abel, Andrew B. and Eberly, Janice C. 1996. Optimal Investment with Costly Reversibility. *Review of Economic Studies* **63** 581-593.

Abel, Andrew B. and Eberly, Janice C. 1997. An exact solution for the investment and value of a firm facing uncertainty, adjustment costs, and irreversibility. *Journal of Economic Dynamics and Control* **21** 831-852.

Bar-Ilan, Avner and Strange, William C. 1996. Investment Lags. *American Economic Review* **86** 610-622.

Bar-Ilan, Avner and Strange, William C. 1996. Urban development with lags. *Journal of Urban Economics* **39** 87-113.

Bar-Ilan, Avner and Strange, William C. 1998. A model of sequential investment. *Journal of Economics Dynamics & Control* **22** 437-463.

Bayraktar, E., and Egami, M 2010. A Unified Treatment of Dividend Payment Problems under Fixed Cost and Implementation Delays. *Mathematical Methods of Operation Research* **71** 325-351.

Bernanke, S. Ben. 1983. Irreversibility, Uncertainty, and Cyclical Investment. *Quarterly Journal of Economics* **98** 85-106.

Bertola, Giuseppe. 1988. Adjustment costs and dynamic factor demands: investment and employment under uncertainty. *Doctoral dissertation, Massachusetts Institute of Technology*

Bruder, Benjamin, and Pham, Huyen. 2009. Impulse Control Problem on Finite Horizon with Execution Delay. *Stochastic Processes and Their Applications* **119** 1436-1439.

Cooper, W. Russell and Haltiwanger, John C. 2006. On the Nature of Capital Adjustment Costs. *Review of Economic Studies* **73** 611-633.

Dai, Min., Huang, Shan and Keppo, Jussi. 2019. Opaque Bank Assets and Optimal Equity Capital. *Journal of Economic Dynamics and Control* **100** 369-394.

Dixit, Avinash. 1991. Irreversible investment with price ceilings. *Journal of Political Economy* **99** 541-557.

Ganuza, Juan-José. 2007. Competition and cost overruns in procurement. *Journal of Industrial Economics* **55** 633–660.

He, Xuedong., Jin, Hanqing., and Zhou, Xunyu. 2015. Dynamic Portfolio Choice when Risk is Measured by Weighted Var. *Mathematics of Operations Research* **40** 773-796.

John, Kose, and Saunders, Anthony. 1983. Asymmetry of information, regulatory lags and optimal incentive contracts: Theory and evidence. *Journal of Finance* **38** 391–404.

Jorgenson, Dale Weldeau. 1963. Investment: Capital theory and investment behavior. *Mit Press*

Kalouptsidi, Myrto. 2014. Time to Build and Fluctuations in Bulk Shipping. *American Economic Review* **104** 564-608.

Kydland, Finn E., and Prescott, Edward C. 1982. Time to build and aggregate fluctuations. *Econometrica* 1345-1370

Lee, B. S., and Ingram, B. F. 1991. Simulation estimation of time-series models. Journal of Econometrics, 47(2-3), 197-205.

Majd, Saman and Pindyck, Robert S. 1987. Time to Build, Option Value, and Investment Decision. *Journal of Financial Economics* **18** 7-27.

McDonald, Robert, and Siegel, Daniel 1986. The Value of Waiting to Invest. *The Quarterly Journal of Economics* **101** 707.

Milne, Alistair and Whally, A. Elizabeth. 2000. A Comment on "Time to Build, Option Value, and Investment Decision". *Journal of Financial Economics* **56** 325-332.

Pindyck, Robert S. 1993. A note on competitive investment under uncertainty. *The American Economic Review* **81** 273-277.

Power, Laura. 1998. The missing link: technology, investment, and productivity. *Review of Economics and Statistics* **80**, 300-313

Sakellaris, Plutarchos. 2004. Patterns of Plant Adjustment. *Journal of Monetary Economics* **51** 425-450.

Shi, Yun, Cui, Xiangy, Yao, Jing and Li, Duan. 2015. Dynamic Trading with Reference Point Adaption and Loss Aversion. *Operations Research* **63** 789-806.

Subramanian, Ajay, and Jarrow, Robert A. 2001. The Liquidity Discount. *Mathematical Finance* **11** 447-474.

Yao, David D., Zhang, Hanqin, and Zhou, Xun Yu. 2002. Stochastic Models and Optimization: With Applications in Queues, Finance and Supply Chains. *Working Paper*.

Zhou, C. 2000. Time-to-build and investment. Review of Economics and Statistics, 82(2), 273-282.

A Appendix

A. Technical Analysis

A.1 Proof of Proposition 1

Proof. We prove that $V(K,X)$ is non-increasing in Δ_I. Denote by $\Pi(\Delta_I)$ the set of admissible strategies with investment delay Δ_I. Assume $\Delta_I > \tilde{\Delta}_I$. For any admissible strategy $\pi = \{(\tau_j^{I,\pi}, I_j^\pi), (\tau_j^{L,\pi}, L_j^\pi)\} \in \Pi(\Delta_I)$ and the associated firm value $V_\pi(\cdot,\cdot;\Delta)$, we construct a strategy $\tilde{\pi} = \{(\tau_j^{I,\tilde{\pi}}, I_j^{\tilde{\pi}}), (\tau_j^{L,\tilde{\pi}}, L_j^{\tilde{\pi}})\}$ as follows:

$$
\begin{aligned}
&I_j^{\tilde{\pi}} = I_j^\pi,\ L_j^{\tilde{\pi}} = L_j^\pi,\\
&\tau_j^{I,\tilde{\pi}} = \tau_j^{I,\pi} + \Delta_I - \tilde{\Delta}_I,\ \tau_j^{L,\tilde{\pi}} = \tau_j^{L,\pi}.
\end{aligned}
$$

It is easy to see that $\tilde{\pi} \in \tilde{\Pi}(\tilde{\Delta}_I)$ and

$$V_\pi(\cdot,\cdot;\Delta) = V_{\tilde{\pi}}(\cdot,\cdot;\tilde{\Delta}) \leqslant V(\cdot,\cdot;\tilde{\Delta}).$$

Maximizing the left-hand side over $\pi \in \Pi(\Delta_I)$, we have

$$V(\cdot,\cdot;\Delta) \leqslant V(\cdot,\cdot;\tilde{\Delta}).$$

33

□

A.2 Definition of parameters in Proposition 2

The parameters in Proposition 2 are given by

$$A = \frac{h}{(1-\gamma)\psi}, \; \psi = r - \gamma\mu + (\delta + \frac{\gamma}{2}\sigma^2)(1-\gamma), \tag{A.1}$$

$$\alpha_{\pm} = \frac{\mu + \delta + \frac{1}{2}\sigma^2 \pm \sqrt{(\mu+\delta+\frac{1}{2}\sigma^2)^2 + 2\sigma^2(r-\mu)}}{\sigma^2}, \tag{A.2}$$

$$C_1 = \frac{(1-b_L)(1-\alpha_-)}{\alpha_+(\alpha_+-\alpha_-)} y_L^{\alpha_+-1} - \frac{A(1-\gamma)(1-\alpha_- - \gamma)}{\alpha_+(\alpha_+-\alpha_-)} y_L^{-1+\alpha_+ +\gamma}, \tag{A.3}$$

$$C_2 = \frac{(1-b_L)(\alpha_+-1)}{\alpha_-(\alpha_+-\alpha_-)} y_L^{\alpha_- -1} - \frac{A(1-\gamma)(\alpha_+ - 1 + \gamma)}{\alpha_-(\alpha_+-\alpha_-)} y_L^{-1+\alpha_- +\gamma}, \tag{A.4}$$

where the boundaries y_I, y_L and target level s_I solves

$$v(y_I) = y_I/s_I v(s_I) - (1+b_I)(y_I/s_I - 1) - c_I, \tag{A.5}$$

$$v(y_I) - y_I v'(y_I) = 1 + b_I - c_I, \tag{A.6}$$

$$v(s_I) - s_I v'(s_I) = 1 + b_I, \tag{A.7}$$

$$v(y_L) - y_L v'(y_L) = 1 - b_L. \tag{A.8}$$

where $v(y) - yv'(y)$ is the firm's marginal q, the marginal valuation of capital $V_K(K, X)$.

A.3 Definition of H, f_1, f_2 in Proposition 3

$$H(y) = Ay^{\gamma} + C_1 y^{1-\alpha_+} + C_2 y^{1-\alpha_-} \tag{A.9}$$

$$f_2(y) = H(y_L)y/y_L + (1-b_L)(1-y/y_L), \tag{A.10}$$

$$\begin{aligned} f_1(y) = \; & e^{-(r-\mu)\Delta_I}(yv(s_I)/s_I - (1+b_I)y/s_I)(1 - \Phi(d_I - \sigma\sqrt{\Delta_I})) \\ & + (1+b_I)e^{-(r+\delta)\Delta_I}(1-\Phi(d_I)) + (1-b_L)e^{-(r+\delta)\Delta_I}\Phi(h_L) \\ & + C_1 y^{1-\alpha_+}(\Phi(d_I - \sigma\sqrt{\Delta_I}(1-\alpha_+)) - \Phi(h_L - \sigma\sqrt{\Delta_I}(1-\alpha_+))) \\ & + C_2 y^{1-\alpha_-}(\Phi(d_I - \sigma\sqrt{\Delta_I}(1-\alpha_-)) - \Phi(h_L - \sigma\sqrt{\Delta_I}(1-\alpha_-))) \\ & + Ay^{\gamma}e^{-\psi\Delta_I}(\Phi(d_I - \sigma\sqrt{\Delta_I}\gamma) - \Phi(h_L - \sigma\sqrt{\Delta_I}\gamma)) \\ & + e^{-(r-\mu)\Delta_I}(yv(y_L)/y_L - (1-b_L)y/y_L)\Phi(y_L - \sigma\sqrt{\Delta_I}) \\ & + (1-\lambda)A(1-e^{-\psi\Delta_I})y^{\gamma} - c_I e^{-(r+\delta)\Delta_I}. \end{aligned} \tag{A.11}$$

34

where $d_I = \frac{\ln(s_I/y)-(\delta+\mu-\frac{1}{2}\sigma^2)\Delta_I}{\sigma\sqrt{\Delta_I}}$, $h_L = \frac{\ln(y_L/y)-(\delta+\mu-\frac{1}{2}\sigma^2)\Delta_I}{\sigma\sqrt{\Delta_I}}$, and y_I, s_I, y_L solves

$$\begin{aligned}
1+b_I &= A(1-\gamma)s_I^{\gamma} + C_1\alpha_+ s_I^{1-\alpha_+} + C_2\alpha_- s_I^{1-\alpha_-}, && \text{(A.12)}\\
f_1(y_I) &= Ay_I^{\gamma} + C_1 y_I^{1-\alpha_+} + C_2 y_I^{1-\alpha_-}, && \text{(A.13)}\\
f_1'(y_I) &= A\gamma y_I^{\gamma-1} + C_1(1-\alpha_+)y_I^{-\alpha_+} + C_2(1-\alpha_-)y_I^{-\alpha_-}, && \text{(A.14)}\\
1-b_L &= A(1-\gamma)y_L^{\gamma} + C_1\alpha_+ y_L^{1-\alpha_+} + C_2\alpha_- y_L^{1-\alpha_-}. && \text{(A.15)}
\end{aligned}$$

B. Proof of Proposition 6

We first solve the value functions V_1, V_2 for the net present value method as follows:

$$V_1 = \mathbb{E}\left[\int_0^{\infty} e^{-rt}\Pi(K_t, X_t)dt\right] = AK_0^{1-\gamma}X_0^{\gamma},$$

and

$$\begin{aligned}
V_2 = \ & E\Big[\max_{I\geqslant 0}\int_0^{\Delta_I} e^{-rt}(1-\lambda)\Pi(K_t,X_t)dt \\
&+ \int_{\Delta_I}^{\infty} e^{-rt}\Pi(K_t+I,X_t)dt - (1+b_I)e^{-r\Delta_I}I - e^{-r\Delta_I}c_I K_{\Delta_I}\Big] \\
&= (1-\lambda)A(1-e^{-\psi\Delta_I})K_0^{1-\gamma}X_0^{\gamma} + AB^{1-\gamma}e^{(\mu-r)\Delta_I}X_0\Phi(-d_I(y)+\sigma\sqrt{\Delta_I}) \\
&\quad + Ae^{-\psi\Delta_I}K_0^{1-\gamma}X_0^{\gamma}\Phi(d_I(y)-\gamma\sigma\sqrt{\Delta_I}) + (1+b_I)e^{-(\delta+r)\Delta_I}K_0\Phi(-d_I(y)) \\
&\quad - (1+b_I)Be^{(\mu-r)\Delta_I}X_0\Phi(-d_I(y)+\sigma\sqrt{\Delta_I}) - c_I e^{-(\delta+r)\Delta_I}K_0,
\end{aligned}$$

where $d_I(y) = \frac{\ln(B/y)-(\delta+\mu-\frac{1}{2}\sigma^2)\Delta_I}{\sigma\sqrt{\Delta_I}}$.

Then we discuss the asymptotic results for two cases: (I) $\lambda = 0$ and (II) $\lambda > 0$.

(I) For the case with $\lambda = 0$, we first investigate the asymptotic behavior of the NPV investment boundary as $\Delta_I \to +\infty$. We have the investment boundary, y_I should satisfy the following equation

$$\begin{aligned}
&- Ae^{-\psi\Delta_I}y^{\gamma-1}\Phi(-d_I(y)+\gamma\sigma\sqrt{\Delta_I}) + [AB^{1-\gamma}-(1+b_I)B]e^{(\mu-r)\Delta_I}\Phi(-d_I(y)+\sigma\sqrt{\Delta_I}) \\
&\quad + (1+b_I)e^{-(\delta+r)\Delta_I}y^{-1}\Phi(-d_I(y)) - c_I e^{-(\delta+r)\Delta_I}y^{-1} = 0.
\end{aligned}$$

It is equivalent to the equation

$$\begin{aligned}
&- Ae^{-\psi_1\Delta_I}y^{\gamma-1}\Phi(-d_I(y)+\gamma\sigma\sqrt{\Delta_I}) + [AB^{1-\gamma}-(1+b_I)B]\Phi(-d_I(y)+\sigma\sqrt{\Delta_I}) \\
&\quad + (1+b_I)e^{-(\delta+\mu)\Delta_I}y^{-1}\Phi(-d_I(y)) - c_I e^{-(\delta+\mu)\Delta_I}y^{-1} = 0, \qquad \text{(A.16)}
\end{aligned}$$

where $\psi_1 = \psi - r + \mu = (1-\gamma)(\frac{1}{2}\sigma^2\gamma + \delta + \mu) > 0$.

Lemma 1. *For any $\Delta_I > 0$, the solution $y_{NPV} > 0$ exists and as $\Delta_I \to +\infty$, $y_{NPV} \to 0$.*

Proof. The existence of solution can be achieved by intermediate value theorem from computing the case $y = 0$ and $y = +\infty$.

To show $y_{\Delta_I} \to 0$, we apply proof by contradiction. Otherwise, we choose a bounded subsequence of $1/y_{NPV}$ with bound M_0, such that

$$
\begin{aligned}
&| - Ae^{-\psi_1\Delta_I}(1/y)^{1-\gamma}\Phi(-d_I(y) + \gamma\sigma\sqrt{\Delta_I}) + (1+b_I)e^{-(\delta+\mu)\Delta_I}\frac{1}{y}\ \Phi(-d_I(y)) - c_I e^{-(\delta+\mu)\Delta_I}\frac{1}{y}\ | \\
\leqslant & e^{-\psi_1\Delta_I}A(1/y)^{1-\gamma} + (1+b_I)e^{-(\delta+\mu)\Delta_I}(1/y) + c_I e^{-(\delta+\mu)\Delta_I}(1/y) \\
\to & 0,
\end{aligned}
$$

while as $\Delta_I \to +\infty$,

$$
\begin{aligned}
&[AB^{1-\gamma} - (1+b_I)B]\Phi(-d_I(y) + \sigma\sqrt{\Delta_I}) \\
=&\gamma AB^{1-\gamma}\Phi(-d_I(y) + \sigma\sqrt{\Delta_I}) \\
\geqslant&\gamma AB^{1-\gamma}\Phi(\frac{-\ln M_0 + \ln B + (\mu + \delta + \frac{1}{2}\sigma^2)\Delta_I}{\sigma\sqrt{\Delta_I}}) \\
\to&\gamma AB^{1-\gamma} > 0.
\end{aligned}
$$

Contradiction. □

By comparing the leading orders in (A.16), we must have $y_{NPV} \sim e^{-(\mu+\delta)\Delta_I}$. And the leading order equation is

$$
AB^{1-\gamma} - (1+b_I)B - c_I e^{-(\delta+\mu)\Delta_I}y^{-1} = 0, \tag{A.17}
$$

The leading term of y_{NPV} is $\frac{c_I}{AB^{1-\gamma}-(1+b_I)B}e^{-(\mu+\delta)\Delta_I} = \frac{c_I}{\gamma AB^{1-\gamma}}e^{-(\mu+\delta)\Delta_I}$.

Then we turn to invetigate the RO method investment boundary when $\lambda = 0$. As $\Delta_I \to +\infty$, the RO method is asymptotically equivalent to an optimal stopping problem with only one investment opportunity. The reason is that the investment lag is too long for the second investment to take effect, then the second and folllowing investments are negligible due to the exponential time discount.

For this optimal stopping time problem, the value function $v(y)$ can be characterized

36

by the following ODE

$$\min\{-\mathcal{L}_1 v,\ v - v_2\} = 0.$$

Since the it is optimal to invest when $y = \frac{X}{K}$ is too high, we guess the solution to satisfy

$$\begin{cases} \mathcal{L}_1 v = 0 & y \leqslant y_{RO} \\ v = v_2 & y \geqslant y_{RO}, \end{cases}$$

which is

$$\begin{cases} v(y) = C_1 y^{1-\alpha_+} + C_2 y^{1-\alpha_-} + Ay^{\gamma} & y \leqslant y_{RO} \\ v = v_2 & y \geqslant y_{RO}, \end{cases}$$

where C_1 and C_2 are constants to be determined. Notice when y is sufficiently small, v/v_1 should be close to 1 since y is unlikely to hit y_{RO} in Δ_I time, then $C_1 = 0$.

By the continuity condition and first order condition, the leading term equations for them are (where the parameters to be determined are y and C_2)

$$\begin{cases} AB^{1-\gamma} - (1+b_I)B - c_I e^{-(\delta+\mu)\Delta_I} y^{-1} - C_2 y^{-\alpha_-} e^{(r-\mu)\Delta_I} = 0, & \text{(A.18)} \\ -c_I e^{-(\delta+\mu)\Delta_I} - \alpha_- C_2 y^{1-\alpha_-} e^{(r-\mu)\Delta_I} = 0. & \text{(A.19)} \end{cases}$$

From (A.19), we have $C_2 y^{-\alpha_-} e^{(r-\mu)\Delta_I} = \frac{c_I}{-\alpha_-} e^{-(\delta+\mu)\Delta_I} y^{-1}$, put it into (A.18), we have

$$\begin{aligned} AB^{1-\gamma} - (1+b_I)B - c_I e^{-(\delta+\mu)\Delta_I} y^{-1} - \frac{c_I}{-\alpha_-} e^{-(\delta+\mu)\Delta_I} y^{-1} &= 0 \\ AB^{1-\gamma} - (1+b_I)B - c_I(1 + \frac{1}{-\alpha_-}) e^{-(\delta+\mu)\Delta_I} y^{-1} &= 0, \end{aligned}$$

which is $y_{RO} = \frac{c_I(1+\frac{1}{-\alpha_-})}{AB^{1-\gamma}-(1+b_I)B} e^{-(\delta+\mu)\Delta_I} = \frac{c_I(1+\frac{1}{-\alpha_-})}{\gamma AB^{1-\gamma}} e^{-(\delta+\mu)\Delta_I} = \frac{1+|\alpha_-|}{|\alpha_-|} y_{NPV}$ as $\Delta_I \to +\infty$.

(II) For the case $\lambda > 0$, the NPV investment boundary should satisfy

$$\begin{aligned} &-\lambda A y^{\gamma-1}(1 - e^{-\psi\Delta_I}) - Ae^{-\psi\Delta_I} y^{\gamma-1} \Phi(-d_I(y) + \gamma\sigma\sqrt{\Delta_I}) \\ &\quad + [AB^{1-\gamma} - (1+b_I)B] e^{(\mu-r)\Delta_I} \Phi(-d_I(y) + \sigma\sqrt{\Delta_I}) \\ &\quad + (1+b_I) e^{-(\delta+r)\Delta_I} y^{-1} \Phi(-d_I(y)) - c_I e^{-(\delta+r)\Delta_I} y^{-1} = 0. \end{aligned} \tag{A.20}$$

The following lemma summarizes the asymptotic behavior of NPV investment boundary given $\lambda > 0$.

37

Lemma 2. *For any $\Delta_I > 0$, the solution $y_{NPV} > 0$ exists and as $\Delta_I \to +\infty$, $y_{NPV} \to +\infty$. Especially, the leading term of y_{NPV} is $\left(\frac{\lambda}{\gamma B^{1-\gamma}}\right)^{\frac{1}{1-\gamma}} e^{\frac{r-\mu}{1-\gamma}\Delta_I}$.*

Proof. When $y \to 0$, left side of A.20 < 0, when $y \to +\infty$, left side of A.20 > 0, by intermediate value theorem, there exists a solution.

As $\Delta_I \to +\infty$, the leading term equation is

$$\begin{aligned}&[AB^{1-\gamma} - (1+b_I)B]e^{(\mu-r)\Delta_I}\Phi(-d_I(y) + \sigma\sqrt{\Delta_I}) + (1+b_I)e^{-(\delta+r)\Delta_I}y^{-1}\Phi(-d_I(y)) \\ =&\lambda A y^{\gamma-1} + c_I e^{-(\delta+r)\Delta_I}y^{-1}.\end{aligned}$$

If the leading order on left side is $(1+b_I)e^{-(\delta+r)\Delta_I}y^{-1}\Phi(-d_I(y))$, then y is at most of order $\frac{e^{-(\mu-r)\Delta_I}}{e^{(\delta+r)\Delta_I}} = e^{-(\mu+\delta)\Delta_I}$, and consequently $y^\gamma \leqslant e^{-\gamma(\mu+\delta)\Delta_I}$, which implies

$$\lim_{\Delta_I \to +\infty} -d_I(y) \to -\infty. \tag{A.21}$$

Then $(1+b_I)e^{-(\delta+r)\Delta_I}y^{-1}\Phi(-d_I(y)) << c_I e^{-(\delta+r)\Delta_I}y^{-1}$, the equation cannot hold, contradiction.

Therefore on left side, the leading order is $[AB^{1-\gamma} - (1+b_I)B]e^{(\mu-r)\Delta_I}\Phi(-d_I(y) + \sigma\sqrt{\Delta_I})$. If right side leading order is $c_I e^{-(\delta+r)\Delta_I}y^{-1}$, y should be of order $e^{-(\mu+\delta)\Delta_I}$, which implies actually $\lambda A y^{\gamma-1}$ is of higher order, contradiction.

In summary, the leading order equation is

$$[AB^{1-\gamma} - (1+b_I)B]e^{(\mu-r)\Delta_I}\Phi(-d_I(y) + \sigma\sqrt{\Delta_I}) = \lambda A y^{\gamma-1},$$

which implies $y \to +\infty$, then $\Phi(-d_I(y) + \sigma\sqrt{\Delta_I}) \to 1$, and

$$\begin{aligned}\lambda A y^{\gamma-1} &= [AB^{1-\gamma} - (1+b_I)B]e^{(\mu-r)\Delta_I} \\ \Rightarrow y &= \left(\frac{\lambda A}{AB^{1-\gamma} - (1+b_I)B}\right)^{\frac{1}{1-\gamma}} e^{-\frac{\mu-r}{1-\gamma}\Delta_I} = \left(\frac{\lambda}{\gamma B^{1-\gamma}}\right)^{\frac{1}{1-\gamma}} e^{-\frac{\mu-r}{1-\gamma}\Delta_I}.\end{aligned} \tag{A.22}$$

□

Let's turn to the RO investment boundary for the case $\lambda > 0$. We first provide the following lemma.

Lemma 3. *In RO case, the leading term of the boundary is $y = \left(\frac{\lambda(1+\frac{1-\gamma}{|\alpha_-|})}{\gamma B^{1-\gamma}}\right)^{\frac{1}{1-\gamma}} e^{\frac{r-\mu}{1-\gamma}\Delta_I}$. Therefore, according to lemma 2, $y_{RO} = (1+\frac{1-\gamma}{|\alpha_-|})^{\frac{1}{1-\gamma}} y_{NPV}$.*

Proof. Similar to the RO method for $\lambda = 0$ case, we can characterize the RO value function as

$$\begin{cases} v = C_2 y^{1-\alpha_-} + A y^{\gamma} & y \leqslant y_{RO} \\ v = v_2 & y \geqslant y_{RO}. \end{cases}$$

According to the smooth-pasting condition,

$$\begin{cases} -\lambda A y^{\gamma-1} + [AB^{1-\gamma} - (1+b_I)B]e^{(\mu-r)\Delta_I} - C_2 y^{-\alpha_-} = 0 \\ -\lambda A(1-\gamma)y^{\gamma} - \alpha_- C_2 y^{1-\alpha_-} = 0. \end{cases}$$

Solve it and we get the result. □

Some may argue that the fixed cost c_I should be imposed at the investment decision time, rather than the end of construction period, in this case, we have the following lemma, which implies the asymptotic ratio y_{RO}/y_{NPV} is the same as Proposition 6.

Lemma 4. *When investment lag $\Delta_I \to +\infty$, the ratio of the optimal investment boundary derived from the real option method and the net present value method satisfies*

$$\lim_{\Delta_I \to \infty} \frac{y_{RO}}{y_{NPV}} \to \begin{cases} 1 + \dfrac{1}{|\alpha_-|} & \text{if } \lambda = 0 \\ \left(1 + \dfrac{1-\gamma}{|\alpha_-|}\right)^{\frac{1}{1-\gamma}}, & \text{if } \lambda > 0. \end{cases}$$

Analysis and modeling of client order flow in limit order markets

Rama Cont[1], Mihai Cucuringu[1,2,3], Vacslav Glukhov[4], and Felix Prenzel*[1,4]

[1] *Mathematical Institute, University of Oxford*
[2] *Department of Statistics, University of Oxford*
[3] *The Alan Turing Institute*
[4] *JP Morgan*†

March 15, 2022

Abstract

Orders in major electronic stock markets are executed through centralised limit order books (LOBs). The availability of historical data have led to extensive research modelling LOBs. Better understanding the dynamics of LOBs and building simulators as a framework for controlled experiments, when testing trading algorithms or execution strategies are among the aims in this area. Most work in the literature models the aggregate view of the limit order book, which focuses on the volume of orders at a given price level using a point process. In addition to this aggregate view, brokers and exchanges also have information on the identity of the agents submitting the order to them. This leads to a more complicated representation of limit order book dynamics, which we attempt to model using a heterogeneous model of order flow.

We present a granular representation of the limit order book, that allows to account for the origins of different orders. Using client order flow from a large broker, we analyze the properties of variables in this representation. The heterogeneity of the order flow is modeled by segmenting clients into different *clusters*, for which we identify representative prototypes. This segmentation appears to be stable both over time, as well as over different stocks. Our findings can be leveraged to build more realistic order flow models that account for the diversity of market participants.

*Corresponding author: prenzel@maths.ox.ac.uk
†Opinions expressed in this paper are those of the authors, and do not necessarily reflect the view of JP Morgan.

Applications of Hawkes Processes in Finance and Insurance

Anatoliy Swishchuk
University of Calgary
2500 University Drive NW
Calgary, Alberta, Canada T2N 1N4
E-mail: aswish@ucalgary.ca
Tel.: +1(403)220-3274

This talk is devoted to stochastic modelling and analyzing of big data arising in finance and insurance, including models with Hawkes processes (HP). We first review some existing stochastic models, such as Markov. semi-Markov ones, etc., and then describe new models based on HP. They include so-called general compound Hawkes processes (GCHP) for modelling the dynamics of limit order books (LOBs), and risk processes with claim arrivals based on GCHP. In finance, quantitative and comparative analyses are performed to find out which model is the best in describing the real dynamics of different LOBs. In insurance, we derive net profit condition and calculate finite time and infinite time ruin probabilities, and we study an optimal investment strategies for an insurer in an incomplete market. We use real data from finance and insurance to justify and implement our results. As a generalization of our one-dimensional results, we also consider multivariate general compound Hawkes processes describing the dynamics of the mid-prices of many stocks. We will talk about multidimensional risk process based on GCHP as well to model a capital for several insurance companies.

Arbitrage-Free Implied Volatility Surface Generation with Variational Autoencoders*

Brian (Xin) Ning†, Sebastian Jaimungal†, Xiaorong Zhang†, and Maxime Bergeron‡

Modelling implied volatility (IV) surfaces in a manner that reflects historical dynamics while remaining arbitrage-free is a challenging open problem in finance. There are numerous approaches driven by stochastic differential equations (SDEs) that aim to do just so, including local volatility models [8], stochastic volatility models [11, 10], stochastic local volatility models [15], jump-diffusion models [7], and regime switching models [3], among many others. Such approaches make specific assumptions on the dynamics of the underlying asset and a choice of an equivalent martingale measure in order to avoid arbitrage. While these assumptions are not necessarily dynamically consistent with historical data, they do allow, e.g., pricing exotic derivatives via Monte Carlo or PDE methods.

An alternative to the SDE approach is to use non-parametric models to approximate IV surfaces directly without making assumptions on the underlying dynamics. For example, ML models such as support vector machines (SVMs) have been used to model such surfaces [17]. The issue of ensuring arbitrage-free surfaces is often tackled jointly during model fitting [2] either through penalisation of arbitrage constraints [1] or by directly encoding them into the network architecture [18]. These approaches, however, typically do not provide any guarantees and may not be arbitrage-free across the entire surface. A recent intriguing approach [6] is to reduce surfaces to arbitrage-free 'factors' – learned, e.g., through principal component analysis (PCA) – which can then be modeled using neural SDEs [14]. This approach, while very promising, relies on the quality of the 'factors' which are often complicated to compute. Another recent approach is that of [5] where the authors use Gaussian processes under shape constraints to generate surfaces and illustrate good fits to S&P data. Here, however, we are interested in the setting of sparse FX data and in generating the distribution over surfaces in a manner that is consistent with the historical data. The construction of arbitrage-free models based on ML approaches for stochastic interest rates has been tackled in [13]. In contrast, our focus is on European options and, more specifically, our application setting is to FX options.

In this paper, we develop a hybrid approach to resolve these issues by using SDE models that are by construction arbitrage-free yet flexible enough to fit arbitrary IV surfaces. One immediate dividend of this approach lies in its ability to produce realistic synthetic training data that can be used to leverage deep learning pricing methods in downstream tasks [9, 12]. The class of SDE models we consider include several time-varying regime switching models and Lévy additive processes which are detailed in the paper. We avoid overfitting by incorporating a Wasserstein penalty to keep the SDE model's risk-neutral density from deviating too far from the candidate one. The SDE model parameters, once fitted to data, represent a parameter subspace reflecting the features embedded in the data. The distribution on the subspace depends on the characteristics of the underlying asset and can be complex. We "learn" this distribution by using Variational Autoencoders (VAEs) which also allows for disentanglement of the subspace in an interpretable manner. SDE model parameters may be generated from the VAE model and used to create IV surfaces that are both faithful to the historical data but also strictly risk-neutral. This is similar in spirit, but distinct from, the tangent Lévy model approach introduced in [4] where a Lévy density is used to generate arbitrage-free prices while, here, the VAE generates parameters of the SDE model.

*The authors thank Ivan Sergienko for his comments on earlier versions of this work. S.J. acknowledges the support of the Natural Sciences & Engineering Research council of Canada [ALLRP 550308 - 20].

†Department of Statistical Sciences, University of Toronto (brian.ning@mail.utoronto.ca, sebastian.jaimungal@utoronto.ca; http://sebastian.statistics.utoronto.ca, xiaorong.zhang@mail.utoronto.ca)

‡Riskfuel Analytics (mb@riskfuel.com; http://riskfuel.com)

The overall approach may be summarised as: (i) fit a rich arbitrage-free SDE model to historical market data to obtain a collection of parameters, (ii) train a generative model, in particular a VAE model, on the collection of SDE model parameters, (iii) sample from the latent space of the generative VAE model, (iv) decode the samples to obtain a collection of SDE model parameters, and (v) use said SDE model and parameters to obtain arbitrage-free surfaces faithful to the historical data. We further refine the VAE model by including conditioning features into the encoding and decoding architectures. This results in a conditional VAE (CVAE) model, first introduced in [16] in a very different setting, for the arbitrage-free model parameter embeddings. We find that the CVAE model outperforms all others when comparing out-of-sample performance.

Our contributions may be summarized as follows: (i) we propose a hybrid approach for generating arbitrage-free synthetic IV surfaces by first calibrating SDE model parameters to historical data and then training a rich VAE model to learn the distribution on the space of SDE model parameters; (ii) we demonstrate that the distribution of IV surfaces from the VAE model is capable of generating surfaces as close to the testing data's distribution as the training set itself, and validate its empirical performance in comparison with several benchmarks; (iii) we extend this approach by including conditional features into the encoding and decoding architectures via the use of CVAE; and (iv) we demonstrate that the CVAE model significantly outperforms all other methods when comparing out-of-sample performance.

REFERENCES

[1] D. Ackerer, N. Tagasovska, and T. Vatter, *Deep smoothing of the implied volatility surface*, in Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020), 2020.

[2] M. Bergeron, N. Fung, Z. Poulos, J. C. Hull, and A. Veneris, *Variational autoencoders: A hands-off approach to volatility*, Available at SSRN 3827447, (2021).

[3] J. Buffington and R. J. Elliott, *Regime switching and european options*, in Stochastic Theory and Control, Springer, 2002, pp. 73–82.

[4] R. Carmona and S. Nadtochiy, *Tangent lévy market models*, Finance and Stochastics, 16 (2012), pp. 63–104.

[5] M. Chataigner, A. Cousin, S. Crépey, M. Dixon, and D. Gueye, *Short communication: Beyond surrogate modeling: Learning the local volatility via shape constraints*, SIAM Journal on Financial Mathematics, 12 (2021), pp. SC58–SC69, https://doi.org/10.1137/20M1381538, https://doi.org/10.1137/20M1381538, https://arxiv.org/abs/https://doi.org/10.1137/20M1381538.

[6] S. N. Cohen, C. Reisinger, and S. Wang, *Arbitrage-free neural-sde market models*, arXiv e-prints, (2021), pp. arXiv–2105.

[7] R. Cont and P. Tankov, *Calibration of jump-diffusion option pricing models: a robust non-parametric approach*, https://ssrn.com/abstract=332400, (2002).

[8] B. Dupire et al., *Pricing with a smile*, Risk, 7 (1994), pp. 18–20.

[9] R. Ferguson and A. Green, *Deeply learning derivatives*, arXiv preprint arXiv:1809.02233, (2018).

[10] P. S. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward, *Managing smile risk*, Wilmott Magazine, 1 (2002), pp. 249–296.

[11] S. L. Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, The review of financial studies, 6 (1993), pp. 327–343.

[12] B. Horvath, A. Muguruza, and M. Tomas, *Deep learning volatility*, arXiv preprint arXiv:1901.09647, (2019).

[13] A. Kratsios and C. Hyndman, *Deep arbitrage-free learning in a generalized hjm framework via arbitrage-regularization*, Risks, 8 (2020), p. 40.

[14] X. Li, T.-K. L. Wong, R. T. Chen, and D. Duvenaud, *Scalable gradients for stochastic differential equations*, in International Conference on Artificial Intelligence and Statistics, PMLR, 2020, pp. 3870–3882.

[15] K. Said, *Pricing exotics under the smile*, RISK, 12 (1999), pp. 72–75.

[16] K. Sohn, H. Lee, and X. Yan, *Learning structured output representation using deep conditional generative models*, Advances in neural information processing systems, 28 (2015), pp. 3483–3491.

[17] Y. Zeng and D. Klabjan, *Online adaptive machine learning based algorithm for implied volatility surface modeling*, Knowledge-Based Systems, 163 (2019), pp. 376–391.

[18] Y. Zheng, Y. Yang, and B. Chen, *Gated deep neural networks for implied volatility surfaces*, arXiv preprint arXiv:1904.12834, (2019).

March 1, 2022

ASSET PRICING AND CORPORATE GOVERNANCE

ARASH FAHIM

Florida State University, Department of Mathematics

VIJAY R. KRISHNA

Florida State University, Department of Economics

In this paper, we establish a asset pricing model for the stock of a firm from its capital structure, which is determined by principal-agent framework similar to DeMarzo and Sannikov [5] or Biais et al. [3]. A principal (shareholders of a company) hire an agent to operate their firm and to report back the cashflow. In return, the principal pays the agent according to a contract set at the time of hiring. The payment process is adapted to the filtration generated by the reported cashflow and not the actual cashflow. The principal does not have information to determine whether the agent reports truthfully or not, and his goal is to find a contract that maximizes his discounted cashflow from the company, subject to individual rationality constraint. As seen in DeMarzo and Sannikov [5] or Biais et al. [3], implementing of an optimal contract dictates a capital structure to the firm and determine the price of the shares. What is different in our paper is that we allow the principal to obtain more accurate information about the action of the agent through running a costly audit process, for example, by reducing the magnitude of noise term in the cashflow process. We show that some features of the price process of the asset of the firm are determined by the capital structure coming from an optimal solution to the principal-agent problem. Our results theoretically justify the empirical evidence that efficient corporate governance by shareholders have positive impact on the stock performance. See for example, Gompers et al. [6].

Our result uses the comparison principle for fully nonlinear partial differential equations through the theory of viscosity solutions. Although we show that the value function is C^2, due to existence of a free boundary, comparison principle is easier to establish through viscosity solution theory. The Hamilton-Jacobi-Bellman equation comes from a mixed singular-continuous control problem, which has an unusual feature that allows us to readily show the concavity of any supersolution. In particular, the concavity of the value function allows us to simplify the HJB equation and to modify arguments in Davis et al. [4], Alvarez [1], and Benth et al. [2] to establish proper comparison principle for the simplified PDE.

The C^2-solution of the PDE yields an optimal contract and an optimal governance strategy for the principal, which is fully determined by the second derivative of the value function. The optimal compensation of the agent is a local time at a

E-mail addresses: afahim@fsu.edu, vkrishna@fsu.edu.

point, i.e., payment boundary. Through comparison principle, we study sensitivity analysis of the value function, governance strategy, and payment boundary with respect to the model parameters.

Since the implementation of optimal contract includes compensating the agent with some dividend-paying shares of the firm, the price of the shares of the company have cash value of the contract for the agent. Another result of this paper shows that optimal contracts shapes the volatility of the stock, the distribution of default, and dividend payment strategy. In other words, this study provides some clue on how capital structure of a company is related to the market risk and credit risk of the stock.

References

[1] O Alvarez. A singular stochastic control problem in an unbounded domain. *Communications in Partial Differential Equations*, 19(11-12):2075–2089, 1994.

[2] Fred Espen Benth, Kenneth Hvistendahl Karlsen, and Kristin Reikvam. Optimal portfolio selection with consumption and nonlinear integro-differential equations with gradient constraint: a viscosity solution approach. *Finance and Stochastics*, 5(3):275–303, 2001.

[3] Bruno Biais, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet. Dynamic security design: Convergence to continuous time and asset pricing implications. *The Review of Economic Studies*, 74(2):345–390, 2007.

[4] Mark HA Davis, Vassilios G Panas, and Thaleia Zariphopoulou. European option pricing with transaction costs. *SIAM Journal on Control and Optimization*, 31(2):470–493, 1993.

[5] Peter M DeMarzo and Yuliy Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *The Journal of Finance*, 61(6):2681–2724, 2006.

[6] Paul Gompers, Joy Ishii, and Andrew Metrick. Corporate governance and equity prices. *The quarterly journal of economics*, 118(1):107–156, 2003.

ASSET SELECTION VIA CORRELATION BLOCKMODEL CLUSTERING

WENPIN TANG, XIAO XU, AND XUN YU ZHOU

ABSTRACT. We aim to cluster financial assets in order to identify a small set of stocks to approximate the level of diversification of the whole universe of stocks. We develop a data-driven approach to clustering based on a correlation blockmodel, in which assets in the same cluster are highly correlated with each other and, at the same time, have the same correlations with all other assets. We devise an algorithm to detect the clusters, with theoretical analysis and practical guidance. Finally, we conduct an empirical analysis to verify the performance of the algorithm.

1. Introduction: Selecting Representative Assets

The modern portfolio theory was pioneered by Markowitz (1952, 1959), in which the key insights are diversification and risk-return tradeoff. One drawback of applying Markowitz's mean-variance portfolio selection approach naïvely is to include *all* the available assets for allocation, which is difficult for small investors or small fund managers and creates technical issues of increased overfitting and difficulty in computing efficient allocation strategies (e.g., DeMiguel et al., 2009).

In this work, we focus on the following natural question: How can we select a *much* smaller subset of the whole stock universe that achieves a *sufficient* level of diversification? In other words, we can decompose the Markowitz model into two stages: asset *selection* and asset *allocation*. Specifically, one first groups or clusters all the assets in a correlation network and then selects one or a few "representative" assets in each group, resulting in a subset of a much smaller number of assets. Asset allocation can then be performed on these representative assets to construct a portfolio. In this work, we study a correlation network (Mantegna, 1999) between stocks and propose a clustering algorithm that groups stocks based on our two criteria, rooted in the problem of asset selection. We provide rigorous theoretical analysis which validates the algorithm and also guides its practical application. Finally, we conduct empirical analysis to verify the performance of the algorithm in portfolio construction.

2. Methods: Two Criteria and the ACC Algorithm

Most of the existing clustering methods applied to portfolio selection are heuristics, hence often difficult to interpret (see Marathe and Shawky, 1999; Das, 2003; Kaufman and Rousseeuw, 1990 with k-means and k-medoids algorithms, Gavrilov et al., 2000 with hierarchical clustering, and Mantegna, 1999; Tumminello et al., 2005 with network filtering). In this work, we propose a new, interpretable, data-driven approach to correlation network clustering method and provide a systematic solution for selecting well-diversified stocks. Our clustering method is based on the following two criteria:

Criterion 1. *Financial assets in the same group have high correlations.*

Criterion 2. *Financial assets in the same group have similar correlations with all other assets.*

Criterion 1 is self-evident: Assets with high correlations may rise and fall simultaneously and should be clustered into the same group. Through Criterion 2, we propose to also cluster in such a way that any two assets in the same group have similar correlations to *all* others in the stock universe. This way, any two assets in the same group are *interchangeable* in terms of their correlations with other assets, and one needs to only choose some idiosyncratic characteristics to select which asset to be included in the portfolio. This makes the choice of representative assets from each cluster simple and transparent.

To develop a new financial clustering approach, taking both criteria into account, we propose a *correlation blockmodel* to capture Criterion 2. This formulation is inspired by the problem of community detection in stochastic blockmodel (Abbe, 2017) and block covariance model (Bunea et al., 2016, 2020). Criterion 1 is then used to calibrate a threshold hyperparameter that controls how variables are grouped together. We devise an algorithm – called ACC (Asset Clustering through Correlation) – to recover the clusters of the blockmodel in polynomial time. To our best knowledge, we are the first to implement both criteria (especially Criterion 2) in financial asset clustering to capture the notion of diversification and the first to utilize the correlation blockmodel to formalize the implementation. This provides interpretability of our clustering approach from the portfolio theory point of view. Our theoretical analysis provides a statistical guarantee for the algorithm which can account for the possible heavy-tailed data intrinsic to financial time series. The information limit from our theoretical results narrows down the search for the hyperparameter, while Criterion 1 is used to cross-validate.

3. Empirical Results: Simple and Performant Portfolio Selection

We conduct an extensive empirical study on the S&P 500 stocks by selecting 15 to 25 stocks at a time via clustering and constructing portfolios using the selected stocks. For comparison, we select stocks from clusters created by the popular k-medoids clustering algorithm and clusters based on S&P's sector and industry classification.

We also consider the set of all S&P 500 sector ETFs, each of which represents a different sector in the S&P 500 Index. For all these groups of stocks, we employ and compare three asset allocation strategies: risk parity, minimum-variance, and Markowitz's mean-variance optimal allocation. The results show that the portfolios constructed using our ACC algorithm outperform the benchmark – the S&P 500 ETF – significantly. The portfolios based on ACC clusters also perform favorably compared to all other portfolios, especially when portfolios are readjusted infrequently.

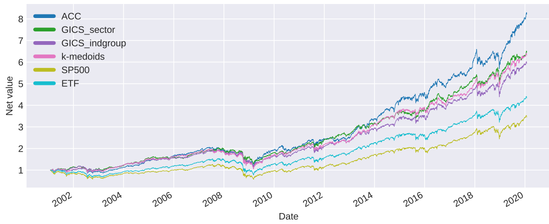

Figure 1. Portfolio from ACC algorithm outperforms other clustering algorithms

4. Conclusion

This work aims to identify a smaller set of stocks that attains an adequate level of diversification compared to the whole universe of stocks. We achieve this by clustering financial assets via exploring the correlation structure. We cluster the assets in a group according to the joint correlation with all other assets. The idea is formalized by the correlation blockmodel, and the ACC algorithm is devised to cluster the model. We conduct rigorous analysis of the ACC algorithm and give practical guidance based on the theoretical results. Numerical experiments show that portfolios constructed based on the ACC clustering algorithm achieve good performance compared to the market benchmark and also other portfolios consisting of similar numbers of assets.

References

E. Abbe. Community detection and stochastic block models: Recent developments. *J. Mach. Learn. Res.*, 18:Paper No. 177, 86, 2017.

F. Bunea, C. Giraud, and X. Luo. Minimax optimal variable clustering in G-models via Cord. *arXiv:1508.01939v2*, 2016. URL `https://arxiv.org/abs/1508.01939v2`.

F. Bunea, C. Giraud, X. Luo, M. Royer, and N. Verzelen. Model assisted variable clustering: Minimax-optimal recovery and algorithms. *Ann. Statist.*, 48(1):111–137, 2020.

N. Das. Hedge fund classification using k-means clustering method. In *9th International Conference on Computing in Economics and Finance*, pages 1–27, 2003. URL `https://depts.washington.edu/sce2003/Papers/284.pdf`.

V. DeMiguel, L. Garlappi, and R. Uppal. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Rev. Financ. Stud.*, 22(5):1915–1953, 2009.

M. Gavrilov, D. Anguelov, P. Indyk, and R. Motwani. Mining the stock market: Which measure is best. In *Proceedings of the 6th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 487–496, 2000.

L. Kaufman and P. J. Rousseeuw. *Partitioning Around Medoids (Program PAM)*, chapter 2, pages 68–125. John Wiley & Sons, Ltd, 1990.

R. N. Mantegna. Hierarchical structure in financial markets. *Eur. Phys. J. B*, 11(1):193–197, 1999.

A. Marathe and H. A. Shawky. Categorizing mutual funds using clusters. *Advances in Quantitative analysis of Finance and Accounting*, 7(1):199–204, 1999.

H. M. Markowitz. Portfolio selection. *J. Finance*, 7(1):77–91, 1952.

H. M. Markowitz. *Portfolio selection: Efficient diversification of investments.* Cowles Foundation for Research, Monograph 16. John Wiley & Sons, Inc., New York, 1959.

M. Tumminello, T. Aste, T. Di Matteo, and R. N. Mantegna. A tool for filtering information in complex systems. *Proc. Natl. Acad. Sci. U.S.A.*, 102(30):10421–10426, 2005.

ASYMPTOTIC ANALYSIS OF LONG-TERM INVESTMENT WITH TWO ILLIQUID AND CORRELATED ASSETS

XINFU CHEN MIN DAI WEI JIANG CONG QIN

Abstract. We use asymptotic analysis to study a long-run portfolio choice problem with two illiquid and correlated assets. Different from most of extant literature in this field, we perform asymptotic expansion in terms of correlation coefficient. Our expansion provides an efficient way to compute the optimal trading boundaries that characterize optimal strategy. Moreover, the leading order expansion implies that the optimal trading boundaries are orthogonal to each other at four corners where two illiquid assets may be traded simultaneously. This orthogonal property suggests that the optimal trading boundaries have certain regularity even at corners.

Keywords: Transaction costs, Asymptotic expansion, Multiple assets

AMS Classification: 91G10, 93E20, 41A60

Department of Mathematics,University of Pittsburgh. (xfc@pitt.edu). Xinfu Chen acknowledges support from NSF grant DMS-1008905.

Department of Mathematics and Risk Management Institute, National University of Singapore (mindai@nus.edu.sg).

Department of Industrial Engineering and Decision Analytics, The Hong Kong University of Science and Technology (weijiang@ust.hk).

Center for Financial Engineering, Soochow University (congqin@suda.edu.cn). Qin acknowledges support from the Start-up Fund at Soochow University.

Title: **Axioms and Properties of Automated Market Makers**
Authors: Maxim Bichuch & Zachary Feinstein

Automated market makers [AMMs] are a decentralized approach to creating a financial market. In contrast to a centralized exchange, AMMs are able to operate without the need for a trust-based relationship and they do not require custodial services. This is vital due to several high-profile exploits at centralized exchanges (e.g., the collapse of the bitcoin exchange Mt. Gox in 2014 and the hacks of CoinCheck and KuCoin). By virtue of being decentralized, AMMs never hold the investors' assets, and therefore are not susceptible to a centralized asset theft. However, AMMs operate by using mathematical formulas to balance reserves to automatically quote market prices and execute swaps, and therefore are susceptible to malicious algorithmic manipulations. Beyond the simple structure of AMMs, a primary benefit is that it allows individuals to pool their assets to become liquidity providers and collect the proceeds associated with serving that vital market function. This democratizes the market making process and supports decentralized finance.

The mathematical structures for AMMs are, currently, based upon an invariance relation for a function – called constant function market makers [CFMMs]. These CFMMs are able to adjust prices and present price impacts for any traders; market efficiency and the law of one price is guaranteed by the actions of arbitrageurs who take advantage of different exchange rates offered by different AMMs and exchanges. Within this talk, we will consider an axiomatic construction of CFMMs so as to provide mathematical properties that all market makers should satisfy. These properties include: no round-trip arbitrage, monotonicity, decreasing marginal returns, infinite liquidity, and marginal invariance and stability to liquidity pooling. This allows for detailed understanding of how the construction of a CFMM impacts both investors and liquidity providers. These properties are compared with AMMs that exist in practice.

Beta and Coskewness Pricing: Perspective from Probability Weighting

Yun Shi, Xiangyu Cui, and Xunyu Zhou

The security market line is found to be often flat or downward-slopping. Our hypothesis is that probability weighting plays a key role and one ought to differentiate the periods when agents overweight both tails of the market return from those when they underweight the two tails. During an overweighting period, agents are more concerned about the extremely bad events and hence demand more risk compensation for taking the beta risk. During an underweighting period, agents deflate the probability of extremely bad events and become less risk-averse or even become risk-loving. Unconditional on probability weighting, these two effects offset each other, resulting in a flat or slightly negative return-beta relationship. We can similarly use probability weighting to investigate coskewness pricing. Overweighting the tails enhances the negative relation between return and coskewness, while underweighting the tails reduces such relation. We support our theory by an extensive empirical study.

Beyond Smart Beta: A Dynamic Statistical Risk Budgeting Approach in Portfolio Construction

Financial markets turmoil over the recent past and uncertainty about future asset returns has seen the emergence of investment strategies that disregard expected returns and focus on risk diversification, such as the well-established risk parity concept and more recently Smart Beta strategies. We present an innovative statistical approach, relying on a factor model that dynamically identify the current risk factors driving asset returns and decompose the portfolio risk into a systematic and specific component.

We propose a dynamic investment strategy allowing investors to diversify away unrewarded risks without restricting their preference of factor risk exposure. Our results display an outperformance of this dynamic statistical framework over traditional asset-based risk budgeting strategies.

Keywords: risk-budgeting, risk diversification, statistical factors, smart beta strategies.

Key messages:

- Dynamic risk-budgeting approach
- Flexible framework allowing customized factor exposure
- Diversify away unrewarded risks
- Outperformance of statistical factor-based strategies over traditional asset-based risk budgeting strategies.

Extended Abstract

Bitcoin miners compete to validate users' transactions via a proof-of-work mechanism using computing power. Thanks to their efforts, miners are compensated with a predetermined number of newly minted Bitcoins plus transaction fees whenever the verified users' transactions are successfully added to the blockchain.

One of the main concerns about Bitcoin mining is its electricity consumption. Many researchers have reported that nowadays, the electricity consumed annually by Bitcoin mining is even higher than that of medium-sized countries. The trend of electricity consumption is even more worrisome. As we have witnessed a tremendous increase in Bitcoin price and mining revenue in terms of the US dollar in the past decade, substantial new miners are lured into joining the mining business. Consequently, both the electricity consumption and the Bitcoin computing power measured by hashrate have grown up rapidly.

There is relatively little literature on the economic modeling of computing power in the Bitcoin network. An exception is Prat and Walter (2021), in which a novel industry equilibrium model is developed by first introducing technology innovation to the model in Caballero and Pindyck (1996) and then calibrating the model to capture the evolution of miners' computing power. However, they mainly focus on examining the entry rule of the miners, and the hashrate in their primary model cannot decrease over time. In addition, for an entry-only model, the resulting electricity consumption must be non-decreasing, and the ratio of electricity consumption to revenue may grow to infinity.

We propose a dynamic industry equilibrium model for Bitcoin electricity consumption in a general framework, including Bitcoin miners' optimal entry and exit with technology innovation. By adopting average operating costs as an approximation to the actual operating costs, we overcome the difficulty of strong path-dependency due to the interaction among entry, exit, and technology innovation. We also formulate the problem from a social planner's perspective and show that his optimal strategy is an industry equilibrium under a competitive market. The penalty method is then applied to find the optimal strategy numerically.

By calibrating the model, we can capture both the upside and downside co-movements of miners' computing power, electricity consumption, and mining revenue. Our model shows that the Bitcoin electricity consumption will not grow indefinitely, with the ratio of Bitcoin electricity consumption to the miners' revenue fluctuating within a range.

CAPITAL DISTRIBUTION CURVES IN POLYNOMIAL SPT MODELS

CHRISTA CUCHIERO, † FLORIAN HUBER †

Abstract

We introduce polynomial processes, in the sense of [1] in the context of stochastic portfolio theory to model the evolution of companies' market capitalizations. Extending the volatility stabilized market models studied by Fernholz and co-authors in [2], we allow correlation between individual stocks via a multiplicative common noise. We investigate the behaviour as the number of companies tends to infinity and show that the empirical measure of the system converges to the solution of a degenerate, non-linear stochastic partial differential equation

$$\mathrm{d}\rho(t) = \frac{\alpha}{2}\langle\rho(t), x\rangle\partial_x^2(x\rho) + \frac{1}{2}\partial_x^2(x^2\rho) - \langle\rho(t), x\rangle\partial_x\rho\,\mathrm{d}t + \partial_x(x\rho)\,\mathrm{d}W_t^0.$$

We obtain that $\rho = \mathbb{P}(Y_t \in \cdot|\sigma(W_s^0, s \leq t))$, with $(Y_t)_{t\geq 0}$ solving the (conditional) McKean Vlasov SDE

$$\mathrm{d}Y_t = \beta\mathbb{E}[Y_t|\sigma(W_s^0, s \leq t)]\,\mathrm{d}t + \sqrt{\alpha Y_t \mathbb{E}[Y_t|\sigma(W_s^0, s \leq t)]}\,\mathrm{d}B_t + Y_t\,\mathrm{d}W_t^0,$$

where $(B_t)_{t\geq 0}$ is a one dimensional Brownian motion independent of $(W_t^0)_{t\geq 0}$ and observe that the pair $(Y_t, \mathbb{E}[Y_t|\sigma(W_s^0, s \leq t)])_{t\geq 0}$ is a joint polynomial process.

† University of Vienna, Department of Statistics and Operations Research, Kolingasse 14-16, Vienna , Austria .

1

References

[1] Christa Cuchiero, Martin Keller-Ressel, and Josef Teichmann. Polynomial processes and their applications to mathematical finance. *Finance and Stochastics*, 16(4):711–740, 2012.

[2] Robert Fernholz and Ioannis Karatzas. Relative arbitrage in volatility-stabilized markets. *Annals of Finance*, 1(2):149–177, 2005.

CBI-TIME-CHANGED LÉVY PROCESSES FOR MULTI-CURRENCY MODELING

CLAUDIO FONTANA[†], ALESSANDRO GNOATTO[‡], AND GUILLAUME SZULDA[†,*]

The Foreign-Exchange (FX) market is one of the largest in the world (see e.g. [fIS19, Woo19]). From the perspective of quantitative finance, modeling the FX market poses several challenges. First, multi-currency models must respect the symmetric structure of FX rates. To illustrate this aspect, let $S^{d,f}$ represent the value of one unit of a foreign currency f measured in units of the domestic currency d. In a multi-currency model, the following symmetric relations must hold:

- $S^{f,d} = 1/S^{d,f}$: the reciprocal of $S^{d,f}$ must coincide with $S^{f,d}$, representing the value of one unit of currency d measured in units of currency f. This is referred to as *inversion*;
- $S^{d,f} = S^{d,e} \times S^{e,f}$, for any other foreign currency e. In other words, the FX rate $S^{d,f}$ must be inferred from $S^{d,e}$ and $S^{e,f}$ through multiplication. This is referred to as *triangulation*.

Besides these symmetric relations, the FX market presents some specific risk characteristics that should be properly reflected in a multi-currency model. First, FX markets are affected by stochastic volatility and jump risk, similarly to equity markets. Second, the dependence between FX rates is typically stochastic and, in particular, shows evidence of unpredictable changes over time, thus generating correlation risk. Third, the skew of the FX volatility smile exhibits a stochastic behavior. This fact has been documented in [CW07] by analyzing the time series of risk-reversals, showing that their values vary significantly over time and exhibit repeated sign changes.

In this paper (see [FGS21a]), we develop a modeling framework for multiple currencies that is consistent with the symmetric structure of FX rates and captures all risk characteristics described above, including stochastic dependence among FX rates as well as between FX rates and their respective volatility. We consider models driven by *CBI-time-changed Lévy processes* (CBITCL processes), a broad and flexible class of stochastic processes that allows for self-exciting jump dynamics, together with stochastic volatility and mean-reversion (we refer the reader to [FGS22, Szu21] for a thorough analysis of CBITCL processes).

[†]Department of Mathematics "Tullio Levi Civita", University of Padova (Italy)

[‡]Department of Economics, University of Verona (Italy)

[*]Laboratoire de Probabilités, Statistique et Modélisation, Université de Paris (France)

E-mail addresses: fontana@math.unipd.it, alessandro.gnoatto@univr.it, szulda@math.unipd.it.

Date: February 28, 2022.

Key words and phrases. FX market; multi-currency market; branching process; self-exciting process; time change; stochastic volatility; deep calibration; affine process.

Acknowledgements: C.F. is grateful to the Europlace Institute of Finance for financial support to this work. G.S. acknowledges hospitality and financial support from the University of Verona, where part of this work has been conducted. This work is part of the project BIRD 2019 "Term structure dynamics in interest rate and energy markets: modelling and numerics" funded by the University of Padova.

The proposed approach is fully analytically tractable, due to the fact that CBITCL processes are affine processes and, therefore, their characteristic function can be explicitly characterized. Moreover, CBITCL processes are *coherent* in the sense of [Gno17], meaning that if an FX rate is modeled by a CBITCL process, then its reciprocal also belongs to the same model class.

We construct our modeling framework by adopting an *artificial currency approach*, which consists in modeling each FX rate as the ratio of two primitive processes, with one primitive process associated to each currency. FX rates then satisfy the inversion and triangulation symmetries by construction, and the model formulation reduces to modeling all primitive processes by means of a common family of CBITCL processes.

By relying on a Girsanov-type result for CBITCL processes, we characterize a class of risk-neutral measures leaving invariant the structure of the model. In particular, this allows preserving the CBITCL property under equivalent changes of probability, permitting us to derive an efficient pricing formula for currency options by means of Fourier techniques.

We analyze the empirical performance of a two-dimensional specification of our framework, driven by tempered-stable CBI processes (as recently introduced in [FGS21b]) and CGMY processes (see [CGMY02]). We perform a calibration of the model with respect to an FX triangle consisting of three major currency pairs (USD-JPY, EUR-JPY, EUR-USD). We propose two calibration methods: a *standard* calibration algorithm and a *deep* calibration algorithm, inspired by the deep learning techniques recently developed in [HMT21]. In particular, we mention that deep calibration is here applied for the first time in a multi-currency setting.

References

[CGMY02] P. Carr, H. Geman, D. B. Madan, and M. Yor. The fine structure of asset returns: an empirical investigation. *The Journal of Business*, 75(2):305–332, 2002.

[CW07] P. Carr and L. Wu. Stochastic skew in currency options. *Journal of Financial Economics*, 86(1):213–247, 2007.

[FGS21a] C. Fontana, A. Gnoatto, and G. Szulda. CBI-time-changed Lévy processes for multi-currency modeling. ArXiv working paper available at `https://arxiv.org/abs/2112.02440`, 2021.

[FGS21b] C. Fontana, A. Gnoatto, and G. Szulda. Multiple yield curve modeling with CBI processes. *Mathematics and Financial Economics*, 15(2):579–610, 2021.

[FGS22] C. Fontana, A. Gnoatto, and G. Szulda. CBI-time-changed Lévy processes. In preparation, 2022.

[fIS19] Bank for International Settlements. BIS Triennial Central Bank Survey: foreign exchange turnover in April 2019. Technical report, BIS, Monetary and Economic Department, April 2019.

[Gno17] A. Gnoatto. Coherent foreign exchange market models. *International Journal of Theoretical and Applied Finance*, 20(1):1750007, 2017.

[HMT21] B. Horvath, A. Muguruza, and M. Tomas. Deep learning volatility. *Quantitative Finance*, 21(1):11–27, 2021.

[Szu21] G. Szulda. *Branching Processes and Multiple Term Structure Modeling.* PhD thesis, Université de Paris, 2021.

[Woo19] P. Wooldridge. FX and OTC derivatives markets through the lens of the Triennial Survey. *BIS Quarterly Review*, December 2019.

Chaotic Hedging with Iterated Integrals

Ariel Neufeld* Philipp Schmocker†

February 25, 2022

Abstract

In this paper, we extend the Wiener-Ito chaos decomposition to the class of local martingale diffusions, whose diffusion coefficient is of linear growth. In addition, by omitting the orthogonality in the chaos expansion, we are able to show that every p-integrable functional, for $1 < p < \infty$, can be represented as sum of iterated or multiple integrals of the underlying process. Using a truncated sum of this expansion and neural networks for the integrands, whose parameters are learned in a machine learning setting, we show that every financial option can be approximated arbitrarily well. In this case, the hedging strategy of the approximating option can be calculated in closed form by using functional Ito calculus.

Summary

The Wiener-Ito chaos expansion can be used to express every square-integrable functional $G \in L^2(\Omega, \mathcal{F}_T, \mathbb{P})$ of Brownian motion $B = (B_t)_{t \in [0,T]}$ in terms of multiple integrals of B, where $\mathbb{F} = (\mathcal{F}_t)_{t \in [0,T]}$ denotes the filtration generated by B (see [5, Section 4]). More precisely, by following the notation of [3], let $(\mathbb{R}^d)^{\otimes n} = \mathbb{R}^d \otimes \cdots \otimes \mathbb{R}^d$ (with convention $(\mathbb{R}^d)^{\otimes 0} := \mathbb{R}$) and $S_n = \{(t_1, ..., t_n) \in [0,T]^n;\, 0 \leq t_1 \leq ... \leq t_n \leq T\}$, from which we define $L^2(S_n)$ as the space of Borel-measurable functions $g : S_n \to (\mathbb{R}^d)^{\otimes n}$ such that the norm

$$\|g\|_{L^2(S_n)} := \left(\sum_{i_1,...,i_n=1}^{d} \int_{S_n} g_{i_1,...,i_n}(t_1, ..., t_n)^2 dt_1 \cdots dt_n \right)^{\frac{1}{2}} \tag{1}$$

is finite. From this, we introduce the *iterated integral* as linear operator $J_n : L^2(S_n) \to L^2(\mathbb{P})$ defined by

$$J_n(g) = \sum_{i_1,...,i_n=1}^{d} \int_0^T \cdots \int_0^{t_2} g_{i_1,...,i_n}(t_1, ..., t_n)\, dB_{t_1}^{i_1} \cdots dB_{t_n}^{i_n},$$

for $g \in L^2(S_n)$, see [3, Definition 1.3]. Hereby, the orthogonality $\mathbb{E}\left[J_n(g) J_m(h)\right] = \delta_{nm} \langle g, h \rangle_{L^2(S_n)}$ holds true, where the inner product $\langle \cdot, \cdot \rangle_{L^2(S_n)}$ induces the norm in (1). In this case, there exists for every $G \in L^2(\Omega, \mathcal{F}_T, \mathbb{P})$ a sequence $(g_n)_{n \in \mathbb{N}_0}$, with $g_n \in L^2(S_n)$, for all $n \in \mathbb{N}$, such that

$$G = \sum_{n=0}^{\infty} J_n(g_n), \qquad \text{in } L^2(\mathbb{P}), \tag{2}$$

see [5, Theorem 4.2], whereas [3, Theorem 1.10] describes the version with multiple integrals.

In our paper, we extend the chaos expansion in (2) to continuous local martingale diffusions, whose diffusion coefficient is of linear growth, and to p-integrable functionals $G \in L^p(\Omega, \mathcal{F}_T, \mathbb{P})$,

*Nanyang Technological University, Division of Mathematical Sciences, 21 Nanyang Link, Singapore 637371, Singapore, ariel.neufeld@ntu.edu.sg

†Nanyang Technological University, Division of Mathematical Sciences, 21 Nanyang Link, Singapore 637371, Singapore, philippt001@e.ntu.edu.sg

for $1 < p < \infty$. As immediate consequence, we obtain a martingale representation theorem in this setting, which generalizes the well-known martingale representation result for Brownian motion (see e.g. [6, Theorem IV.43]).

For this purpose, we use the time-space harmonic Hermite polynomials $(H_n)_{n\in\mathbb{N}_0}$, which are obtained from the (probabilist's) Hermite polynomials $(h_n)_{n\in\mathbb{N}_0}$ defined by

$$h_n(x) = (-1)^n e^{\frac{1}{2}x^2} \frac{d^n}{dx^n} \left(e^{-\frac{1}{2}x^2} \right)$$

such that $H_n(x,t) = t^{\frac{n}{2}} h_n(x/\sqrt{t})$, for $n \in \mathbb{N}_0$ and $(x,t) \in \mathbb{R} \times [0,\infty[$. These time-space harmonic Hermite polynomials $(H_n)_{n\in\mathbb{N}_0}$ can be used to compute iterated integrals $J_n(g_0^{\otimes n})$ of a product function $g_0^{\otimes n} := g_0 \otimes \cdots \otimes g_0 : S_n \to (\mathbb{R}^d)^{\otimes n}$ with respect to the continuous local martingale diffusion. In the case of Brownian motion, this relationship between the Hermite polynomials and the normal distribution can even be extended to more general integrands $g \in L^2(S_n)$ satisfying an orthonormality condition, see [5, Theorem 3.1] with multiple integrals.

In this setting, we approximate a financial option $G \in L^p(\mathbb{P})$ by a truncated sum of iterated integrals, whose integrands are replaced by specific neural networks $(\varphi_0, ..., \varphi_N) \in \bigoplus_{n=0}^N \mathcal{NN}_{n,d}^\rho$, which are learned in a machine learning setting. Hereby, we obtain by using functional Ito calculus developed in [4, 1, 2] a closed-form expression for the hedging strategy of the approximating financial option $G^{(\varphi_0,...,\varphi_N)} \in L^p(\mathbb{P})$. In addition, by applying the Girsanov transformation, we extend the numerical application to semimartingales whose drift process is of linear growth.

References

[1] R. Cont and D.-A. Fournié. Change of variable formulas for non-anticipative functionals on path space. *Journal of Functional Analysis*, 259(4):1043–1072, 2010.

[2] R. Cont and D.-A. Fournié. Functional Ito calculus and stochastic integral representation of martingales. *Annals of Probability*, 41(1):109–133, 01 2013.

[3] G. Di Nunno, B. Øksendal, and F. Proske. *Malliavin Calculus for Lévy Processes with Applications to Finance*. Springer Science + Business Media, Berlin Heidelberg, 2008.

[4] B. Dupire. Functional Itô calculus. Technical report, Bloomberg, 2009. Bloomberg Portfolio Research Paper No. 2009-04-FRONTIERS.

[5] K. Ito. Multiple Wiener integral. *Journal of the Mathematical Society of Japan*, 3(1):157–169, 05 1951.

[6] P. Protter. *Stochastic integration and differential equations*, volume 21 of *Stochastic modelling and applied probability*. Springer, Berlin, 2nd ed., 3rd print edition, 2005.

Cheaper by the Bundle: The Interaction of Frictions and Option Exercise in Variable Annuities

Thorsten Moenig
Department of Risk, Insurance, and Healthcare Management; Temple University
1801 Liacouras Walk, Alter Hall 611, Philadelphia, PA 19122. USA

Daniel Bauer*
Department of Risk and Insurance; University of Wisconsin-Madison
975 University Avenue, Madison, WI 53706. USA

December 2019

Abstract

Variable Annuity (VA) providers sell complex baseline products at considerable fees, and then offer a variety of optional guarantees that further compound contract complexity. This is in contrast to mutual funds or investment-only VAs, where companies compete over low baseline fees. We argue that this is driven by benefits in bundling the baseline option features with the more advanced riders, to the extent that the baseline features commonly have *negative* marginal option values. This is possible due to market frictions, and particularly specific taxation rules for VA benefits, affecting policyholder exercise. In this paper, we document the relevance of this mechanism, both theoretically in a simple binomial model and empirically by evaluating existing VA guarantees in a more advanced continuous-time option pricing framework.

VAs are investment vehicles offered by U.S. life insurers since the 1970s. They are equipped with certain option features that can be exercised upon survival to a certain age or contract year—so-called Guaranteed Living Benefits (GLBs)—or upon death—so-called Guaranteed Minimum Death Benefits (GMDBs). The VA market is large. With nearly USD 2 trillion in net assets, VAs account for almost a quarter of the US insurance industry's total assets. And despite a recent downtrend in sales and GLB election rates, annual new premiums still well exceed $100bn, and over 75% of all new VAs include GLBs. In particular, various versions of relatively complex withdrawal guarantees—so-called Guaranteed Minimum Withdrawal Benefits (GMWBs) and Guaranteed Lifetime Withdrawal Benefits (GLWBs)—have been the most popular living benefits since the mid-2000s.

Most VA products include various option features at baseline, particularly a return-of-premium GMDB, financed via a so-called *mortality and expense* (M&E) fee that can be quite substantial. GLBs are then offered on top of the baseline contracts, with the

*Corresponding author. Phone: +1-608-265-4119. E-mail addresses: moenig@temple.edu (T. Moenig); daniel.bauer@wisc.edu (D. Bauer).

product details—including guarantee levels, step-up and surrender features, and associated rider fees—varying tremendously across issuers/products. This is distinct from other competing investment opportunities, where baseline products and fees are lean. In this paper, we investigate this aspect of VA products, particularly asking whether there are advantages to *bundling* the different option layers.

We focus on GMWB riders, one of the most popular class of VA riders. Under a GMWB rider policyholders have the right (but not the obligation) to make fixed periodic withdrawals irrespective of the VA's performance. We commence by analyzing a basic VA plus GMWB product in a two-period binomial setting, where the policyholder's payoff is subject to personal taxes. It is not surprising that the marginal option values in the bundle differ from their stand-alone values, which is also true in frictionless settings. What is surprising, however, is that the impact of the tax friction on optimal exercise behavior can have the perverse consequence that including the baseline GMDB will decrease the total option value from the perspective of the option writer (the insurance company). Withdrawals are costly for insurers as they move the option further in-the-money and reduce future fee income. Adding on the GMDB will avert or delay costly withdrawals, since withdrawing also decreases the guaranteed death benefit amount. In addition, withdrawing results in frictional costs (e.g. due to foregone tax benefits) to the policyholder, but not to the insurer. Therefore, the present value of the change in exercise behavior (to the insurer) can exceed the marginal value of adding on the death benefit option—leading to a *negative* marginal option value of the GMDB.

We evaluate the validity and relevance of this mechanism by analyzing eleven VA plus GMWB products offered in the marketplace, accounting for product differences and details in the contracts, in a generalized Black-Scholes model. These products roughly reflect the cross section of products in the market. We find that while adding on a GMDB rider to a VA+GWMB policy with a conventional proportional fee structure increases the policyholder's valuation by an average 1.7% of the investment amount, the insurer itself *profits* by nearly 1.9%—thus the marginal value of producing the option is considerably *negative*.[1] Notably, we obtain negative option values for all considered VA+GMWB products with a conventional fee structure, although a less common modification to the fee structure appears to invert the result. Our insights also do not extend to *lifetime* withdrawal guarantees.

Our findings help explain the prevalence and design of existing VA products, and they have implications for (tax) regulation. We document that accounting for policyholder incentives affects the cost of producing VA option features, to the extent that adding on a rider may reduce total value. Exploiting such situations may present a way to enhance product appeal and further differentiate a certain VA product without impeding the firm's bottom line.[2] At the same time, of course the benefits to policyholders and insurance companies will be financed by reduced tax revenue, and it is up to the regulator to decide to what extent firms economizing on tax advantages is desirable.

[1] This possibility of both option writers and buyers profiting at the expense of the Government when accounting for taxation was already raised by Myron Scholes (1976), one of the fathers of modern option pricing theory.

[2] While taxes and regulation are familiar drivers of financial innovation, typically the focus is on innovations that circumvent—rather than exploit—(tax) regulations.

Two financial markets should be considered equivalent if they contain the same investment opportunities at the same costs. We formalize this by defining a notion of isomorphism between markets, and hence defining the category of financial markets.

We prove a general theorem relating the symmetries of a market (i.e. the isomorphisms of that market to itself) to mutual fund theorems in that market.

Given a notion of isomorphism, one may attempt to classify objects up to isomorphism: that is to identify the isomorphism equivalence classes. One typically proceeds by defining invariants that remain unchanged under isomorphism, and one aims to show that these invariants are sufficiently powerful to identify the market. We will examine three classification theorems and the related mutual fund theorems.

First, the markets considered by Markowitz are classified by their efficient frontier which gives the basic invariant of a Markowitz market. The classical mutual fund theorems may be reproved by symmetry arguments. This gives a significant generalization of the theorems. Where the classical mutual fund theorem shows that one specific optimization problem may be solved using only convex combinations of two portfolios, the symmetry argument shows the same theorem applies to any invariantly defined problem.

Second, we consider one period complete markets where the key invariant is the distribution function of the Radon-Nikodym derivative of the pricing measure with respect to the physical measure. This is not a sufficiently powerful invariant to classify such markets. However, given a one period complete market we may extend it to a two-period market by allowing a trader to gamble their market winnings at a casino. The distribution function of the Radon-Nikodym derivative fully classifies the two period markets arising in this way. We will see how this allows one to reduce investment problems in one period complete markets to a calculus of variations problem and how this result may be interpreted as a mutual fund theorem.

Finally, we consider continuous time complete markets driven by diffusion processes, where the basic invariant is the absolute value of the market price of risk. We show that markets with a deterministic absolute value for the market price of risk are classified by this market price of risk function and dimension. This implies that they are isomorphic to Black-Scholes Merton markets. For example, Bachelier markets are isomorphic to Black-Scholes Merton markets and one can find many other examples such as local volatility models and stochastic volatility models which are isomorphic to Black-Scholes Merton markets. A corollary of this result is that all such markets have a large symmetry group and so admit a powerful mutual fund theorem that applies to all invariant convex investment problems for that market.

11^{th} World Congress of
The Bachelier Finance Society 2020

Computation of greeks in rough Volterra stochastic volatility models

Jan Pospíšil
University of West Bohemia, FAV NTIS, Univerzitní 2732/8, 301 00 Plzeň, Czech Republic
Email: honik@kma.zcu.cz

This is a joint work with

Josep Vives
Universitat de Barcelona, Facultat de Matemàtiques, Gran Via 585, 08007 Barcelona, Spain
Email: josep.vives@ub.edu

Abstract. In this talk we demonstrate how to use Malliavin calculus techniques to compute several financial derivatives sensitivities that are generally referred to as greeks. Models considered are rough volatility models, type of options are European. We find particular Malliavin weights for several rough Volterra models and numerically demonstrate the convergence.

Keywords. stochastic volatility models; rough volatility; option pricing
JEL classification. C63; G12
AMS classification. 60G22; 91G20; 91G60

Conditional Importance Sampling for Expectations over Convex Polyhedra

Dohyun Ahn and Lewen Zheng

We consider the problem of estimating the expectation over a convex polyhedron specified by a set of linear inequalities. This problem encompasses a multitude of financial applications including portfolio management, systemic risk quantification, and exotic option pricing. We particularly focus on the case where the target event is rare which corresponds to, for example, large portfolio losses, extreme systemic failures, and deep out of the money options in the aforementioned applications, respectively. This rare event setting renders the naive Monte Carlo method inefficient and requires variance reduction techniques. Assuming that random vectors follow normal variance mixture distributions, which include multivariate normal distribution and multivariate t distribution, we develop a novel conditional importance sampling (CIS) method for the computation of the said expectation by exploiting the geometry of the target polyhedron and concentrating the sampling density almost exclusively inside the target set. We prove the bounded relative error of the CIS estimator under a general rare event regime which embraces a wide range of asymptotic regimes considered in the literature. The proposed method significantly outperforms the existing approaches in various numerical experiments in terms of accuracy and computational burden.

Deep calibration for consistent recalibration models

Long Abstract - BFS, Hong Kong 2022

Matteo Gambara* & Josef Teichmann†

15th March 2022

Keywords – Consistent recalibration models, neural networks, implied volatility

1 Motivation

Calibration is an essential feature of financial models that aspire to be used in real life and, often, a model is chosen instead of another based not only on the possibility of reproducing the observed data, but also on the reliability and speed of the calibration process. Traditionally, the first step consists in specifying the dynamics of the underlying price process and then to select parameters and state variables in order to match liquid derivatives' prices daily observed in the market. While pursuing this approach, one will most of the time encounter at least the following two problems:

1. To reproduce all prices, we might need to add other parameters or to increase the dimension of the state space (e.g. adding hidden variables), but sometimes these can become redundant or meaningless.

2. After a certain amount of time, the model needs to be recalibrated, because the current configuration is not able to reproduce anymore observed data. Acceptable for state variables, but unsatisfactory for parameters (since these are assumed to be deterministic and constant).

2 Possible solutions

A possible fix to these problems is represented by consistent recalibration models as introduced by Richter and Teichmann in [3], which adopt a HJM philosophy, because a *codebook* is prescribed (deterministic 1-to-1 map between the reference asset and derivatives' prices), which instantaneously look like finite factor models and in which, thanks to a revision of Hull-White extension, parameters can be treated as state variables (able to vary over time). If we denote with X the log-return price process, the codebook, as already chosen in [1] and [2], consists in the so called *forward characteristic process* η, for which we generally assume a decomposition with respect to a d-dimensional semimartingale M: $\eta_t(u,T) = \eta_0(u,T) + \int_0^t \alpha_s(u,T)\,ds + \sum_{i=1}^d \int_0^t \beta_s^i(u,T)\,dM_s^i$. It can be proven the equivalence between the conditional expectation condition

$$\mathbb{E}\left[\exp(\mathrm{i}\,\langle u, X_t\rangle)\right|\mathcal{F}_s] = \exp\left(\mathrm{i}\,\langle u, X_s\rangle + \int_s^t \eta_s(u,r)\,dr\right) \tag{2.1}$$

and the following two: 1. Short-end condition: existence of predictable characteristics $\kappa_t^X(u) = \eta_{t-}(u,t)$ for $t \geq 0$ and $u \in \mathbb{R}^n$ and 2. HJM drift condition (ensuring the model is arbitrage-free): $\int_t^T \alpha_t(u,r)\,dr = \eta_{t-}(u,t) - \kappa_t^{(X,M)}\left(u, -\mathrm{i}\int_t^T \beta_t(u,r)\,dr\right)$.

*matteo.gambara@math.ethz.ch, ETH Zürich, Rämistrasse 101, Zürich.
†josef.teichmann@math.ethz.ch, ETH Zürich, Rämistrasse 101, Zürich.

3 Affine stochastic volatility models

It is possible to use the previous framework for affine stochastic volatility models. Let $C \subset \mathbb{R}^m$ be a proper convex cone, (X, Y) a homogenous affine process taking values in $\mathbb{R}^n \times C$, i.e. (X, Y) is a time-homogeneous Markov process relative to some filtration $(\mathcal{F}_t)$ and with state space $D = \mathbb{R}^n \times C$ such that it is stochastically continuous and its Fourier-Laplace transform has exponential affine dependence on the initial state. By the *Riccati ODEs*, we have $\phi(u,0,t) = \int_0^t F(u, \psi_C(u,0,s))ds$ and $\psi_C(u,0,t) = \int_0^t R_C(u, \psi_C(u,0,s))ds$ from which the characteristic process can be modified as

$$\eta_t(u,T) = F_T(\mathrm{i}u, \psi_C(\mathrm{i}u, 0, T-t)) + \langle R_C(\mathrm{i}u, \psi_C(\mathrm{i}u, 0, T-t)), Y_t \rangle, \tag{3.1}$$

where F is now time-dependent by adding a Lévy process to X (time-inhomogeneous).

3.1 An example: consistent recalibration for Bates model

If we denote with L the new Lévy process, then the equations are

$$\begin{aligned}
&dX(t) = \Big(r - q - {}^1\!/\!{}_2\, v(t)\Big)dt + \sqrt{v(t)}dW_1(t) + dL(t), \quad X(0) = x_0 \\
&dv(t) = k\left[\theta_t - v(t)\right]dt + \sigma_t\sqrt{v(t)}dW_2(t), \quad v(0) = v_0, \\
&dW_1(t)\, dW_2(t) = \rho_t\, dt, \qquad \rho_t \in [-1, 1] \\
&dL(t) = Y_t\, dN(t) - \bar{\mu}_t \lambda\, dt,
\end{aligned}$$

where $Y \sim \mathcal{N}(\nu_L, \delta_L^2)$, $\bar{\mu} = \mathbb{E}[e^Y - 1] = \exp(\nu_L + \frac{1}{2}\delta_L^2) - 1$, $\widetilde{\mu}_L$ is the compensated Lévy random measure, λ is the instantaneous intensity of the Poisson process $(N_t)_{t\geq 0}$, Y the normally distributed jump size. The Hull-White extension modifies F, so it becomes $T \mapsto F_T(u_1, u_2)$. In particular,

$$F_T(u_1, u_2) = k\theta_T u_2 + (r-q)u_1 + \widetilde{\kappa}_T(u_1),$$

where $\widetilde{\kappa}$ is the compensated cumulant generating function of the Lévy measure.

4 Calibration with Neural Networks

Exploit the direct map (from parameters to volatilities) to create a neural network (NN) that maps volatilities onto volatilities, thus learning the identity. In the training process only the first part of the (total) network (volatilities to parameters) is trained. In other words, $NN_3 := (NN_2 \circ NN_1)$, where NN_1 was already trained before and NN_3 has (basically) to learn the identity function. In some sense, we could see NN_2 as the inverse Neural Network of NN_1.
This method allows to solve an inverse problem which was intractable before and to calibrate online the Bates model. Thus, we are able to show the evolution of an arbitrage-free implied volatility surface for indefinite time without breaking arbitrage conditions for the first time.

References

[1] René Carmona and Sergey Nadtochiy. Tangent lévy market models. *Finance and Stochastics*, 16(1):63–104, Jan 2012.

[2] Jan Kallsen and Paul Krühner. On a heath–jarrow–morton approach for stock options. *Finance and Stochastics*, 19(3):583–615, Jul 2015.

[3] Anja Richter and Josef Teichmann. Discrete time term structure theory and consistent recalibration models. *SIAM Journal on Financial Mathematics*, 8(1):504–531, 2017.

Continuous-time incentives in hierarchies

Emma Hubert*

February 21, 2022

Abstract

In this talk, we will study continuous–time optimal contracting in a hierarchy model which generalizes the model of Sung (2015) [4]. The hierarchy is modeled by a series of interlinked Principal–Agent problems, leading to a sequence of Stackelberg equilibria. More precisely, the Principal can contract with the Managers to incentivise them to act in her best interest, despite only observing the net benefits of the total hierarchy. Managers in turn subcontracts the Agents below them. Both Agents and Managers each independently control a stochastic process representing their outcome. We will see through a simple example that even if the Agents only control the drift of their outcome, the Managers control the volatility of the Agents' continuation utility. Even this first simple example justifies the use of recent results on optimal contracting for drift and volatility control, and therefore the theory on 2BSDEs, developed in the second more theoretical part. We will also discuss some possible extensions of this model.

Starting from a simple model. The framework of this present work is inspired by the model developed by Sung (2015) [4], where a manager is hired by a principal to subcontract with N agents. Each worker (manager and agents) controls his own output process, and all outputs are assumed to be independent. His model includes a *bi-level* moral hazard: first, the manager does not observe the effort of the agents, but only the resulting outputs; second, the principal observes only the total benefit of the hierarchy, *i.e.*, the difference between the sum of the outputs of all workers and the sum of the contracts paid to the agents. However, instead of studying a continuous-time version of the model, Sung states that: '[f]or ease of exposition and without loss of generality, we formulate a discrete-time model which is analogous to its continuous-time counterpart' ([4, p. 2]). Extending the reasoning of Holmström and Milgrom (1987) [2], he therefore restricts the study to *linear contracts*, in the sense that they are linear with respect to the outcome, and states that '[t]his assumption is without loss of generality, as long as our results are interpreted in the context of continuous-time models' ([4, p. 3]).

However, while the restriction to linear contracts can be justified in Sung's framework for the first Stackelberg equilibrium, this is no longer the case for the contract offered by the principal to the (top) manager. More precisely, although the workers control only the drift of their outcomes, the manager controls both the drift and the volatility of the net benefit. Therefore, according to Cvitanić, Possamaï, and Touzi (2018) [1], it appears that the type of contracts considered by Sung is sub-optimal. Indeed, in continuous time, it is not sufficient to limit oneself to linear contracts (in the sense of [2]) when the volatility of the state variable is controlled: the optimal form of contracts should contain an additional part indexed on the quadratic variation of the net benefit. However, in the one-period model

*Department of Operations Research and Financial Engineering (ORFE), Princeton University, NJ 08544; eh3988@princeton.edu.

of Sung, this controlled quadratic variation cannot be estimated (unlike in continuous time), which leads to a fundamental gap between these two frameworks. From our point of view, this gap motivates a full study of Sung's model in continuous time.

Main contributions. In this paper, we provide a systematic method to solve any hierarchy problems of this sort, including those in which workers can also control the volatility of the output process, and not just the drift. The search for the optimal contract therefore requires the application of the theory of second-order backward stochastic differential equations (2BSDEs for short), subject to a slight extension to take into account the plurality of agents in the hierarchy. Furthermore, we show that, in a general way, the contract offered by the manager to one of his subordinate agents must be indexed not only on the performance of this particular agent, but also on the performance of other workers.

The model we develop allows us to determine the optimal form of incentives for a particular hierarchical structure, which can be extended in a straightforward way to a larger scale hierarchy. Although theoretical, the results we obtain give intuitions based on solid conceptual considerations to know which levers could be activated to incentivise workers within a hierarchy. In particular, the indexation of the contract on the quadratic variation of net profits for the managers argues in favour of remunerating them through stock options. These results can be applied to problems of incentives within a firm with a hierarchical structure, but also and above all as soon as work is delegated to an external entity. For example, these multi-layered incentive problems can be used for insurance-related applications, to model the relationships between a policyholder, an insurance company, and a reinsurer.

An overview of the paper. This work consists of two parts. In the first part, we study Sung's model [4] and some extensions. This opening example highlights the differences between the discrete-time model and its continuous-time equivalent, concerning the volatility control and the form of the contracts. In particular, this example leads to the conclusion that in order to rigorously study a continuous-time hierarchy problem, it is not possible to consider the associated discrete-time model with linear contracts. In our opinion, these conclusions justify the use of the theory of 2BSDEs to tackle problems of moral hazard within a hierarchy. The second part of this paper is devoted to the study of a more general model. In particular, we consider a more comprehensive structure of hierarchy since the principal can contract with m managers, who in turns subcontract with the agents in their teams. Moreover, the workers (managers and agents) can now control the volatility of their output, in addition to the drift. Finally, we consider general preferences, allowing us to recover the exponential utility functions (CARA) of Sung's model [4], but also other cases such as the one of risk-neutral workers.

Keywords. Principal–agent problems, moral hazard, hierarchical contracting, 2BSDEs.

AMS 2020 subject classifications. Primary: 91A65; secondary: 91B41, 60H30, 93E20.

References

[1] J. Cvitanić, D. Possamaï, and N. Touzi. Dynamic programming approach to principal–agent problems. *Finance and Stochastics*, 22(1):1–37, 2018.

[2] B. Holmström and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303–328, 1987.

[3] E. Hubert. *Interactions and incentives: between contract theory and mean–field games.* PhD thesis, Université Gustave Eiffel, 2020.

[4] J. Sung. Pay for performance under hierarchical contracting. *Mathematics and Financial Economics*, 9(3):195–213, 2015.

In this talk, we first propose a new class of metrics and show that under such metrics, the convergence of empirical measures in high dimensions is free of the curse of dimensionality, in contrast to Wasserstein distance. Proposed metrics originate from the maximum mean discrepancy, which we generalize by proposing criteria for test function spaces. Examples include RKHS, Barron space, and flow-induced function spaces. One application studies the construction of Nash equilibrium for the homogeneous n-player game by its mean-field limit (mean-field game). Then we discuss mean-field games with common noise and propose a deep learning algorithm based on fictitious play and signatures in rough path theory. The first part of the work collaborates with Jiequn Han and Jihao Long (arXiv:2104.12036); the second part is the joint work with Ming Min (ICML2021).

Convergence rate for pricing of American Option by penalisation method and its Monte-Carlo approximation

Wanqing WANG*, Emmanuel GOBET†, Mingyu XU‡

December 30, 2021

In this paper, we study the pricing of American Option taking advantage of the link with Reflected Backward Stochastic Differential equation (RBSDE in short). If we set the payoff as the reflecting obstacle S, a RBSDE generally writes as

$$\begin{cases} Y_t = S_T + \int_t^T f(s, Y_s, Z_s)\mathbf{d}s + K_T - K_t - \int_t^T Z_s \mathbf{d}B_s, \quad 0 \le t \le T, \\ Y_t \ge S_t, \quad 0 \le t \le T, \\ \int_0^T (Y_t - S_t)\mathbf{d}K_t = 0, \end{cases} \tag{0.0.1}$$

where the driver f stands for the self-financing properties and the possible market fictions. At time t, the price of American option $\operatorname*{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}^* \left[e^{-\int_t^\tau r_s \mathbf{d}s} S_\tau | \mathcal{F}_t \right]$ coincides with Y_t for some specific driver f when we deal with Itô market [1]. The solution to (0.0.1) can be approximated by a penalised BSDE (PBSDE in short) under the form

$$Y_t^\lambda = S_T + \int_t^T f(s, Y_s^\lambda, Z_s^\lambda)\mathbf{d}s + \lambda \int_t^T (Y_s^\lambda - S_s)^- \mathbf{d}s - \int_t^T Z_s^\lambda \mathbf{d}B_s.$$

Our main contribution is to prove that the convergence of $(Y^\lambda, Z^\lambda, K^\lambda)$ to (Y, Z, K) holds at rate $\frac{1}{\lambda}$ or $\frac{1}{\sqrt{\lambda}}$ as $\lambda \to +\infty$ according to some appropriate convex properties of S.
In addition, we design a discrete-time numerical scheme for PBSDE and prove its rate of convergence in term of the time-step h and λ. Results are illustrated using Monte Carlo simulations. We also conduct comparison studies with [1–4].

*This research is supported by the *Chinese Scholarship council(CSC)* and *the Chair Stress Test, RISK Management and Financial Steering of the Foundation Ecole Polytechnique.*

*CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, Route de Saclay, 91128 Palaiseau Cedex, France. email: **wanqing.wang@polytechnique.edu**

†CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, Route de Saclay, 91128 Palaiseau Cedex, France. email: **emmanuel.gobet@polytechnique.edu**

‡School of Mathematics, Fudan University, Shanghai, China. email: **xumy@fudan.edu.cn**

References

[1] N. El Karoui, C. Kapoudjian, E. Pardoux, S. Peng, and M. Quenez. Reflected solutions of backward SDE's and related obstacle problems for PDE's. *Ann. Probab.*, 25(2):702–737, 1997.

[2] N. El Karoui, S. Peng, and M. C. Quenez. Backward stochastic differential equations in finance. *Mathematical finance*, 7(1):1–71, 1997.

[3] E. Gobet, J. Lemor, and X. Warin. A regression-based monte carlo method to solve backward stochastic differential equations. *The Annals of Applied Probability*, 15(3):2172–2202, 2005.

[4] J. Mémin, S. Peng, and M. Xu. Convergence of solutions of discrete reflected backward SDE's and simulations. *Acta Math. Appl. Sin., Engl. Ser.*, 24(1):1–18, 2008.

Special Sessions in memory of Prof. Peter Carr

Convex duality in continuous pricing models

Speaker: Lorenzo Torricelli

Abstract: We provide an alternative description of diffusive asset pricing models using the theory of convex duality. Instead of specifying an underlying martingale security process and deriving option price dynamics, we directly specify a stochastic differential equation for the dual delta, i.e. the option delta as a function of strike, and attain a process describing the option convex conjugate. For valuation, the convex conjugate of an option price is seen to satisfy a certain initial value problem dual to Dupire (1994) equation, and the option price can be derived by inverting the Legendre transform. We discuss in detail the primal and dual specifications of two known cases, the Normal Bachelier (1900) model and Carr and Torricelli (2021) logistic price model, and show that the dynamics of the latter retain a much simpler expression when the dual formulation is used.

Cost-efficient payoffs under model ambiguity

Carole Bernard,[*]
Gero Junike[†]
Thibaut Lux[‡]
Steven Vanduffel[§¶]

June 11, 2021

Abstract

In a market without ambiguity all payoffs have a know distribution. A payoff that is cheapest possible in reaching a given target distribution is called cost-efficient and investors with law-invariant increasing preferences only purchase cost-efficient payoffs. In the presence of ambiguity the distribution of a payoff is no longer known. A payoff is called robust cost-efficient if its worst-case distribution stochastically dominates a target distribution and is cheapest possible in doing so. We study the link between this notion of "robust cost-effciency" and the maxmin expected utility setting of Gilboa and Schmeidler,and more generally with robust preferences in a possibly non-expected utility setting. We illustrate our study with examples in a log-normal market with uncertainty on the drift and on the volatility of the risky asset.

KEYWORDS: Cost-efficient payoffs, model ambiguity, maxmin utility, robust preferences, drift and volatility uncertainty
JEL classification: C02, C63, D80
AMS classification: 91B30, 62E17

[*]Corresponding author : Carole Bernard, Department of Accounting, Law and Finance at Grenoble Ecole de Management (email: carole.bernard@grenoble-em.com) and Department of Economics and Political Sciences at Vrije Universiteit Brussel (VUB).

[†]Institute of Mathematics, Carl von Ossietzky University, 26111 Oldenburg, Germany (e-mail: gero.junike@uol.de).

[‡]Helvetia Insurance Group (e-mail: thibaut.lux@helvetia.ch). *The views expressed in this manuscript do not represent the opinions of Helvetia Insurance Group.

[§]Steven Vanduffel, Department of Economics and Political Sciences at Vrije Universiteit Brussel (VUB). (e-mail: steven.vanduffel@vub.ac.be).

[¶]This work was supported by the FWO Odysseus project FWOODYS11 of Carole Bernard at VUB.

Credit Valuation Adjustment with Replacement Closeout: Theory and Algorithms

Abstract The replacement closeout convention has drawn more and more attention since the 2008 financial crisis. Compared with the conventional risk-free closeout, the replacement closeout convention incorporates the creditworthiness of the counterparty and thus providing a more accurate estimate of the Mark-to-market value of a financial claim. In contrast to the risk-free closeout, the replacement closeout renders a nonlinear valuation system, which constitutes the major difficulty in the valuation of the counterparty credit risk.

In this paper, we show how to tame the nonlinearity attributed to the replacement closeout in the theoretical and computational analysis. In the theoretical part, we prove the unique solvability of the nonlinear valuation system and study the impact of the replacement closeout on the credit valuation adjustment. In the computational part, we propose a neural network-based algorithm for solving the (high dimensional) nonlinear valuation system. We numerically compare the computational cost for the valuations with risk-free and replacement closeouts. The numerical tests indicate that the application of the replacement closeout would not burden the valuation.

Abstract for Decomposition of Optimal Dynamic Portfolio Choice in Incomplete Market

Yiwen Shen, Chenxu Li, Olivier Scaillet

This paper establishes a new structural decomposition of optimal dynamic portfolio choice under general incomplete-market diffusion models with flexible utilities. We derive explicit dynamics of the components for the optimal policy, and obtain an equation system for solving the shadow price of market incompleteness. We show the shadow price depends on both market state and wealth level. Consequently, a new hedge component emerges for non-myopic investors with wealth-dependent utilities. Our structural decomposition reveals the fundamental impact on optimal policy from market incompleteness and wealth dependent utilities. As an important application, we establish the decomposition of optimal policy under general models with HARA and CRRA utilities, and reveal the fundamental connection between the optimal policy under these two utilities. Finally, we develop a closed-form maturity decomposition approach to implement the sophisticated representations of optimal policy. This maturity decomposition approach can be applied to general incomplete market models with flexible utilities, where existing numerical approaches do not generally apply.

Deep Learning Algorithms for Hedging with Frictions

Xiaofei Shi Daran Xu Zhanhao Zhang

As observed in many empirical papers, markets are imperfect, meaning that arbitrary quantities cannot be traded immediately at the quoted market price because of taxes, regulations, and the limited liquidity of the assets. Typical examples include linear transaction taxes as well as fixed transaction costs. As reported and studied in [Almgren, 03, Almgren and Chriss, 01, Almgren et al., 05, Lillo et al., 03], empirical estimates of actual transaction costs typically correspond to a 3/2-th power of the order flow. Accordingly, the large trading volume quickly impacts the market liquidity, which, in turn, drastically changes the agents' behaviors. Hence, optimally scheduling the order flow in anticipation of market liquidity shortage is crucial.

There is a large body of literature on optimal execution strategies as well as dynamic portfolio optimization models with market illiquidity and price impacts, see [De Lataillade et al., 12, Amihud et al., 06, Dumas and Luciano, 91, Liu, 04, Janeček and Shreve, 10, Shreve and Soner, 94, Kohlmann and Tang, 02], and with recent results in [Bank et al., 17, Garleanu and Pedersen, 13, Guasoni and Weber, 20, Soner and Touzi, 13]. By assuming the transaction costs correspond to the square of the size of the order flow, the optimal trading policy can be given in a closed-form, as shown in [Garleanu and Pedersen, 13, Kohlmann and Tang, 02]. However, no closed-form solution is available under general nonlinear transaction costs with general market dynamics. To obtain tractable results, researchers then focus on the *small* transaction costs limit, just as considered in [Almgren and Li, 16, Bayraktar et al., 18, Guasoni and Weber, 20, Kallsen and Muhle-Karbe, 17, Moreau et al., 17, Soner and Touzi, 13]. These elegant asymptotic formulas were proved rigorously by [Ahrens, 15, Herdegen and Muhle-Karbe, 18, Kallsen and Li, 13, Shreve and Soner, 94]. The quality of the leading-order approximation therefore arises as an important follow-up question, especially in empirical applications.

From the numerical side, recent advances in highly accurate machine learning models have introduced powerful new tools for studying high-dimensional optimization problems, such as hedging with frictions. The FBSDE solver, developed by Han, Jentzen, and E in [Han et al., 18], can solve a dynamic portfolio optimization problem by finding the solution to a BSDE system. For the first time, this algorithm overcame the curse of dimensionality in numerical solutions to high-dimensional SDE and associated PDE, as pointed out by [Beck et al., 20, Grohs et al., 18]. The convergence analysis is established by [Han and Long, 20] under the same short-term existence assumptions in [Delarue, 02]. At a higher level, the FBSDE solvers find the solution of the FBSDEs through a *supervised learning* framework, since only the accuracy of the terminal value of the backward components is served as the goal functional of the algorithm. However, this algorithm does not scale well with the trading time horizon and the time discretization. For calibrated trading parameters as in [Gonon et al., 21], the time horizon for the algorithm to work is required to be unreasonably small. In the meantime, with the development of modern model-free techniques, *reinforcement learning* algorithms are also widely used in single-agent optimization problems. Indeed, as shown in the groundbreaking papers [Han and Weinan, 16, Becker et al., 19, Buehler et al., 19a, Buehler et al., 19b, Casgrain et al., 19, Huré et al., 18, Moallemi and Wang, 21, Mulvey et al., 20, Hu, 19, Reppen and Soner, 20, Ruf and Wang, 21], we treat the utility functions as targets and directly parametrize and learn the optimal trading policy. Moreover, reinforcement learning frameworks are analyzed rigorously as in [Wang et al., 20, Wang and Zhou, 20]. Therefore, just like in [Moallemi and Wang, 21], a natural question is how these methods compare in practice, and our goal is to understand and compare the two types of methods, especially document the advantages and drawbacks of supervised and reinforcement learning algorithms.

This paper studies the deep learning-based numerical algorithms for hedging problems in frictional markets in a general setting. First, we explore different machine learning architectures, including the FBSDE solver and the deep hedging algorithms. We implement both methods, document the tuning procedures, and discuss their advantages and disadvantages. Under the smallness assumption on the magnitude of the transaction costs, the leading-order approximation of the optimal trading speed can be identified through the solution to a nonlinear ODE. With calibrated parameters from real-world time-series data as in [Gonon et al., 21], we compare our numerical results with the asymptotic formulas in a Bachelier model for the price dynamics. Implementation details can be found here: https://github.com/InnerPeas/ML-for-Transaction-Costs. We summarize the our empirical observations (High/Medium/Low or Yes/No) in the following table:

	FBSDE Solver	Deep Hedging
Scalability wrt time	Low	Medium
Scalability wrt dimension	Medium	High
Convergence speed	High (if algorithm converges)	Medium
Hardness of tuning	Low (if algorithm converges)	Medium
Sensitivity to calculation precision	High	Medium
Multiple stocks	Yes	Yes
General utility functions	Limited	Yes
General market dynamics	Yes	Yes

These comparison results help explain the leading-order asymptotic formula and provide us with ideas on the usage of each method in practice. In summary, the FBSDE only performs well under short trading horizons, while it fails to work in long trading horizons. In contrast, the deep hedging algorithm has stable and reliable performance for short and moderately long trading horizons, while it starts to get unstable when the trading horizon is getting too large. In addition, the hyperparameter tuning of the deep hedging algorithm becomes inefficient when the model has more sophisticated dynamics. However, the leading-order approximation works well under long trading horizons, while it deviates from the ground truth by a significant amount only shortly before maturity. In fact, we relate the smallness assumption on the transaction costs level λ and with the trading time horizon T; namely, the smallness assumption of λ and the definition of long trading horizon are both relative quantities, depending on whether $\sqrt{\lambda}/T$ is of higher order of 1. Even more surprisingly, $O(\sqrt{\lambda}/T)$ is the approximation accuracy for the leading-order approximation, and it is *independent* of the choice of the transaction costs. Finally, based on these observations, we propose our "pasting" algorithm to aggregate the leading-order approximations and the deep learning-based algorithms together, hence fully utilizing both of their advantages. At the initial time, we start with the leading-order formula to approximate the optimal trading strategy, then switch to deep hedging algorithm to learn the optimal trading strategy starting at $T - O(\sqrt{\lambda})$ to maturity T, in order to have the scalability in both trading time horizon and dimension to work in cases where the trading time horizon is long and there are multiple stocks in the market. Based on the needs, other deep learning-based algorithms can also be applied starting $O(\sqrt{\lambda})$ before maturity T. In addition, this is preliminary documentation for the numerical methods that could be generalized to solve Radner and Nash equilibrium models, where several agents interact strategically, alongside with [Casgrain et al., 19, Gonon et al., 21].

Deep Learning Statistical Arbitrage*

Jorge Guijarro-Ordonez[†] Markus Pelger[‡] Greg Zanotti[§]

October 25, 2021

Abstract

Statistical arbitrage identifies and exploits temporal price differences between similar assets. We propose a unifying conceptual framework for statistical arbitrage and develop a novel deep learning solution, which finds commonality and time-series patterns from large panels in a data-driven and flexible way. First, we construct arbitrage portfolios of similar assets as residual portfolios from conditional latent asset pricing factors. Second, we extract the time series signals of these residual portfolios with one of the most powerful machine learning time-series solutions, a convolutional transformer. Last, we use these signals to form an optimal trading policy, that maximizes risk-adjusted returns under constraints. We conduct a comprehensive empirical comparison study with daily large cap U.S. stocks. Our optimal trading strategy obtains a consistently high out-of-sample Sharpe ratio and substantially outperforms all benchmark approaches. It is orthogonal to common risk factors, and exploits asymmetric local trend and reversion patterns. Our strategies remain profitable after taking into account trading frictions and costs. Our findings suggest a high compensation for arbitrageurs to enforce the law of one price.

Keywords: statistical arbitrage, pairs trading, machine learning, deep learning, big data, stock returns, convolutional neural network, transformer, attention, factor model, market efficiency, investment.

JEL classification: C14, C38, C55, G12

*We thank Robert Anderson, Jose Blanchet, Marcelo Fernandes, Kay Giesecke, Lisa Goldberg, Robert Korajczyk, Martin Lettau, Marcelo Medeiros, George Papanicolaou and seminar and conference participants at Stanford, UC Berkeley, the Econometric Research in Finance Workshop, the Meeting of the Brazilian Finance Society, World Online Seminars on Machine Learning in Finance, NBER-NSF Time-Series Conference, NVIDIA AI Webinar, Vanguard Academic Seminar, INFORMS and the Western Conference on Mathematical Finance for helpful comments. We thank MSCI for generous research support.

[†]Stanford University, Department of Mathematics, Email: jguiord@stanford.edu.

[‡]Stanford University, Department of Management Science & Engineering, Email: mpelger@stanford.edu.

[§]Stanford University, Department of Management Science & Engineering, Email: gzanotti@stanford.edu.

I. Introduction

Statistical arbitrage is one of the pillars of quantitative trading, and has long been used by hedge funds and investment banks. The term statistical arbitrage encompasses a wide variety of investment strategies, which identify and exploit temporal price differences between similar assets using statistical methods. Its simplest form is known as "pairs trading". Two stocks are selected that are "similar", usually based on historical co-movement in their price time-series. When the spread between their prices widens, the arbitrageur sells the winner and buys the loser. If their prices move back together, the arbitrageur will profit. While Wall Street has developed a plethora of proprietary tools for sophisticated arbitrage trading, there is still a lack of understanding of how much arbitrage opportunity is actually left in financial markets. In this paper we answer the two key questions around statistical arbitrage: What are the important elements of a successful arbitrage strategy and how much realistic arbitrage is in financial markets?

Every statistical arbitrage strategy needs to solve the following three fundamental problems: Given a large universe of assets, what are long-short portfolios of similar assets? Given these portfolios, what are time series signals that indicate the presence of temporary price deviations? Last, but not least, given these signals, how should an arbitrageur trade them to optimize a trading objective while taking into account possible constraints and market frictions? Each of these three questions poses substantial challenges, that prior work has only partly addressed. First, it is a hard problem to find long-short portfolios for all stocks as it is a priori unknown what constitutes "similarity". This problem requires considering all the big data available for a large number of assets and times, including not just conventional return data but also exogenous information like asset characteristics. Second, extracting the right signals requires detecting flexibly all the relevant patterns in the noisy, complex, low-sample-size time series of the portfolio prices. Last but not least, optimal trading rules on a multitude of signals and assets are complicated and depend on the trading objective. All of these challenges fundamentally require flexible estimation tools that can deal with many variables. It is a natural idea to use machine learning techniques like deep neural networks to deal with the high dimensionality and complex functional dependencies of the problem. However, our problem is different from the usual prediction task, where machine learning tools excel. We show how to optimally design a machine learning solution to our problem that leverages the economic structure and objective.

In this paper, we propose a unifying conceptual framework that generalizes common approaches to statistical arbitrage. Statistical arbitrage can be decomposed into three fundamental elements: (1) arbitrage portfolio generation, (2) arbitrage signal extraction and (3) the arbitrage allocation decision given the signal. By decomposing different methods into their arbitrage portfolio, signal and allocation element, we can compare different methods and study which components are the most relevant for successful trading. For each step we develop a novel machine learning implementation, which we compare with conventional methods. As a result, we construct a new deep learning statistical arbitrage approach. Our new approach constructs arbitrage portfolios with a conditional latent factor model, extracts the signals with the currently most successful machine learning time-series method and maps them into a trading allocation with a flexible neural network. These components are integrated and optimized over a global economic objective, which maximizes the risk-adjusted return under constraints. Empirically, our general model outperforms out-of-sample the leading benchmark approaches and provides a clear insight into the structure of statistical arbitrage.

To construct arbitrage portfolios, we introduce the economically motivated asset pricing perspective to create them as residuals relative to asset pricing models. This perspective allows us to take advantage of the recent developments in asset pricing and to also include a large set of firm characteristics in the construction of the arbitrage portfolios. We use fundamental risk factors and conditional and unconditional statistical factors for our asset pricing models. Similarity between assets is captured by similar exposure to those factors. Arbitrage Pricing Theory implies that, with an appropriate model, the corresponding factor portfolios represent the "fair price" of each of the assets. Therefore, the residual portfolios relative to the asset pricing factors capture the temporary deviations from the fair price of each of the assets and should only temporally deviate from their long-term mean. Importantly, the residuals are tradeable portfolios, which are only weakly cross-sectionally correlated, and close to orthogonal to firm characteristics and systematic factors. These properties allow us to extract a stationary time-series model for the signal.

To detect time series patterns and signals in the residual portfolios, we introduce a filter perspective and estimate them with a flexible data-driven filter based on convolutional networks combined with transformers. In this way, we do not prescribe a potentially misspecified function to extract the time series structure, for example, by estimating the parameters of a given parametric time-series model, or the coefficients of a decomposition into given basis functions, as in conventional methods. Instead, we directly learn in a data-driven way what the optimal pattern extraction function is for our trading objective. The convolutional transformer is the ideal method for this purpose. Convolutional neural networks are the state-of-the-art AI method for pattern recognition, in particular in computer vision. In our case they identify the local patterns in the data and may be thought as a nonlinear and learnable generalization of conventional kernel-based data filters. Transformer networks are the most successful AI model for time series in natural language processing. In our model, they combine the local patterns to global time-series patterns. Their combination results in a data-driven flexible time-series filter that can essentially extract any complex time-series signal, while providing an interpretable model.

To find the optimal trading allocation, we propose neural networks to map the arbitrage signals into a complex trading allocation. This generalizes conventional parametric rules, for example fixed rules based on thresholds, which are only valid under strong model assumptions and a small signal dimension. Importantly, these components are integrated and optimized over a global economic objective, which maximizes the risk-adjusted return under constraints. This allows our model to learn the optimal signals and allocation for the actual trading objective, which is different from a prediction objective. The trading objective can maximize the Sharpe ratio or expected return subject to a risk penalty, while taking into account constraints important to real investment managers, such as restricting turnover, leverage, or proportion of short trades.

Our comprehensive empirical out-of-sample analysis is based on the daily returns of roughly the 550 largest and most liquid stocks in the U.S. from 1998 to 2016. We estimate the out-of-sample residuals on a rolling window relative to the empirically most important factor models. These are observed fundamental factors, for example the Fama-French 5 factors and price trend factors, locally estimated latent factors based on principal component analysis (PCA) or locally estimated conditional latent factors that include the information in 46 firm-specific characteristics and are based on the Instrumented PCA (IPCA) of Kelly et al. (2019). We extract the trading signal with one of the most successful parametric models, based on the mean-reverting Ornstein-Uhlenbeck process, a frequency decomposition of the time-series with a Fourier transformation and our novel convolutional network with transformer. Finally, we compare the trading allocations based on parametric or nonparametric rules estimated with different risk-adjusted trading objectives.

Our empirical main findings are five-fold. First, our deep learning statistical arbitrage model substantially outperforms all benchmark approaches out-of-sample. In fact, our model can achieve an impressive annual Sharpe ratio larger than four. While respecting short-selling constraints we can obtain annual out-of-sample mean returns of 20%. This performance is four times better than one of the best parametric arbitrage models, and twice as good as an alternative deep learning model without the convolutional transformer filter. These results are particularly impressive as we only trade the largest and most liquid stocks. Hence, our model establishes a new standard for arbitrage trading.

Second, the performance of our deep learning model suggests that there is a substantial amount of short-term arbitrage in financial markets. The profitability of our strategies is orthogonal to market movements and conventional risk factors including momentum and reversal factors and does not constitute a risk-premium. Our strategy performs consistently well over the full time horizon. The model is extremely robust to the choice of tuning parameters, and the period when it is estimated. Importantly, our arbitrage strategy remains profitable in the presence of realistic transaction and holdings costs. Assessing the amount of arbitrage in financial markets with unconditional pricing errors relative to factor models or with parametric statistical arbitrage models, severely underestimates this quantity.

Third, the trading signal extraction is the most challenging and separating element among different arbitrage models. Surprisingly, the choice of asset pricing factors has only a minor effect on the overall performance. Residuals relative to the five Fama-French factors and five locally estimated principal component factors perform very well with out-of-sample Sharpe ratios above 3.2 for our deep learning model. Five conditional IPCA factors increase the out-of-sample Sharpe ratio to 4.2, which suggests that asset characteristics provide additional useful information. Increasing the number of risk factors beyond five has only a marginal effect. Similarly, the other benchmark models are robust to the choice of factor model as long as it contains sufficiently many factors. The distinguishing element is the time-series model to extract the arbitrage signal. The convolutional transformer doubles the performance relative to an identical deep learning model with a pre-specified frequency filter. Importantly, we highlight that time-series modeling requires a time-series machine learning approach, which takes temporal dependency into account. An off-the-shelf nonparametric machine learning method like conventional neural networks, that estimates an arbitrage allocation directly from residuals, performs substantially worse.

Fourth, successful arbitrage trading is based on local asymmetric trend and reversion patterns. Our convolutional transformer framework provides an interpretable representation of the underlying patterns, based on local basic patterns and global "dependency factors". The building blocks of arbitrage trading are smooth trend and reversion patterns. The arbitrage trading is short-term and the last 30 trading days seem to capture the relevant information. Interestingly, the direction of policies is asymmetric. The model reacts quickly on downturn movements, but more cautiously on uptrends. More specifically, the "dependency factors" which are the most active in downturn movements focus only on the most recent 10 days, while those for upward movements focus on the first 20 days in a 30-day window.

Fifth, time-series-based trading patterns should be extracted from residuals and not directly from returns. For an appropriate factor model, the residuals are only weakly correlated and close to stationary in both, the time and cross-sectional dimension. Hence, it is meaningful to extract a uniform trading pattern, that is based only on the past time-series information, from the residuals. In contrast stock returns are dominated by a few factors, which severely limits the actual independent time-series information, and are strongly heterogenous due to their variation in firm characteristics. While the level of stock returns is extremely hard to predict, even with flexible machine learning methods, residuals capture relative movements and remove the level component. These properties make residuals analyzable from a purely time-series based perspective and, unlike the existing literature, they allow us to incorporate alternative data into the portfolio construction process. This also highlights a fundamental difference with most of the existing financial machine learning literature: We do not use characteristics to get features for prediction, but rather to generate new data orthogonal to these features.

Related Literature

Our paper builds on the classical statistical arbitrage literature, in which the three main problems of portfolio generation, pattern extraction, and allocation decision have traditionally been considered independently. Classical statistical methods of generating arbitrage portfolios have mostly focused on obtaining multiple pairs or small portfolios of assets, using techniques like the distance method of Gatev et al. (2006), the cointegration approach of Vidyamurthy (2004), or copulas as in Rad et al. (2016). In contrast, more general methods that exploit large panels of stock returns include the use of PCA factor models, as in Avellaneda and Lee (2010) and its extension in Yeo and Papanicolaou (2017), and the maximization of mean-reversion and sparsity statistics as in d'Aspremont (2011). We include the model of Yeo and Papanicolaou (2017) as the parametric benchmark model in our study as it has one of the best empirical performances among the class of parametric models. Our paper paper contributes to this literature by introducing a general asset pricing perspective to obtain the arbitrage portfolios as residuals. This allows us to take advantage of conditional asset pricing models, that include time-varying firm characteristics in addition to the return time-series, and provides a more disciplined, economically motivated approach. The signal extraction step for these models assumes parametric time series models for the arbitrage portfolios, whereas the allocations are often decided from the estimated parameters by using stochastic control methods or given threshold rules and one-period optimizations. Some representative papers of the first approach include Jurek and Yang (2007), Mudchanatongsuk et al. (2008), Cartea and Jaimungal (2016), Lintilhac and Tourin (2016) and Leung and Li (2015), whereas the second one is illustrated by Elliott et al. (2005) and Yeo and Papanicolaou (2017). Both approaches are special cases of our more general framework. Mulvey et al. (2020) and Kim and Kim (2019) are examples of including machine learning elements within the parametric statistical arbitrage framework, by either solving a stochastic control problem with neural networks or estimating a time-varying threshold rule with reinforcement learning.

Our paper is complementary to the emerging literature that uses machine learning methods for asset pricing. While the asset pricing literature aims to explain the risk premia of assets, our focus is on the residual component which is not explained by the asset pricing models. Chen et al. (2019), Bryzgalova et al. (2019) and Kozak et al. (2020) estimate the stochastic discount factor (SDF), which explains the risk premia of assets, with deep neural networks, decision trees or elastic net regularization. These papers employ advanced statistical methods to solve a conditional method of moment problem in the presence of many variables. The workhorse models in equity asset pricing are based on linear factor models exemplified by Fama and French (1993, 2015). Recently, new methods have been developed to extract statistical asset pricing factors from large panels with various versions of principal component analysis (PCA). The Risk-Premium PCA in Lettau and Pelger (2020a,b) includes a pricing error penalty to detect weak factors that explain the cross-section of returns. The high-frequency PCA in Pelger (2020) uses high-frequency data to estimate local time-varying latent risk factors and the Instrumented PCA (IPCA) of Kelly et al. (2019) estimates conditional latent factors by allowing the loadings to be functions of time-varying asset characteristics. Gu et al. (2021) generalize IPCA to allow the loadings to be nonlinear functions of characteristics.

Our paper is related to the growing literature on return prediction with machine learning methods, which has shown the benefits of regularized flexible methods. In their pioneering work Gu et al. (2020) conduct a comparison of machine learning methods for predicting the panel of individual U.S. stock returns based on the asset-specific characteristics and economic conditions in the previous period. In a similar spirit, Bianchi et al. (2019) predict bond returns and Freyberger et al. (2020) use different methods for predicting stock returns. This literature is fundamentally estimating the risk premia of assets, while our focus is on understanding and exploiting the temporal deviations thereof. This different goal is reflected in the different methods that are needed. These return predictions estimate a nonparametric model between current returns and large set of covariates from the last period, but do not estimate a time-series model. In contrast, the important challenge that we solve is to extract a complex time-series pattern. A related stream of this literature forecasts returns using past returns, generally followed by some long-short investment policy based on the prediction. For example, Krauss et al. (2017) use various machine learning methods for this type of prediction.[1] However, they use general nonparametric function estimates, which are not specifically designed for time-series data. Lim and Zohren (2020) show that it is important for machine learning solutions to explicitly account for temporal dependence when they are applied to time-series data. Forecasting returns and building a long-short portfolio based on the prediction is different from statistical arbitrage trading as it combines a risk premium and potential arbitrage component. It is not based on temporary price differences and also in general not orthogonal to common risk factors and market movements. In this paper we highlight the challenge of inferring complex time-series information and argue that using returns directly as an input to a time-series machine learning method, is suboptimal as returns are dominated by a few factor time-series and heterogeneous due to cross-sectionally and time-varying characteristics. In contrast, appropriate residuals are locally stationary and hence allow the extraction of a complex time-series pattern.

Naturally, our work overlaps with the literature on using machine learning tools for investment. The SDF estimated by asset pricing models, like in Chen et al. (2019) and Bryzgalova et al. (2019), directly maps into a conditionally mean-variance efficient portfolio and hence an attractive investment opportunity. However, by construction this investment portfolio is not orthogonal but fully exposed to systematic risk, which is exactly the opposite for an arbitrage portfolio. Prediction approaches also imply investment strategies, typically long-short portfolios based on the prediction signal. However, estimating a signal with a prediction objective, is not necessarily providing an optimal signal for investment. Bryzgalova et al. (2019) and Chen et al. (2019) illustrate that machine learning models that use a trading objective can result in a substantially more profitable investment than models that estimate a signal with a prediction objective, while using the same information as input and having the same flexibility. This is also confirmed in Cong et al. (2020), who use an investment objective and reinforcement learning to construct machine learning investment portfolios. Our paper contributes to this literature by estimating investment strategies, that are orthogonal to systematic risk and are based on a trading objective with constraints.

Finally, our approach is also informed by the recent deep learning for time series literature. The transformer method was first introduced in the groundbreaking paper by Vaswani et al. (2017). We are the first to bring this idea into the context of statistical arbitrage and adopt it to the economic problem.

II. Model

The fundamental problem of statistical arbitrage consists of three elements: (1) The identification of similar assets to generate arbitrage portfolios, (2) the extraction of time-series signals for the temporary deviations of the similarity between assets and (3) a trading policy in the arbitrage

[1] Similar studies include Fischer et al. (2019), Chen et al. (2018), Huck (2009), and Dunis et al. (2006).

portfolios based on the time-series signals. We discuss each element separately.

A. *Arbitrage portfolios*

We consider a panel of excess returns $R_{n,t}$, that is the return minus risk free rate of stock $n = 1, ..., N_t$ at time $t = 1, ..., T$. The number of available assets at time t can be time-varying. The excess return vector of all assets at time t is denoted as $R_t = \begin{pmatrix} R_{1,t} & \cdots & R_{N_t,t} \end{pmatrix}^\top$.

We use a general asset pricing model to identify similar assets. In this context, similarity is defined as the same exposure to systematic risk, which implies that assets with the same risk exposure should have the same fundamental value. We assume that asset returns can be modeled by a conditional factor model:

$$R_{n,t} = \beta_{n,t-1}^\top F_t + \epsilon_{n,t}.$$

The K factors $F \in \mathbb{R}^{T \times K}$ capture the systematic risk, while the risk loadings $\beta_{t-1} \in \mathbb{R}^{N_t \times K}$ can be general functions of the information set at time $t-1$ and hence can be time-varying. This general formulation includes the empirically most successful factor models. In our empirical analysis we will include observed traded factors, e.g. the Fama-French 5 factor model, latent factors based on the principal components analysis (PCA) of stock returns and conditional latent factors estimated with Instrumented Principal Component Analysis (IPCA).

Without loss of generality, we can treat the factors as excess returns of traded assets. Either the factors are traded, for example a market factor, in which case we include them in the returns R_t. Otherwise, we can generate factor mimicking portfolios by projecting them on the asset space, as for example with latent factors:

$$F_t = {w_{t-1}^F}^\top R_t.$$

We define *arbitrage portfolio* as residual portfolios $\epsilon_{n,t} = R_{n,t} - \beta_{n,t-1}^\top F_t$. As factors are traded assets, the arbitrage portfolio is itself a traded portfolio: Hence, the residual portfolio equals

$$\epsilon_t = R_t - \beta_{t-1} {w_{t-1}^F}^\top R_t = \underbrace{\left(I_{N_t} - \beta_{t-1} {w_{t-1}^F}^\top \right)}_{\Phi_{t-1}} R_t = \Phi_{t-1} R_t. \tag{1}$$

Arbitrage portfolios are projections on the return space that annihilate systematic asset risk. For an appropriate asset pricing model, the residual portfolios should not earn a risk premium. This is the fundamental assumption behind any arbitrage argument. As deviations from a mean of zero have to be temporary, arbitrage trading bets on the mean revision of the residuals. In particular, for an appropriate factor model the residuals will have the following properties:

1. The unconditional mean of the arbitrage portfolios is zero: $\mathbb{E}[\epsilon_{n,t}] = 0$.
2. The arbitrage portfolios are only weakly cross-sectionally dependent.

We denote by $\mathcal{F}_t$ the filtration generated by the returns R_t, which include the factors, and the information set that captures the risk exposure β_t, which is typically based on asset specific characteristics or past returns.

B. *Arbitrage signal*

The *arbitrage signal* is extracted from the time-series of the arbitrage portfolios. These time-series signals are the input for a trading policy. An example for an arbitrage signal would be a parametric model for mean reversion that is estimated for each arbitrage portfolio. The trading strategy for each arbitrage portfolio would depend on its speed of mean reversion and its deviation from the long run mean. More generally, the arbitrage signal is the estimation of a time-series model, which can be parametric or nonparametric. An important class of models are filtering approaches. Conceptually, time-series models are multivariate functional mappings between sequences which take into account the temporal order of the elements and potentially complex dependencies between the elements of the input sequence.

We apply the signal extraction to the time-series of the last L lagged residuals, which we denote in vector notation as

$$\epsilon^L_{n,t-1} := \begin{pmatrix} \epsilon_{n,t-L} & \cdots & \epsilon_{n,t-1} \end{pmatrix}.$$

The arbitrage signal function is a mapping $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ from $\mathbb{R}^L$ to $\mathbb{R}^p$, where $\boldsymbol{\Theta}$ defines an appropriate function space:

$$\boldsymbol{\theta}(\cdot) : \epsilon^L_{n,t-1} \to \theta_{n,t-1}.$$

The signals $\theta_{n,t-1} \in \mathbb{R}^p$ for the arbitrage portfolio n at time t only depend on the time-series of lagged returns $\epsilon^L_{n,t-1}$. Note that the dimensionality of the signal can be the same as for the input sequence. Formally, the function $\boldsymbol{\theta}$ is a mapping from the filtration $\mathcal{F}^{\epsilon,L}_{n,t-1}$ generated by $\epsilon^L_{n,t-1}$ into the filtration $\mathcal{F}^{\theta}_{n,t-1}$ generated by $\theta_{n,t-1}$ and $\mathcal{F}^{\theta}_{n,t-1} \subseteq \mathcal{F}^{\epsilon,L}_{n,t-1}$. We use the notation of evaluating functions elementwise, that is $\boldsymbol{\theta}(\epsilon^L_{t-1}) = \begin{pmatrix} \theta_{1,t-1} & \cdots & \theta_{N_t,t-1} \end{pmatrix} = \theta_{t-1} \in \mathbb{R}^{N_t}$ with $\epsilon^L_{t-1} = \begin{pmatrix} \epsilon_{1,t-1} & \cdots & \epsilon_{N_t,t-1} \end{pmatrix}$.

The arbitrage signal $\theta_{n,t-1}$ is a sufficient statistic for the trading policy; that is, all relevant information for trading decisions is summarized in it. This also implies that two arbitrage portfolios with the same signal get the same weight in the trading strategy. More formally, this means that the arbitrage signal defines equivalence classes for the arbitrage portfolios. The most relevant signals summarize reversal patterns and their direction with a small number of parameters. A potential trading policy could be to hold long positions in residuals with a predicted upward movement and go short in residuals that are in a downward cycle.

This problem formulation makes two implicit assumptions. First, the residual time-series follow a stationary distribution conditioned on its lagged returns. This is a very general framework that includes the most important models for financial time-series. Second, the first L lagged returns are a sufficient statistic to obtain the arbitrage signal $\theta_{n,t-1}$. This reflects the motivation that arbitrage is a temporary deviation of the fair price. The lookback window can be chosen to be arbitrarily large, but in practice it is limited by the availability of lagged returns.

C. *Arbitrage trading*

The trading policy assigns an investment weight to each arbitrage portfolio based on its signal. The allocation weight is the solution to an optimization problem, which models a general risk-return tradeoff and can also include trading frictions and constraints. An important case are mean-variance efficient portfolios with transaction costs and short sale constraints.

An arbitrage allocation is a mapping from $\mathbb{R}^p$ to $\mathbb{R}$ in a function class $\boldsymbol{W}$, that assigns a weight $w^{\epsilon}_{n,t-1}$ for the arbitrage portfolio $\epsilon_{n,t-1}$ in the investment strategy using only the arbitrage signal $\theta_{n,t}$:

$$\boldsymbol{w}^{\boldsymbol{\epsilon}} : \theta_{n,t-1} \to w^{\epsilon}_{n,t-1}.$$

Given a concave utility function $U(\cdot)$, the allocation function is the solution to

$$\max_{\boldsymbol{w}^{\boldsymbol{\epsilon}} \in \boldsymbol{W}, \boldsymbol{\theta} \in \boldsymbol{\Theta}} \mathbb{E}_{t-1} \left[U \left(w^R_{t-1} R_t \right) \right] \tag{2}$$

$$\text{s.t.} \qquad w^R_{t-1} = \frac{{w^{\epsilon}_{t-1}}^{\top} \Phi_{t-1}}{\| {w^{\epsilon}_{t-1}}^{\top} \Phi_{t-1} \|_1} \qquad \text{and} \qquad w^{\epsilon}_{t-1} = \boldsymbol{w}^{\boldsymbol{\epsilon}}(\boldsymbol{\theta}(\epsilon^L_{t-1})). \tag{3}$$

In the presence of trading costs, we calculate the expected utility of the portfolio net return by subtracting from the portfolio return $w^R_{t-1} R_t$ the trading costs that are associated with the stock allocation w^R_{t-1}. The trading costs can capture the transaction costs from frequent rebalancing and the higher costs of short selling compared to long positions. The stock weights w^R_{t-1} are normalized to add up to one in absolute value, which implicitly imposes a leverage constraint. The conditional expectation uses the general filtration $\mathcal{F}_{t-1}$.

This is a combined optimization problem, which simultaneously solves for the optimal allocation function and arbitrage signal function. As the weight is a composition of the two functions, i.e. $w^{\epsilon}_{t-1} = \boldsymbol{w}^{\boldsymbol{\epsilon}}(\boldsymbol{\theta}(\epsilon^L_{t-1}))$, the decomposition into a signal and allocation function is in general not uniquely identified. This means there can be multiple representations of $\boldsymbol{\theta}$ and $\boldsymbol{w}^{\boldsymbol{\epsilon}}$, that will result in the same trading policy. We use a decomposition that allows us to compare the problem to classical arbitrage approaches, for which this separation is uniquely identified. The key feature of the signal function $\boldsymbol{\theta}$ is that it models a time-series, that means it is a mapping that explicitly models the temporal order and the dependency between the elements of ϵ^L_{t-1}. The allocation function $\boldsymbol{w}^{\boldsymbol{\epsilon}}$ can be a complex nonlinear function, but does not explicitly model time-series behavior. This means that $\boldsymbol{w}^{\boldsymbol{\epsilon}}$ is implicitly limited in the dependency patterns of its input elements that it can capture.

We will consider arbitrage trading that maximizes the Sharpe ratio or achieves the highest average return for a given level of variance risk. More specifically we will solve for

$$\max_{\boldsymbol{w}^{\boldsymbol{\epsilon}}\in\boldsymbol{W},\boldsymbol{\theta}\in\boldsymbol{\Theta}} \frac{\mathbb{E}\left[{w_{t-1}^{R}}^{\top} R_t\right]}{\sqrt{\mathrm{Var}({w_{t-1}^{R}}^{\top} R_t)}} \qquad \text{or} \qquad \max_{\boldsymbol{w}^{\boldsymbol{\epsilon}}\in\boldsymbol{W},\boldsymbol{\theta}\in\boldsymbol{\Theta}} \mathbb{E}[{w_{t-1}^{R}}^{\top} R_t] - \gamma \mathrm{Var}({w_{t-1}^{R}}^{\top} R_t) \tag{4}$$

$$\text{s.t.} \qquad w_{t-1}^{R} = \frac{{w_{t-1}^{\epsilon}}^{\top}\Phi_{t-1}}{\|{w_{t-1}^{\epsilon}}^{\top}\Phi_{t-1}\|_1} \qquad \text{and} \qquad w_{t-1}^{\epsilon} = \boldsymbol{w}^{\boldsymbol{\epsilon}}(\boldsymbol{\theta}(\epsilon_{t-1}^{L})). \tag{5}$$

for some risk aversion parameter γ. We will consider this formulation with and without trading costs.[2]

Many relevant models estimate the signal and allocation function separately. The arbitrage signals can be estimated as the parameters of a parametric time-series model, the serial moments for a given stationary distribution or a time-series filter. In these cases, the signal estimation solves a separate optimization problem as part of the estimation. Given the signals, the allocation function is the solution of

$$\max_{\boldsymbol{w}^{\boldsymbol{\epsilon}}\in\boldsymbol{W}} \mathbb{E}_{t-1}\left[U\left(w_{t-1}^{R} R_t\right)\right] \qquad \text{s.t.} \qquad w_{t-1}^{R} = \frac{w^{\epsilon}(\theta_{t-1})^{\top}\Phi_{t-1}}{\|w^{\epsilon}(\theta_{t-1})^{\top}\Phi_{t-1}\|_1}. \tag{6}$$

We provide an extensive study of the importance of the different elements in statistical arbitrage trading. We find that the most important driver for profitable portfolios is the arbitrage signal function; that is, a good model to extract time-series behavior and to time the predictable mean reversion patterns is essential. The arbitrage portfolios of asset pricing models that are sufficiently rich result in roughly the same performance. Once an informative signal is extracted, parametric and nonparametric allocation functions can take advantage of it. We find that the key element is to consider a sufficiently general class of functions $\boldsymbol{\Theta}$ for the arbitrage signal and to estimate the signal that is the most relevant for trading. In other words, the largest gains in statistical arbitrage come from flexible time-series signals $\boldsymbol{\theta}$ and a joint optimization problem.

D. Models for Arbitrage Signal and Allocation Functions

In this section we introduce different functional models for the signal and allocation functions. They range from the most restrictive assumptions for simple parametric models to the most flexible model, which is our sophisticated deep neural network architecture. The general problem is the estimation of a signal and allocation function given the residual time-series. Here, we take the residual returns as given, i.e. we have selected an asset pricing model. In order to illustrate the key elements of the allocation functions, we consider trading the residuals directly. Projecting the residuals back into the original return space is identical for the different methods and discussed in the empirical part. The conceptual steps are illustrated in Figure 1.

[2] While the conventional mean-variance optimization applies a penalty to the variance, we found that it is numerically beneficial to use the standard deviation instead. Both formulations describe the same fundamental trade-off.

Figure 1: Conceptual Arbitrage Model

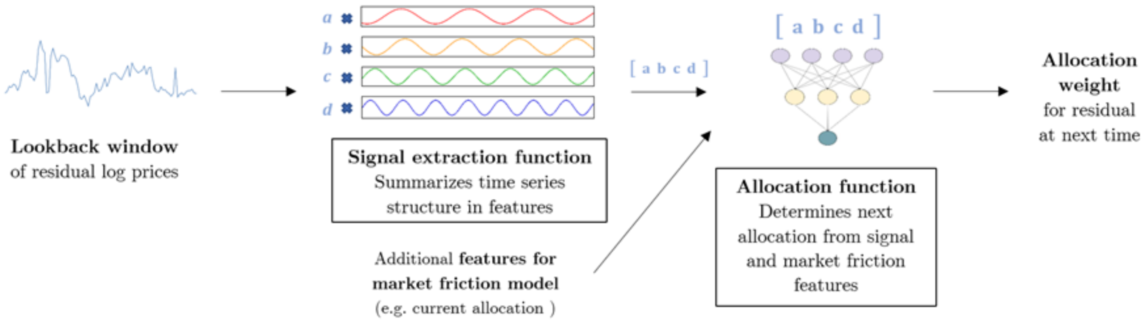

This figure illustrates the structure of our model for trading residuals. The model takes as input a lookback window of the last L cumulative returns or log prices of a residual portfolio at a given time and outputs the predicted optimal allocation weight for that residual for the next time. The model is composed of a signal extraction function and an allocation function, whose purposes are explained in the figure.

The input to the signal extraction functions are the last L cumulative residuals. We simplify the notation by dropping the time index $t-1$ and the asset index n and define the generic input vector

$$x := \mathbf{Int}\left(\epsilon_{n,t-1}^{L}\right) = \begin{pmatrix} \epsilon_{n,t-L} & \sum_{l=1}^{2} \epsilon_{n,t-L-1+l} & \cdots & \sum_{l=1}^{L} \epsilon_{n,t-L-1+l} \end{pmatrix}.$$

Here the operation $\mathbf{Int}(\cdot)$ simply integrates a discrete time-series. We can view the cumulative residuals as the residual "price" process. We discuss three different classes of models for the signal function $\boldsymbol{\theta}$ that vary in the degree of flexibility of the type of patterns that they can capture. Similarly, we consider different classes of models for the allocation function $\boldsymbol{w}^{\boldsymbol{\epsilon}}$.

D.1. Parametric Models

Our first benchmark method is a parametric model and corresponds to classical mean-reversion trading. In this framework, the cumulative residuals x are assumed to be the discrete realizations of continuous time model:

$$x = \begin{pmatrix} X_1 & \cdots & X_L \end{pmatrix}.$$

Following Avellaneda and Lee (2010) and Yeo and Papanicolaou (2017) we model X_t as an Ornstein-Uhlenbeck (OU) process

$$dX_t = \kappa\left(\mu - X_t\right)dt + \sigma dB_t$$

for a Brownian motion B_t. These are the standard models for mean-reversion trading and Avellaneda and Lee (2010) among others have shown their good empirical performance.

The parameters of this model are estimated from the moments of the discretized time-series, as described in detail along with the other implementation details in Appendix B.B. The parameters for each residual process, the last cumulative sum and a goodness of fit measure form the signals for the Ornstein-Uhlenbeck model:

$$\theta^{\text{OU}} = \begin{pmatrix} \hat{\kappa} & \hat{\mu} & \hat{\sigma} & X_L & R^2 \end{pmatrix}.$$

Following Yeo and Papanicolaou (2017) we also include the goodness of fit parameter R^2 as part of the signal. R^2 has the conventional definition of the ratio of squared values explained by the model normalized by total squared values. If the R^2 value is too low, the predictions of the model seem to be unreliable, which can be taken into account in a trading policy. Hence, for each cumulative residual vector $\epsilon^L_{n,t-1}$ we obtain the signal

$$\theta^{\text{OU}}_{n,t-1} = \begin{pmatrix} \hat{\kappa}_{n,t-1} & \hat{\mu}_{n,t-1} & \hat{\sigma}_{n,t-1} & \sum_{l=1}^{L} \epsilon_{n,t-1+l} & R^2_{n,t-1} \end{pmatrix}.$$

Avellaneda and Lee (2010) and Yeo and Papanicolaou (2017) advocate a classical mean-reversion thresholding rule, which implies the following allocation function[3]:

$$\boldsymbol{w}^{\epsilon|\text{OU}}\left(\theta^{\text{OU}}\right) = \begin{cases} -1 & \text{if } \frac{X_L-\mu}{\sigma/\sqrt{2\kappa}} > c_{\text{thresh}} \text{ and } R^2 > c_{\text{crit}} \\ 1 & \text{if } \frac{X_L-\mu}{\sigma/\sqrt{2\kappa}} < -c_{\text{thresh}} \text{ and } R^2 > c_{\text{crit}} \\ 0 & \text{otherwise} \end{cases}$$

The threshold parameters c_{thresh} and c_{crit} are tuning parameters. The strategy suggests to buy or sell residuals based on the ratio $\frac{X_L-\mu}{\sigma/\sqrt{2\kappa}}$. If this ratio exceeds a threshold, it is likely that the process reverts back to its long term mean, which starts the trading. If the R^2 value is too low, the predictions of the model seem to be unreliable, which stops the trading. This will be our parametric benchmark model. It has a parametric model for both the signal and allocation function.

Figure 2 illustrates this model with an empirical example. In this figure we show the allocation weights and signals of the Ornstein-Uhlenbeck with threshold model as well as the more flexible models that we are going to discuss next. The models are estimated on the empirical data, and the residual is a representative empirical example. In more detail, we consider the residuals from five IPCA factors and estimate the benchmark models as explained in the empirical Section III.J. The left subplots display the cumulative residual process along with the out-of-sample allocation weights w^{ϵ}_l that each model assigns to this specific residual. The evaluation of this illustrative example is a simplification of the general model that we use in our empirical main analysis. In this example, we consider trading only this specific residual and hence normalize the weights to $\{-1, 0, 1\}$. In our empirical analysis we trade all residuals and map them back into the original stock returns. The middle column shows the time-series of estimated out-of-sample signals for each model, by applying the θ_l arbitrage signal function to the previous $L = 30$ cumulative returns of the residual. The right column displays the out-of-sample cumulative returns of trading this particular residual

[3]The allocation function is derived by maximizing an expected trading profit. This deviates slightly from our objective of either maximizing the Sharpe ratio or the expected return subject to a variance penalty. As this is the most common arbitrage trading rule, we include it as a natural benchmark.

based on the corresponding allocation weights.

Figure 2: Illustrative Example of Allocation Weights and Signals for Different Methods

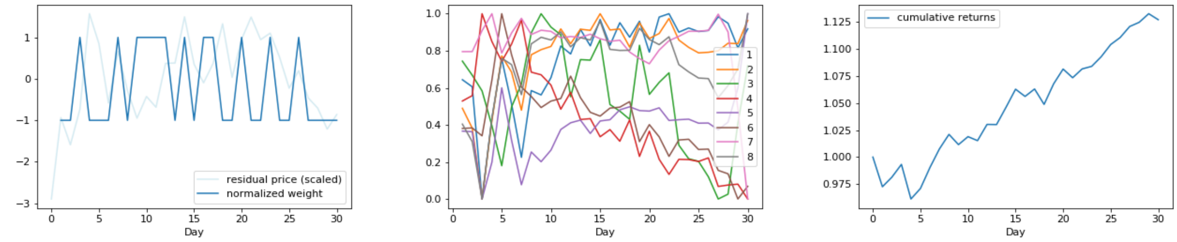

(a) Cumulative residuals x_l and allocation weight $w_l^{\epsilon|\text{CNN+Trans}}$

(b) Signal $\theta_l^{\text{CNN+Trans}}$

(c) Cumulative returns of CNN+Trans strategy

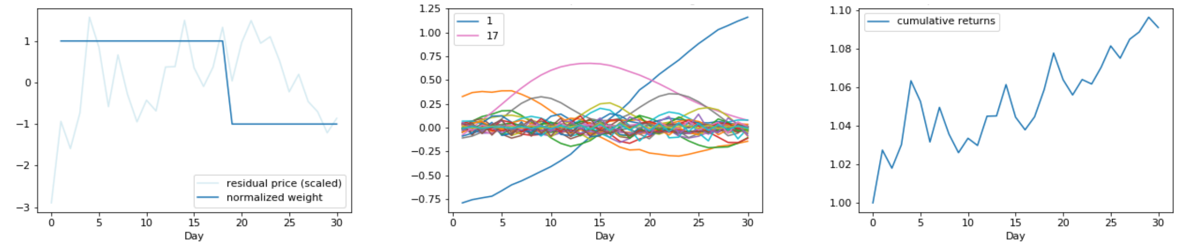

(d) Cumulative residual x_l and allocation weight $w_l^{\epsilon|\text{FFT}}$

(e) Signal θ_l^{FFT}

(f) Cumulative returns of Fourier+FFN strategy

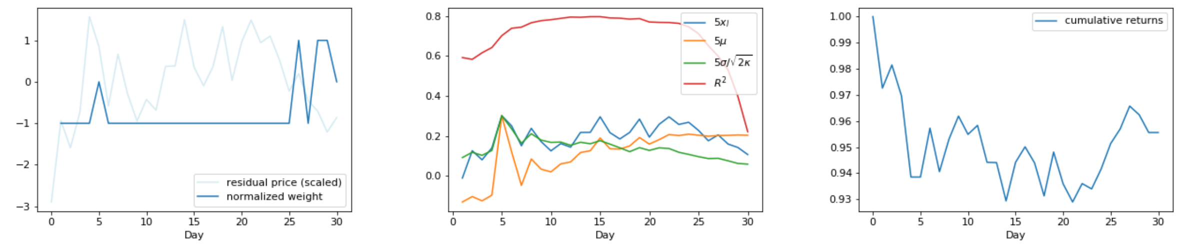

(g) Cumulative residual x_l and allocation weight $w_l^{\epsilon|\text{OU}}$

(h) Signal θ_l^{OU}

(i) Cumulative returns of OU+Thresh strategy

These plots are an illustrative example of the allocation weights and signals of the Ornstein-Uhlenbeck with Threshold (OU+Thres), Fast Fourier Transform (FFT) with Feedforward Neural Network (FFN), and Convolutional Neural Network (CNN) with Transformer models for a specific cumulative residual. The models are estimated on the empirical data, and the residual is a representative empirical example. In more detail, we consider the residuals from five IPCA factors and estimate the benchmark models as explained in Section III.J. The left subplots display the cumulative residual process along with the out-of-sample allocation weights $w_l^{\epsilon|\cdot}$ that each model assigns to this specific residual. In this example, we consider trading only this specific residual and hence normalize the weights to $\{-1, 0, 1\}$. The middle column plots show the time-series of estimated out-of-sample signals for each model, by applying the $\theta_l^{\cdot}$ arbitrage signal function to the previous L cumulative returns of the residual. The right column plots display the out-of-sample cumulative returns of trading this particular residual based on the corresponding allocation weights. We use a rolling lookback window of $L = 30$ days to estimate the signal and allocation, which we evaluate for the out-of-sample on the next 30 days. The plots only show the out-of-sample period. The evaluation of this illustrative example is a simplification of the general model that we use in our empirical main analysis, where we trade all residuals and map them back into the original stock returns.

The last row in Figure 2 shows the results for the OU+Threshold model. The cumulative return of trading this residual is negative, suggesting that the parametric model fails. The residual time-series with the corresponding allocation weights in subplot (g) explain why. The trading allocation does not assign a positive weight during the uptrend and wrongly assigns a constant negative weight, when the residual price process follows a mean-reversion pattern with positive and negative returns. A parametric model can break down if it is misspecified. This is not only the case for trend patterns, but also if there are multiple mean reversion patterns of different frequencies. Subplot (h) shows the signal.[4] We see that changes in the allocation function are related to sharp changes in at least one of the signals, but overall, the signal does not seem to represent the complex price patterns of the residual.

A natural generalization is to allow for a more flexible allocation function given the same time-series signals. We will consider for all our models also a general feedforward neural network (FFN) to map the signal into an allocation weight. FFNs are nonparametric estimators that can capture very general functional relationships.[5] Hence, we also consider the additional model that restricts the signal function, but allows for a flexible allocation function:

$$\boldsymbol{w}^{\epsilon|\text{OU-FFN}}\left(\theta^{\text{OU}}\right) = \boldsymbol{g}^{\text{FFN}}\left(\theta^{\text{OU}}\right).$$

We will show empirically that the gains of a flexible allocation function are minor relative to the very simple parametric model.

D.2. Pre-Specified Filters

As a generalization of the restrictive parametric model of the last subsection, we consider more general time-series models. Many relevant time-series models can be formulated as filtering problems. Filters are transformations of time-series that provide an alternative representation of the original time-series which emphasizes certain dynamic patterns.

A time-invariant linear filter can be formulated as

$$\theta_l = \sum_{j=1}^{L} W_j^{\text{filter}} x_j,$$

which is a linear mapping from $\mathbb{R}^L$ into $\mathbb{R}^L$ with the matrix $W^{\text{filter}} \in \mathbb{R}^{L\times L}$. The estimation of causal ARMA processes is an example for such filters. A spectral decomposition based on a frequency filter is the most relevant filter for our problem of finding mean reversion patterns.

A Fast Fourier Transform (FFT) provides a frequency decomposition of the original time-series and separates the movements into mean reverting processes of different frequencies. FFT applies the filter $W_j^{\text{FFT}} = e^{\frac{2\pi i}{L}j}$ in the complex plane, but for real-valued time-series it is equivalent to

[4]For better readability we have scaled the parameters of the OU process by a factor of five, but this still represents the same model as the scaling cancels out in the allocation function. As a minor modification, we use the ratio $\sigma/\sqrt{2\kappa}$ as a signal instead of two individual parameters, as the conventional regression estimator of the OU process directly provides the ratio, but requires additional moments for the individual parameters. However, this results in an equivalent presentation of the model as only the ratio enters the allocation function.

[5]Appendix B.A provides the details for estimating a FFN as a functional mapping $\boldsymbol{g}^{\text{FFN}} : \mathbb{R}^p \to \mathbb{R}$.

fitting the following model:

$$x_l = a_0 + \sum_{j=1}^{L/2-1} \left(a_j \cdot \cos\left(\frac{2\pi j}{L} l\right) + b_j \cdot \sin\left(\frac{2\pi j}{L} l\right)\right) + a_{L/2} \cos(\pi l).$$

The FFT representation is given by coefficients of the trigonometric representation

$$\theta^{\mathrm{FFT}} = \begin{pmatrix} a_0 & \cdots & a_{L/2} & b_1 & \cdots & b_{L/2-1} \end{pmatrix} \in \mathbb{R}^L.$$

The coefficients a_l and b_l can be interpreted as "loadings" or exposure to long or short-term reversal patterns. Note that the FFT is an invertible transformation. Hence, it simply represents the original time-series in a different form without losing any information. It is based on the insight that not the magnitude of the original data but the relative relationship in a time-series matters.

We use a flexible feedforward neural network for the allocation function

$$\boldsymbol{w}^{\epsilon|\mathrm{FFT}}\left(\theta^{\mathrm{FFT}}\right) = \boldsymbol{g}^{\mathrm{FFN}}\left(\theta^{\mathrm{FFT}}\right).$$

The usual intuition behind filtering is to use the frequency representation to cut off frequencies that have low coefficients and therefore remove noise in the representation. The FFN is essentially implementing this filtering step of removing less important frequencies.

We illustrate the model within our running example in Figure 2. The middle row shows the results for the FFT+FFN model. The cumulative residual in subplot (d) seems to be a combination of low and high-frequency movements with an initial trend component. The signal in subplot (e) suggests that the FFT filter seems to capture the low frequency reversal pattern. However, it misses the high-frequency components as indicated by the simplistic allocation function. The trading strategy takes a long position for the first half and a short position for the second part. While this simple allocation results in a positive cumulative return, in this example it neglects the more complex local reversal patterns.

While the FFT framework is an improvement over the simple OU model as it can deal with multiple combined mean-reversion patterns of different frequencies, it fails if the data follows a pattern that cannot be well approximated by a small number of the prespecified basis functions.

For completeness, our empirical analysis will also report the case of a trivial filter, which simply takes the residuals as signals, and combines them with a general allocation function:

$$\boldsymbol{\theta}^{\mathrm{ident}}(x) = x = \theta^{\mathrm{ident}}$$
$$\boldsymbol{w}^{\epsilon|\mathrm{FFN}}\left(\theta^{\mathrm{ident}}\right) = \boldsymbol{g}^{\mathrm{FFN}}(x).$$

This is a good example to emphasize the importance of a time-series model. While FFNs are flexible in learning low dimensional functional relationships, they are limited in learning a complex dependency model. For example, the FFN architecture we consider is not sufficiently flexible to learn the FFT transformation and hence has a worse performance on the original time-series compared to frequency-transformed time-series. While Xiaohong Chen and White (1999) have shown that FFNs are "universal approximators" of low-dimensional functional relationships, they also show that FFN can suffer from a curse of dimensionality when capturing complex dependencies between the input. Although the time domain and frequency domain representations of the input are equivalent under the Fourier transform, clearly the time-series model implied by the frequency domain representation allows for a more effective learning of an arbitrage policy. However, the choice of the pre-specified filter limits the time-series patterns that can be exploited. The solution is our data driven filter presented in the next section.

D.3. Convolutional Neural Network with Transformer

Our benchmark machine learning model is a Convolutional Neural Network (CNN) combined with a Transformer. It uses the most advanced state-of-the art machine learning tools tailored to our problem. Convolutional networks are in fact the most successful networks for computer vision, i.e. for pattern detection. Transformers have rapidly become the model of choice for sequence modeling such as Natural Language Processing (NLP) problems, replacing older recurrent neural network models such as the Long Short-Term Memory (LSTM) network.

The CNN and transformer framework has two key elements: (1) Local filters and (2) the temporal combination of these local filters. The CNN can be interpreted as a set of data driven flexible local filters. A transformer can be viewed as a data driven flexible time-series model to capture complex dependencies between local patterns. We use the CNN+Transformer to generate the time-series signal. The allocation function is then modeled as a flexible data driven allocation with an FFN.

The CNN estimates D local filters of size D_{size}:

$$y_l^{(0)} = \sum_{m=1}^{D_{\text{size}}} W_m^{(0)} x_{l-m+1}$$

for a matrix $W^{(0)} \in \mathbb{R}^{D_{\text{size}} \times D}$. The local filters are a mapping from $x \in \mathbb{R}^L$ to $y^{(0)} \in \mathbb{R}^{L \times D}$ given by the convolution $y^{(0)} = W^{(0)} * x$. Figure 3 shows examples of these local filters for $D_{\text{size}} = 3$. The values of $y^{(0)}$ can be interpreted as the "loadings" or exposure to local basis patterns. For example, if x represents a global upward trend, its filtered representation should have mainly large values for the local upward trend filter.

The convolutional mapping can be repeated in multiple layers to obtain a multi-layer CNN. First, the output of the first layer of the CNN is transformed nonlinearly by applying the ReLU$(\cdot)$ function:

$$x_{l,d}^{(1)} = \text{ReLU}\left(y_{l,d}^{(0)}\right) := \max(y_{l,d}^{(0)}, 0).$$

Figure 3: Examples of Local Filters

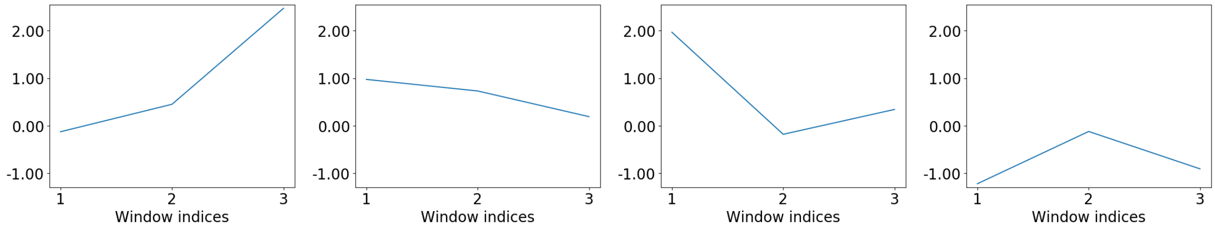

(a) Upward trend **(b)** Downward trend **(c)** Up reversal **(d)** Down reversal

These figures show the most important local filters estimated for the benchmark model in our empirical analysis. These are projections of our higher dimensional nonlinear filter from a 2-layer CNN into two-dimensional linear filters.

The second layer is given by a higher dimensional filtering projection:

$$y_{l,d}^{(2)} = \sum_{m=1}^{D_{\text{size}}} \sum_{j=1}^{D} W_{d,j,m}^{(1)} x_{l-m+1,j}^{(1)},$$

$$x_{l,d}^{(2)} = \text{ReLU}\left(y_{l,d}^{(1)}\right).$$

The final output of the CNN is $\tilde{x} \in \mathbb{R}^{L \times D}$. Our benchmark model is a 2-layered convolutional neural network. The number of layers is a hyperparameter selected on the validation data. Figure 4 illustrates the structure of the 2-layer CNN. While this description captures all the conceptual elements, the actual implementation includes additional details, such as bias terms, instance normalization and residual connection to improve the implementation as explained in Appendix B.C.

Figure 4: Convolutional Network Architecture

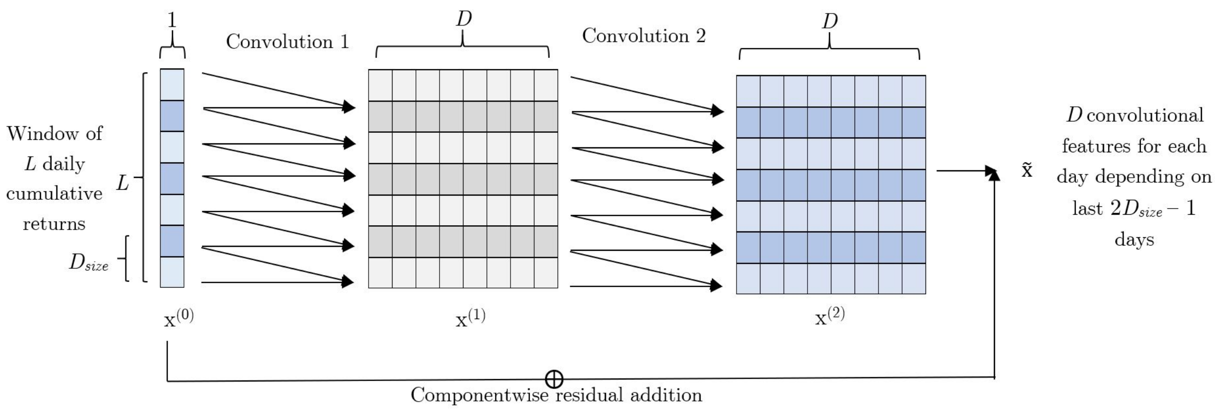

This figure shows the structure of our convolutional network. The network takes as input a window of L consecutive daily cumulative returns a residual, and outputs D features for each block of D_{size} days. Each of the features is a nonlinear function of the observations in the block, and captures a common pattern.

For a 1-layer CNN without the final nonlinear transformation, i.e. for a simple local linear filter, the patterns can be visualized by the vectors $W_m^{(0)}$. In our case of a 2-layer CNN the local filter can capture more complex patterns as it applies a 3-dimensional weighting scheme in the array $W^{(1)}$ and nonlinear transformations. In order to visualize the type of patterns, we project the local

17

filter into a simple local linear filter. We want to find the basic patterns that activate only one of the D filters, but none of the others, i.e. we are looking for an orthogonal representation of the projection.[6]

The example plots for local filters in Figure 3 are projections of our higher dimensional nonlinear filter into two-dimensional linear filters. The examples show some of the most important local filters for our empirical benchmark model. While these projections are of course not complete representations of the nonlinear filters of the CNN, they provide an intuition for the type of patterns which are activated by specific filters. Our 2-layer CNN network has a local window size of $D_{\text{size}} = 2$, but because of the 2-layer structure it captures information from two neighboring points. Hence, the projection on a one-dimensional linear filter has a local window size of three as depicted in Figure 3.

The output of the CNN $\tilde{x} \in \mathbb{R}^{L \times D}$ is used as an input to the transformer. The CNN projection provides a more informative representation of the dynamics than the original time-series as it captures the relative local dependencies between data points. However, by construction the CNN is only a local representation, and we need the transformer network to detect the global patterns. A transformer network is a model of temporal dependencies between local filters. Given the local structure $\tilde{x}$ the transformer estimates the temporal interactions between the L different blocks by computing a "global pattern projection":

Assume there are H different global patterns. The transformer will calculate projections on these H patterns with the "attention weights". We first introduce a simplified linear projection model before extending it to the actual transformer. For each of the $i = 1, ..., H$ patterns we have projections defined by $\alpha_i \in \mathbb{R}^{L \times L}$:

$$h_{i,l}^{\text{simple}} = \sum_{j=1}^{L} \alpha_{i,l,j} \tilde{x}_j \qquad \text{for } l = 1, ..., L \text{ and } i = 1, .., H.$$

The "attention function" $\boldsymbol{\alpha}_i(.,.) \in [0, 1]$ captures dependencies between the local patterns:

$$\alpha_{i,l,j} = \boldsymbol{\alpha}_i \left(\tilde{x}_l, \tilde{x}_j\right) \qquad \text{for } i = 1, ..., H.$$

Each projection $h_{i,l}^{\text{simple}}$ is called an "attention head". These attention heads $h_{i,l}^{\text{simple}}$ could be interpreted as "loadings" or "exposure" for a specific "pattern factor" α_i. For example, a global upward trend can be captured by an attention function that puts weight on subsequent local upward trends. Another example would be sinusoidal mean reversion patterns which would put weights on

[6] A local filter can be formalized as a mapping from the local D_{size} points of a sequence to the activation of the D filters: $\phi : \mathbb{R}^{D_{\text{size}}} \to \mathbb{R}^D$. Denote by $e_d \in \mathbb{R}^D$ a vector that is 0 everywhere except for the value 1 at position d, i.e. $e_d = \begin{pmatrix} 0 & \cdots & 1 & 0 & \cdots 0 \end{pmatrix}$. Fundamentally, we want to invert the local filter to obtain $\phi^{-1}(e_d)$ to find the local sequences that only activates filter d. In general, the inverse is a set and not unique. Our example basic patterns in Figure 3 solve

$$\operatorname{argmin}_{x_{\text{loc},d} \in \mathbb{R}^{D_{\text{size}}}} \|\phi(x_{\text{loc},d}) - e_d\|_2 \qquad \text{for } d = 1, ..., D.$$

alternating "curved" local basis patterns. The projection on these weights captures how much a specific time-series $\tilde{x}$ is exposed to this global pattern. Hence, $h_{i,l}^{\text{simple}}$ measures the exposure to the global pattern i at time l of the time-series $\tilde{x}$. Each attention head can focus on a specific global pattern, which we then combine to obtain our signal.

The fundamental challenge is to learn attention functions that can model complex dependencies. The crucial innovation in transformers is their modeling of the attention functions $\boldsymbol{\alpha}_i$ and attention heads h_i. In order to deal with the high dimensionality of the problem, transformers consider lower dimensional projections of $\tilde{x}$ into $\mathbb{R}^{D/H}$ and use the lower-dimensional scaled dot product attention mechanism for $\boldsymbol{\alpha}_i$ as explained in Appendix B.C. More specifically, each attention head $h_i \in \mathbb{R}^{L\times D/H}$ is based on[7] the projected input $\tilde{x}W_i^V$ with $W_i^V \in \mathbb{R}^{D\times D/H}$ and $\alpha_i \in \mathbb{R}^{L\times L}$:

$$h_i = \alpha_i \tilde{x} W_i^V \qquad \text{for } i = 1, .., H.$$

The projection on all global basis patterns $h \in \mathbb{R}^{L\times D/H}$ is given by a weighted linear combination of the different attention heads

$$h^{\text{proj}} = \begin{pmatrix} h_1 & \cdots & h_H \end{pmatrix} W^O$$

with $W^O \in \mathbb{R}^{D\times D}$. This final projection can, for example, model a combination of a global trend and mean reversion patterns. In conclusion, h^{proj} represents the time-series in terms of the H global patterns. This is analogous to a Fourier filter, but without pre-specifying the global patterns a priori. All parts of the CNN+Transformer network, i.e. the local patterns, the attention functions and the projections on global patterns, are estimated from the data.

The trading signal $\theta^{\text{CNN+Trans}}$ equals the global pattern projection for the final cumulative return projection[8] h_L^{proj}:

$$\theta^{\text{CNN+Trans}} = h_L^{\text{proj}} \in \mathbb{R}^H,$$

which is then used as input to a time-wise feedforward network allocation function

$$\boldsymbol{w}^{\epsilon|\text{CNN+Trans}}\left(\theta^{\text{CNN+Trans}}\right) = \boldsymbol{g}^{\text{FFN}}\left(\theta^{\text{CNN+Trans}}\right).$$

The separation between signal and allocation is not uniquely identified as we use a joint optimization problem. We have chosen a separation that maps naturally into the classical examples considered in the previous subsections. Figure 5 illustrates the transformer network architecture. We have presented a 1-layer transformer network, which is our benchmark model. The transformed data

[7]The actual implementation also includes bias terms which we neglect here for simplicity. The Appendix provides the implementation details.

[8]In principle, we can use the complete matrix $h \in \mathbb{R}^{L\times D}$ as the signal. However, conceptually the global pattern at the end of the time period should be the most relevant for the next realization of the process. We have also implemented a transformer that uses the full matrix, with similar results and the variable importance rankings suggest that only h_L is selected in the allocation function.

Figure 5: Transformer Network Architecture

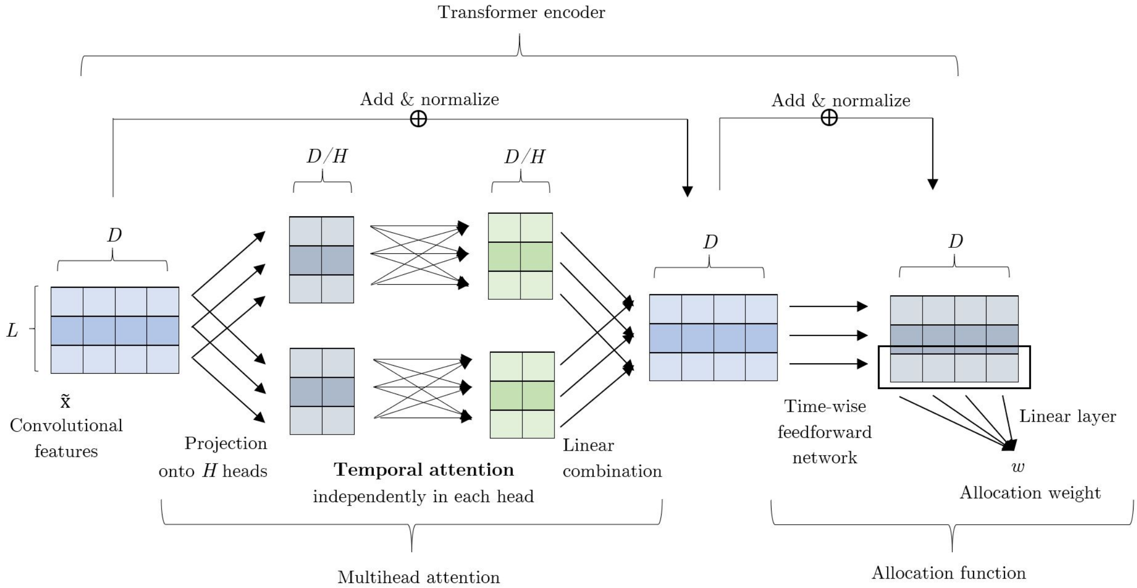

This figure shows the structure of our transformer network. The model takes as input the matrix $\tilde{x} \in \mathbb{R}^{L \times D}$ that we obtain as output of the convolutional network depicted in Figure 4, which contains D features for each of the L blocks of the original time series. These features are projected onto H attention heads, which independently quantify the temporal relation between the blocks and aggregate them into hidden states. These hidden states are finally combined by a feedforward network and a linear layer to predict the optimal allocation for the residual in the next day.

can be used as input in more iterations of the transformer to obtain a multi-layer transformer.

We illustrate the CNN+Transformer model in the first row of Figure 2 for an empirical residual example. First, it is apparent that the cumulative returns of the strategy in subplot (c) outperforms the previous two models. This is because the allocation weights in subplot (a) capture not only the low frequency reversal patterns, but also the high-frequency cycles and trend components. This also implies that the allocation weights change more frequently to capture the higher frequency components. This more sophisticated allocation function requires a more complex signal as illustrated in subplot (b). Each change in the allocation can be traced back to changes in at least one of the signals. While the signals themselves are hard to interpret, we will leverage the transformer structure to extract interpretable "global dependency factors" in our main analysis. Figure A.2 in the Appendix provides another example to illustrate the differences between the three models. This example has a strong negative trend component with a superimposed mean-reversion. Only the CNN+Transformer captures both type of patterns.

III. Empirical Analysis

A. Data

We collect daily equity return data for the securities on CRSP from January 1978 through December 2016. We use the first part of the sample to estimate the various factor models, which gives us the residuals for the time period from January 1998 to December 2016 for the arbitrage trading. The arbitrage strategies trade on a daily frequency at the close of each day. We use the daily adjusted returns to account for dividends and splits and the one-month Treasury bill rates from the Kenneth French Data Library as the risk-free rate. In addition, we complement the stock returns with the 46 firm-specific characteristics from Chen et al. (2019), which are listed in Table A.I. All these variables are constructed either from accounting variables from the CRSP/Compustat database or from past returns from CRSP. The full details on the construction of these variables are in the Internet Appendix of Chen et al. (2019).

Our analysis uses only the most liquid stocks in order to avoid trading and market friction issues. More specifically, we consider only the stocks whose market capitalization at the previous month was larger than 0.01% of the total market capitalization at that previous month, which is the same selection criterion as in Kozak et al. (2020). On average this leaves us with approximately the largest 550 stocks, which correspond roughly to the S&P 500 index. This is an unbalanced dataset, as the stocks that we consider each month need not be the same as in the next month, but it is essentially balanced on a daily frequency in rolling windows of up to one year in our trading period from 1998 through 2016. For each stock we have its cross-sectionally centered and rank-transformed characteristics of the previous month. This is a standard transformation to deal with the different scales which is robust to outliers and time-variation, and has also been used in Chen et al. (2019), Kozak et al. (2020), Kelly et al. (2019), and Freyberger et al. (2020).

Our daily residual time-series start in 1998 as we have a large number of missing values in daily individual stock returns prior to this date, but almost no missing daily values in our sample.[9] We want to point out that the time period after 1998 also seems to be more challenging for arbitrage trading or factor trading, and hence our results can be viewed as conservative lower bounds.

B. *Factor model estimation*

As discussed in Section II.A, we construct the statistical arbitrage portfolios by using the residuals of a general factor model for the daily excess returns of a collection of stocks. In particular, we consider the three empirically most successful families of factor models in our implementation. For each family, we conduct a rolling window estimation to obtain daily residuals out of sample from 1998 through 2016. This means that the residual composition matrix Φ_{t-1} of equation (1) depends only on the information up to time $t-1$, and hence there is no look-ahead bias in trading the residuals. The rolling window estimation is necessary because of the time-variation in risk exposure of individual stocks and the unbalanced nature of a panel of individuals stock returns.

The three classes of factor models consists of pre-specified factors, latent unconditional factors and latent conditional factors:

[9]Of all the stocks that have daily returns observed in a the local lockback window of $L = 30$ days, only 0.1% have a missing return the next day for the out-of-sample trading, in which case we do not trade this stock. Hence, our data set of stocks with market capitalization higher than 0.01% of the total market capitalization, has essentially no missing daily values on a local window for the time period after 1998.

1. **Fama-French factors:** We consider 1, 3, 5 and 8 factors based on various versions and extensions of the Fama-French factor models and downloaded from the Kenneth French Data Library. We consider them as tradeable assets in our universe. Each model includes the previous one and adds additional characteristic-based risk factors:
   (a) $K = 1$: CAPM model with the excess return of a market factor
   (b) $K = 3$: Fama-French 3 factor model includes a market, size and value factor
   (c) $K = 5$: Fama-French 3 factor model + investment and profitability factors
   (d) $K = 8$: Fama-French 5 factor model + momentum, short-term reversal and long-term reversal factors.

   We estimate the loadings of the individual stock returns daily with a linear regression on the factors with a rolling window on the previous 60 days and compute the residual for the current day out-of-sample. This is the same procedure as in Carhart (1997). At each day we only consider the stocks with no missing observations in the daily returns within the rolling window, which in any window removes at most 2% of the stocks given our market capitalization filter.
2. **PCA factors:** We consider 1, 3, 5, 8, 10, and 15 latent factors, which are estimated daily on a rolling window. At each time $t-1$, we use the last 252 days, or roughly one trading year, to estimate the correlation matrix from which we extract the PCA factors.[10] Then, we use the last 60 days to estimate the loadings on the latent factors using linear regressions, and compute residuals for the current day out-of-sample. At each day we only consider the stocks with no missing observations in the daily returns during the rolling window, which in any window removes at most 2% of the stocks given our market capitalization filter.
3. **IPCA factors:** We consider 1, 3, 5, 8, 10, and 15 factors in the Instrumented PCA (IPCA) model of Kelly et al. (2019). This is a conditional latent factor model, in which the loadings β_{t-1} are a linear function of the asset characteristics at time $t-1$. As the characteristics change at most each month, we reestimate the IPCA model on rolling window every year using the monthly returns and characteristics of the last 240 months. The IPCA provides the factor weights and loadings for each stock as a function of the stock characteristics. Hence, we do not need to estimate the loadings for individual stocks with an additional time-series regression, but use the loading function and the characteristics at time $t-1$ to obtain the out-of-sample residuals at time t. The other details of the estimation process are carried out in the way outlined in Kelly et al. (2019).

In addition to the factor models above, we also include the excess returns of the individual stocks without projecting out the factors. This “zero-factor model” simply consists of the original excess returns of stocks in our universe and is denoted as $K = 0$. For each factor model, in our empirical analysis we observe that the cumulative residuals exhibit consistent and relatively regular mean-reverting behavior, with some occasional jumps. After taking out sufficiently many factors, the

[10]This is the same procedure as in Avellaneda and Lee (2010).

residuals of different stocks are only weakly correlated.

C. Implementation

Given the daily out-of-sample residuals from 1998 through 2016 we estimate the trading signal and policy on a rolling window to obtain the out-of-sample returns of the strategy. For each strategy we calculate the annualized sample mean μ, annualized volatility σ and annualized Sharpe ratio[11] $\text{SR} = \frac{\mu}{\sigma}$. The Sharpe ratio represents a risk-adjusted average return. Our main models estimate arbitrage strategies to maximize the Sharpe ratio without transaction costs. In Section III.E we also consider a mean-variance objective and in Section III.K we include transaction costs in the estimation and evaluation.

Our strategies trade the residuals of all stocks, which are mapped back into positions of the original stocks. We use the standard normalization that the absolute values of the individual stock portfolio weights sums up to one, i.e. we use the normalization $\|\omega^R_{t-1}\|_1 = 1$. This normalization implies a leverage constraint as short positions are bounded by one. The trading signal is based on a local lookback window of $L = 30$ days. We show in Section III.I, that the results are robust to this choice and are very comparable for a lookback window of $L = 60$ days. Our main results use a rolling window of 1,000 days to estimate the deep learning models. For computational reasons we re-estimate the network only every 125 days using the previous 1,000 days. Section III.I shows that our results are robust to this choice. Our main results show the out-of-sample trading performance from January 2002 to December 2016 as we use the first four years to estimate the signal and allocation function.

The hyperparameters for the deep learning models are based on the validation results summarized in Appendix C.A. Our benchmark model is a 2-layer CNN with $D = 8$ local convolutional filters and local window size of $D_{size} = 2$ days. The transformer has $H = 4$ attention heads, which can be interpreted as capturing four different global patterns. The results are extremely robust to the choice of hyperparameters. Appendix B.D includes all the technical details for implementing the deep learning models.

D. Main Results

Table I displays the main results for various arbitrage models. It reports the annualized Sharpe ratio, mean return and volatility for our principal deep trading strategy CNN+Transformer and the two benchmark models, Fourier+FFN and OU+Threshold, for every factor model described in Section III.B. The CNN+Transformer model and Fourier+FFN model are estimated with a Sharpe ratio objective. We obtain the daily out-of-sample residuals for different number of factors K for the time period January 1998 to December 2016. The daily returns of the out-of-sample arbitrage trading is then evaluated from January 2002 to December 2016, as we use a rolling window of four

[11]We obtain the annualized metrics from the daily returns using the standard calculations $\mu = \frac{252}{T}\sum_{t=1}^{T} R_t$ and $\sigma = \sqrt{\frac{252}{T}\sum_{t=1}^{T}(R_t - \mu)^2}$.

years to estimate the deep learning models.

Table I: OOS Annualized Performance Based on Sharpe Ratio Objective

	Factors	Fama-French			PCA			IPCA		
Model	K	SR	μ	σ	SR	μ	σ	SR	μ	σ
	0	1.64	13.7%	8.4%	1.64	13.7%	8.4%	1.64	13.7%	8.4%
	1	3.68	7.2%	2.0%	2.74	15.2%	5.5%	3.22	8.7%	2.7%
CNN	3	3.13	5.5%	1.8%	3.56	16.0%	4.5%	3.93	8.6%	2.2%
+	5	3.21	4.6%	1.4%	3.36	14.3%	4.2%	4.16	8.7%	2.1%
Trans	8	2.49	3.4%	1.4%	3.02	12.2%	4.0%	3.95	8.2%	2.1%
	10	-	-	-	2.81	10.7%	3.8%	3.97	8.0%	2.0%
	15	-	-	-	2.30	7.6%	3.3%	4.17	8.4%	2.0%
	0	0.36	4.9%	13.6%	0.36	4.9%	13.6%	0.36	4.9%	13.6%
	1	0.89	3.2%	3.5%	0.80	8.4%	10.6%	1.24	6.3%	5.0%
Fourier	3	1.32	3.5%	2.7%	1.66	11.2%	6.7%	1.77	7.8%	4.4%
+	5	1.66	3.1%	1.8%	1.98	12.4%	6.3%	1.90	7.7%	4.1%
FFN	8	1.90	3.1%	1.6%	1.95	10.1%	5.2%	1.94	7.8%	4.0%
	10	-	-	-	1.71	8.2%	4.8%	1.93	7.6%	3.9%
	15	-	-	-	1.14	4.8%	4.2%	2.06	7.9%	3.8%
	0	-0.18	-2.4%	13.3%	-0.18	-2.4%	13.3%	-0.18	-2.4%	13.3%
	1	0.16	0.6%	3.8%	0.21	2.1%	10.4%	0.60	3.0%	5.1%
OU	3	0.54	1.6%	3.0%	0.77	5.2%	6.8%	0.88	3.8%	4.3%
+	5	0.38	0.9%	2.3%	0.73	4.4%	6.1%	0.97	3.8%	4.0%
Thresh	8	1.16	2.8%	2.4%	0.87	4.4%	5.1%	0.91	3.5%	3.8%
	10	-	-	-	0.63	2.9%	4.6%	0.86	3.1%	3.6%
	15	-	-	-	0.62	2.4%	3.8%	0.93	3.2%	3.5%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) of our three statistical arbitrage models for different numbers of risk factors K, that we use to obtain the residuals. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. CNN+Trans denotes the convolutional network with transformer model, Fourier+FFN estimates the signal with a FFT and the policy with a feedforward neural network and lastly, OU+Thres is the parametric Ornstein-Uhlenbeck model with thresholding trading policy. The two deep learning models are calibrated on a rolling window of four years and use the Sharpe ratio objective function. The signals are extracted from a rolling window of $L = 30$ days. The $K = 0$ factor model corresponds to directly using stock returns instead of residuals for the signal and trading policy.

First, we confirm that it is crucial to apply arbitrage trading to residuals and not individual stock returns. The stock returns, denoted as the $K = 0$ model, perform substantially worse than any type of residual within the same model and factor family. This is not surprising as residuals for an appropriate factor model are expected to be better described by a model that captures mean reversion. Importantly, individual stock returns are highly correlated and a substantial part of the returns is driven by the low dimensional factor component.[12] Hence, the complex nonparametric models are actually not estimated on many weakly dependent residual time-series, but most time-series have redundant information. In other words, the models are essentially calibrated on only a few factor time-series, which severely limits the structure that can be estimated. However, once

[12] Pelger (2020) shows that around one third of the individual stock returns is explained by a latent four-factor model.

Table II: Significance of Arbitrage Alphas based on Sharpe Ratio Objective

CNN+Trans model

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	11.6%	6.4***	30.3%	13.7%	6.3***	11.6%	6.4***	30.3%	13.7%	6.3***	11.6%	6.4***	30.3%	13.7%	6.3***
1	7.0%	14***	2.4%	7.2%	14***	14.9%	10***	0.6%	15.2%	11***	8.1%	12***	9.5%	8.7%	12***
3	5.5%	12***	1.2%	5.5%	12***	15.8%	14***	1.7%	16.0%	14***	8.2%	15***	6.0%	8.6%	15***
5	4.5%	12***	2.3%	4.6%	12***	14.1%	13***	1.3%	14.3%	13***	8.3%	16***	3.9%	8.7%	16***
8	3.3%	9.4***	2.1%	3.4%	9.6***	12.0%	12***	0.9%	12.2%	12***	7.8%	15***	5.0%	8.2%	15***
10	-	-	-	-	-	10.5%	11***	0.7%	10.7%	11***	7.7%	15***	4.0%	8.0%	15***
15	-	-	-	-	-	7.5%	8.8***	0.5%	7.6%	8.9***	8.1%	16***	4.2%	8.4%	16***

Fourier+FFN model

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	2.7%	0.8	8.6%	4.9%	1.4	2.7%	0.8	8.6%	4.9%	1.4	2.7%	0.8	8.6%	4.9%	1.4
1	3.0%	3.3**	3.3%	3.2%	3.5***	7.4%	2.7**	3.3%	8.4%	3.1**	4.8%	4.0***	16.4%	6.3%	4.8***
3	3.2%	4.7***	4.2%	3.5%	5.1***	10.9%	6.3***	2.2%	11.2%	6.4***	6.8%	6.4***	13.0%	7.8%	6.9***
5	2.9%	6.1***	3.5%	3.1%	6.4***	12.1%	7.5***	1.5%	12.4%	7.6***	6.7%	6.9***	13.3%	7.7%	7.4***
8	3.0%	7.2***	3.2%	3.1%	7.4***	10.0%	7.5***	0.9%	10.1%	7.6***	6.8%	7.0***	13.3%	7.8%	7.5***
10	-	-	-	-	-	8.0%	6.5***	1.0%	8.2%	6.6***	6.8%	7.1***	12.7%	7.6%	7.5***
15	-	-	-	-	-	4.7%	4.3***	0.4%	4.8%	4.4***	7.1%	7.6***	12.2%	7.9%	8.0***

OU+Thresh model

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	-4.5%	-1.4	13.4%	-2.4%	-0.7	-4.5%	-1.4	13.4%	-2.4%	-0.7	-4.5%	-1.4	13.4%	-2.4%	-0.7
1	-0.2%	-0.2	13.5%	0.6%	0.6	0.7%	0.3	6.3%	2.1%	0.8	1.7%	1.4	18.9%	3.0%	2.3*
3	0.9%	1.2	10.4%	1.6%	2.1*	4.3%	2.5*	4.3%	5.2%	3.0**	2.6%	2.6**	18.8%	3.8%	3.4***
5	0.5%	0.9	6.8%	0.9%	1.5	3.7%	2.4*	3.2%	4.4%	2.8**	2.8%	3.0**	17.7%	3.8%	3.8***
8	0.6%	1.2	5.5%	1.0%	1.9	3.9%	3.0**	1.9%	4.4%	3.4***	2.3%	2.6**	17.6%	3.5%	3.6***
10	-	-	-	-	-	2.6%	2.2*	1.4%	2.9%	2.4*	2.1%	2.5*	17.6%	3.1%	3.3***
15	-	-	-	-	-	2.1%	2.1*	0.7%	2.4%	2.4*	2.3%	2.8**	18.1%	3.2%	3.6***

This table shows the out-of-sample pricing errors α of the arbitrage strategies relative of the Fama-French 8 factor model and their mean returns μ for the different arbitrage models and different number of factors K that we use to obtain the residuals. We run a time-series regression of the out-of-sample returns of the arbitrage strategies on the 8-factor model (Fama-French 5 factors + momentum + short-term reversal + long-term reversal) and report the annualized α, accompanying t-statistic value t_α, and the R^2 of the regression. In addition, we report the annualized mean return μ along with its accompanying t-statistic t_μ. The hypothesis test are two-sided and stars indicate p-values of 5% (*), 1% (**), and 0.1% (***). All results use the out-of-sample daily returns from January 2002 to December 2016 and the deep learning models are based on a Sharpe ratio objective.

we extract around $K = 5$ factors with any of the different factor models, the performance does not substantially increase by adding more factors. This suggests that most of commonality is explained by a small number of factors.

Second, the CNN+Transformer model strongly dominates the other benchmark models in terms of Sharpe ratio and average return. The Sharpe ratio is approximately twice as large as for a comparable Fourier+FFN model and four times higher for the corresponding parametric OU+Threshold model. Using IPCA residuals, the CNN+Transformer achieves the impressive out-of-sample Sharpe ratio of around 4, in spite of trading only the most liquid large cap stocks and the time period after 2002. The mean returns of the CNN+Transformer are similar to the Fourier+FFN model, but have substantially smaller volatilities, which results in the higher Sharpe ratios. The parametric

Figure 6. Cumulative OOS Returns of Different Arbitrage Strategies

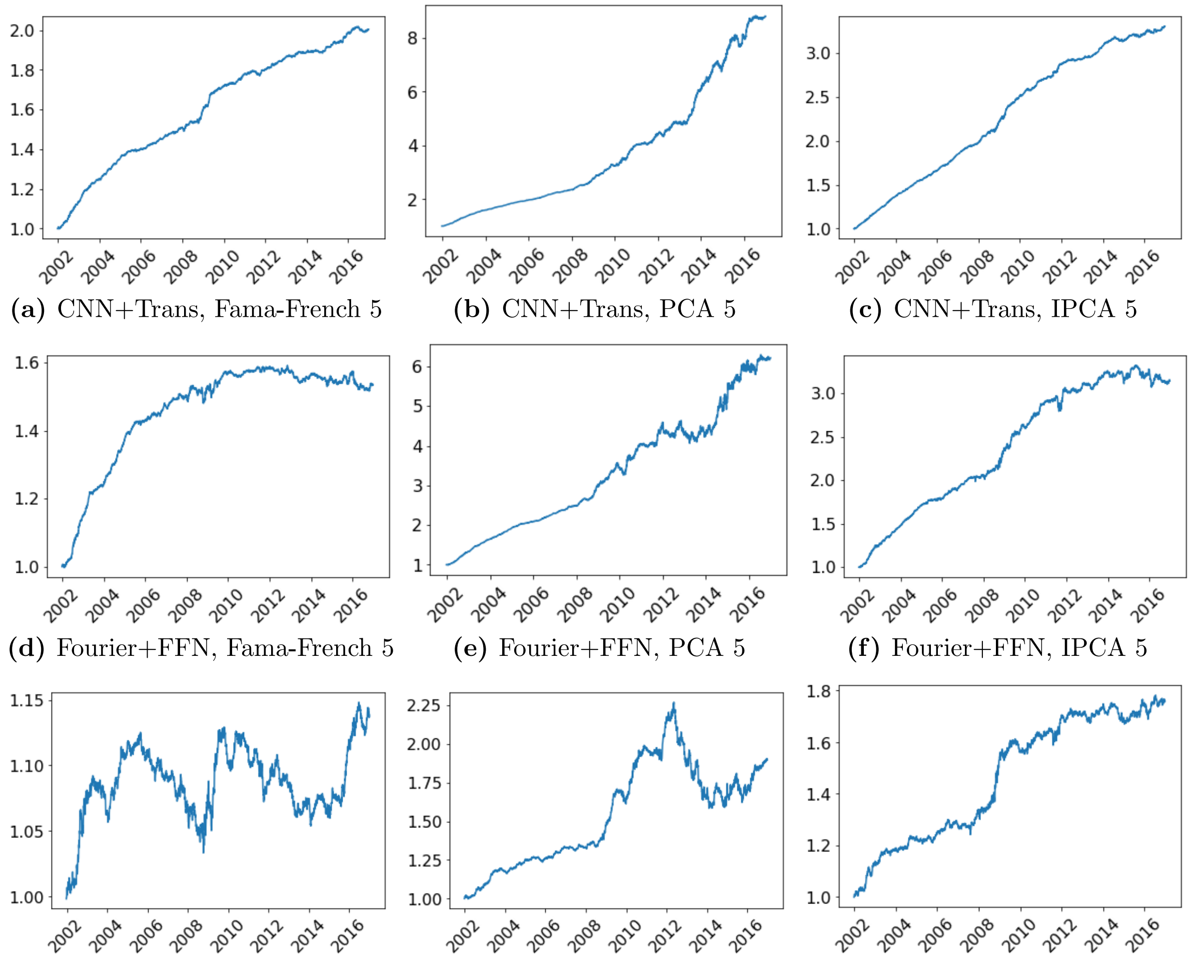

(a) CNN+Trans, Fama-French 5 **(b)** CNN+Trans, PCA 5 **(c)** CNN+Trans, IPCA 5

(d) Fourier+FFN, Fama-French 5 **(e)** Fourier+FFN, PCA 5 **(f)** Fourier+FFN, IPCA 5

(g) OU+Thresh Fama-French 5 **(h)** OU+Thresh PCA 5 **(i)** OU+Thresh IPCA 5

These figures show the cumulative daily returns of the arbitrage strategies for our representative models on the out-of-sample trading period between January 2002 and December 2016. We estimate the optimal arbitrage trading strategies for our three benchmark models based on the out-of-sample residuals of the Fama-French, PCA and IPCA 5-factor models. The deep learning models use the Sharpe ratio objective.

mean-reversion model achieves positive mean returns with Sharpe ratios close to one for the IPCA residuals, but as expected is too restrictive relative to the flexible models. The Fourier+FFN has the same flexibility as the CNN+Transformer in its allocation function, but is restricted to a pre-specified signal structure. The difference in performance quantifies the importance of extracting the complex time-series signals.

Third, the average return of the arbitrage strategies is large in spite of the leverage constraints. Normalizing the individual stock weights w^R_{t-1} so sum up in absolute value to one limits the short-selling. The CNN+Transformer with a five-factor PCA residual achieves an attractive annual mean return of around 14%. This means that the strategies do not require an infeasible amount of leverage to yield an average return that might be required by investors. In other words, the high

Sharpe ratios are not the results of vanishing volatility but a combination of high average returns with moderate volatility.

Fourth, the arbitrage strategies are qualitatively robust to the choice of factor models to obtain residuals. The Fama-French and PCA factor lead to very similar Sharpe ratio results, suggesting that they explain a similar amount of co-movement in the data. However, as the mean returns of PCA factors are usually higher than the mean returns of the Fama-French factors, the risk factors are different. This confirms the findings of Pelger (2020) and Lettau and Pelger (2020b), who show that PCA factors do not coincide with Fama-French type factors and explain different mean returns. The IPCA factors use the additional firm-specific characteristic information. The resulting residuals achieve the highest Sharpe ratios, which illustrates that conditional factor models can capture more information than unconditional models. Including the momentum and reversal factors in the Fama-French 8 factor models to obtain residuals still results in profitable arbitrage strategies. Hence, the arbitrage strategies are not simply capturing a price trend risk premium.

The returns of the CNN+Transformer arbitrage strategies are statistically significant and not subsumed by conventional risk factors. Table II reports the out-of-sample pricing errors α of the arbitrage strategies relative of the Fama-French 8 factor model and their mean returns μ. We run a time-series regression of the out-of-sample returns of the arbitrage strategies on the 8-factor model (Fama-French 5 factors + momentum + short-term reversal + long-term reversal) and report the annualized α, accompanying t-statistic value t_α, and the R^2 of the regression. In addition, we report the annualized mean return μ along with its accompanying t-statistic t_μ. The arbitrage strategies for the CNN+Transformer model for $K \geq 1$ are all statistically significant and not explained by the Fama-French 5 factors or price trend factors. Importantly, the pricing errors are essentially as large as the mean returns, which implies that the returns of the CNN+Transformer arbitrage strategies do not carry any risk premium of these eight factors. This is supported by the R^2 values, which are close to zero for the Fama-French or PCA residuals, and hence confirm that these arbitrage portfolios are essentially orthogonal to the Fama-French 8 factors. In contrast, one third of the individual return variation for $K = 0$ is explained by those risk factors. However, even in that case the pricing errors are significant. The residuals of IPCA factors have a higher correlation with the Fama-French 8 factors, suggesting that the conditional IPCA factor model extracts factors that are inherently different from the conventional risk factors. Note that the residuals of a Fama-French 8 factor model are not mechanically orthogonal in the time-series regression on the Fama-French 8 factors, as we construct out-of-sample residuals based on rolling window estimates. The parametric arbitrage strategies are largely explained by conventional risk factors. The third subtable in Table II shows that the residuals of the OU+Threshold model for Fama-French of PCA residuals do not have statistically significant pricing errors on a 1% level.

The CNN+Transformer has a consistent out-of-sample performance and is not affected by various negative events. Figure 6 shows the cumulative out-of-sample returns of the arbitrage strategies for our representative models. We select the residuals of the various five-factor models as including additional factors has only minor effects on the performance. Note that the CNN+Transformer model has consistently almost always positive returns, while maintaining a low volatility and avoiding any large losses. Importantly, the performance of the CNN+Transformer is nearly completely immune to both the "quant quake" which affected quantitative trading groups and funds engaging in statistical arbitrage in August 2007 (Barr (2007)), and the period of poor performance in quant funds during 2011–2012 (Ebrahimi (2013)). The Fourier+FFN model also performs similarly well until the financial crisis, but its risk increases afterwards as displayed by the larger volatility and larger drawdowns. The performance of the parametric model is visibly inferior. This illustrates that although all strategies trade the same residuals, which should be orthogonal to common market risk, profitable arbitrage trading requires an appropriate signal and allocation policy.

Table III: OOS Annualized Performance Based on Mean-Variance Objective

CNN+Trans strategy, mean-variance objective function

	Fama-French			PCA			IPCA		
K	SR	μ	σ	SR	μ	σ	SR	μ	σ
0	0.83	9.5%	11.4%	0.83	9.5%	11.4%	0.83	9.5%	11.4%
1	3.15	10.5%	3.3%	2.21	27.3%	12.3%	2.83	15.9%	5.6%
3	2.95	7.8%	2.6%	2.38	22.6%	9.5%	3.13	17.9%	5.7%
5	3.03	5.9%	2.0%	2.75	19.6%	7.1%	3.21	18.2%	5.7%
8	2.96	4.2%	1.4%	2.68	16.6%	6.2%	3.18	17.0%	5.4%
10	-	-	-	2.67	15.3%	5.7%	3.21	16.6%	5.2%
15	-	-	-	2.20	8.7%	4.0%	3.34	16.3%	4.9%

Fourier+FFN strategy, mean-variance objective function

	Fama-French			PCA			IPCA		
K	SR	μ	σ	SR	μ	σ	SR	μ	σ
0	0.28	5.5%	19.3%	0.28	5.5%	19.3%	0.28	5.5%	19.3%
1	0.38	2.5%	6.7%	0.48	16.6%	34.8%	0.56	9.7%	17.2%
3	1.16	4.3%	3.7%	0.34	32.1%	93.1%	1.06	17.6%	16.7%
5	1.30	3.1%	2.4%	0.37	22.5%	61.2%	1.17	17.0%	14.5%
8	1.73	3.6%	2.0%	0.67	17.4%	25.9%	1.21	14.4%	11.9%
10	-	-	-	0.45	7.4%	16.4%	1.06	12.6%	11.9%
15	-	-	-	0.56	5.7%	10.2%	1.17	12.1%	10.4%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) of our CNN+Transformer and Fourier+FFN models for different numbers of risk factors K, that we use to obtain the residuals. We use a mean-variance objective function with risk aversion $\gamma = 1$. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. The two deep learning models are calibrated on a rolling window of four years. The signals are extracted from a rolling window of $L = 30$ days. The $K = 0$ factor model corresponds to directly using stock returns instead of residuals for the signal and trading policy.

E. Mean-Variance Objective

The deep learning statistical arbitrage strategies can achieve high average returns in spite of leverage constraints. Our main deep learning models are estimated with a Sharpe ratio objective. As the sum of absolute stock weights is normalized to one, the arbitrage strategies impose an implicit leverage constraint. We show that the average return can be increased while maintaining this leverage constraint. For this purpose we change the objective for the deep learning model to a mean-variance objective. In order to illustrate the effect of the different objective function, we set the risk aversion parameter to $\gamma = 1$.

Table IV: Significance of Arbitrage Alphas based on Mean-Variance Objective

CNN+Trans model

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	5.8%	2.2*	19.6%	9.5%	3.2**	5.8%	2.2*	19.6%	9.5%	3.2**	5.8%	2.2*	19.6%	9.5%	3.2**
1	9.9%	12***	7.1%	10.5%	12***	26.3%	8.3***	1.6%	27.3%	8.6***	14.0%	11***	23.5%	15.9%	11***
3	7.5%	11***	5.3%	7.8%	11***	22.1%	9.1***	2.2%	22.6%	9.2***	16.6%	12***	17.6%	17.9%	12***
5	5.7%	11***	5.3%	5.9%	12***	19.0%	10***	3.2%	19.6%	11***	16.7%	12***	16.0%	18.2%	12***
8	4.4%	9.8***	3.6%	4.6%	10***	16.3%	10***	1.6%	16.6%	10***	15.5%	12***	18.3%	17.0%	12***
10	-	-	-	-	-	14.8%	10***	1.7%	15.3%	10***	15.2%	13***	20.6%	16.6%	12***
15	-	-	-	-	-	8.5%	8.4***	0.9%	8.7%	8.5***	14.8%	13***	21.6%	16.3%	13***

Fourier+FFN model

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	3.2%	0.7	8.4%	5.5%	1.1	3.2%	0.7	8.4%	5.5%	1.1	3.2%	0.7	8.4%	5.5%	1.1
1	2.8%	1.6	1.8%	2.5%	1.5	15.4%	1.7	1.3%	16.6%	1.9	7.9%	1.8	2.6%	9.7%	2.2*
3	4.1%	4.4***	3.4%	4.3%	4.5***	30.3%	1.3	0.1%	32.1%	1.3	17.4%	4.1***	1.9%	17.6%	4.1***
5	2.9%	4.8***	3.1%	3.1%	5.0***	21.0%	1.3	0.1%	22.5%	1.4	15.9%	4.3***	2.6%	17.0%	4.5***
8	3.5%	6.8***	2.3%	3.6%	7.0***	17.4%	2.6**	0.3%	17.2%	2.6**	12.9%	4.3***	4.4%	14.4%	4.7***
10	-	-	-	-	-	7.1%	1.7	0.3%	7.4%	1.8	11.7%	3.9***	3.5%	12.6%	4.1***
15	-	-	-	-	-	5.5%	2.1*	0.1%	5.7%	2.2*	11.3%	4.3***	4.0%	12.1%	4.5***

This table shows the out-of-sample pricing errors α of the arbitrage strategies relative of the Fama-French 8 factor model and their mean returns μ for the different arbitrage models and different number of factors K that we use to obtain the residuals. We use a mean-variance objective function with risk aversion $\gamma = 1$. We run a time-series regression of the out-of-sample returns of the arbitrage strategies on the 8-factor model (Fama-French 5 factors + momentum + short-term reversal + long-term reversal) and report the annualized α, accompanying t-statistic value t_α, and the R^2 of the regression. In addition, we report the annualized mean return μ along with its accompanying t-statistic t_μ. The hypothesis test are two-sided and stars indicate p-values of 5% (*), 1% (**), and 0.1% (***). All results use the out-of-sample daily returns from January 2002 to December 2016.

Tables III and IV collect the results for the Sharpe ratio, mean, volatility and significance tests. As expected the Sharpe ratios are slightly lower compared to the corresponding model with Sharpe ratio objective, but the mean returns are substantially increased. The CNN+Transformer model achieves average annual returns around 20% with PCA and IPCA residuals while the volatility is only around half as large as the one of a market portfolio. The mean returns are statistically highly significant and not spanned by conventional risk factors or a price trend risk premium. The Fourier+FFN model can also obtain high average returns, but those come at the cost of a substantial volatility. Overall, we confirm that the more flexible signal extraction function of the CNN+Transformer is crucial for the superior performance.

F. Unconditional Residual Means

The unconditional average of residuals is not a profitable strategy and does not provide information about the potential arbitrage profitability contained in the residuals. A natural question to ask is if the residuals themselves have a risk premium component and if trading an equally weighted portfolio of residuals could be profitable. Table V shows the performance of this simple strategy. If we do not project out any factors ($K = 0$), this strategy essentially trades an equally weighted market portfolio. Table VI reports the test statistics relative to the Fama-French 8 factor model, which completely subsumes the market risk premium. Once we regress out at least 3 factors, the equally weighted residuals have a mean return of around 1% or lower. The low volatility confirms that the residuals are only weakly cross-sectionally dependent and are largely diversified away. The moderately large Sharpe ratios for PCA residuals is a consequence of the near zero volatility. Scaling up the mean returns to a meaningful magnitude would potentially require an unreasonable amount of leverage. Overall, we confirm that residuals need to be combined with a signal and trading policy that takes advantage of the time series patterns in order to achieve a profitable strategy.

Table V: OOS Annualized Performance of Unconditional Average Residuals

Equally Weighted Residuals									
	Fama-French			PCA			IPCA		
K	SR	μ	σ	SR	μ	σ	SR	μ	σ
0	0.52	11.2%	21.4%	0.52	11.2%	21.4%	0.52	11.2%	21.4%
1	0.39	1.9%	4.8%	-0.23	-0.4%	1.5%	0.76	3.2%	4.2%
3	0.18	0.7%	3.7%	0.34	0.3%	0.9%	0.76	2.0%	2.7%
5	0.22	0.8%	3.5%	0.93	0.7%	0.7%	0.63	1.4%	2.3%
8	-0.17	-0.5%	2.9%	1.04	0.6%	0.5%	0.66	1.4%	2.2%
10	-	-	-	0.90	0.4%	0.5%	0.65	1.3%	2.1%
15	-	-	-	1.08	0.4%	0.4%	0.62	1.3%	2.0%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) of equally weighted residuals. We evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. The $K = 0$ factor model corresponds to directly using stock returns instead of residuals for the signal and trading policy.

IPCA factors are close to uncorrelated with conventional risk factors. The R^2 values in Table VI are as expected for the Fama-French factors and, not surprisingly, after regressing out all of those factors, the cross-sectional average of the residuals is essentially orthogonal to those factors. The PCA residuals show a very similar behavior. However, the conditional IPCA model leaves a component in the residuals that it is highly correlated with conventional risk factors. In this sense, the IPCA factors extract a factor model that is quite different from the Fama-French factors.

Importantly, unconditional means and alphas of asset pricing residuals are a poor measure of arbitrage opportunities. The mean and alphas of residuals that are optimally traded based on their time series patterns have mean returns that can be larger by a factor of 50. This implies more generally, that the unconditional perspective of evaluating asset pricing models could potentially overstate the efficiency of markets and the pricing ability of asset pricing models.

Table VI: Significance of Arbitrage Alphas Based on Unconditional Average Residuals

Equally Weighted Residuals

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	1.4%	1.4	97.0%	11.2%	2.0*	1.4%	1.4	97.0%	11.2%	2.0*	1.4%	1.4	97.0%	11.2%	2.0*
1	0.4%	0.4	36.6%	1.9%	1.5	0.0%	0.0	25.8%	-0.4%	-0.9	0.4%	1.1	85.0%	3.2%	2.9**
3	0.4%	0.4	9.6%	0.7%	0.7	0.4%	1.9	13.1%	0.3%	1.3	0.9%	3.3**	84.1%	2.0%	2.9**
5	0.2%	0.2	7.0%	0.8%	0.9	0.7%	4.2***	5.9%	0.7%	3.6***	0.4%	2.0*	89.4%	1.4%	2.4*
8	-0.6%	-0.8	0.7%	-0.5%	-0.7	0.6%	4.5***	4.1%	0.6%	4.0***	0.4%	2.1*	89.3%	1.4%	2.5*
10	-	-	-	-	-	0.5%	3.8***	3.0%	0.4%	3.5***	0.3%	1.9	89.4%	1.3%	2.5*
15	-	-	-	-	-	0.4%	4.3***	2.0%	0.4%	4.2***	0.3%	1.6	89.0%	1.3%	2.4*

This table shows the out-of-sample pricing errors α of cross-sectionally equally weighted residuals relative of the Fama-French 8 factor model and their mean returns μ for the different arbitrage models and different number of factors K that we use to obtain the residuals. We run a time-series regression of the out-of-sample returns of the arbitrage strategies on the 8-factor model (Fama-French 5 factors + momentum + short-term reversal + long-term reversal) and report the annualized α, accompanying t-statistic value t_α, and the R^2 of the regression. In addition, we report the annualized mean return μ along with its accompanying t-statistic t_μ. The hypothesis test are two-sided and stars indicate p-values of 5% (*), 1% (**), and 0.1% (***). All results use the out-of-sample daily returns from January 2002 to December 2016.

G. *Importance of Time-Series Signal*

How important is the flexibility in the signal extraction function relative to the allocation function? So far, we have considered a rigid parametric model for the signal and allocation function and a flexible allocation function but either a pre-specified time-series filter or a data-driven flexible filter. In A.VII in Appendix we also report the results for two additional model variations, which serve as ablation tests emphasizing the central importance of applying a time-series model to extract a signal extraction function from the data.

The first model, OU+FFN, uses the same 4-dimensional OU signal as the OU+Threshold policy, but replaces the threshold allocation function with an FFN allocation function. This FFN allocation function has the same architecture as that of the Fourier+FFN policy, except the input is 4-dimensional instead of 30-dimensional. The results show that even despite using a very flexible allocation function, the results are similar or even worse than the simple parametric thresholding rule. This points to the weakness of the OU signal representation: although the allocation function is a powerful universal approximator, it cannot accomplish much with an information-poor input. If the optimal allocation function given the simple OU signal is well approximated by the parametric thresholding rule, then the nonparametric FFN offers too much flexibility without comparable efficiency, which leads to a noisier estimate of a simple function and hence worse out-of-sample performance.

The second model does not extract a time-series signal from the residuals, but uses the residuals themselves as signal to a flexible FFN allocation function. As the allocation function uses the same type of network as for the CNN+Transformer, Fourier+FFN or OU+FFN, this setup directly assesses the relevance of using a time-series model for the signal. The FFN model also performs worse than the deep learning models that apply a time-series filter to the residuals. This is a good example to emphasize the importance of a time-series model. While FFNs are flexible in learning low dimensional functional relationships, they are limited in learning a complex dependency model if the training data is limited. For example, the FFN is not sufficiently efficient to learn an FFT-like transformation and hence has a substantially worse performance on the original time-series compared to frequency-transformed time-series.

In summary, the flexible data-driven signal extraction function of the CNN+Transformer model seems to be the critical element for statistical arbitrage. A flexible allocation function is not sufficient to compensate for an uninformative signal.

H. Dependency between Arbitrage Strategies

The trading strategies for different factor models are only weakly correlated. In Table A.VI, we report the correlations of the returns of our CNN+Transformer strategies across factor models with 3, 5, and 10 factors, based on the Sharpe ratio objective function strategy.[13] Notably, the correlations between strategies from different factor model families range from roughly 0.2 to 0.45, indicating that strategies for different factor model families can be used as part of a diversification strategy. While the performance of the arbitrage trading for the residuals obtained with different families of factor models is comparable, the factors themselves are different. Hence, even if the arbitrage signal and allocation functions are similar, the resulting strategies can be weakly correlated. The within-family correlations range from 0.4 to 0.85, indicating that the residuals from the same class of factor model capture similar patterns.

I. Stability over Time

Our results are robust to length of the local window to extract the trading signal. We re-estimate the CNN+Transformer model on an extended rolling lookback window of $L = 60$ days, while keeping the rest of the model structure the same. Tables VII and VIII show that the results are robust to the choice of lookback window. Extending the local window to 60 trading days, which is close to three months, leads to essentially the same performance as using only the most recent $L = 30$ trading days to infer the signal. This is further evidence that the arbitrage signal is different from conventional momentum or reversal strategies that incorporate information from longer time periods. As the signal can be inferred from the most recent past, it implies that either the arbitrage signal depends only on the most recent days or that those days are sufficient to infer the relevant time-series structure. In the next section, we provide evidence that the arbitrage trading signals put strong emphasis on most recent two weeks before the trading.

A constant-in-time signal and allocation function captures a large fraction of the arbitrage information. We re-estimate the CNN+Transformer model with a constant model instead of the rolling window calibration. Our main models are estimated on a rolling window of four years, which allows the models to adopt to changing economic conditions. Here we use either the first $T_{\text{train}} = 4$ years (1,000 trading days) or $T_{\text{train}} = 8$ years (2,000 trading days) to estimate the signal and allocation

[13] The correlations for the mean-variance objective function are similar.

Table VII: OOS Annualized Performance of CNN+Trans for 60 Days Lookback Window

	Fama-French			PCA			IPCA		
K	SR	μ	σ	SR	μ	σ	SR	μ	σ
0	1.50	13.5%	9.0%	1.50	13.5%	9.0%	1.50	13.5%	9.0%
1	2.95	9.6%	3.2%	2.68	15.8%	5.9%	3.14	8.8%	2.8%
3	3.21	8.7%	2.7%	3.49	16.8%	4.8%	3.84	9.6%	2.5%
5	3.23	6.8%	2.1%	3.54	16.0%	4.5%	3.90	9.2%	2.4%
8	2.96	4.2%	1.4%	3.02	12.5%	4.2%	3.93	8.7%	2.2%
10	-	-	-	2.67	9.9%	3.7%	3.98	9.2%	2.3%
15	-	-	-	2.36	8.1%	3.4%	4.24	9.6%	2.3%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) of the CNN+Transformer model for different numbers of risk factors K, that we use to obtain the residuals. The signals are extracted from a rolling window of $L = 60$ days. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. The model is calibrated on a rolling window of four years and uses the Sharpe ratio objective function. The $K = 0$ factor model corresponds to directly using stock returns instead of residuals for the signal and trading policy.

Table VIII: Significance of Arbitrage Alphas for 60 Days Lookback Window

CNN+Trans Model , Sharpe objective function, $L = 60$ days lookback window															
	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	11.8%	5.6***	19.5%	13.5%	5.8***	11.8%	5.6***	19.5%	13.5%	5.8***	11.8%	5.6***	19.5%	13.5%	5.8***
1	9.1%	11***	7.2%	9.6%	11***	15.5%	10***	1.2%	15.8%	10***	8.2%	12***	10.1%	8.8%	12***
3	8.3%	12***	7.1%	8.7%	12***	16.5%	13***	2.5%	16.8%	14***	9.2%	15***	9.3%	9.6%	15***
5	6.5%	12***	6.0%	6.8%	13***	15.6%	13***	2.2%	16.0%	14***	8.8%	15***	10.3%	9.2%	15***
8	4.1%	11***	3.2%	4.2%	11***	12.2%	11***	1.6%	12.5%	12***	8.3%	15***	8.9%	8.7%	15***
10	-	-	-	-	-	9.7%	10***	1.0%	9.9%	10***	8.8%	15***	8.3%	9.2%	15***
15	-	-	-	-	-	8.1%	9.1***	0.7%	8.1%	9.1***	9.2%	16***	9.3%	9.6%	16***

This table shows the out-of-sample pricing errors α of the arbitrage strategies relative of the Fama-French 8 factor model and their mean returns μ for the CNN+Transformer model and different number of factors K that we use to obtain the residuals. The signals are extracted from a rolling window of $L = 60$ days. We run a time-series regression of the out-of-sample returns of the arbitrage strategies on the 8-factor model (Fama-French 5 factors + momentum + short-term reversal + long-term reversal) and report the annualized α, accompanying t-statistic value t_α, and the R^2 of the regression. In addition, we report the annualized mean return μ along with its accompanying t-statistic t_μ. The hypothesis test are two-sided and stars indicate p-values of 5% (*), 1% (**), and 0.1% (***). All results use the out-of-sample daily returns from January 2002 to December 2016 and are based on a Sharpe ratio objective.

function, and then keep those functions constant for the remaining out-of-sample trading period. The results are reported in Tables IX and X. As expected the performance decreases relative to a time-varying model with re-estimation, which suggests that there is some degree of time-variation in the signal and allocation function. The longer training window of 8 years results in slightly higher Sharpe ratios than the 4 year window, as the model has more data and more variety in the market environment to learn the arbitrage information. Importantly, the constant CNN+Transformer still substantially outperforms the other benchmark models, Fourier+FFN and OU+Threshold, even if those are estimated on a rolling window. We conclude that the constant signal and allocation function for the CNN+Transformer model already capture a substantial amount of arbitrage information. Therefore, the constant functions serve as a meaningful model to analyze in more detail in the next section.

Table IX: OOS Annualized Performance of CNN+Trans for Constant Model

$T_{\text{train}} = 4$ years

	Fama-French			PCA			IPCA		
K	SR	μ	σ	SR	μ	σ	SR	μ	σ
0	1.10	8.5%	7.8%	1.10	8.5%	7.8%	1.10	8.5%	7.8%
1	1.90	4.5%	2.3%	0.66	5.2%	7.9%	0.94	3.1%	3.3%
3	1.60	3.6%	2.2%	1.65	8.7%	5.3%	1.82	5.3%	2.9%
5	1.81	3.0%	1.7%	1.93	9.8%	5.1%	2.09	5.4%	2.6%
8	1.70	2.5%	1.5%	2.04	9.6%	4.7%	1.89	5.0%	2.6%
10	-	-	-	2.06	9.1%	4.4%	1.77	4.7%	2.7%
15	-	-	-	1.82	7.0%	3.9%	2.09	5.5%	2.7%

$T_{\text{train}} = 8$ years

	Fama-French			PCA			IPCA		
K	SR	μ	σ	SR	μ	σ	SR	μ	σ
0	1.33	12.0%	9.0%	1.33	12.0%	9.0%	1.33	12.0%	9.0%
1	2.06	5.0%	2.4%	1.81	15.2%	8.4%	2.02	8.5%	4.2%
3	2.46	5.3%	2.2%	2.04	13.1%	6.4%	2.47	7.5%	3.0%
5	1.82	3.2%	1.8%	1.91	11.9%	6.2%	2.64	7.6%	2.9%
8	1.48	2.5%	1.7%	1.89	10.8%	5.7%	2.71	8.3%	3.1%
10	-	-	-	1.82	10.0%	5.5%	2.68	8.2%	3.1%
15	-	-	-	1.38	6.2%	4.5%	2.70	7.8%	2.9%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) of the CNN+Transformer model for different numbers of risk factors K. We estimate the model on only once on the first T_{train} days and keep it constant on the remaining test set. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 1998 $+T_{\text{train}}$ to December 2016. The signals are extracted from a rolling window of $L = 30$ days and we use the Sharpe ratio objective function.

J. Estimated Structure

What are the patterns that our CNN+Transformer model can learn and exploit? In order to answer this question, we analyze the different building blocks of our benchmark model and show their structure on representative and informative residuals inputs. Our goal is to ascertain, characterize, and explain the role that the convolutional features and attention heads play in the determination of the final allocation weight and recognition of salient time series patterns. The benchmark model for this section is the CNN+Transformer based on IPCA 5-factor residuals and a Sharpe ratio objective. The model is calibrated on the first 8 years of data and kept constant, which allows us to study the signal and allocation function.

As an illustrative example, Figure 7 shows the allocation and return on representative residuals. The left subplot displays an out-of-sample time-series of cumulative returns of a randomly selected residual and its value in the allocation function. We normalize the allocation weight to have an absolute value of one, that is, for this illustrative example we only trade this particular residual. The right subplots depicts the payoff of trading the specific residual with the displayed allocation function. The first residual shows mean-reversion patterns, which are successfully detected and exploited by the function $\boldsymbol{w}^{\epsilon|\text{CNN+Trans}}$. The second residual has a downward trend, which is

Table X: Significance of Arbitrage Alphas for Constant Model

CNN+Trans model, Sharpe objective function, $T_{\text{train}} = 4$ years

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	8.4%	4.2***	3.0%	8.5%	4.3***	8.4%	4.2***	3.0%	8.5%	4.3***	8.4%	4.2***	3.0%	8.5%	4.3***
1	4.0%	6.8***	5.9%	4.5%	7.3***	4.1%	2.0*	4.5%	5.2%	2.5*	3.1%	3.7***	1.6%	3.1%	3.6***
3	3.2%	5.7***	4.9%	3.6%	6.2***	8.2%	6.1***	2.7%	8.7%	6.4***	5.3%	7.4***	11.7%	5.3%	7.0***
5	2.8%	6.6***	4.3%	3.0%	7.0***	9.3%	7.1***	1.8%	9.8%	7.5***	5.5%	8.6***	8.3%	5.4%	8.1***
8	2.3%	6.1***	5.1%	2.5%	6.6***	9.0%	7.5***	2.2%	9.6%	7.9***	5.0%	7.7***	8.2%	5.0%	7.3***
10	-	-	-	-	-	8.6%	7.5***	1.9%	9.1%	8.0***	5.1%	8.0***	16.6%	4.7%	6.9***
15	-	-	-	-	-	6.8%	6.8***	1.0%	7.0%	7.1***	5.8%	9.3***	17.6%	5.5%	8.1***

CNN+Trans model, Sharpe objective function, $T_{\text{train}} = 8$ years

	Fama-French					PCA					IPCA				
K	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ	α	t_α	R^2	μ	t_μ
0	10.1%	4.1***	18.1%	12.0%	4.4***	10.1%	4.1***	18.1%	12.0%	4.4***	10.1%	4.1***	18.1%	12.0%	4.4***
1	4.4%	6.5***	14.3%	5.0%	6.8***	14.5%	5.8***	2.5%	15.2%	6.0***	7.0%	6.6***	30.6%	8.5%	6.7***
3	4.9%	7.9***	11.6%	5.3%	8.2***	12.8%	6.7***	2.7%	13.1%	6.8***	7.0%	7.9***	8.2%	7.5%	8.2***
5	2.9%	5.8***	12.3%	3.2%	6.0***	11.6%	6.2***	1.6%	11.9%	6.3***	7.1%	8.7***	12.1%	7.6%	8.7***
8	2.3%	4.7***	5.4%	2.5%	4.9***	10.2%	6.0***	3.1%	10.8%	6.3***	7.7%	9.0***	14.6%	8.3%	9.0***
10	-	-	-	-	-	9.4%	5.7***	2.6%	10.0%	6.0***	7.7%	8.9***	11.3%	8.2%	8.9***
15	-	-	-	-	-	6.0%	4.4***	0.9%	6.2%	4.6***	7.4%	8.9***	11.2%	7.8%	8.9***

This table shows the out-of-sample pricing errors α of the arbitrage strategies relative of the Fama-French 8 factor model and their mean returns μ for the CNN+Transformer model and different number of factors K. We estimate the model on only once on the first T_{train} days and keep it constant on the remaining test set. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 1998 $+T_{\text{train}}$ to December 2016. The signals are extracted from a rolling window of $L = 30$ days and we use the Sharpe ratio objective function. We run a time-series regression of the out-of-sample returns of the arbitrage strategies on the 8-factor model (Fama-French 5 factors + momentum + short-term reversal + long-term reversal) and report the annualized α, accompanying t-statistic value t_α, and the R^2 of the regression. In addition, we report the annualized mean return μ along with its accompanying t-statistic t_μ. The hypothesis test are two-sided and stars indicate p-values of 5% (*), 1% (**), and 0.1% (***).

also correctly detected and taken advantage of by the model.[14] These examples suggest that the CNN+Transformer model can learn mean-reversion and trend patterns. Figures 2 and A.2 are further examples with the same takeaways.

Next, we "dissect" the CNN+Transformer model to understand what type of functions it can estimate. Our analysis begins with the eight basic convolutional patterns learned by the convolutional layers of our network, which are displayed in Figure 8. The CNN represents a given time-series as a matrix of exposures to local basic patterns. As explained in Section II.D.3, these local filters are more complicated than simple local linear filters, but we can project our CNN filters into two-dimensional orthogonal linear filters, which are more interpretable. These local patterns are the building blocks to construct global patterns. We see that these basic patterns display a variety of salient price behavior which are considered to be important. Basic patterns 4 and 6 capture local upward trends, basic patterns 3 and 7 track local downward trends and basic patterns 1, 5 and 8 learn reversion patterns. However, the basis patterns do not include very spiked, sharp changes. Overall, the building blocks seem to be sufficient to construct any smooth trend and mean-reversion pattern.

[14]Note that our in our empirical study the model trades in all residuals and is not limited to trade only in one residual. Hence, the empirical performance is substantially better as shown in Figure 6.

Figure 7: Examples of Allocation and Returns of CNN+Transformer Strategy

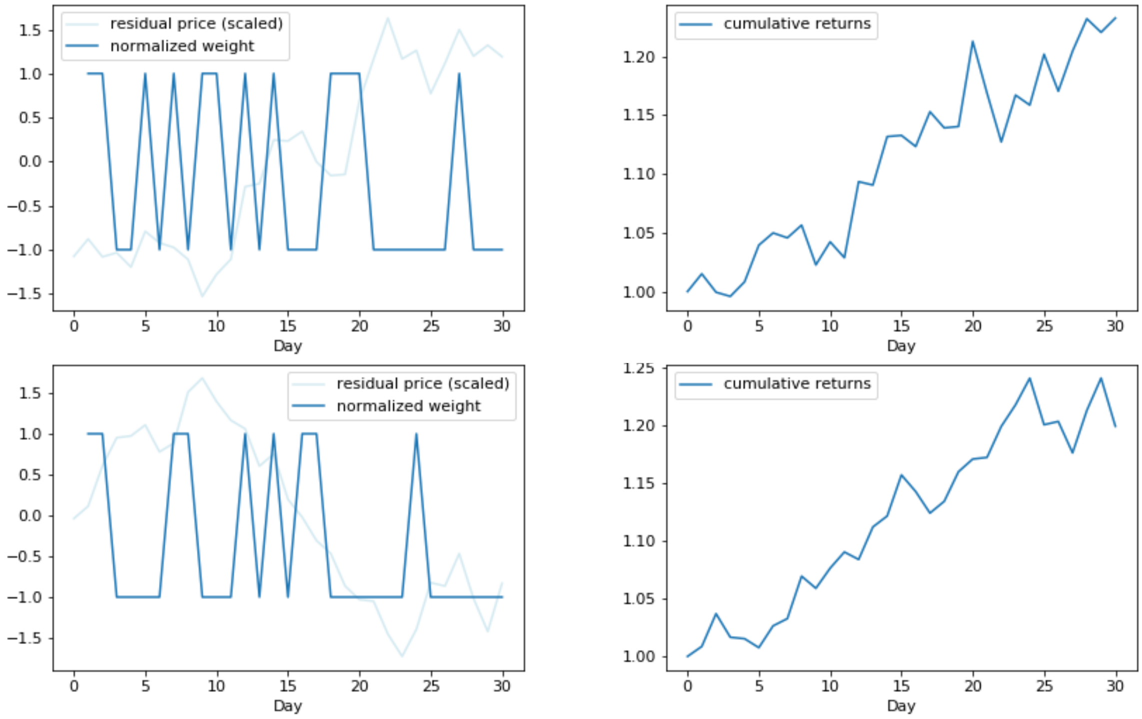

These plots display representative examples of the CNN+Transformer out-of-sample arbitrage trading on a sample of residuals from the IPCA 5-factor model. The left subplots show the normalized cumulative returns of the residuals and the normalized allocation weight, which the specific residual has in the trading strategy. The right subplots illustrates the payoff of trading the specific residual with the displayed allocation function. The model is calibrated on the first 8 years of data and kept constant.

We can understand the global patterns learned by the transformer by studying the attention function. The attention functions $\boldsymbol{\alpha}_i(.,.)$ of each attention head $i = 1, ..., H = 4$ capture the dependencies between the local patterns. Our arbitrage signal can be interpreted as "loadings" to these "attention factors". We use the same H attention functions for all residuals, but in order to visualize them, we evaluate them for a given residual time-series, which yield the attention weights per head:

$$\alpha_{i,l,j} = \boldsymbol{\alpha}_i\left(\tilde{x}_l, \tilde{x}_j\right) \qquad \text{for } i = 1, ..., H.$$

As our signal only depends on the final cumulative return projection h_L^{proj}, the attention weights $\alpha_{i,L,j}$ for $i = 1, .., H$ and $j = 1, ..., L$ contain all the "global factor" information. Hence, we will plot the $H \times L$ dimensional attention weights of the attention heads to understand which global patterns are activated by specific time-series.

Figure 9 plots the attention head weights for simulated sinusoidal residual input time-series. Note that the attention head weights discover the sinusoidal pattern although the model was estimated on the empirical data and not specifically trained for this simulated input. The different attention heads capture different patterns. The fourth attention head displayed in red has the

Figure 8: Local Basic Patterns of Benchmark Model

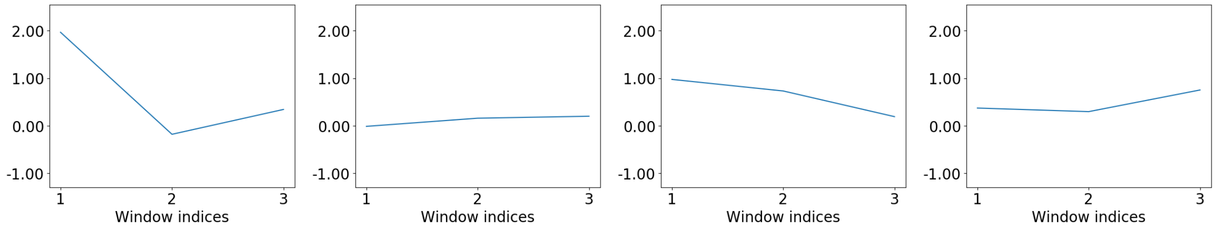

(a) Basic pattern 1 **(b)** Basic pattern 2 **(c)** Basic pattern 3 **(d)** Basic pattern 4

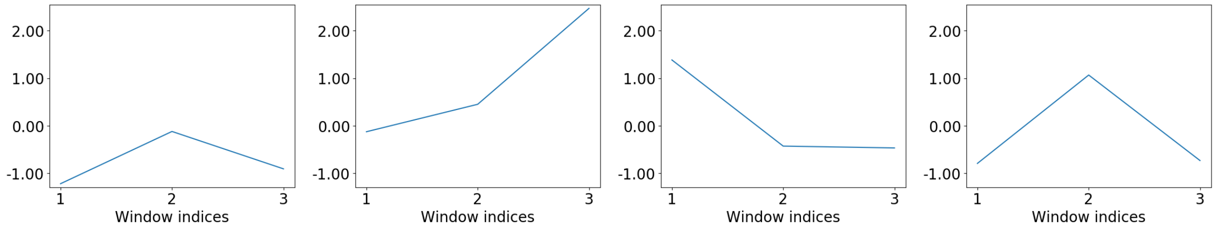

(e) Basic pattern 5 **(f)** Basic pattern 6 **(g)** Basic pattern 7 **(h)** Basic pattern 8

These figures show the $D = 8$ local filters of the CNN estimated for the benchmark model in our empirical analysis. These are projections of our higher dimensional nonlinear filter from a 2-layer CNN into two-dimensional linear filters. These building blocks are labeled "basic patterns". The benchmark model is the CNN+Transformer model based on IPCA 5-factor residuals. We estimate the model on only once on the first T_{train}=8 years based on the Sharpe ratio objective.

strongest activation and capture high-frequency mean reversion patterns. These attention head weights are positive for negative realizations. We will label these fourth attention head weights a "negative reversal" pattern. The third attention head weights depicted by the green curve co-move with the mean-reversion patterns of the original time-series, that is they are positive for high values, but seem to be only activated if this positive "hilltop" appears at the beginning of the time-series. If the mean-reversion cycle achieves its positive values at the end, the third attention head is not activated. We will label these third attention head weights the "early reversal" pattern. The first attention head in blue seems to be a "dampened" version of the fourth red attention head. Figure A.3 in the Appendix shows additional simulated input time-series that confirm this interpretation. In summary, the different attention head weights can be assigned to specific global patterns.

In Figure 10, we plot the different components of the CNN+Transformer model evaluated on a randomly selected, representative 30-day empirical residual. The cumulative residual in subfigure (a) is the input to the CNN. This $L = 30$ dimensional vector is represented by the CNN in terms of its "exposure" to local basic pattern. The subfigures (d)–(k) show this $D \times L$ dimensional representation, which is the output of the CNN. As we have $D = 8$ local filters, we obtain eight time-series that display the activation to each filter. For example, basic pattern 1 is associated with a "reversal kink" in subfigure 8(a) and hence has the strongest activation to this basic pattern on day 28, when the residual has a downward spike. This $D \times L$ matrix of exposures to local patterns is the input to the transformer. The attention head weights in subfigure (b) connect the local patterns to a global pattern. The fourth attention head weight in red has its highest values during the "bottom" of the residual movements, confirming our previous intuition that this

Figure 9: Example Attention Weights for Sinusoidal Residual Inputs

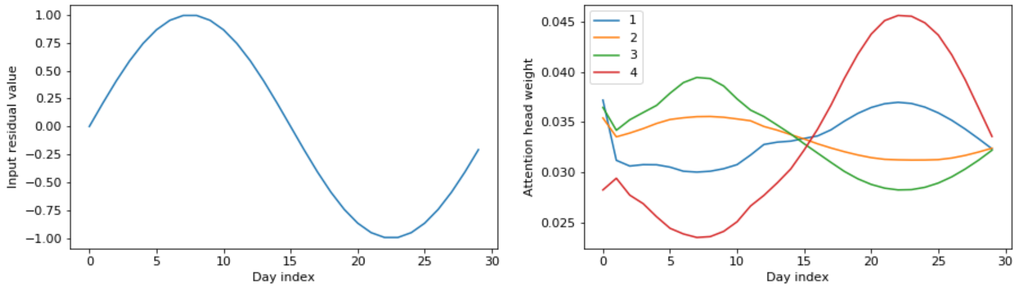

(a) Input residual and attention head weights for $x_l = \sin\left(2\pi \frac{l}{30}\right)$

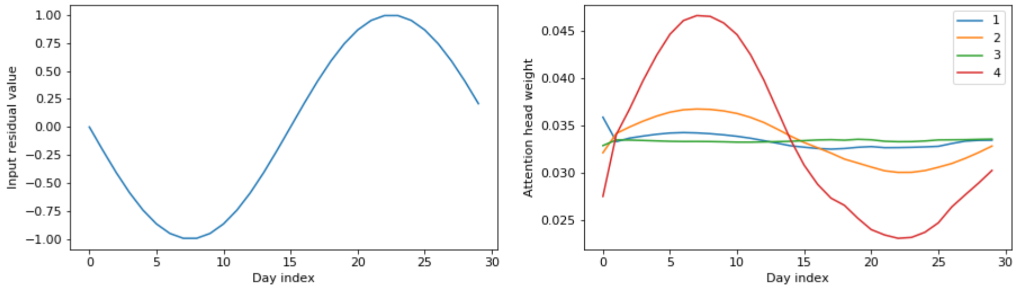

(b) Input residual and attention head weights for for $x_l = \sin\left(2\pi \frac{l+15}{30}\right)$

These plots show the attention head weights of the CNN+Transformer benchmark model for simulated sinusoidal residual input time series. Both sine functions have a cycle of 30 days and the second is shifted by 15 days. The right subplot shows the attention weights for the $H = 4$ attention heads for the specific residuals. The empirical benchmark model is the CNN+Transformer model based on IPCA 5-factor residuals. We estimate the model on only once on the first T_{train}=8 years based on the Sharpe ratio objective.

attention head activates during bad times. The third attention head weight in green spikes during the "top" at the beginning of the residual time-series, which is in line with our interpretation as an early reversal pattern. The first attention head in blue is a dampened version of the red attention head. The average over the four attention head weights depicted in subfigure (c) suggests that the heads attend on average more closely to the latter half of the time series.

In Figure 11 we generalize the analysis of Figure 10 to study the model structure of the benchmark model over time. While Figure 10 represents a "snapshot" for one point in time, we now display the allocation weights and attention head weights of a single representative residual for different times. Subfigure 11(a) shows the cumulative residual time-series. For a specific point in time we use the lagged $L = 30$ days as an input to obtain the allocation weights and attention head weights for that time. The out-of-sample allocation weights correctly change the directions to exploit the patterns in the residual time-series. The attention head weights over time offer additional insights into the structure of the global patterns. Each vertical slice from window index 1 to window index 30 displays the normalized attention weights for the time point under the slice. The third attention head, which was displayed in green in Figures 9 and 10 has the largest values during "up-patterns" of the residual, for example for 2007, 2010 and 2012. Importantly, the attention weights focus on the early days within the 30-day window. This confirms our previous interpretation as an "early reversal" pattern. Attention head four, which was previously represented as a red line, has

Figure 10: CNN+Transformer Model Structure for Representative Residual

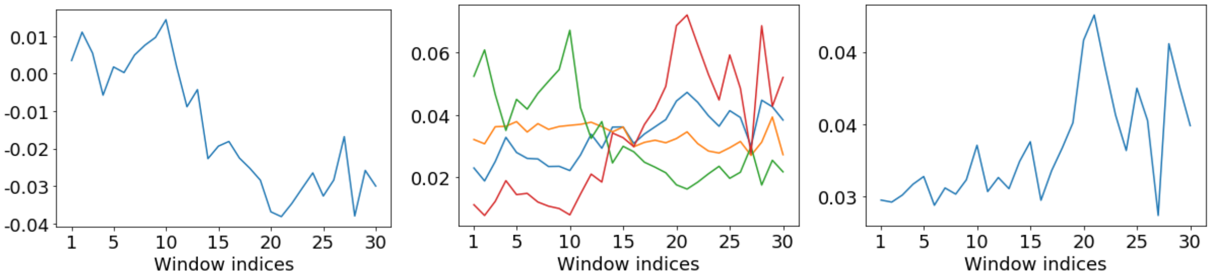

(a) Cumulative residual **(b)** Attention weights per head **(c)** Average attention weights

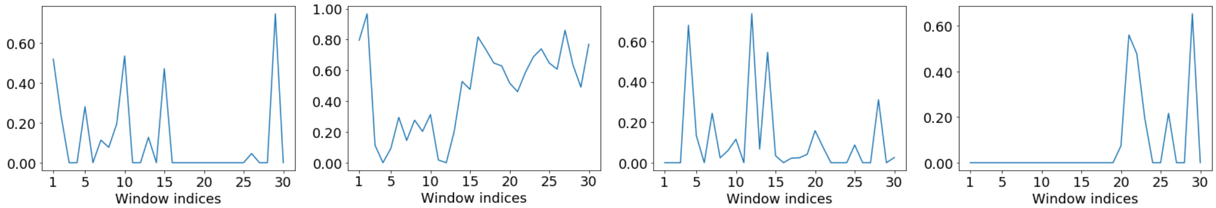

(d) CNN activations for basic pattern 1 **(e)** CNN activations for basic pattern 2 **(f)** CNN activations for basic pattern 3 **(g)** CNN activations for basic pattern 4

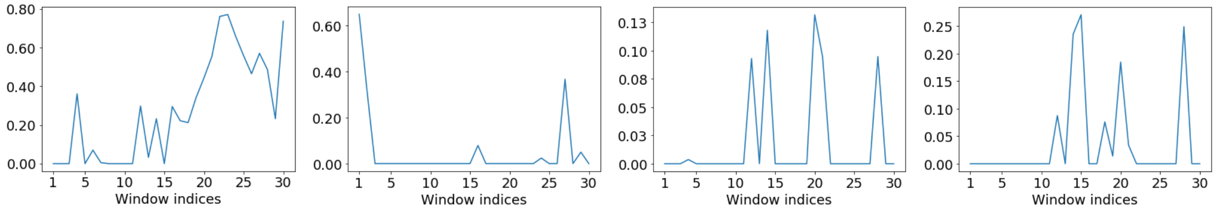

(h) CNN activations for basic pattern 5 **(i)** CNN activations for basic pattern 6 **(j)** CNN activations for basic pattern 7 **(k)** CNN activations for basic pattern 8

These figures illustrate the different components of the CNN+Transformer benchmark model evaluated for a randomly selected, representative empirical residual. The cumulative residual returns, which are the input to the model, are plotted in (a). The convolutional activations (d)–(k) quantify the exposure of the residual time-series to local basis filters. Subplot (b) displays the attention weights for the $H = 4$ attention heads, which represent global dependency patterns. Subplot (c) shows the average of these four attention head weights. The empirical benchmark model is the CNN+Transformer model based on IPCA 5-factor residuals. We estimate the model on only once on the first T_{train}=8 years based on the Sharpe ratio objective.

the highest values during down-times, such as 2009, 2014 and the middle of 2016. In contrast to attention head 3, this head focuses on the immediate past within the local window. Attention head 1, which is a dampened down-version, focuses more uniformly on all the values within the local window.

The average attention weights in (b) illustrate the asymmetric response of the transformer network. During uptrends, it focuses on the residual prices which are further in the past part of the window to decide which to position to take; however, during downtrends, it focuses on the most recent cumulative residual prices in the lookback window, which indicates that it is taking into account the latest data in order to decide what position to take. This indicates that our CNN+Transformer policy network has learned to act swiftly during downtrends, and more slowly during uptrends. This shows that our model learns in particular the commonly repeated wisdom that "markets take

Figure 11: CNN+Transformer Model Structure for Representative Residual Over Time

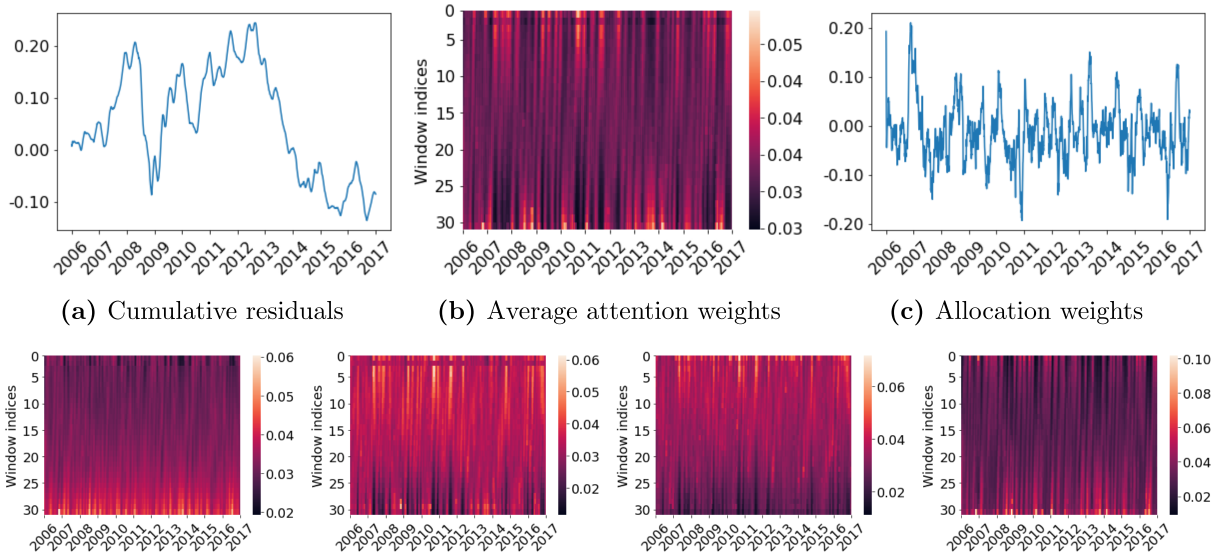

(a) Cumulative residuals **(b)** Average attention weights **(c)** Allocation weights

(d) Attention weights for head 1 **(e)** Attention weights for head 2 **(f)** Attention weights for head 3 **(g)** Attention weights for head 4

These figures illustrate the out-of-sample behavior from 2006 to 2016 of the CNN+Transformer benchmark model for a single residual time-series. The cumulative residual returns are plotted in (a), and the suggested allocation weights before cross-sectional normalization are plotted in (c). The attention head weights (d)–(g) quantify the activation for each attention head over time. Subplot (b) shows the average of these weights over the four heads for different times. All time-series have been smoothed using a simple moving average with a 30-day window for better presentation. The empirical benchmark model is the CNN+Transformer model based on IPCA 5-factor residuals. We estimate the model on only once on the first T_{train}=8 years based on the Sharpe ratio objective.

Figure 12: Variable Importance for Allocation Weight

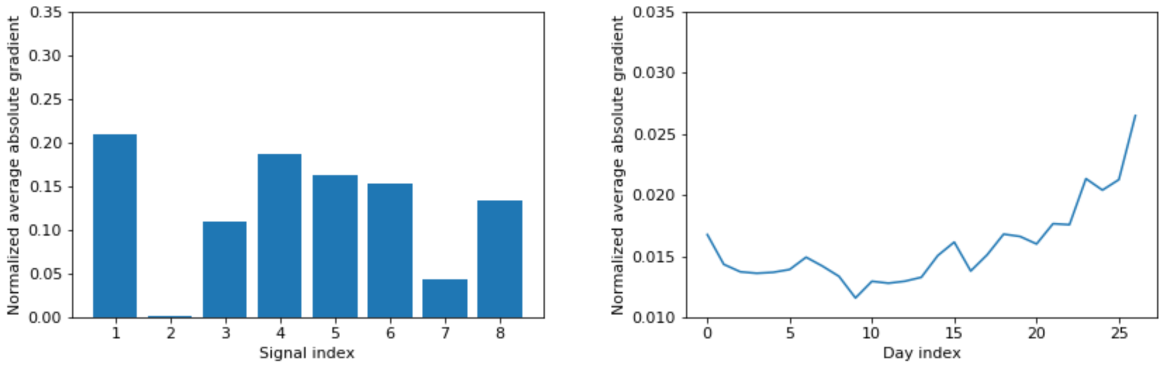

(a) Importance of Local Basic Patterns **(b)** Importance of Residual Days

These figures show the normalized average absolute gradient (NAAG) of the allocation weight with respect to various inputs to intermediate layers in the CNN+Transformer benchmark network. A higher NAAG indicates a higher importance. Subplot (a) quantifies the importance of the $D = 8$ different convolutional filters, that is, we display the gradient with respect to the output of the convolutional network, which is the input to the self-attention layer. In (b), we report the importance of the first 27 days of the input residual time series. Each average absolute gradient is normalized by dividing each element by the sum of all elements. The empirical benchmark model is the CNN+Transformer model based on IPCA 5-factor residuals. We estimate the model on only once on the first T_{train}=8 years based on the Sharpe ratio objective.

escalators up and elevators down". This asymmetric policy is a key benefit of the attention-based

40

model, which cannot easily be replicated by the parametric Ornstein-Uhlenbeck or fixed basis pattern benchmark models we compare against. The convolutional subnetwork's patterns provide translation invariant information about what kind of trend is present within each 3-day subwindow of the 30-day cumulative residual price lookback window, which allows the transformer subnetwork to form a stable attention function that results in this unique policy.

Figure 12 sheds further light on which days and patterns are important. The figure shows the normalized average absolute gradient (NAAG) of the allocation weight with respect to various inputs to intermediate layers in the CNN+Transformer benchmark network. A higher NAAG indicates a higher importance. Subplot (a) quantifies the importance of the $D = 8$ different basic patterns. We observe that the flat basic pattern 2 has a negligible weight, while basis patterns that are needed for trend or reversal patterns have high importance. In (b), we report the importance of the first 27 days of the input residual time series.[15] Crucially, all previous days matter, which emphasizes that the trading allocation depends on the past dynamics. The most recent 14 days seem to get on average more attention for the trading decisions. However, as indicated in Figure 11, the importance of the days seems to be asymmetric for different global patterns.

K. Market Frictions and Transaction Costs

Our deep learning arbitrage strategies remain profitable in the presence of realistic trading frictions. In practice, trading costs associated with high turnover or large short-selling positions can diminish the profitability of arbitrage trading. In order to ensure that our model produces economically meaningful results, we extend it to the setting in which both transaction costs and holding costs are accounted for. We do not model market frictions associated with market impact, as in our empirical analysis we restrict the asset universe to stocks with large market capitalization, which are especially liquid.

Table XI: OOS Performance of CNN+Trans with Trading Frictions

IPCA factor model						
	Sharpe ratio			Mean-variance		
K	SR	μ	σ	SR	μ	σ
0	0.52	8.5%	16.3%	0.22	2.6%	11.9%
1	0.85	5.9%	6.9%	0.86	5.5%	6.4%
3	1.24	6.6%	5.4%	1.16	6.9%	5.9%
5	1.11	5.5%	5.0%	1.02	5.3%	5.3%
10	0.98	5.1%	5.2%	1.04	5.4%	5.2%
15	0.94	4.8%	5.1%	1.02	5.1%	5.0%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) for the CNN+Transformer model with trading frictions on IPCA residuals. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. The models are calibrated on a rolling window of four years and use either the Sharpe ratio or mean-variance objective function with trading costs ($\text{cost}(w^R_{t-1}, w^R_{t-2}) = 0.0005\|w^R_{t-1} - w^R_{t-2}\|_{L^1} + 0.0001\|\min(w^R_{t-1}, 0)\|_{L^1}$). The signals are extracted from a rolling window of $L = 30$ days.

[15] As the attention head weights are determined relative to the last local 3-day window, that subwindow has mechanically a larger weight and is not comparable to the other 27 days.

Figure 13: Turnover of CNN+Transformer Model with and without Trading Friction Objective

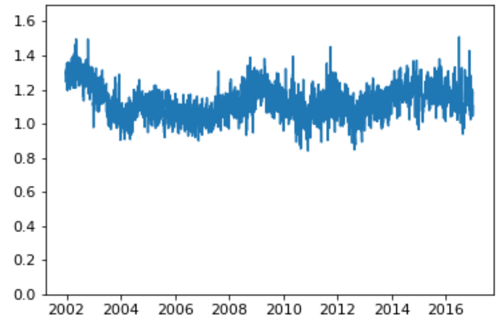

(a) No Trading Friction Penalty

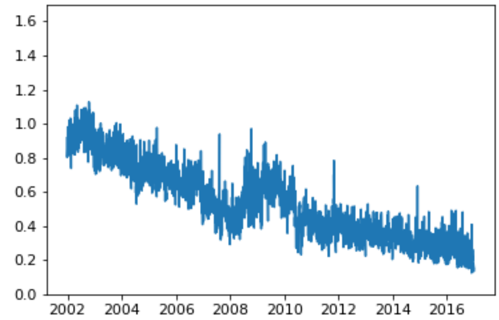

(b) With Trading Friction Penalty

These figures show the daily turnover of CNN+Transformer model with and without trading friction objective on the representative IPCA 5-factor residuals for the out-of-sample trading period between January 2002 and December 2016. The models are calibrated on a rolling window of four years and use the Sharpe ratio objective function with or without trading costs ($\text{cost}(w^R_{t-1}, w^R_{t-2}) = 0.0005\|w^R_{t-1} - w^R_{t-2}\|_{L^1} + 0.0001\|\min(w^R_{t-1}, 0)\|_{L^1}$). We define turnover as the ℓ_1 norm of the difference between allocation weight vectors at consecutive times, i.e. $||w^R_{t-1} - w^R_{t-2}||_{L^1}$.

Figure 14: Proportion of Short Allocation Weights of CNN+Transformer Model with and without Trading Friction Objective

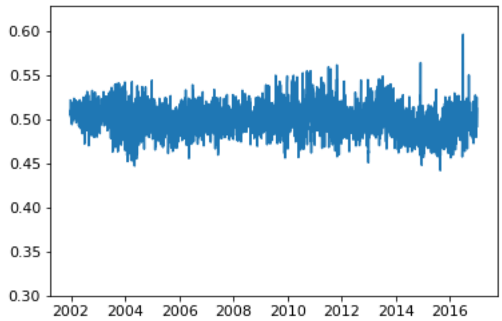

(a) No Trading Friction Penalty

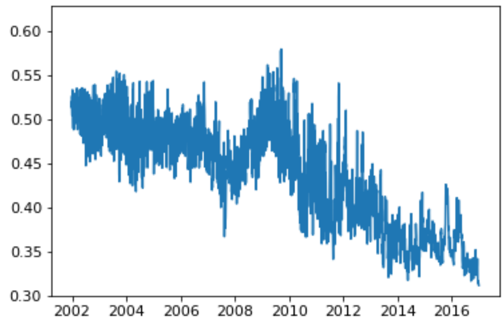

(b) With Trading Friction Penalty

These figures show the daily fraction of short trades of the CNN+Transformer strategies with and without trading friction objective on the representative IPCA 5-factor residuals for the out-of-sample trading period between January 2002 and December 2016. Each plot shows the absolute value of the sum of negative weights $\|\min(w^R_{t-1}, 0)\|_{L^1}$ relative to the sum of absolute values of all weights, which is normalized to $\|w^R_{t-1}\|_1 = 1$. The models are calibrated on a rolling window of four years and use the Sharpe ratio objective function with or without trading costs ($\text{cost}(w^R_{t-1}, w^R_{t-2}) = 0.0005\|w^R_{t-1} - w^R_{t-2}\|_{L^1} + 0.0001\|\min(w^R_{t-1}, 0)\|_{L^1}$).

In our market-friction extension, the daily returns R_t of the strategy now have constant linear penalties associated with the daily turnover and the proportion of short trades. These penalties quantify proportional transaction costs, which are used to model trading fees, size of the bid-ask spread, etc., and holding costs, which are used to model short borrow rate fees charged by a brokerage. In particular, we incorporate a subset of the market friction models proposed by Boyd et al. (2017), which are commonly used in the statistical arbitrage literature.[16] Mathematically, we

[16] See for example Avellaneda and Lee (2010), Yeo and Papanicolaou (2017) and Krauss et al. (2017).

subtract the market-friction costs

$$\text{cost}(w_{t-1}^R, w_{t-2}^R) = 0.0005\|w_{t-1}^R - w_{t-2}^R\|_{L^1} + 0.0001\|\min(w_{t-1}^R, 0)\|_{L^1}$$

from the portfolio returns and use these net portfolio returns in the optimization problem of section II.C, where $w_{t-1}^R \in \mathbb{R}^{N_{t-1}}$ is the strategy's allocation weight vector at time $t-1$. The first penalty term represents a slippage/transaction cost of 5 basis points per transaction, whereas the second one is a holding cost of 1 basis point per short position. Both costs are universal for all times and all stocks. This corresponds to a modification of the objective function in the training and evaluation parts of our algorithm. We use this model for the sake of illustration and simplicity given that in our empirical study we trade a universe of highly liquid US stocks, but more complicated models[17] may be included in the computations without any significant structural changes.

Table XI displays the Sharpe ratios, average returns, and volatility of our CNN+Transformer model under market frictions for IPCA residuals. The results for PCA residuals are collected in the Appendix in Table XI with very similar findings. We exclude the Fama-French factor model from the analysis with market frictions, as we take the traded factors from Kenneth French Data Library as given, which are based on a larger stock universe with different trading costs and, hence, would not be directly comparable to the IPCA and PCA results.[18] As expected the Sharpe ratios are lower and range from 0.94 to 1.24 for a reasonable number of IPCA factors. The Sharpe ratio and mean-variance objective have the desired effects, but lead to overall very similar results. Importantly, the arbitrage strategies retain their economic significance even in the presence of trading costs.

These results present a lower bound on the profitability under trading frictions, as we have made four simplifying assumptions. First, in the current implementation the factor composition cannot be changed due to trading costs. A possible extension could construct the latent risk factors by including the trading friction objective. For example, the sparse representation of latent factors as in Pelger and Xiong (2021) would reduce trading costs. Second, because the policy with frictions is recursive, we are conducting an approximate training process to maintain parallelization given our computational resources and the large volume of data, but this may lead to suboptimal optimization results. However, it would be possible to conduct an exact sequential training process at the cost of more computation. Third, our modified architecture with the market-friction objective is given by the simplest modification to our architecture without frictions, but it is possible that the optimal transaction and holding cost-minimizing strategy has a more complicated functional form or is not Markovian and requires additional previous allocations. Last but not least, we keep the hyperparameters of our main analysis, but we could potentially improve the performance by

[17]For example, those considering time and stock-dependent transaction costs or market impact of the trades on the stock prices.

[18]Regardless, each factor corresponds to a portfolio of traded assets, and thus the residuals of this model could be traded in a number of a number of ways under suitable extensions. For example, we could include ETFs which try to track a value or size premium, project these latent factors onto our asset universe, or approximate each factor with a number of sparse subset of assets in our asset universe as in Pelger and Xiong (2021). However, these changes constitute differences that would make the results incomparable to the PCA and IPCA results.

employing hyperparameter tuning.

The effect of trading frictions is time-varying and our model can exploit particularly profitable arbitrage time periods by increasing trading and short positions. In Figure 13 we analyze the daily turnover of a representative CNN+Transformer strategy based on IPCA 5-factor residuals and a Sharpe ratio objective. Broadly, we see that our model with trading friction penalty is able to adapt by decreasing daily turnover. However, our model seems to reduce turnover based on trading opportunities. During the times of high market volatility such as 2007–2009, arbitrage trading could be potentially be more profitable, which our model takes advantage of. On the other hand, during the later years of the calm bull market from 2011–2015, strategies with less turnover could maintain profitability. This pattern is confirmed in Figure 14 which shows the daily proportion of allocation weights, which are short stocks in our universe. As expected the holding cost friction model reduces the overall proportion of short trades. Interestingly, our model is able to intelligently choose time periods during which it can maximize performance by taking positions with higher short proportion, such as the market turmoil at the end of 2015 and the financial crisis of 2008. Effectively, this indicates that the CNN+Transformer trading policy has learned to avoid holding and transaction costs by generally modifying the original strategy's allocations to be less short-biased on average, and to more appropriately enter short-dominant positions during relevant subperiods.

IV. Conclusion

In this paper, we introduce a unifying conceptual framework to compare different statistical arbitrage approaches based on the decomposition into (1) arbitrage portfolio generation, (2) signal extraction and (3) allocation decision. We develop a novel deep learning statistical arbitrage approach. It uses conditional latent factors to generate arbitrage portfolios. The signal is estimated with a CNN+Transformer, which combines global dependency patterns with local filters. The allocation is estimated with a nonparametric FFN based on a global trading objective.

We conduct a comprehensive empirical out-of-sample study on U.S. equities and demonstrate the potential of machine learning methods in arbitrage trading. Our CNN+Transformer substantially outperforms all benchmark approaches. The implied trading strategies are not spanned by conventional risk factors, including price trend factors, and survive realistic transaction and holding costs. Our model provides insights into optimal trading policies which are based on asymmetric trend and reversion patterns. In particular, our study shows that the trading signal extraction is the most challenging and separating element among different statistical arbitrage approaches.

Our findings contribute to the debate on efficiency of markets. We quantify the scope of profits that arbitrageurs can achieve in equity markets. Importantly, the substantial profitability of our arbitrage strategies is not inconsistent with equilibrium asset pricing, following similar arguments as in Gatev et al. (2006). It could rather be viewed as empirical evidence about how efficiency is maintained in practice. We document non-declining profitability of arbitrage trading over time, which suggests that the profits are compensation for arbitrageurs to enforce the law of one price. Our findings also suggest that unconditional means of asset pricing residuals as a measure of alpha might not correctly reflect the amount of arbitrage left in financial markets.

REFERENCES

Avellaneda, M. and J.-H. Lee (2010): "Statistical arbitrage in the US equities market," *Quantitative Finance*, 10, 761–782.

Barr, A. (2007): "Quant quake shakes hedge-fund giants," *Marketwatch.*

Bianchi, D., M. Büchner, and A. Tamoni (2019): "Bond risk premia with machine learning," *Working paper.*

Boyd, S., E. Busseti, S. Diamond, R. N. Kahn, K. Koh, P. Nystrup, J. Speth, et al. (2017): "Multi-period trading via convex optimization," *Foundations and Trends in Optimization*, 3, 1–76.

Bryzgalova, S., M. Pelger, and J. Zhu (2019): "Forest through the Trees: Building Cross-Sections of Stock Returns," *Working paper.*

Carhart, M. M. (1997): "On persistence in mutual fund performance." *Journal of Finance*, 52, 57–82.

Cartea, A. and S. Jaimungal (2016): "Algorithmic trading of co-integrated assets," *International Journal of Theoretical and Applied Finance*, 19, 165038.

Chen, L., M. Pelger, and J. Zhu (2019): "Deep learning in asset pricing," *Working paper.*

Chen, Y., W. Chen, and S. Huang (2018): "Developing Arbitrage Strategy in High-frequency Pairs Trading with Filterbank CNN Algorithm," in *2018 IEEE International Conference on Agents (ICA)*, 113–116.

Cong, L., K. Tang, J. Wang, and Y. Zhang (2020): "Single-Step Portfolio Construction through Reinforcement Learning and Economically Interpretable AI," *Working paper.*

d'Aspremont, A. (2011): "Identifying small mean-reverting portfolios," *Quantitative Finance*, 11, 351–364.

Dunis, C., J. Laws, and B. Evans (2006): "Modelling and trading the soybean-oil crush spread with recurrent and higher order networks: a comparative analysis," *Neural Network World*, 16, 193–213.

Ebrahimi, H. (2013): "Hedge fund quants lose money in 2012," *The Telegraph.*

Elliott, R., J. Van Der Hoek, and W. Malcolm (2005): "Pairs trading," *Quantitative Finance*, 5, 513–545.

Fama, E. F. and K. R. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of.*

——— (2015): "A five-factor asset pricing model," *Journal of financial economics*, 116, 1–22.

Fischer, T. G., C. Krauss, and A. Deinert (2019): "Statistical arbitrage in cryptocurrency markets," *Journal of Risk and Financial Management*, 12, 31.

Freyberger, J., A. Neuhierl, and M. Weber (2020): "Dissecting characteristics nonparametrically," *Review of Financial Studies*, 33, 2326–2377.

Gatev, E., W. N. Goetzmann, and K. Rouwenhorst (2006): "Pairs trading: performance of a relative-value arbitrage rule," *Review of Financial Studies*, 19, 797–827.

Gu, S., B. Kelly, and D. Xiu (2021): “Autoencoder Asset Pricing Models,” *Journal of Econometrics*, 222, 429–450.

Gu, S., B. T. Kelly, and D. Xiu (2020): “Empirical Asset Pricing Via Machine Learning,” *Review of Financial Studies*, 33, 2223–2273.

Huck, N. (2009): “Pairs selection and outranking: An application to the S&P 100 index,” *European Journal of Operational Research*, 196, 819–825.

Jurek, J. W. and H. Yang (2007): “Dynamic portfolio selection in arbitrage,” in *EFA 2006 Meetings Paper*.

Kelly, B., S. Pruitt, and Y. Su (2019): “Characteristics Are Covariances: A Unified Model of Risk and Return,” *Journal of Financial Economics*, 134, 501–524.

Kim, T. and H. Y. Kim (2019): “Optimizing the Pairs-Trading Strategy Using Deep Reinforcement Learning with Trading and Stop-Loss Boundaries,” *Complexity*, 2019.

Kozak, S., S. Nagel, and S. Santosh (2020): “Shrinking the Cross Section,” *Journal of Financial Economics*, 135, 271–292.

Krauss, C., X. A. Doa, and N. Huck (2017): “Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500,” *European Journal of Operational Research*, 259, 689–702.

Lettau, M. and M. Pelger (2020a): “Estimating Latent Asset Pricing Factors,” *Journal of Econometrics*, 218, 1–31.

——— (2020b): “Factors that Fit the Time-Series and Cross-Section of Stock Returns,” *Review of Financial Studies*, 33, 2274–2325.

Leung, T. and X. Li (2015): “Optimal mean reversion trading with transaction costs and stop-loss exit,” *International Journal of Theoretical and Applied Finance*, 18, 1550020.

Lim, B. and S. Zohren (2020): “Time Series Forecasting With Deep Learning: A Survey,” *Working paper.*

Lintilhac, P. S. and A. Tourin (2016): “Model-based pairs trading in the bitcoin markets,” *Quantitative Finance*, 17, 703–716.

Mudchanatongsuk, S., J. A. Primbs, and W. Wong (2008): “Optimal pairs trading: A stochastic control approach,” in *2008 American Control Conference*, IEEE, 1035–1039.

Mulvey, J. M., Y. Sun, M. Wang, and J. Ye (2020): “Optimizing a portfolio of mean-reverting assets with transaction costs via a feedforward neural network,” *Quantitative Finance*, forthcoming.

Pelger, M. (2020): “Understanding Systematic Risk: A High-Frequency Approach,” *Journal of Finance*, 74, 2179–2220.

Pelger, M. and R. Xiong (2021): “Interpretable sparse proximate factors for large dimensions,” *Journal of Business & Economic Statistics, conditionally accepted.*

Rad, H., R. K. Y. Low, and R. Faff (2016): “The profitability of pairs trading strategies: distance, cointegration and copula methods,” *Quantitative Finance*, 16, 1541–1558.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin (2017): “Attention Is All You Need,” *NIPS Conference on Neural Information Processing Systems.*

Vidyamurthy, G. (2004): *Pairs Trading: Quantitative Methods and Analysis*, Wiley.

Xiaohong Chen and H. White (1999): “Improved rates and asymptotic normality for nonparametric neural network estimators,” *IEEE Transactions on Information Theory*, 45, 682–691.

Yeo, J. and G. Papanicolaou (2017): “Risk control of mean-reversion time in statistical arbitrage,” *Risk and Decision Analysis*, 6, 263–290.

Appendix A. Data

Appendix A. List of the Firm-Specific Characteristics

Table A.I: Firm Characteristics by Category

	Past Returns	
(1)	r2_1	Short-term momentum
(2)	r12_2	Momentum
(3)	r12_7	Intermediate momentum
(4)	r36_13	Long-term momentum
(5)	ST_Rev	Short-term reversal
(6)	LT_Rev	Long-term reversal
	Investment	
(7)	Investment	Investment
(8)	NOA	Net operating assets
(9)	DPI2A	Change in property, plants, and equipment
(10)	NI	Net Share Issues
	Profitability	
(11)	PROF	Profitability
(12)	ATO	Net sales over lagged net operating assets
(13)	CTO	Capital turnover
(14)	FC2Y	Fixed costs to sales
(15)	OP	Operating profitability
(16)	PM	Profit margin
(17)	RNA	Return on net operating assets
(18)	ROA	Return on assets
(19)	ROE	Return on equity
(20)	SGA2S	Selling, general and administrative expenses to sales
(21)	D2A	Capital intensity
	Intangibles	
(22)	AC	Accrual
(23)	OA	Operating accruals
(24)	OL	Operating leverage
(25)	PCM	Price to cost margin
	Value	
(26)	A2ME	Assets to market cap
(27)	BEME	Book to Market Ratio
(28)	C	Ratio of cash and short-term investments to total assets
(29)	CF	Free Cash Flow to Book Value
(30)	CF2P	Cashflow to price
(31)	D2P	Dividend Yield
(32)	E2P	Earnings to price
(33)	Q	Tobin's Q
(34)	S2P	Sales to price
(35)	Lev	Leverage
	Trading Frictions	
(36)	AT	Total Assets
(37)	Beta	CAPM Beta
(38)	IdioVol	Idiosyncratic volatility
(39)	LME	Size
(40)	LTurnover	Turnover
(41)	MktBeta	Market Beta
(42)	Rel2High	Closeness to past year high
(43)	Resid_Var	Residual Variance
(44)	Spread	Bid-ask spread
(45)	SUV	Standard unexplained volume
(46)	Variance	Variance

This table shows the 46 firm-specific characteristics sorted into six categories. More details on the construction are in the Internet Appendix of Chen et al. (2019).

Appendix B. Implementation of Different Models

Appendix A. Feedforward Neural Network (FFN)

In the Fourier+FFN and FFN models, we utilize a feedforward network with L^{FFN} layers as illustrated in Figure A.1. Each hidden layer takes the output from the previous layer and transforms it into an output as

$$x^{(l)} = \text{ReLU}\left(W^{\text{FFN},(l-1)\top}x^{(l-1)} + w_0^{\text{FFN},(l-1)}\right) = \text{ReLU}\left(w_0^{\text{FFN},(l-1)} + \sum_{k=1}^{K^{(l-1)}} w_k^{\text{FFN},(l-1)} x_k^{(l-1)}\right)$$

$$y = W^{\text{FFN},(L^{\text{FFN}})\top}x^{(L)} + w_0^{\text{FFN},(L^{\text{FFN}})}$$

with hidden layer outputs $x^{(l)} = (x_1^{(l)}, ..., x_{K^{(l)}}^{(l)}) \in \mathbb{R}^{K^{(l)}}$, parameters $W^{(\text{FFN},l)} = (w_1^{\text{FFN},(l)}, ..., w_{K^{(l)}}^{\text{FFN},(l)}) \in \mathbb{R}^{K^{(l)} \times K^{(l-1)}}$ for $l = 0, ..., L^{\text{FFN}} - 1$ and $W^{\text{FFN},(L^{\text{FFN}})} \in \mathbb{R}^{K^{(L)}}$, and where $\text{ReLU}(x_k) = \max(x_k, 0)$.

Figure A.1. Feedforward Network Architecture

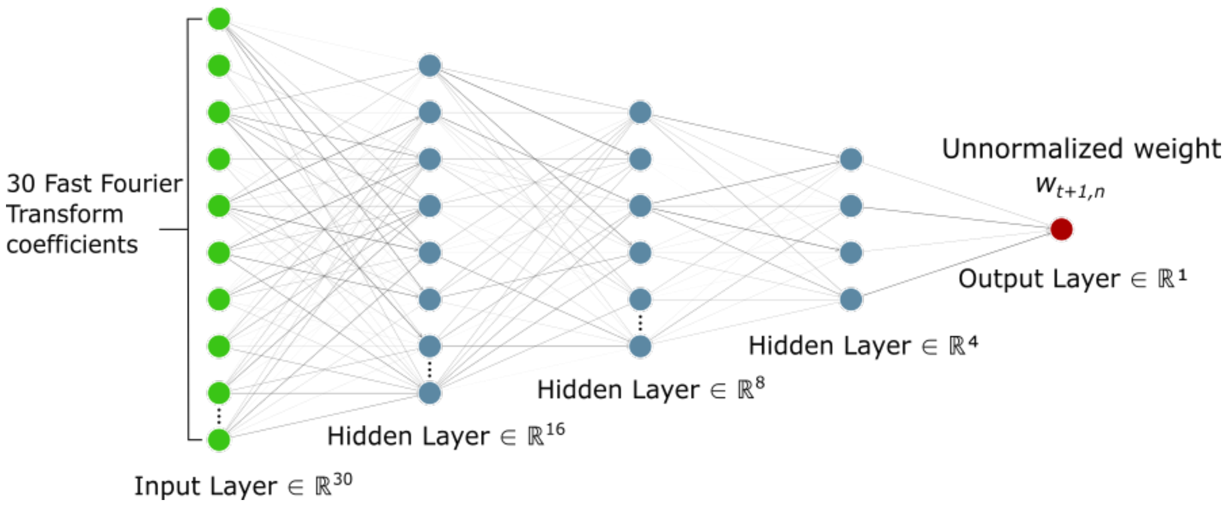

Appendix B. Ornstein-Uhlenbeck Model

Following Avellaneda and Lee (2010) and Yeo and Papanicolaou (2017) we model X_t as an Ornstein-Uhlenbeck (OU) process

$$dX_t = \kappa\,(\mu - X_t)\,dt + \sigma dB_t$$

for a Brownian motion B_t. As the analytical solution of the above stochastic differential equation is

$$X_{t+\Delta t} = (1 - e^{-\kappa\Delta t})\mu + e^{-\kappa\Delta t}X_t + \sigma \int_t^{t+\Delta t} e^{-\kappa(t+\Delta t-s)}dB_s$$

for any Δt, we can without loss of generality set $\Delta t = 1$, and estimate the parameters κ, μ and σ from the AR(1) model

$$X_{t+1} = a + bX_t + e_t,$$

where each e_t is a normal, independent and identically distributed random variable with mean 0. The parameters are estimated with a standard linear regression, which yields

$$\hat{\kappa} = -\frac{\log(\hat{b})}{\Delta t}, \quad \hat{\mu} = \frac{\hat{a}}{1-\hat{b}}, \quad \frac{\hat{\sigma}}{\sqrt{2\hat{\kappa}}} = \sqrt{\frac{\hat{\sigma}_e^2}{1-\hat{b}^2}}.$$

The strategy depends on the ratio $\frac{X_L - \hat{\mu}}{\hat{\sigma}\sqrt{2\hat{\kappa}}}$. Note that this is only defined for $b < 1$ which is equivalent to parameter restrictions that the OU process is mean-reverting. The trading policy depends on the thresholds c_{thresh} and c_{crit}, which are hyperparameters. These hyperparameters are selected on the validation data from the candidate values $c_{\text{thresh}} \in \{1, 1.25, 1.5\}$ and $c_{\text{crit}} \in \{0.25, 0.5, 0.75\}$. Our benchmark model has the values $c_{\text{thresh}} = 1.25$ and $c_{\text{crit}} = 0.25$, which coincides with the optimal values in Avellaneda and Lee (2010) and Yeo and Papanicolaou (2017).

Appendix C. Convolutional Neural Network with Transformer

Appendix C.1. Convolutional Neural Network

In our empirical application, we consider a 2-layered convolutional network with some standard technical additions. The network takes as input a window $x^{(0)} = x \in \mathbb{R}^L$ of L consecutive daily cumulative returns or log prices of a residual, and outputs the feature matrix $\tilde{x} \in \mathbb{R}^{L\times F}$ given by computing the following quantities for $l = 1, \dots, L, d = 1 \dots, D$

$$y_{l,d}^{(0)} = b_d^{(0)} + \sum_{m=1}^{D_{\text{size}}} W_{d,m}^{(0)} x_{l-m+1}^{(0)}, \quad x_{l,d}^{(1)} = \text{ReLU}\left(\frac{y_{1,d}^{(0)} - \mu_d^{(0)}}{\sigma_d^{(0)}}\right). \tag{A.1}$$

$$y_{l,d}^{(1)} = b_d^{(1)} + \sum_{m=1}^{D_{\text{size}}} \sum_{j=1}^{D} W_{d,j,m}^{(1)} x_{l-m+1,j}^{(1)}, \quad x_{l,d}^{(2)} = \text{ReLU}\left(\frac{y_{l,d}^{(1)} - \mu_d^{(1)}}{\sigma_d^{(1)}}\right), \tag{A.2}$$

$$\tilde{x}_{l,d} = x_{l,d}^{(2)} + x_l^{(0)}, \tag{A.3}$$

where

$$\mu_k^{(i)} = \frac{1}{L}\sum_{l=1}^{L} y_{l,k}^{(i)}, \quad \sigma_k^{(i)} = \sqrt{\frac{1}{L}\sum_{l=1}^{L}\left(y_{l,k}^{(i)} - \mu_k^{(i)}\right)^2}.$$

and $b^{(0)}, b^{(1)} \in \mathbb{R}^D$, $W^{(0)} \in \mathbb{R}^{D\times D_{\text{size}}}$ and $W^{(1)} \in \mathbb{R}^{D\times D\times D_{\text{size}}}$ are parameters to be estimated. Compared with the simple convolutional network introduced in the main text, the previous equations incorporate three standard technical improvements commonly used in deep learning practice. First, they include "bias terms" $b^{(i)}$ in the first part of equations A.1 and A.2 to allow for more flexible modeling. Second, they include so-called "instance normalization" before each activation function to speed up the optimization and avoid vanishing gradients caused by the saturation of the ReLU activations. Third, they include a "residual connection" in equation A.3 to facilitate gradient propagation during training.

Appendix C.2. Transformer Network

The benchmark model in our empirical application is a one-layer transformer, following the implementation of the seminal paper of Vaswani et al. (2017). First, the sequence of features $\tilde{x} \in \mathbb{R}^{L\times D}$ is projected onto D/H-dimensional subspaces (called the "attention heads") for an integer H dividing D, obtaining, for $1 \leq i \leq H$,

$$V_i = \tilde{x}W_i^V + b_i^V \in \mathbb{R}^{L\times D/H}, \quad K_i = \tilde{x}W_i^K + b_i^K \in \mathbb{R}^{L\times D/H}, \quad Q_i = \tilde{x}W_i^Q + b_i^Q \in \mathbb{R}^{L\times D/H},$$

where $W_i^V, W_i^K, W_i^Q \in \mathbb{R}^{D\times F/H}$, $b_i^V, b_i^K, b_i^Q \in \mathbb{R}^{D/H}$ are parameters to be estimated. Next, each projection V_i is processed temporally obtaining the hidden states $h_i \in \mathbb{R}^{L\times D/H}$, with

$$h_{i,l} = \sum_{j=1}^{L} w_{l,j,i} V_{i,j} \in \mathbb{R}^{D/H}, \quad w_{l,j,i} = \frac{\exp(K_{i,l} \cdot Q_{i,j})}{\sum_{m=1}^{L} \exp(K_{i,l} \cdot Q_{i,m})} \in [0,1].$$

These states are then concatenated and linearly combined, obtaining the last hidden state

$$h = \text{Concat}(h_1, ..., h_h)W^O + b^O \in \mathbb{R}^{L\times D},$$

where $W^O \in \mathbb{R}^{F\times F}, b^O \in \mathbb{R}^F$ are parameters to be estimated.

Finally, h is normalized and processed time-wise though a 2-layered feedforward network as described in detail in the original paper (Vaswani et al. (2017)). The number of hidden units in the intermediate layer is a technical hyperparameter that we call HDN in section C.A. This network also has dropout regularization with hyperparameter called DRP in section C.A.

Appendix D. Network Estimation Details

As explained in Section II.C, we estimate the parameters of the models with neural networks by solving the optimization problems introduced in equation (4) or in equation (6) of section II.C, depending on the model and the objective function. In all cases, this is done by replacing the mean and variance by their annualized sample counterpart over a training set, and by finding the optimal network parameters with stochastic gradient descent using PyTorch's Adam optimizer and the optimization hyperparameters learning rate and number of optimization epochs described in detail in section C.A.

As mentioned in Section III.C, our main results use rolling windows of 1,000 days as training sets. The networks are reestimated every 125 days to strike a balance between computational efficiency and adaptation to changing economic conditions, and the strategies' returns are always obtained out-of-sample. Additionally, to be able to train our model over these long windows without running into memory issues, we split each training window into temporal "batches", as is commonly done in deep learning applications. Each batch contains the returns and residuals for all the stocks in a subwindow of 125 days of the original training window, with the subwindows being consecutive and non-overlapping (i.e., for a training window of 1000 days, we split it into the subwindow containing the first 125 days, the subwindow containing the days between the 126th day and the 250th day, etc.). The optimization process is applied successively to each batch, completing the full sequence of batches before starting a new optimization iteration or epoch.

In the implementation of our optimization procedure under market frictions, we found it useful to include the last allocation as an additional input to the allocation function $\boldsymbol{w}^{\boldsymbol{\epsilon}}$, as the inclusion of the cost term makes the objective function depend on it. However, the inclusion of the previous allocations in either the objective function or the architecture of the model complicates the parallelization of the training and evaluation computations, because after this change the model requires the output of previous lookback window in order to compute the output of the current window. To allow training to remain parallelized, which is desirable for reasonable computational speed given the volume of data of our empirical application, in our implementation of the training function in each epoch, we take the previous allocations from the output of the previous epoch and use them as a pre-computed approximation of the allocations for the current epoch. This approximation converges in our empirical experiments and allows us to maintain parallelization, but may produce suboptimal results. For evaluation purposes, however, everything is computed exactly and with no approximations using a sequential approach.

Throughout section III, all presented results have been computed with PyTorch 1.5 and have been parallelized across 8 NVIDIA GeForce GTX Titan V GPUs, on a server with two Intel Xeon E5-2698v3 32-core CPUs and 1 TB of RAM. The full rationale for the hyperparameter choices are described in detail in section C.A, but for a CNN+Transformer model with a lookback window of 30 days, 8 convolutional filters with a filter size of 2, 4 attention heads, 125-day reestimation using a rolling lookback window of 1000, it takes our deep model approximately 7 hours to be periodically estimated and run in our 19 years of daily out-of-sample data with our universe of on average $\sim$550 stocks per month.

Appendix C. Additional Empirical Results

Appendix A. Robustness to Hyperparameter Selection

In this subsection, we describe our hyperparameter selection procedure and explore additional hyperparameter choices to show that the performance of our strategies is extremely robust to our choices. These results complement the time stability checks we exhibited in Section III.I. To decide which hyperparameters we would select for use in our network, we fixed a validation dataset as follows: we took the first 1000 trading days of our data set of residuals (all trading days from January 1, 1998 through December 31, 2001) of the 5-factor IPCA-based model, which is estimated with a 20-year rolling window. Because it is solely used for training in our rolling train/test procedures used to compute strategy returns, this data is completely in-sample, and thus completely avoids look-ahead bias which would influence any of our out-of-sample trading results in the main text. We started with a reasonable set for our hyperparameters, and tested also additional points adjacent to these sets.[19] For each model represented by a point on the grid, we trained the model using the Sharpe ratio objective on the first 750 days of the 1000 trading days, and evaluated it by its out-of-sample Sharpe ratio on the last 250 days of the 1000 trading days. We tested 16 combinations of hyperparameters, which are illustrated in Table A.II. The results of our test on the last 250 days of our validation data are displayed in Table A.III.

The results in Table A.III show that all Sharpe ratios fall within a tight range of values, which is roughly $[3.5, 4.2]$. Means and volatilities concentrate similarly, falling within $[13\%, 17.8\%]$ and $[3.6\%, 4.3\%]$. Computation of 95% bootstrapped confidence intervals for mean return shows that all models' confidence intervals contain the interval $[10\%, 20\%]$, with volatilies similarly contained. Hence, these models are statistically not distinguishable. Given the statistical insignificance of the differences in performance of these models, we chose the model displayed in Table A.II, which is the most parsimonious one, that is it has the smallest number of parameters, and hence benefits low GPU memory usage and ease of interpretability.

Table A.II: Hyperparameter options for the network in the empirical analysis

Notation	Hyperparameters	Candidates	Chosen
D	Number of filters in the convolutional network	8, 16	8
ATT	Number of attention heads	2, 4	4
HDN	Number of hidden units in the transformer's linear layer	2D, 3D	2D
DRP	Dropout rate in the transformer	0.25, 0.5	0.25
D_{size}	Filter size in the convolutional network	2	2
LKB	Number of days in the residual lookback window	30	30
WDW	Number of days in the rolling training window	1000	1000
RTFQ	Number of days of the retraining frequency	125	125
BTCH	Batch size, in days	125	125
LR	Learning rate	0.001	0.001
EPCH	Number of optimization epochs	100	100
OPT	Optimization method	Adam	Adam

This table shows the parameters for our network architecture with respect to the Sharpe ratio on our validation data and the candidates we tried In DRP, we follow the convention that the dropout rate p is the proportion of units which are removed.

To ensure that our results are stable across several choices of hyperparameters, we study the results of four additional models with perturbed hyperparameters. This complements our robustness results of Section III.I regarding the size of the lookback window and the retraining frequency. The four additional networks and their hyperparameter configurations are listed in Table A.IV, with Network 1 being the network studied throughout this empirical section. Network 2 corresponds to more filters and commensurately more hidden units to consume them, and higher dropout rates to more strongly regularize the additional parameters. Network 3 halves the number of attention heads from our original specification. Networks 4 and 5 modify the size of the rolling training window from 1000 trading days to 1250 and 750 trading days, respectively, which corresponds closely to three and five calendar years. These additional hyperparameter configurations constitute local

[19] Note the computation over a large set of hyperparameters is computationally infeasible, which requires us to restrict the set to reasonable values.

Table A.III: Performance of candidate models on the last year of the validation data set

D	ATT	HDN	DRP	SR	μ	σ
8	2	2	0.25	3.81	16.3%	4.3%
8	2	2	0.50	3.92	16.0%	4.1%
8	2	3	0.25	3.79	16.2%	4.3%
8	2	3	0.50	4.00	16.4%	4.1%
8	4	2	0.25	3.81	15.6%	4.1%
8	4	2	0.50	4.13	17.8%	4.3%
8	4	3	0.25	3.82	15.6%	4.1%
8	4	3	0.50	4.16	17.4%	4.2%
16	2	2	0.25	4.00	14.8%	3.7%
16	2	2	0.50	4.06	16.2%	4.0%
16	2	3	0.25	4.11	14.9%	3.6%
16	2	3	0.50	4.06	16.6%	4.1%
16	4	2	0.25	3.93	15.6%	4.0%
16	4	2	0.50	3.66	13.9%	3.8%
16	4	3	0.25	4.18	16.8%	4.0%
16	4	3	0.50	3.51	13.0%	3.7%

This table shows the model performance with respect to the Sharpe ratio, mean, and volatility on our validation data set for the candidate models implied by Table A.II. The models are trained on the first three years of the validation data set (1998–2000) and tested on the last year (2001). In DRP, we follow the convention that the dropout rate p is the proportion of units which are removed.

perturbations in hyperparameter space, to which our strategies' performance should be relatively robust.

Table A.IV: Alternative best performing models on the data from 2002–2016

Model	FLNB	FLSZ	ATT	HDN	DRP	LKB	WDW
Network 1	[1,8]	2	4	16	0.25	30	1000
Network 2	[1,16]	2	4	32	0.5	30	1000
Network 3	[1,8]	2	2	16	0.25	30	1000
Network 4	[1,8]	2	4	16	0.25	30	1250
Network 5	[1,8]	2	4	16	0.25	30	750

This table reports four of the best performing models for our network architecture with respect to the Sharpe ratio on our data from 2002–2016 and the candidates described in Table A.II. Our original network, which is studied throughout this section, is labeled as Network 1.

In Table A.V, we report the results of these models on a representative subset of 5-factor models, which are now evaluated on the full out-of-sample data. We see that the Sharpe ratios are broadly similar across all three different perturbations of network architecture hyperparameters (i.e., number of filters, number of attention heads, and dropout rate). The small range of values induced by these choices shows that our network performs similarly over a variety of sensible network parameters, and highlights the efficacy of our reasonable choice of convolutional, attentional, and feedforward subnetworks which specialize in finding small temporal patterns, arranging these patterns throughout time, and deciding on allocations based on these arranged patterns. The only

Table A.V: Performance of the alternative models on our benchmark residual datasets, 2002–2016

	Fama-French 5			PCA 5			IPCA 5		
Model	SR	μ	σ	SR	μ	σ	SR	μ	σ
Network 1	3.21	4.6%	1.4%	3.36	14.3%	4.2%	4.16	8.7%	2.1%
Network 2	3.16	4.6%	1.4%	3.26	13.9%	4.3%	4.35	8.4%	1.9%
Network 3	3.30	4.8%	1.4%	3.17	13.4%	4.2%	4.00	8.4%	2.1%
Network 4	2.93	4.1%	1.4%	2.74	11.7%	4.3%	3.96	7.9%	2.0%
Network 5	3.13	4.9%	1.6%	3.52	15.0%	4.3%	3.77	8.6%	2.3%

This table shows the average annualized returns, volatilities and Sharpe ratios of our alternative models from Table A.IV on our three benchmark residual datasets, trained with the Sharpe ratio objective function.

hyperparameter choice which can be considered to be slightly lower than the rest is the choice of using a rolling training window of 1250 trading days. This suggests that statistical arbitrage policies are time-varying to adopt to changing economic conditions. Our results show that using the latest 3 or 4 years of data for the estimation of the parameters of this model achieves an adequate balance between capturing changing conditions and using sufficient data.

For the allocation function feedforward network (FFN) utilized for the Fourier+FFN model, we choose a reasonable architecture based on deep learning conventions and have verified that the results are robust to this choice. Because the input of the network are the $L = 30$ coefficients of the Fourier decomposition of each residual window $(X_l^{(n,t)})_{1\leq l\leq L}$ and the output is the corresponding allocation weight $w_{n,t} \in \mathbb{R}$, we follow standard deep learning practices and consider 3 hidden layers with dimensions 16,8,4 regularized with a dropout rate of 0.25. We use the ReLU activation function, and train using the same procedure outlined in B.D, with the same batch size, learning rate, number of optimization epochs, and optimization method as in Table A.II.

Appendix B. Interpretation

Figure A.2: Additional Examples of Allocation Weights and Signals

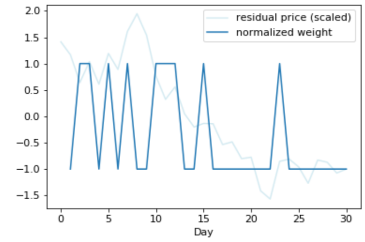

(a) Cumulative residuals x_l and allocation weight $w_l^{\epsilon|\text{CNN+Trans}}$

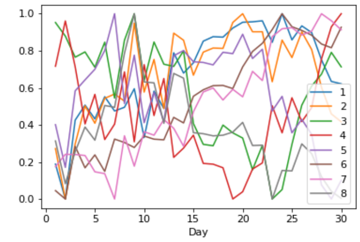

(b) Signal $\theta_l^{\text{CNN+Trans}}$

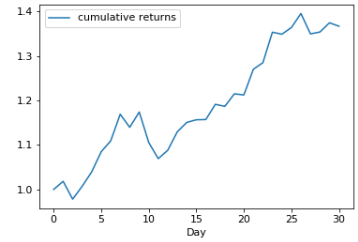

(c) Cumulative returns of CNN+Trans strategy

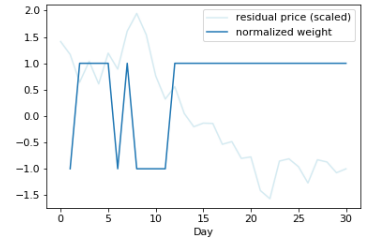

(d) Cumulative residual x_l and allocation weight $w_l^{\epsilon|\text{FFT}}$

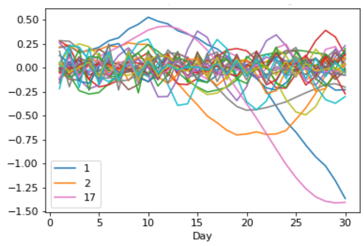

(e) Signal θ_l^{FFT}

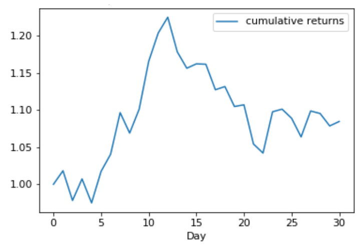

(f) Cumulative returns of Fourier+FFN strategy

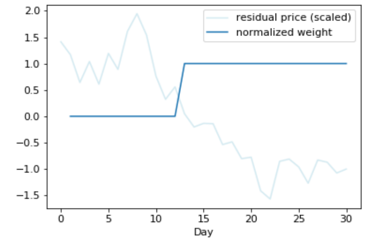

(g) Cumulative residual x_l and allocation weight $w_l^{\epsilon|\text{OU}}$

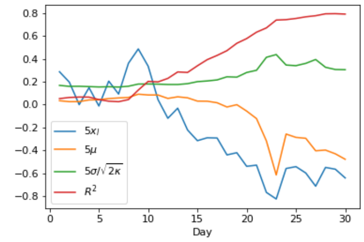

(h) Signal θ_l^{OU}

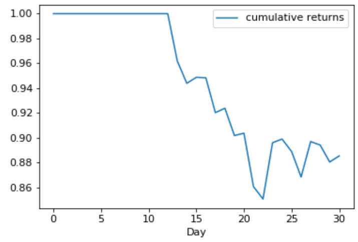

(i) Cumulative returns of OU+Thres strategy

These plots are an illustrative example of the allocation weights and signals of the Ornstein-Uhlenbeck with Threshold (OU+Thres), Fast Fourier Transform (FFT) with Feedforward Neural Network (FFN), and Convolutional Neural Network (CNN) with Transformer models for a specific cumulative residual. The models are estimated on the empirical data, and the residual is a representative empirical example. In more detail, we consider the residuals from five IPCA factors and estimate the benchmark models as explained in Section III.J. The left subplots display the cumulative residual process along with the out-of-sample allocation weights $w_l^{\epsilon|\cdot}$ that each model assigns to this specific residual. In this example, we consider trading only this specific residual and hence normalize the weights to $\{-1, 0, 1\}$. The middle column plots show the time-series of estimated out-of-sample signals for each model, by applying the $\theta_l^{\cdot}$ arbitrage signal function to the previous L cumulative returns of the residual. The right column plots display the out-of-sample cumulative returns of trading this particular residual based on the corresponding allocation weights. We use a rolling lookback window of $L = 30$ days to estimate the signal and allocation, which we evaluate for the out-of-sample on the next 30 days. The plots only show the out-of-sample period. The evaluation of this illustrative example is a simplification of the general model that we use in our empirical main analysis, where we trade all residuals and map them back into the original stock returns.

Figure A.3: Example Attention Weights for Sinusoidal Residual Inputs

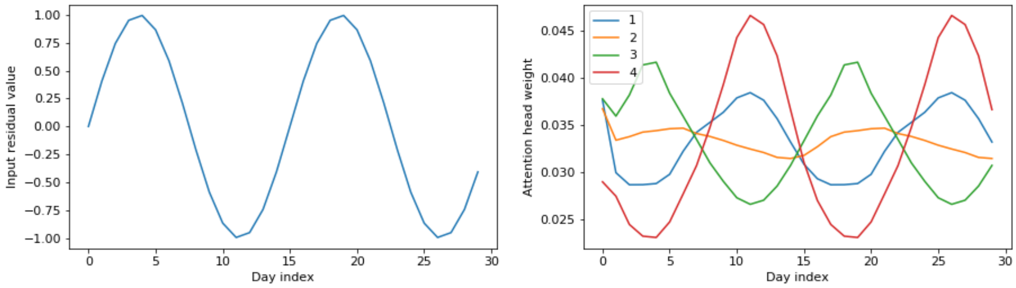

(a) Input residual and attention head weights for $\omega_{\text{sin}} = 2/30$

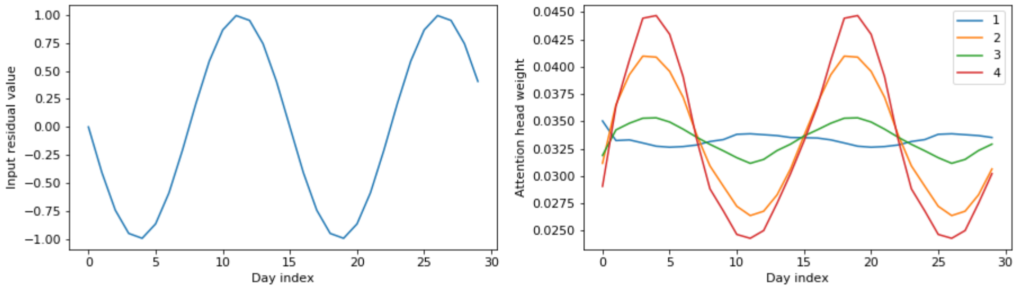

(b) Input residual and attention head weights for $\omega_{\text{sin}} = 28/30$

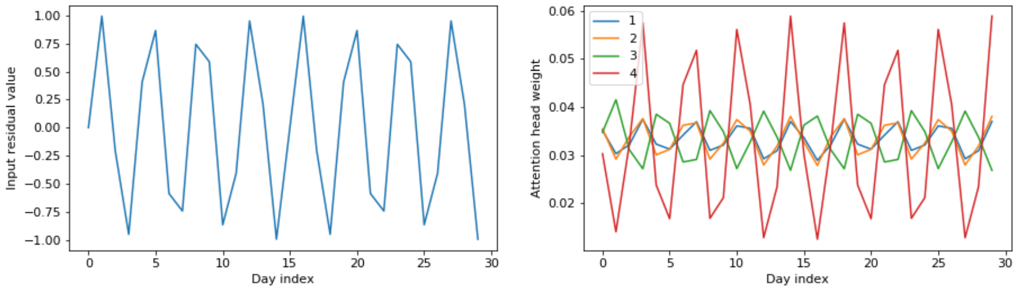

(c) Input residual and attention head weights for $\omega_{\text{sin}} = 8/30$

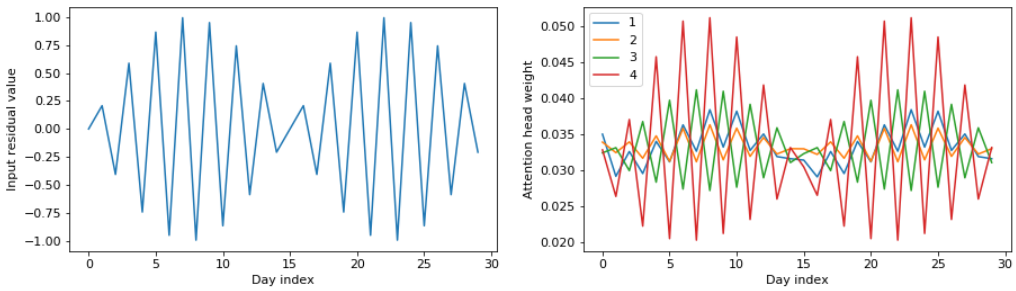

(d) Input residual and attention head weights for $\omega_{\text{sin}} = 14/30$

These plots show the attention head weights of the CNN+Transformer benchmark model for simulated sinusoidal residual input time series. The inputs are $x_l = \sin(2\pi\omega_{\text{sin}} l)$, for various ω_{sin} and $l \in \{0, ..., 29\}$. The right subplot shows the attention weights for the $H = 4$ attention heads for the specific residuals. The empirical benchmark model is the CNN+Transformer model based on IPCA 5-factor residuals. We estimate the model on only once on the first T_{train}=8 years based on the Sharpe ratio objective.

Appendix C. Dependency between Arbitrage Strategies

Table A.VI: Correlations between the Returns of the CNN+Transformer Arbitrage Strategies

	Fama-French 3	PCA 3	IPCA 3	Fama-French 5	PCA 5	IPCA 5	PCA 10	IPCA 10
Fama-French 3	1.00	0.32	0.44	0.62	0.25	0.43	0.21	0.44
PCA 3	0.32	1.00	0.32	0.34	0.62	0.35	0.41	0.36
IPCA 3	0.44	0.32	1.00	0.37	0.28	0.81	0.21	0.75
Fama-French 5	0.62	0.34	0.37	1.00	0.28	0.39	0.23	0.40
PCA 5	0.25	0.62	0.28	0.28	1.00	0.29	0.47	0.31
IPCA 5	0.43	0.35	0.81	0.39	0.29	1.00	0.23	0.84
PCA 10	0.21	0.41	0.21	0.23	0.47	0.23	1.00	0.25
IPCA 10	0.44	0.36	0.75	0.40	0.31	0.84	0.25	1.00

This table reports the correlations of our CNN+Transformer strategies for some representative choices of the factor models. The correlations are calculated with returns of the out-of-sample arbitrage trading from January 2002 to December 2016. The models are calibrated on a rolling window of four years and use the Sharpe ratio objective function. The signals are extracted from a rolling window of $L = 30$ days.

Appendix D. Time-Series Signal

In this appendix, we report the OOS returns of strategies using alternative models for the ablation tests in Section III. For the FFN feedforward network, we use the same architecture, hyperparameters, optimization settings, etc. as in the Fourier+FFN model utilized throughout the empirical results section and described in Appendix C.A. For the OU+FFN model, because the input is the low-dimensional OU signal in $\mathbb{R}^4$, we consider a 3 hidden layer with dimensions 4,4,4 regularized with a dropout rate of 0.25. We use the sigmoid activation function, and estimate it using the same procedure outlined in section B.D, with the same batch size, learning rate, number of optimization epochs, and optimization method as in Table A.II.

Table A.VII: OOS Annualized Performance Based on Sharpe Ratio Objective

	Factors	Fama-French			PCA			IPCA		
Model	K	SR	μ	σ	SR	μ	σ	SR	μ	σ
OU + FFN	0	0.50	10.6%	21.3%	0.50	10.6%	21.3%	0.50	10.6%	21.3%
	1	0.34	0.8%	2.3%	0.05	0.7%	11.9%	0.60	4.8%	8.0%
	3	0.16	0.2%	1.4%	0.44	3.4%	7.8%	0.70	4.6%	6.6%
	5	0.17	0.2%	1.2%	0.68	4.7%	7.0%	0.66	4.2%	6.3%
	8	-0.34	-0.3%	1.0%	0.51	3.1%	6.0%	0.60	3.9%	6.2%
	10	-	-	-	0.26	1.3%	5.0%	0.56	3.5%	6.2%
	15	-	-	-	0.31	1.4%	4.3%	0.54	3.3%	6.1%
FFN	0	0.57	8.8%	15.3%	0.57	8.8%	15.3%	0.57	8.8%	15.3%
	1	0.60	2.0%	3.3%	0.53	6.2%	11.7%	1.07	6.5%	6.1%
	3	1.02	2.6%	2.6%	1.15	8.2%	7.2%	1.50	7.6%	5.0%
	5	1.32	2.3%	1.7%	1.42	9.8%	6.9%	1.55	7.3%	4.7%
	8	1.31	2.1%	1.6%	1.05	6.4%	6.2%	1.52	7.2%	4.7%
	10	-	-	-	0.70	3.5%	5.0%	1.48	7.0%	4.7%
	15	-	-	-	0.51	2.4%	4.8%	1.68	7.5%	4.5%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) of our three statistical arbitrage models for different numbers of risk factors K, that we use to obtain the residuals. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. OU+FFN denotes a parametric Ornstein-Uhlenbeck model to extract the signal, but a flexible feedforward neural network to estimate the allocation function. RawFFN takes the residuals directly as signals and estimates an allocation function with a feedforward neural network. The deep learning models are calibrated on a rolling window of four years and use the Sharpe ratio objective function. The signals are extracted from a rolling window of $L = 30$ days. The $K = 0$ factor model corresponds to directly using stock returns instead of residuals for the signal and trading policy.

Appendix E. Trading Friction Results for PCA Residuals

Table A.VIII: OOS Performance of CNN+Trans with Trading Frictions

PCA factor model						
	Sharpe ratio			Mean-variance		
K	SR	μ	σ	SR	μ	σ
0	0.52	8.5%	16.3%	0.22	2.6%	11.9%
1	0.88	7.3%	8.4%	0.79	9.0%	11.4%
3	0.90	5.7%	6.3%	0.62	4.7%	7.6%
5	0.81	4.5%	5.6%	0.68	4.4%	6.4%
10	-0.08	-0.4%	4.8%	-0.08	-0.4%	4.6%
15	-0.87	-3.7%	4.3%	-0.96	-3.5%	3.7%

This table shows the out-of-sample annualized Sharpe ratio (SR), mean return (μ), and volatility (σ) for the CNN+Transformer model with trading frictions on PCA residuals. We use the daily out-of-sample residuals from January 1998 to December 2016 and evaluate the out-of-sample arbitrage trading from January 2002 to December 2016. The models are calibrated on a rolling window of four years and use either the Sharpe ratio or mean-variance objective function with trading costs $\mathrm{cost}(w^R_{t-1}, w^R_{t-2}) = 0.0005\|w^R_{t-1} - w^R_{t-2}\|_{L^1} + 0.0001\|\min(w^R_{t-1}, 0)\|_{L^1}$. The signals are extracted from a rolling window of $L = 30$ days.

Deep Limit Order Book Trading: Half a second please!

Presenter: Jie Yin

Joint work with Hoi Ying Wong

Department of Statistics, The Chinese University of Hong Kong, Shatin, Hong Kong

Extended Abstract

We introduce an expert deep-learning system for limit order book (LOB) trading for markets in which the stock tick frequency is longer than or close to 0.5 seconds, such as the Chinese A-share market. This half a second enables our system, which is trained with a deep-learning architecture, to integrate price prediction, trading signal generation, and optimization for capital allocation on trading signals altogether. It also leaves sufficient time to submit and execute orders in exchanges before the next tick-report.

The key ingredients to enable our system to be realized and profitable in practice are followed. The basic features of the LOB are normalized and analyzed using a deep convolutional neural network (DCNN) model to classify current trading actions: long, short, and none. This DCNN model contains multiple convolutional layers, inception module and LSTM module for training on very large LOB datasets. As the universal features can be extracted by deep learning techniques, group training is conducted. Besides, we find that the number of signals generated from the system can be used to rank stocks for the preference of LOB trading. Other contributions of this novel system are the application of the focal loss function to training to deal with imbalanced classification and optimization using the fractional Kelly growth criterion under risk control, which is further achieved by the risk measure, value at risk (VaR). These mechanisms greatly enhance returns and reduce risks. Furthermore, our system considers factors in real trading circumstances, such as order choices, transaction costs, and extreme losses.

Finally, we test our data-driven LOB trading system with two types of data: simulated data generated using the zero-intelligence agent-based model and real LOB data from the Chinese A-share market. The simulation experiments demonstrate the characteristics of the trading system in different market sentiments, while the empirical study with real data confirms significant profits after factoring in transaction costs and risk requirements. This innovative system is inspirational and will serve as a reference for financial firms or organizations engaging in LOB trading to better use machine learning and optimization.

Abstract. We present an approach, based on deep neural networks, that allows identifying robust statistical arbitrage strategies in financial markets. Robust statistical arbitrage strategies refer to self-financing trading strategies that allow profitable trading under model ambiguity. The presented novel methodology does not suffer from the curse of dimensionality nor does it depend on the identification of cointegrated pairs of assets and is therefore applicable even on high-dimensional financial markets or in markets where classical pairs trading approaches fail, for what we also provide evidence in empirical investigations. In these investigations we further provide a routine how the ambiguity set of admissible measures can be derived from market data. Thus, the approach can be considered as being entirely data-driven.

1. Introduction

The term *statistical arbitrage* is commonly used in finance to describe trading strategies which are profitable on average, but, in contrast to pure arbitrage strategies, not necessarily in every market scenario that is deemed to be possible. This generalized notion therefore forms the foundation for a systematic approach enabling to trade profitably even in markets in which it is difficult or impossible to detect pure arbitrage strategies.

A popular class of trading strategies that are often referred to as statistical arbitrage strategies are *pairs trading* strategies, which all rely on the fundamental idea that, given two financial assets are strongly related, e.g., through a cointegration property (compare e.g. [3]) or through a low level of the variance of the spread between the two assets (compare e.g. [8]), then deviations of the spread are assumed to last only for a short period of time and, thus, eventually the spread of the asset pair will return to its long-term equilibrium. This mean-reversion property is exploited by trading in the opposite direction of the deviation after the deviation has exceeded a certain threshold and by clearing the position when the spread has reached again a level close to its long-term equilibrium, see also [1], [6], [7], [8], [10], [11], and [13].

The apparent drawback of using pairs trading strategies is the strong dependence on the underlying mean-reversion property of the spread process. Indeed, if the mean-reversion relation breaks down and thus the spread does not converge to its equilibrium, then the pairs trading approach will in general not be profitable. Moreover, in an empirical study, [5] provide evidence that the profitability of pairs trading strategies has declined over the recent years, mainly as it has become increasingly difficult to identify profitable asset pairs.

Our contribution is based on another class of statistical arbitrage strategies possessing the advantage not to rely on a mean-reversion property and which are therefore applicable independent of the prevalent market scenario and, as we will show, even in periods where the pairs trading approach fails. Our approach relies on the idea of Bondarenko ([2]), who introduced and characterized statistical arbitrage strategies as self-financing strategies which are profitable on average given any terminal value of the underlying securities at a fixed maturity $t_n > 0$, i.e., if $\Phi(S)$ denotes the profit of a trading strategy investing in underlying securities $S = (S)_{t_0 \leq t \leq t_n}$, then statistical arbitrage strategies fulfil $\mathbb{E}[\Phi(S) \mid S_{t_n}] \geq 0$. Based on this idea, [9] generalized the notion of statistical arbitrage from [2] by introducing $\mathcal{G}$-arbitrage defined through zero-cost payoffs Y (which are not necessarily the payoffs of self-financing trading strategies) fulfilling

$$\mathbb{E}[Y \mid \mathcal{G}] \geq 0 \tag{1.1}$$

for $\mathcal{G}$ being a σ-algebra $\mathcal{G} \subseteq \sigma(S)$, which allows, in particular, to take into account more flexible choices of trading strategies, possibly adjusted to available information. Moreover, the results from [2, Proposition 1], [9, Proposition 6], and [14, Theorem 3.3] characterize the existence of $\mathcal{G}$-arbitrage strategies by relating the absence of self-financing strategies fulfilling (1.1) to the existence of $\mathcal{G}$-measurable Radon-Nikodym densities, a result which can be considered as an extension of the fundamental theorem of asset-pricing ([4]) which connects the absence of arbitrage with the existence of pricing measures. The authors from [14] further propose and validate empirically an embedding-methodology to exploit statistical arbitrage on financial markets. All these mentioned contributions assume that the underlying securities behave according to a previously fixed underlying probability measure $\mathbb{P}$. To account for ambiguity w.r.t. the choice of an underlying probability measure [12] recently introduced the notion of $\mathcal{P}$-robust $\mathcal{G}$-arbitrage referring to self-financing trading strategies that allow for $\mathcal{G}$-arbitrage independent which measure $\mathbb{P}$ from the considered ambiguity set $\mathcal{P}$ is the "correct" underlying probability measure, as the strategy is required to be profitable for all measures from the ambiguity set. Hence, this notion allows, in particular, to take into account probability measures that would incur losses for conventional pairs trading strategies, and thus to determine trading strategies that are profitable even in scenarios where pairs trading fails.

We build on the notion of $\mathcal{P}$-robust $\mathcal{G}$-arbitrage and establish a numerical routine involving deep neural networks to determine $\mathcal{P}$-robust statistical arbitrage strategies, i.e., $\mathcal{P}$-robust $\mathcal{G}$-arbitrage strategies for the choice $\mathcal{G} = \sigma(S_{t_n})$, where S_{t_n} denotes the terminal values of the underling securities. To this end, we introduce a functional which penalizes self-financing trading strategies that are not statistical arbitrage strategies for each of the considered measures from the ambiguity set $\mathcal{P}$ and therefore, by construction, the functional prefers $\mathcal{P}$-robust statistical arbitrage strategies as optimal solution. We show as a main result that minimizing this functional among strategies that can be represented as the outputs of deep neural networks is, for a sufficiently large penalization parameter, equivalent to the determination of $\mathcal{P}$-robust statistical arbitrage strategies. In particular, through the representation of trading strategies by neural networks, our approach does, in contrast to approaches relying on linear programming, not suffer from the *curse of dimensionality* and is hence tractable even when the trading strategy depends on a large amount of underlying assets.

In an extensive study we show how it is possible to determine based on historical financial market data an ambiguity set $\mathcal{P}$. In contrast to the determination of statistical arbitrage strategies w.r.t. a fixed underlying probability measure, taking into account the ambiguity set allows, particularly, that the future developments of the underlying process differ from the historical evolution and still computed $\mathcal{P}$-robust statistical arbitrage strategies turn out to be profitable. Indeed, by using this data-driven ambiguity set we provide empirical evidence for the profitability of our approach in overall bad market scenarios, in high-dimensional financial markets, and in scenarios where classical pairs trading approaches fail.

References

[1] Marco Avellaneda and Jeong-Hyun Lee. Statistical arbitrage in the us equities market. *Quantitative Finance*, 10(7):761–782, 2010.

[2] Oleg Bondarenko. Statistical arbitrage and securities prices. *The Review of Financial Studies*, 16(3):875–919, 2003.

[3] João Caldeira and Guilherme V Moura. Selection of a portfolio of pairs based on cointegration: A statistical arbitrage strategy. *Available at SSRN 2196391*, 2013.

[4] Freddy Delbaen and Walter Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische annalen*, 300(1):463–520, 1994.

[5] Binh Do and Robert Faff. Does simple pairs trading still work? *Financial Analysts Journal*, 66(4):83–95, 2010.

[6] Binh Do, Robert Faff, and Kais Hamza. A new approach to modeling and estimation for pairs trading. In *Proceedings of 2006 financial management association European conference*, volume 1, pages 87–99. Citeseer, 2006.

[7] Robert J Elliott, John Van Der Hoek*, and William P Malcolm. Pairs trading. *Quantitative Finance*, 5(3):271–276, 2005.

[8] Evan Gatev, William N Goetzmann, and K Geert Rouwenhorst. Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, 19(3):797–827, 2006.

[9] Stefan Kassberger and Thomas Liebmann. Additive portfolio improvement and utility-efficient payoffs. *Mathematics and Financial Economics*, 11(2):241–262, 2017.

[10] Christopher Krauss. Statistical arbitrage pairs trading strategies: Review and outlook. *Journal of Economic Surveys*, 31(2):513–545, 2017.

[11] Christopher Krauss, Xuan Anh Do, and Nicolas Huck. Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the s&p 500. *European Journal of Operational Research*, 259(2):689–702, 2017.

[12] Eva Lütkebohmert and Julian Sester. Robust statistical arbitrage strategies. *Quantitative Finance*, 21(3):379–402, 2021.

[13] Hossein Rad, Rand Kwong Yew Low, and Robert Faff. The profitability of pairs trading strategies: distance, cointegration and copula methods. *Quantitative Finance*, 16(10):1541–1558, 2016.

[14] Christian Rein, Ludger Rüschendorf, and Thorsten Schmidt. Generalized statistical arbitrage concepts and related gain strategies. *Mathematical Finance*, 31(2):563–594, 2021.

Dispersion-Constrained Martingale Schrödinger Problems and the Exact Joint S&P 500/VIX Smile Calibration Puzzle

Julien Guyon[1,2,3]

[1] Quantitative Research, Bloomberg L.P., New York

[2] Department of Mathematics, Columbia University

[3] Courant Institute of Mathematical Sciences, NYU

We solve for the first time[1] a longstanding puzzle of quantitative finance that has often been described as the Holy Grail of volatility modeling: build a model that jointly and exactly calibrates to the prices of S&P 500 (SPX) options, VIX futures, and VIX options. So far the best attempts, which used parametric continuous-time (jump-)diffusion models on the SPX, only produced approximate fits. We use a very different, nonparametric and discrete-time approach. Given a VIX future maturity T_1, we consider the set $\mathcal{P}$ of all joint probability measures on the SPX at T_1, the VIX at T_1, and the SPX at $T_2 = T_1 + 30$ days which are perfectly calibrated to the full SPX smiles at T_1 and T_2, and the full VIX smile at T_1, and which also satisfy the martingality constraint on the SPX as well as the requirement that the VIX is the implied volatility of the 30-day log-contract on the SPX.

We first consider robust hedging in this setting. By casting the superreplication problem as what we call a *dispersion-constrained martingale optimal transport problem*, we establish a strong duality theorem and, as a result, prove that the absence of joint SPX/VIX arbitrage is equivalent to the set $\mathcal{P}$ of jointly calibrating models being nonempty. Should they arise, joint arbitrages are identified using classical linear programming. In the absence of joint arbitrage, we then provide a solution to the joint calibration puzzle by solving what we call a *dispersion-constrained martingale Schrödinger problem*: we choose a reference measure and build the unique jointly calibrating model that minimizes the relative entropy. We establish several dual versions of the problem, one of which has a natural financial interpretation in terms of exponential utility indifference pricing, and prove absence of duality gaps. The minimum entropy jointly calibrating model is explicit in terms of what we call the dual *Schrödinger portfolio*, i.e., the maximizer of the dual problems, should it exist. We numerically compute this Schrödinger portfolio using an extension of the Sinkhorn algorithm, in the spirit of De March and Henry-Labordère (2019). Our numerical experiments show that the algorithm performs very well in both low and high volatility regimes.

Along the way, we provide new variants, as well as a new proof, of strong duality theorems for the classical Schrödinger problem and for a mixed *Schrödinger-Monge-Kantorovich problem* (also known as *entropic optimal transport problem*) that has recently attracted a lot of attention in the optimal transport community, which are interesting in themselves. Our methodology applies not only to the VIX, but also to any index computed as a function of the price of an option on some underlying asset.

Finally, we explain how our technique of building a jointly calibrating model extends to continuous-time models, (a) using a martingale interpolation of the discrete-time model, and (b) using Schrödinger bridges, in the spirit of [19].

[1] A first, much shorter version of this work was published in *Risk* in April 2020 [14]. This version includes new theorems, proofs, analysis, and numerical tests. In particular it develops a theory of dispersion-constrained martingale Schrödinger problems and proves strong duality results for them.

References

[1] Avellaneda, M.: *Minimum-relative-entropy calibration of asset pricing models*, International Journal of Theoretical and Applied Finance, 1(4):447–472, 1998.

[2] Beiglböck, M., Henry-Labordère, P., Penkner, F.: *Model-independent bounds for option prices: A mass-transport approach*, Finance Stoch., 17(3):477–501, 2013.

[3] Cuturi, M.: *Sinkhorn distances: Lightspeed computation of optimal transport*, Advances in neural information processing systems, 2292–2300, 2013.

[4] De March, A., Henry-Labordère, P.: *Building arbitrage-free implied volatility: Sinkhorn's algorithm and variants*, SSRN preprint 3326486, 2019.

[5] De Marco, S., Henry-Labordère, P.: *Linking vanillas and VIX options: A constrained martingale optimal transport problem*, SIAM Journal on Financial Mathematics, 6:1171–1194, 2015.

[6] Föllmer, H.: *Random fields and diffusion processes*, Hennequin PL. (eds) École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87, Lecture Notes in Mathematics, vol 1362, Springer, Berlin, Heidelberg, 1988.

[7] Föllmer, H., Gantert, N.: *Entropy Minimization and Schrödinger Processes in Infinite Dimensions*, Annals of Probability, 25(2):901–926, 1997.

[8] Frittelli, M.: *The minimal entropy martingale measure and the valuation problem in incomplete markets*, Mathematical Finance, 10(1):39–52, 2000.

[9] Gatheral, J.: *Consistent modeling of SPX and VIX options*, presentation at the Bachelier congress, London, July 18, 2008.

[10] Guo, I., Loeper, G., Obłój, J., Wang, S.: *Joint Modeling and Calibration of SPX and VIX by Optimal Transport*, SIAM J. Finan. Math. 13(1), 2022.

[11] Guyon, J., Menegaux, R., Nutz, M.: *Bounds for VIX futures given S&P 500 smiles*, Finan. & Stoch., 21(3):593–630, 2017.

[12] Guyon, J.: *On the joint calibration of SPX and VIX options*, presentations at Jim Gatheral's 60th birthday conference, New York, October 14, 2017; at Imperial College London, March 28, 2018; at QuantMinds 2018, Lisbon, May 16, 2018.

[13] Guyon J.: *Inversion of Convex Ordering in the VIX Market*, Quantitative Finance 20(10):1597–1623, 2020.

[14] Guyon, J.: *The joint SPX/VIX smile calibration puzzle solved*, Risk, April 2020.

[15] Guyon, J.: *The joint SPX/VIX smile calibration puzzle solved in continuous time with Schrödinger bridges*, talk at the Bachelier Finance Society One World Seminar, July 16, 2020.

[16] Guyon, J.: *The VIX future in Bergomi models: fast approximation formulas and joint calibration with S&P 500 skew*, preprint, 2020.

[17] Guyon, J.: *Dispersion-Constrained Martingale Schrödinger Problems and the Exact Joint S&P 500/VIX Smile Calibration Puzzle*, preprint, available at `ssrn.com/abstract=3853237`, 2021.

[18] Henry-Labordère, P.: *Model-free Hedging: A Martingale Optimal Transport Viewpoint*, Chapman & Hall/CRC Financial Mathematics Series, 2017.

[19] Henry-Labordère, P.: *From (Martingale) Schrodinger Bridges to a New Class of Stochastic Volatility Models*, SSRN, 2019.

[20] Léonard, C.: *A survey of the Schrödinger problem and some of its connections with optimal transport*, Discrete Contin. Dyn. Syst., 34(4):1533–1574, 2014.

[21] Pennanen, T., Perkkiö, A.-P.: *Convex duality in nonlinear optimal transport*, Journal of Functional Analysis 277:1029–1060, 2019.

[22] Peyré, G., Cuturi, M.: *Computational Optimal Transport: With Applications to Data Science*, Foundations and Trends in Machine Learning, 11(5-6):355–607, 2019.

[23] Rüschendorf, L., Thomsen, W.: *Note on the Schrodinger equation and I-projections*, Statist. Probab. Lett. 17:369–375, 1993.

[24] Sinkhorn, R.: *Diagonal equivalence to matrices with prescribed row and column sums*, The American Mathematical Monthly, 74(4):402–405, 1967.

[25] Sion, M.: *On general minimax theorems*, Pacific Journal of Mathematics, 8(1):171–176, 1958.

[26] Villani, C.: *Topics in optimal transportation*, American Mathematical Society, Graduate Studies in Mathematics 58, 2003.

Dual Formulation of the Optimal Consumption Problem with Multiplicative Habit Formation*

Thijs Kamma†
Dept. of Quantitative Economics
Maastricht University
NETSPAR

Antoon Pelsser
Dept. of Quantitative Economics
Maastricht University & University of Amsterdam
NETSPAR

January 29, 2022

Abstract

This paper provides a dual formulation of the optimal consumption problem with internal multiplicative habit formation. In this problem, the agent derives utility from the ratio of consumption to the internal habit component. Due to this multiplicative specification of the habit model, the optimal consumption problem is not strictly concave and incorporates irremovable path-dependency. As a consequence, standard Lagrangian techniques fail to supply a candidate for the corresponding dual formulation. Using Fenchel's Duality Theorem, we manage to identify a candidate formulation and prove that it satisfies strong duality. On the basis of this strong duality result, we develop an evaluation mechanism to measure the accuracy of analytical or numerical approximations to the optimal solutions.

Keywords: Fenchel duality, habit formation, life-cycle investment, stochastic optimal control, utility maximisation

JEL Classification: C61, D15, D53, D81, G11

*We are very grateful to An Chen, Michel Vellekoop, Bas Werker and to the conference and seminar participants at the Netspar Pension Day (2021), Tilburg University (2021), Ulm University (2021), the International Pension Workshop (2022) and the Winter Seminar on Mathematical Finance (2022), for their helpful comments and suggestions.

†Corresponding author. Mailing address: t.kamma@maastrichtuniversity.nl. PO Box 616, 6200 MD Maastricht, The Netherlands. Phone: +31 (0) 43 388 3695.

Dynamic Mean-Variance Efficient Fractional Kelly Portfolios in a Stochastic Volatility Model

Xue Dong He*and Zhaoli Jiang†

Kelly portfolio, which was proposed by Kelly (1956), is a popular investment strategy in the market. It is the strategy that maximizes the logarithmic growth rate of wealth in the long run. Breiman et al. (1961), Thorp (1975), Hakansson (1975), Hakansson and Miller (1975), Algoet and Cover (1988), Bell and Cover (1988, 1980), Ethier (2004), and many other works show that Kelly strategy possesses some good properties. For example, in some specific settings or among a certain class of portfolio strategies, the Kelly portfolio asymptotically maximizes the rate of asset growth and the median of wealth. For other good properties possessed by the Kelly portfolio, see for instance Maclean et al. (2010).

Samuelson (1971, 1975) and Ziemba et al. (2016) show that it is possible for the Kelly portfolio to lose a lot of money. In addition, Kelly strategy can be too risky for many investors. For example, for a stock with annual mean excess return rate 4% and volatility 20%, the Kelly strategy is to invest 100% of wealth in the stock, while many investors believe that this investment strategy is too risky. To address these issues, MacLean et al. (1992) proposes fractional Kelly portfolios, which suggest to invest a certain percentage of wealth in Kelly portfolio and hold the remaining as cash or in a risk-free asset. As argued by MacLean et al. (1992), if one perceives Kelly strategy as one that maximizes the growth side of an investment and holding the risk-free asset as the strategy that maximizes the security side of an investment, then a fractional Kelly portfolio trades off growth and security of an investment. MacLean et al. (1992), MacLean and Ziemba (1999), Maclean et al. (2010), and Dohi et al. (1995) study fractional Kelly portfolios in various settings and find some nice properties for these portfolios.

Fractional Kelly portfolios, however, are constructed heuristically rather than as the con-

*Corresponding Author. Room 505, William M.W. Mong Engineering Building, Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong, Telphone: +852-39438336, Email: xdhe@se.cuhk.edu.hk.

†Risk Management Institute, National University of Singapore, Singapore, Email: rmijz@nus.edu.sg.

sequence of optimality. Thus, these portfolios do not necessarily optimize certain investment criteria. MacLean et al. (1992) and MacLean and Ziemba (1999) show that in general fractional Kelly portfolios do not optimize the growth or security of investment under various measures of growth and security. In particular, fractional Kelly portfolios are not necessarily mean-variance efficient.

In the present paper, we attempt to improve the mean-variance efficiency of fractional Kelly portfolios in a market with stochastic volatility. There is abundant evidence of stochastic volatility (Campbell and Hentschel, 1992; Pagan and Schwert, 1990), and stochastic volatility models have been extensively used in option pricing and portfolio selection. In the present paper, we consider a multi-stock version of the stochastic volatility model used in Liu (2001), Chacko and Viceira (2005), and Kraft (2005). We then propose a mean-variance portfolio selection problem in which an agent, at any time, wants to achieve as high expected return as a fractional Kelly portfolio and to minimize the variance of portfolio return. Because of the expected return target and the objective of minimizing return variance, the problem is time inconsistent in that the agent's plan that is optimal at certain time is not necessarily optimal at future time. Consequently, the plan that is optimal today cannot be consistently implemented by the agent in the future. To address this issue, we follow the literature on time-inconsistent problems to consider equilibrium portfolio strategies for the agent: In this notion, the agent consider her selves at different time to be different players in a game and an equilibrium in the game is sought. As a result, the agent is not willing to deviate from an equilibrium portfolio strategy at any time, so this strategy can be consistently implemented by the agent throughout the entire investment horizon.

Next, we derive an equilibrium portfolio strategy in closed form, up to solving a one-dimensional Ricatti equation. Under this strategy, the agent invests a wealth-independent percentage of her wealth in the stocks, and this percentage is proportional to a power transformation of the market state that represents the stochastic volatility in the market. Moreover, we can decompose the strategy into three parts. The first part is the same as the fractional Kelly strategy, which is also myopic in that it becomes optimal when the investment horizon is infinitesimally short. The second part is the hedging demand caused by the stochastic volatility. Because the hedging demand can possibly make the expected return target unachieved, certain adjustment must be made, leading to the third part of the equilibrium portfolio strategy.

The equilibrium portfolio strategy is not global optimal in that at each time, it does not necessarily minimize the portfolio return variance. Thus, it is not automatically implied that the equilibrium strategy leads to a smaller return variance than the fractional Kelly strategy.

We, however, are able to prove that the equilibrium portfolio improves the mean-variance efficiency of the fractional Kelly portfolio in that the former leads to a smaller return variance than the latter while achieving the same expected return at any time, and the improvement is strict as long as the equilibrium portfolio differs from the fractional Kelly portfolio.

Next, we conduct numerical studies by estimating the model parameters, computing the equilibrium portfolio, and comparing it with the fractional Kelly portfolio using the real data of the S&P500 Index, gold futures (ZGZ1), 30-year Treasury bond futures, and Nikkei225 Index. We consider investment in a risk-free asset and in two risky assets, where the first risky asset is chosen to be the S&P500 Index and the other is chosen to be one of gold futures (ZGZ1), 30-year Treasury bond futures, and Nikkei225 Index. We find that the improvement in the mean-variance efficiency, measured in the proportional reduction in return variance by the equilibrium strategy, is economically significantly for the investment in the S&P500 Index and ZGZ1, because these two risky assets are not highly correlated and have large market price of risk. We also find that the longer the investment horizon is, the more the equilibrium strategy improves on the mean-variance efficiency. We also compare the equilibrium portfolio in our model with the portfolio maximizing expected utility of wealth with a constant relative risk aversion degree (Liu, 2001) and with the portfolio in a mean-variance analysis of logarithmic return rates (Dai et al., 2021), and show that they are significantly different.

Finally, we infer a time-varying risk aversion degree from our model. We prove that the implied risk aversion degree is smaller and thus the agent is more risk seeking when the market condition becomes better. We also demonstrate numerically that the implied risk aversion degree is decreasing with respect to the investment horizon and vanishes when the horizon goes to infinity. This shows that the agent is insensitive to risk when she makes investment decisions to optimize the mean-variance criterion of her wealth in very distant future.

References

Algoet, P. H. and Cover, T. M. (1988). Asymptotic optimality and asymptotic equipartition properties of log-optimum investment, *Annals of Probability* **16**(2): 876–898.

Bell, R. and Cover, T. M. (1988). Game-theoretic optimal portfolios, *Management Science* **34**(6): 724–733.

Bell, R. M. and Cover, T. M. (1980). Competitive optimality of logarithmic investment, *Mathematics of Operations Research* **5**(2).

Breiman, L. et al. (1961). Optimal gambling systems for favorable games, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, The Regents of the University of California.

Campbell, J. Y. and Hentschel, L. (1992). No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of financial Economics* **31**(3): 281–318.

Chacko, G. and Viceira, L. M. (2005). Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, *The Review of Financial Studies* **18**(4): 1369–1402.

Dai, M., Jin, H., Kou, S. and Xu, Y. (2021). A dynamic mean-variance analysis for log returns, *Management Science* **67**(2): 1093–1108.

Dohi, T., Tanaka, H., Kaio, N. and Osaki, S. (1995). Alternative growth versus security in continuous dynamic trading, *European Journal of Operational Research* **84**(2): 265–278.

Ethier, S. N. (2004). The Kelly system maximizes median fortune, *Journal of Applied Probability* **41**(4): 1230–1236.

Hakansson, N. H. (1975). On optimal myopic portfolio policies, with and without serial correlation of yields, *Stochastic Optimization Models in Finance*, Elsevier, pp. 401–411.

Hakansson, N. H. and Miller, B. L. (1975). Compound-return mean-variance efficient portfolios never risk ruin, *Management Science* **22**(4): 391–400.

Kelly, J. L. (1956). A new interpretation of information rate, *Bell System Technical Journal* **35**: 917–926.

Kraft, H. (2005). Optimal portfolios and heston's stochastic volatility model: an explicit solution for power utility, *Quantitative Finance* **5**(3): 303–313.

Liu, J. (2001). Dynamic portfolio choice and risk aversion. SSRN287095.

Maclean, L. C., Thorp, E. O. and Ziemba, W. T. (2010). Long-term capital growth: the good and bad properties of the kelly and fractional kelly capital growth criteria, *Quantitative Finance* **10**(7): 681–687.

MacLean, L. C. and Ziemba, W. T. (1999). Growth versus security tradeoffs in dynamic investment analysis, *Annals of Operations Research* **85**: 193–225.

MacLean, L., Ziemba, W. T. and Blazenko, G. (1992). Growth versus security in dynamic investment analysis, *Management Science* **38**(11): 1562–1585.

Pagan, A. R. and Schwert, G. W. (1990). Alternative models for conditional stock volatility, *Journal of econometrics* **45**(1-2): 267–290.

Samuelson, P. A. (1971). The fallacy of maximizing the geometric mean in long sequences of investing or gambling, *Proceedings of the National Academy of Sciences* **68**(10): 2493–2496.

Samuelson, P. A. (1975). Lifetime portfolio selection by dynamic stochastic programming, *Stochastic Optimization Models in Finance*, Elsevier, pp. 517–524.

Thorp, E. O. (1975). Portfolio choice and the Kelly criterion, *Stochastic Optimization Models in Finance*, Elsevier, pp. 599–619.

Ziemba, W. T. et al. (2016). A response to Professor Paul A. Samuelson's objections to Kelly capital growth investing, *World Scientific Book Chapters* pp. 249–274.

Dynamic Mean-Variance Portfolio Selection under Factor Models

Xiangyu Cui, Duan Li, Yun Shi and Lanzhi Yang

Based on the financial literature and the real financial data, we propose a system of factor models for both the return and risk. Then, we consider the multi-period mean-variance portfolio selection problem under the proposed factor models. We derive the semi-analytical optimal portfolio policy, which also takes a linear form of the current wealth level. The corresponding coefficients of the optimal policy are related to a particular stochastic process, which reflects how the investor evaluates the investment opportunity in the future and is called future investment opportunity (FIO). We propose a numerical algorithm for deriving FIO and reveal the relationship between FIO and the variance-optimal signed martingale measure (VSMM) of the market. The empirical analysis based on U.S. market further shows that by incorporating the factor models of the return and risk, the optimal portfolio policy can achieve better out-of-sample performance than the ones without considering the dynamic factor models of return and risk.

Abstract

We study a dynamic mean–variance portfolio selection problem with return predictability and trading frictions from price impact. Applying mean-field type control theory, we provide a characterization of an equilibrium trading strategy for an investor facing stochastic investment opportunities. An explicit equilibrium strategy is derived in terms of the solution to a generalized matrix Riccati differential equation, and a sufficient condition is also provided to ensure the latter's well-posedness. Our solution indicates that the investor should trade gradually towards a target portfolio which accounts for return predictability, price impact and time-consistency. Moreover, an asymptotic analysis around small liquidity costs shows that the investor's target portfolio is an equilibrium portfolio without price impact in the first-order sense, and that her first-order approximated value function does not deteriorate significantly for sufficiently small liquidity costs. Finally, our numerical results demonstrate that the target portfolio is more conservative than an equilibrium portfolio without price impact

Special Sessions in memory of Prof. Mark H.A. Davis

E-backtesting risk measures

Ruodu Wang

Abstract: Expected Shortfall (ES) is the most important risk measure in finance and insurance. One of the most challenging tasks in risk modeling practice is to backtest ES forecasts provided by financial institutions, based only on daily realized portfolio losses without imposing specific models. Recently, the notion of e-values has gained attention as potential alternatives to p-values as measures of uncertainty, significance and evidence. We use e-values and e-processes to construct a model-free backtest of ES, which can be naturally generalized to many other risk measures and statistical quantities.

Effective Algorithms for Optimal Portfolio Deleveraging Problem with Cross Impact

Hezhi Luo[*], Yuanyuan Chen[†], Xianye Zhang[‡], Duan Li[§], Huixian Wu[¶]

February 24, 2022

Introduction

In the financial market, leverage has been widely used by investors to significantly increase their returns of an investment. Although high leverage may be beneficial in boom periods, it may cause serious cash flow problems in recessionary periods. During the recessionary periods, investors often need to unwind their portfolios to reduce leverage. In this paper, we investigate the optimal portfolio deleveraging (OPD) problem with permanent and temporary price impacts, where the objective is to maximize equity while meeting a prescribed debt/equity requirement.

Brown et al. (2010) first study the optimal deleveraging problem with both permanent and temporary price impacts of trading. With a convexity assumption, Brown et al. (2010) derive an analytical optimal trading strategy for the optimal deleveraging problem. Without this convexity assumption, Chen et al. (2014) study the optimal deleveraging problem considered in Brown et al. (2010) and propose an efficient Lagrangian method to solve this non-convex quadratic program. Furthermore, Chen et al. (2015) extend the Lagrangian method to the optimal deleveraging problem with non-linear temporary price impact. However, Chen et al. (2014, 2015) still make the assumption that both permanent and temporary price impact matrices are diagonal in their models. In other words, no cross impact among different assets is considered in Brown et al. (2010) or Chen et al. (2014, 2015).

However, the existence of cross impact among different assets has been both empirically documented and theoretically justified. (See e.g., Fleming et al. (1998), Kyle & Xiong (2001), Andrade et al. (2008), Pasquariello & Vega (2015) and Benzaquen et al. (2017)).

We thus investigate in this paper how to solve the optimal deleveraging problem with cross impact, by considering general permanent and temporary price impact matrices, which could be not only non-diagonal, but also asymmetric. The resulting problem is then a non-separable and non-convex program with a quadratic objective function and non-homogeneous quadratic and box constraints, which is known to be NP-hard.

[*]Department of Mathematics, Zhejiang Sci-Tech University, Hangzhou, Zhejiang 310018, China. Email: hzluo@zstu.edu.cn

[†]Corresponding author. Department of Finance and Insurance, Nanjing University, Nanjing, Jiangsu 210093, China. Email:yychen@nju.edu.cn

[‡]Department of Mathematics, Zhejiang Sci-Tech University, Hangzhou, Zhejiang 310018, China.

[§]School of Data Science, City University of Hong Kong, Hong Kong. Email: dli226@cityu.edu.hk

[¶]Department of Mathematics, Hangzhou Dianzi University, Hangzhou, Zhejiang 310018, China.

Methodology

We reformulate the non-convex quadratic program as a D.C. program via spectral decomposition and simultaneous diagonalization (cf. Newcomb (1961)), and then propose two efficient algorithms for it: the SCO algorithm and the SCOBB algorithm. We show that our SCO algorithm converges to a KKT point of the transformed problem, while the SCOBB algorithm can efficiently identify the global ϵ-optimal solution. We establish the global convergence of the SCOBB algorithm and estimate its complexity. Specifically, we show that the SCOBB algorithm has a worst-case complexity bound $\mathcal{O}\left(N \prod_{i=1}^{r} \left\lceil \frac{\sqrt{r+\rho_1 s}(z_u^i - z_l^i)}{2\sqrt{\epsilon}} \right\rceil \right)$, where N is the complexity to solve the relaxed subproblem (a convex QP).

Numerical experiments

We conduct numerical experiments using historical data from NASDAQ to illustrate the application of our algorithm. In particular, we use NASDAQ TotalView-ITCH data to estimate the temporary and permanent price impact matrices of representative stocks. As shown in our experiments, the impact matrices are not only non-diagonal, but also asymmetric, which is also found in the empirical results documented in Table V by Pasquariello & Vega (2015). This finding demonstrates the importance of considering general impact matrices in our paper.

Randomly generated instances are also considered. According to our numerical experiments, our SCOBB algorithm can effectively identify the global optimal solution to medium- and large-scale instances of the optimal deleveraging problem with limited number of negative eigenvalues of the matrices in the quadratic terms in our optimization problem. Meanwhile, although we cannot prove the global optimality, our SCO algorithm always provides the global optimal solution in our numerical experiments for all the instances within short computational time, and thus also serve the purpose as an efficient algorithm for the optimal deleveraging strategy.

References

Andrade, S. C., Chang, C., & Seasholes, M. S. (2008). Trading imbalances, predictable reversals, and cross-stock price pressure. *Journal of Financial Economics*, *88*(2), 406–423.

Benzaquen, M., Mastromatteo, I., Eisler, Z., & Bouchaud, J.-P. (2017). Dissecting cross-impact on stock markets: An empirical analysis. *Journal of Statistical Mechanics: Theory and Experiment*, *2017(2)*(2), 023406. doi: 10.1088/1742-5468/aa53f7

Brown, D. B., Carlin, B., & Lobo, M. S. (2010). Optimal portfolio liquidation with distress risk. *Management Science*, *56(11)*, 1997–2014.

Chen, J., Feng, L., & Peng, J. (2015). Optimal deleveraging with nonlinear temporary price impact. *European Journal of Operational Research*, *244*, 240–247.

Chen, J., Feng, L., Peng, J., & Ye, Y. (2014). Analytical results and efficient algorithm for optimal portfolio deleveraging with market impact. *Operation Research*, *62(1)*, 195–206.

Fleming, J., Kirby, C., & Ostdiek, B. (1998). Information and volatility linkages in the stock, bond, and money markets. *Journal of Financial Economics*, *49*(1), 111–137.

Kyle, A. S., & Xiong, W. (2001). Contagion as a wealth effect. *The Journal of Finance*, *56*(4), 1401–1440.

Newcomb, R. W. (1961). On the simultaneous diagonalization of two semi-definite matrices. *Quarterly of Applied Mathematics*, *19*(2), 144–146.

Pasquariello, P., & Vega, C. (2015). Strategic cross-trading in the u.s. stock market. *Review of Finance*, *19*(1), 229–282.

Efficient Allocations under Ambiguous Model Uncertainty

Chiaki Hara
Kyoto University

Sujoy Mukerji
Queen Mary University of London

Frank Riedel
University of Bielefeld

Jean Marc Tallon
Paris School of Economics

March 14, 2022

Abstract

We investigate consequences of ambiguity on ex ante efficient allocations in an exchange economy. The ambiguity we consider is embodied in the model uncertainty perceived by the decision maker: they are unsure what would be the appropriate probability measure to apply to evaluate consumption contingent on a state space Ω and keep in consideration a set $\mathcal{P}$ of alternative probabilistic laws p. We study the case where the typical consumer in the economy is ambiguity averse with smooth ambiguity preferences and $\mathcal{P}$ is point identified, i.e., the true law $p \in \mathcal{P}$ can be recovered empirically from events in Ω, a framework axiomatized recently by Denti and Pomatto (Econometrica, forthcoming). Differently from the literature, we allow for the case where the aggregate risk is ambiguous and agents are heterogeneously ambiguity averse. Our analysis addresses, in particular, the full range of set-ups where under expected utility the Pareto efficient consumption sharing rule is a linear function of the aggregate endowment. We identify systematic differences ambiguity aversion introduces to optimal sharing arrangements in these environments and also characterize the representative consumer. Furthermore, we investigate the implications for the state-price function, in particular, the effect of heterogeneity in ambiguity aversion.

Efficient exotic options pricing and hedging under AEPD model with biased control variate method

Yongzeng Lai (ylai@wlu.ca)
Department of Mathematics, Wilfrid Laurier University, Waterloo, Ontario, Canada, N2L 3C5

Abstract

Options and other financial derivatives are widely used in the financial and insurance industries. Empirical studies using worldwide real financial data show that the famous Black-Scholes-Merton's model for options is not that realistic. Various more realistic models are proposed to extend the Black-Scholes-Merton model. However, under these more realistic models for asset prices, no closed-form formulas for option prices exist usually. The Monte Carlo simulation method is used for pricing and hedging of complex exotic options under these models since these pricing and hedging problems are usually high dimensional. The main drawback of the Monte Carlo method is its slow convergence. The control variate method is one of the main variance reductions to accelerate the convergence of the Monte Carlo simulation method.

In this paper, we propose a new more realistic option pricing model, the asymptotic exponential power distribution, for the underlying asset prices. We then use the Monte Carlo combined with the biased control variate method to speed up the convergence. To further improve the convergence, we also combine the randomized quasi-Monte Carlo method with the biased control variate method. Our numerical tests demonstrate that the biased control variate method is very efficient in reducing variances while the option biases are very small: compared with the Monte Carlo method, the variance reduction ratios are up to a few thousand, and when combined with the randomized quasi-Monte Carlo method, the variance reduction ratios are up to tens of millions, in pricing exotic options, such as Asian style options.

Keywords: Asymmetric Exponential Power Distribution; Options Pricing; Monte Carlo and quasi-Monte Carlo Simulation; Variance Reduction Method; Control Variate Method.

Efficient Markets and Contingent Claims Valuation: An Information Theoretic Approach – part II

Jussi Lindgren, Aalto University, 02150 Espoo, Finland
jussi.lindgren@aalto.fi

The aim of the oral presentation is to show how the Efficient Market Hypothesis can be operationalized by considering the financial market as a complex system, which optimizes an information functional. The financial market solves a stochastic optimal control problem, in which the log-price of the market assets follow an Ito diffusion. The target functional is an entropy functional/information functional, for which we construct the Hamilton-Jacobi-Bellman equation (HJB). The HJB-equation can be linearized using the Hopf-Cole transformation to obtain the Black-Scholes partial differential equation. Therefore the Black-Scholes-Merton equation reflects the efficiency criteria of the financial market.

The HJB-equation is then used to derive an optimal transport equation for the market drift, which supports Black-Scholes pricing. The transport equation obtained is the backwards Burgers equation, which has been used in modelling turbulent fluid flow. Finally, it is shown how the market reaches a thermodynamic equilibrium. The key results support earlier findings by Hodges and Carverhill [1], in which it was shown, how the market risk premium obeys the Burgers equation.

We also find a Gibbs/Boltzmann distribution for the equilibrium market, from which we can deduce empirical predictions for the distribution of the prices of financial derivatives in the equilibrium. Finally, we simulate the solutions for the Burgers equation using Monte Carlo methods. This work follows the peer-reviewed published manuscript by the present author [2].

[1] Hodges, S.; Carverhill, A. Quasi Mean Reversion in an Efficient Stock Market: The Characterisation of Economic Equilibria which Support Black–Scholes Option Pricing. *Econ. J.* **1993**, *103*, 395–405.

[2] Lindgren, J. Efficient Markets and Contingent Claims Valuation: An Information Theoretic Approach. *Entropy* **2020**, *22*, 1283.

Evaluation of Asset Pricing Models: Optimal Risk Premia and Goodness-of-Fit Measures

Robert L. Kimmel*

June 30, 2020

Under Revision—This Version not for Distribution

Abstract

We consider linear factor models, in which expected returns are linear functions of beta coefficients on explanatory factors. Such models have been in common use in finance for decades, and many methods for assignment of risk premia to the factors and evaluation of the model's fit have appeared in the literature. We show that there is essentially a unique method for assigning risk premia, and a unique method for assessing the fit of the model, based on only two assumptions: the fit of the model is judged solely by its predictions of the assets' expected returns, and the fit improves when the prediction error for an asset decreases, holding the prediction error of all uncorrelated assets fixed. The unique (to within monotonic transformation and additional "tie-breaking" criteria) goodness-of-fit is based on the maximum Sharpe ratio that can be achieved using the factor mimicking portfolios, and the unique risk premia assigned to factors that are themselves excess returns are simply the expected excess returns of the factors.

*email: robert.l.kimmel@icloud.com.

Evaluation of Deep Learning Algorithms for Quadratic Hedging

Zhiwen Dai* Lingfei Li† Gongqiu Zhang‡

February 22, 2022

Abstract

We solve the quadratic hedging problem by deep learning in discrete time. We consider three deep learning algorithms corresponding to three architectures of neural network approximation: approximating controls of different periods by different feedforward neural networks (FNNs) as proposed by Han and E (2016), using a single FNN with decision time as an input to approximate controls of different periods, and using a recursive neural network (RNN) to utilize historical information. We evaluate these methods under the discrete-time Black-Scholes model and the DCC-GARCH model for hedging basket options on portfolios of up to 100 assets with time to maturity up to one year. We compare them in terms of their hedging error on the test data, the extent of overlearning, the learned hedging strategy, training speed and scalability. Our results favour the single FNN and RNN approximations overall while the multiple FNN approximation can fail for a large portfolio and a long maturity. We also evaluate the performance of the single FNN and RNN algorithms in a data-driven framework, where data is generated by resampling without assuming any parametric model.

*Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. Email: zwdai@se.cuhk.edu.hk.

†Corresponding author. Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR. Email: lfli@se.cuhk.edu.hk.

‡School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China. Email: zhanggongqiu@cuhk.edu.cn.

Execution in an aggregator with multiple traders and informed flow

Dieter Hendricks*[1,2], Michael Harvey[5], Daniel Pesch[1,4] and Stephen J. Roberts[1,3]

[1]Oxford-Man Institute of Quantitative Finance, University of Oxford
[2]Systematic Trading Strategies, Rand Merchant Bank (Global Markets)
[3]Machine Learning Research Group, Department of Engineering Science, University of Oxford
[4]Saïd Business School, University of Oxford
[5]School of Computer Science and Applied Mathematics, University of the Witwatersrand

*Corresponding author: dieter.hendricks@gmail.com

Abstract: In over-the-counter (OTC) markets, price discovery is obscured by *fragmented liquidity* and *information asymmetry* between liquidity providers (LPs) and traders. Traders use quote *aggregators* to consolidate indicative price quotes from a subset of LPs with which they have set up bilateral agreements for conditions of trade. A deal request is then typically sent through to the originating LP with the best indicative price. The LP further has an option to either honour the deal request at the indicative price, or reject the request if an adverse price move is expected (so-called *last look*). Different traders in this market typically source quotes from different subsets of LPs, with some degree of overlap. Individual LPs on the other hand do not typically observe the price quotes of their peers. Oomen (2017a,b) investigated the dynamics of this market in detail, providing a framework for interrogating LP gross revenue and market share, and trader effective spread and deal rejection rate. We develop a flexible Agent-Based Model (ABM) framework to interrogate the macro and micro behaviour of OTC markets with heterogeneous LP and trader types. The ABM framework allows us to recover the results in Oomen (2017a,b) for uninformed (noise) traders, while also permitting increased complexity in agent specification, allowing us to examine a wider range of scenarios to understand the nature of *adverse selection* in OTC liquidity provision. We investigate two extensions of Oomen's work: 1) We consider the impact of *informed trader flow* by allowing traders to infer an *expected true price* from observed quotes in their aggregator, using this as a basis for buying/selling decisions, and 2) We consider the impact of multiple traders with informed flow (coupled via overlapping LP quote visibility) on the LP's gross revenue and market share. We show how an LP may adjust their nominal spread and last look threshold to control revenue or market share in this environment. We also briefly address the implications of the contentious abolishment of last look practices for FX market making at Tier 1 banks, given the rise of non-bank electronic market makers and increased pressure for market-wide firm pricing.

Keywords: adverse selection, informed trading, liquidity provision, over-the-counter markets, agent-based modelling

References

Oomen, R. (2017a) Execution in an aggregator, *Quantitative Finance*, 17(3):383-404.
Oomen, R. (2017b) Last look, *Quantitative Finance*, 17(7):1057-1070.

Existence of an equilibrium with limited participation

Kim Weston

Rutgers University

March 15, 2022

A limited participation economy models the real-world phenomenon that some investors have access to more of the market than others. Basak and Cuoco (RFS'98) introduced a continuous-time, running consumption model of limited participation with two (classes of) investors: an unconstrained investor with access to a complete market, and a constrained investor who cannot trade in the stock market and faces incompleteness. Equilibrium existence results have so far been limited to considering logarithmic constrained investors, in part due to the complications that arise with an endogenously-determined stochastic interest rate.

This work provides an extension of Basak and Cuoco's model to the case of exponential investors. A coupled system of quadratic BSDEs describes the equilibrium. Deriving the characterizing BSDE system requires the presence of a traded annuity plus the Markov property afforded by the underlying Brownian structure. The equilibrium existence proof exploits the triangular-quadratic structure on the drivers and makes use of the BSDE existence result of Xing and Žitković (AP'18).

In addition to extending work on limited participation economies and proving the existence of a solution to a coupled, quadratic BSDE system, this work provides one of the few existence results of an incomplete financial equilibrium with a stochastic interest rate.

Exploratory Control with Tsallis Entropy for Latent Factor Models

Ryan Donnelly [1], Sebastian Jaimungal[2]

[1] King's College London

[2] University of Toronto

We study optimal control in models with latent factors where the agent controls the distribution over actions, rather than actions themselves, in both discrete and continuous time. To encourage exploration of the state space, we reward exploration with Tsallis Entropy and derive the optimal distribution over states – which we prove is q-Gaussian distributed with location characterized through the solution of an FBSΔE and FBSDE in discrete and continuous time, respectively. Finally, we apply the approach to a statistical arbitrage problem, where the asset price has latent drift that may be interpreted as a trading signal which the agent filters.

Extreme Value Theory for Mean-Field Interacting Particle Systems via Propagation of Chaos

Nikolaos Kolliopoulos, Martin Larsson and Zeyu Zhang
Department of Mathematical Sciences, Carnegie Mellon University

March 15, 2022

Abstract

We establish an asymptotic extreme value theory for the interacting particles of a finite system with correlation due to a mean-field term in the drifts of the SDEs that invovles the empirical measure of the system. Since the system is finite, the particles are not necessarily conditionally independent, rendering the techniques of standard extreme value theories inapplicable. We show that at any time t, the largest particle of our system behaves asymptotically exactly as that of the approximate system arising from the propagation of chaos theory when the number of particles is large, where the particles are independent and existing results can be implemented.

1 Extended Abstract

In this paper we consider a mean-field system of N particles, where the dynamics of the i-th particle $X_{\cdot}^{i,N}$ are given by the SDE:

$$\begin{aligned} X_t^{i,N} \quad = \quad & X_0^i + \int_0^t C\left(s, \left(X_r^{i,N}\right)_{0\leq r\leq s}\right) ds + \int_0^t A\left(s, \left(X_r^{i,N}\right)_{0\leq r\leq s}\right) \\ & \times \left[B\left(s, \left(X_r^{i,N}\right)_{0\leq r\leq s}, \int g_s\left(\left(X_r^{i,N}\right)_{0\leq r\leq s}, x\right) d\mu_s^{N,N}(x)\right) ds + dW_s^i\right] \end{aligned} \tag{1.1}$$

for $i \in \{1, 2, ..., N\}$, where W^1, W^2, ... are Brownian motions defined in some filtered probability space $\left(\Omega, \mathcal{F}, \left(\mathcal{F}_t\right)_{t\geq 0}, \mathbb{P}\right)$, adapted to the filtration $\left(\mathcal{F}_t\right)_{t\geq 0}$, and

$$\mu_s^{N,N} = \frac{1}{N}\sum_{j=1}^{N} \delta_{\left(X_r^{j,N}\right)_{0\leq r\leq s}} \tag{1.2}$$

is the empirical measure of the paths of the particles constituting our system. We are interested in the asymptotic behaviour of the distribution of $\max_{1\leq j\leq N} X_t^{j,N}$ as $N \longrightarrow +\infty$, for any positive time t. Our motivation is the rank-based model for a large portfolio of shares

$$X_t^{i,N} \quad = \quad X_0^i + \int_0^t B\left(\frac{1}{N}\sum_{j=1}^{N} 1_{\left\{X_s^{j,N} < X_s^{i,N}\right\}}\right) ds + \sigma W_t^i$$

1

(1.3)

which is a special case of the system (1.1), where we would like to study the asymptotic behaviour (for large N) of the top-performing company. To study this distribution of $\max_{1\leq j\leq N} X_t^{j,N}$ as $N \longrightarrow +\infty$ for large values of N, we consider a normalized maximum of the form

$$Y_t^N = \max_{1\leq j\leq N} \frac{X_t^{j,N} - b(t,N)}{a(t,N)} \tag{1.4}$$

and we want to check whether we can pick the normalizing deterministic numbers $b(t,N)$ and $a(t,N)$ (depending only on t and N) so that Y_t^N converges weakly to some non-trivial distribution as $N \longrightarrow +\infty$.

Considering now the corresponding McKean-Vlasov system of independent particles:

$$\begin{aligned} X_t^{i,\infty} &= X_0^i + \int_0^t C\left(s, \left(X_r^{i,\infty}\right)_{0\leq r\leq s}\right) ds + \int_0^t A\left(s, \left(X_r^{i,\infty}\right)_{0\leq r\leq s}\right) \\ &\quad \times \left[B\left(s, \left(X_r^{i,\infty}\right)_{0\leq r\leq s}, \int g_s\left(\left(X_r^{i,\infty}\right)_{0\leq r\leq s}, x\right) d\mu_s^{\infty,\infty}(x)\right) ds + dW_s^i\right] \end{aligned} \tag{1.5}$$

for $i \in \mathbb{N}$, where $\mu_s^{\infty,\infty}$ is the weak limit of the empirical measure

$$\mu_s^{N,\infty} = \frac{1}{N}\sum_{j=1}^N \delta_{\left(X_r^{j,\infty}\right)_{0\leq r\leq s}} \tag{1.6}$$

as $N \longrightarrow +\infty$, The particles $X_\cdot^{i,\infty}$ of the limiting system are pairwise independent and identically distributed, which means that the standard results of Extreme Value Theory are applicable and it is quite likely that we can choose $a(t,N)$ and $b(t,N)$ so that the corresponding sequence of normalized maxima for the system (1.5), i.e

$$Y_t^N = \max_{1\leq j\leq N} \frac{X_t^{j,\infty} - b(t,N)}{a(t,N)}, \tag{1.7}$$

converges weakly as $N \longrightarrow +\infty$ to some random variable Z with a non-trivial distribution. In this paper we show that when this happens, the sequence of normalized maxima for the system (1.1) converges also Z in distribution as $N \longrightarrow +\infty$, i.e we have that

$$Y_t^N = \max_{1\leq j\leq N} \frac{X_t^{j,N} - b(t,N)}{a(t,N)} \longrightarrow Z \tag{1.8}$$

in distribution as $N \longrightarrow +\infty$ where Z can be determined.

Feynman-Kac formula for BSDEs with jumps and time delayed generators associated to path-dependent nonlinear Kolmogorov equations

Luca Di Persio
Department of Computer Science,
University of Verona, Verona, Italy,
luca.dipersio@univr.it

Matteo Garbelli
Department of Mathematics,
University of Trento, Trento, Italy
matteo.garbelli@unitn.it

Adrian Zalinescu
Department of Faculty of Computer Science,
Alexandru Ioan Cuza University, Iasi, Romania,
adrian.zalinescu@info.uaic.ro

We study a system of forward-backward stochastic differential equations (FBSDEs) in the spirit of the theory introduced by Pardoux and Peng in [8, 9]. The forward process evolves according to a jump-diffusion dynamic as in the framework of [1]. We rely on the theory developed by Delong in, e.g., [5, 6] to study the backward component with a delayed generator and driven by a Lèvy-type noise. Furthermore, we provide a probabilistic representation of a (mild) solution of a path-dependent nonlinear Kolmogorov equation for a fixed time horizon. The concept of mild solutions to delay equations, introduced in [7], generalizes the standard notion of viscosity solutions. Thereafter, we establish a non-linear Feynman–Kac representation formula to associate the FBSDE solution to the solution coming from the path dependent nonlinear Kolmogorov equation, generalizing the results provided in [3, 4]. We also present two financial applications. The first deals with the Large Investor Problem focusing on a stock price with jump-diffusion dynamic, moving from the classical setting defined in [2]. The second describes an insurance problem with a payment dynamics consistent

with a step process and evaluated by means of a dynamic risk measure.

Keywords
BSDEs; Feynman-Kac formula; jump-diffusion; path-dependence.

References

[1] Barles, G., R. Buckdahn, E. Pardoux, Backward stochastic differential equations and integral-partial differential equations. Stochastics and Stochastics Reports 60, 1997.

[2] Cvitanic J. and Ma J., Hedging options for a large investor and forward-backward SDEs. The annals of applied probability, 6(2): 370-398, 1996.

[3] Cordoni, Francesco, Di Persio, Luca, Maticiuc, Lucian and Zălinescu, Adrian, A stochastic approach to path-dependent nonlinear Kolmogorov equations via BSDEs with time-delayed generators and applications to finance. Stochastic Processes and their Applications, 2020.

[4] Cordoni, F., Di Persio, L. Oliva, I. A nonlinear Kolmogorov equation for stochastic functional delay differential equations with jumps. Nonlinear Differ. Equ. Appl. 24, 16, 2017.

[5] Delong, Łukas, Backward Stochastic Differential Equations with Jumps and Their Actuarial and Financial Applications: BSDEs with Jumps. Springer London. EAA Series, 2013.

[6] Delong, Łukasz, Peter Imkeller, Backward stochastic differential equations with time delayed generators - results and counterexamples. Ann. Appl. Probab. 20 (4), 2010.

[7] Fuhrman, M., Tessitore, G. Generalized Directional Gradients, Backward Stochastic Differential Equations and Mild Solutions of Semilinear Parabolic Equations. Appl Math Optim 51, 2005.

[8] Pardoux, Etienne, Shige Peng, Backward SDEs and quasilinear PDEs, in Stochastic partial differential equations and their applications, Springer, 1992.

[9] Peng, Shige, Backward Stochastic Differential Equation, Nonlinear Expectation and Their Applications. Proceedings of the International Congress of Mathematicians Hyderabad, India, 2010.

Special Sessions in memory of Prof. Peter Carr

Financial Interpretation of Feller’ s Factorization

Claudio Tebaldi

Abstract: The infinitessimal generator of a time-homogeneous univariate diffusion process is a secondorder linear ordinary differential operator (SOLODO) with no loading on the the identity operator. Feller[1] famously factorized this generator into successive differentiations with respect to scale and speed measure. Later, Feller[2] also factored an extended generator which is a SOLODO that loads on the identity operator in a particular way. We provide a novel financial interpretation of this latter factorization.

Fire Sales, Default Cascades and Complex Financial Networks
On Limit Theorems

Hamed Amini* Zhongyuan Cao † Agnès Sulem ‡

We present a general tractable framework for understanding the tradeoffs between the fire sales and default cascades in complex financial networks. Our work integrates the fire sales loss into the cascades of insolvencies in interbank networks, see e.g. [3, 5]. This extends our previous central limit theorems in [1] for the pure default cascade process in heterogeneous financial networks. Our limit theorems quantify how price mediated contagion across institutions with common asset holding could worsen cascades of insolvencies in heterogeneous financial networks.

We consider a general tractable model for the fire sales endowed by a quick-reaction default contagion in a heterogeneous financial network, subject to an exogenous macroeconomic shock. By reducing the dimension of the problem by classifying the financial institutions according to different types, in an appropriate type space and under some regularity assumptions, we show the limit theorems of the total final sold shares during the default cascade. In particular, under suitable assumptions on the degree and threshold distributions, we show that the final sold shares has asymptotically Gaussian fluctuations.

We next state the limit theorems for the market prices of assets after shock given by the so called inverse demand function. Following [2, 4], we assume that all the liquidation happens simultaneously and instantaneously. Since the network is quick reacted, there exist some equilibrium for the asset value after the fire sales. Formally, set a interval $[p_{\min}, p_0]$ for the illiquid asset price, we define the *equilibrium price* of the illiquid asset as

$$p_n^\star = \sup\{p \in [p_{\min}, p_0] : p \leqslant g(\Gamma_n(p)/n)\},$$

where g is the inverse demand function and $\Gamma_n(p)$ denotes the final sold shares under price p for a size n network.

We then investigate the value in equilibrium and further give some limit theorems regarding

*Georgia State University, Department of Risk Management and Insurance, Atlanta, GA 30303, USA, email: hamini@gsu.edu.

†INRIA Paris, 2 rue Simone Iff, CS 42112, 75589 Paris Cedex 12, France, and Université Paris-Dauphine, email: zhongyuan.cao@inria.fr.

‡INRIA Paris, 2 rue Simone Iff, CS 42112, 75589 Paris Cedex 12, France, and Université Paris-Est, email: agnes.sulem@inria.fr.

the price in equilibrium. Last, we also extend our framework to the multi-type assets case and the limit theorems for the total sold shares and the price in equilibrium of assets also follows.

References

[1] Hamed Amini, Zhongyuan Cao, and Agnes Sulem. Limit theorems for default contagion and systemic risk. *Available at SSRN 3811107*, 2021.

[2] Hamed Amini, Damir Filipović, and Andreea Minca. Uniqueness of equilibrium in a payment system with liquidation costs. *Operations Research Letters*, 44(1):1–5, 2016.

[3] Maxim Bichuch and Zachary Feinstein. Optimization of fire sales and borrowing in systemic risk. *SIAM Journal on Financial Mathematics*, 10(1):68–88, 2019.

[4] Rodrigo Cifuentes, Gianluigi Ferrucci, and Hyun Song Shin. Liquidity risk and contagion. *Journal of the European Economic association*, 3(2-3):556–566, 2005.

[5] Nils Detering, Thilo Meyer-Brandis, Konstantinos Panagiotou, and Daniel Ritter. An integrated model for fire sales and default contagion. *Mathematics and Financial Economics*, 15(1):59–101, 2021.

Special Sessions in memory of Prof. Mark H.A. Davis

Forward performance criteria and incomplete information

Thaleia Zariphopoulou

Abstract: In this talk, I will introduce forward performance processes in Ito-diffusion markets under incomplete information. They give rise to new filtering problems that are ill-posed but, on the other hand, capture real time changes in information flow. Examples for locally riskless forward processes will be presented.

BFS Congress, June 2020, Hong Kong

GARCH minimum variance portfolio

Linda Chamakh
Ecole Polytechnique

Supervisor(s): Prof. Emmanuel Gobet (Ecole Polytechnique), Prof. Zoltán Szabó (Ecole Polytechnique) and Dr. Jean-Philippe Lemor (BNP Paribas)

Email: linda.chamakh@polytechnique.edu

Abstract:

In this work, we propose a solution for the minimum variance portfolio, e.g. the allocation which will minimize the future *portfolio realized variance* over the investment span, when the asset returns have a conditional distribution characterized by a Markovian and ergodic *instantaneous covariance*. We place ourselves in a discrete time setting (daily frequency).

Usually, the investor seeking for the minimum variance portfolio will search for the allocation minimizing the associated historical portfolio variance. Doing so, she takes the empirical historical covariance matrix as a proxy for the unknown *realized covariance* (e.g. the empirical covariance associated to future returns over the holding period). In our framework, we can explicitly derive the optimal realized covariance estimator at the time of investment, which is simply the conditional expectation of the realized covariance. This estimator coincides with the conditional expectation of the sum of instantaneous covariances over the holding period. Under specific model, such as the linear multivariate GARCH model, explicit formula can be obtained for this estimator.

Using tools from martingale and ergodic processes theories, we provide asymptotic as well as finite sample guarantees of this optimality. We compare our conditional realized variance-based minimum variance portfolio to the long-term average covariance matrix-based portfolio, but our results are valid for any other benchmark, and also if the allocation is subject to static constraints. Our main result is a minoration of the probability that our portfolio outperforms the benchmark, as a function of the number of rebalancing of the portfolio.

- Assuming that the instantaneous covariance process is Markovian and satisfies specific characteristics, applying Fort-Moulines [2, prop. 2], we have a control on the centered moments of sum of functions of the process over a fixed-time trajectory, polynomial in the number of samples to the power half the moment order. Using this control, we can show that the difference between the finite sample conditional variance of the benchmark portfolio and the proposed portfolio remains positive when the number of rebalancing goes to infinity, and that it remains close to this limit with high probability at fixed number of rebalancing.

- Exploiting the fact that that the difference between realized and conditional variance defines a martingale increment, we use standard Burkhölder type of inequality [4, th. 2.10] to get moment control on the realized versus conditional variance of our portfolios.

Each moment control is polynomial in the time horizon (number of rebalancing times time holding period) with a power proportional to the maximal moment of the covariance if the variance is finite, and linear in the time horizon if it is infinite. The bound on the probability that the finite cumulated variance of our portfolio is smaller than the benchmark is obtained via standard Markov inequalities.

We elaborate our results with the specific model of the GARCH Constant Conditional Correlation (GARCH-CCC) [1, 3]. This model satisfies the required Markovian and ergodicity assumptions under explicit conditions on its parameters values. In this model, the covariance process only has few finite moments, which motivated our restriction to moment control -rather than exponential one for example. Numerical illustrations are given under this model, with GARCH-CCC parameters fitted on real financial returns from various asset classes. The conditional GARCH minimum variance portfolio outperforms the stationary covariance-based minimum variance portfolio with a probability of about 70% for one rebalancing at a monthly frequency. It goes up to 90% after three rebalancing, and 100% after a year. Illustration of the bound tightness is given by comparing the theoretical polynomial rate with the distribution tails exponent via Hill estimation.

Keywords: minimum variance portfolio, conditional covariance, GARCH, concentration inequalities.

References

[1] Tim Bollerslev et al. Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *Review of Economics and statistics*, 72(3):498–505, 1990.

[2] Gersende Fort, Eric Moulines, et al. Convergence of the Monte Carlo expectation maximization for curved exponential families. *The Annals of Statistics*, 31(4):1220–1259, 2003.

[3] Christian Francq and Jean-Michel Zakoian. *GARCH models: structure, statistical inference and financial applications*. John Wiley & Sons, 2019.

[4] Peter Hall and Christopher C Heyde. *Martingale limit theory and its application*. Academic press, 2014.

GARCH option pricing based on convolutional time series neural network with ensemble empirical mode decomposition

Zheng Fang *, Jianying Xie †

December 25, 2021

1 Extended Abstract

The traditional Black–Scholes model has challenges in measuring the real volatility in the market price of the underlying assets. Option pricing by GARCH addresses these challenges by taking into account implied volatility [1]. However, there are noises from different frequencies in the time series, which cannot be measured by the GARCH model, and these noises will affect the volatility measurement. The ensemble empirical mode decomposition (EEMD) in depicting the fluctuation features of financial time series, and Hua [2] proposed EEMD-GARCH model, but it used past data to obtain the model parameters, which is not indicative of future results. In this paper, we are motivated by those issues, we adapted EEMD in GARCH model. In addition, we built our neural network model in extracting the information from EEMD's results shown in Figure 1. To address challenges of scale insensitive in the neural networks, we leveraged the autoregressive component in training our model, and we adapted the convolutional neural network in capturing the dependencies between multiple IMFs and capturing long and short-term dependencies by Recurrent Neural Network (RNN). We also used locally risk-neutral valuation relationship (LRNVR) and Monte Carlo simulations, we embedded the forecasting results of EEMD from the neural network into the simulation to obtain the martingale measure for the price of the underlying assets. Other GARCH option pricing models preset the unit risk premium (λ), but our model optimized λ with all model parameters by Nelder-Mead algorithm. We conduct the empirical analysis in S&P500 Option, and found that our hybrid GARCH model showed better performance than the traditional Black-Scholes model, GARCH model, and hybrid EEMD-GARCH

*Department of Data Science and Artificial Intelligence, Monash University, Clayton, Victoria 3800, Australia. Corresponding author. Email: zfan51@student.monash.edu

†Department of Econometrics and Business Statistics, Monash University, Melbourne, 3800, Australia.

model in pricing the European option. We will provide more detailed methodology in the full-length paper.

Keywords – `EEMD, neural network, time series, GARCH, option pricing`

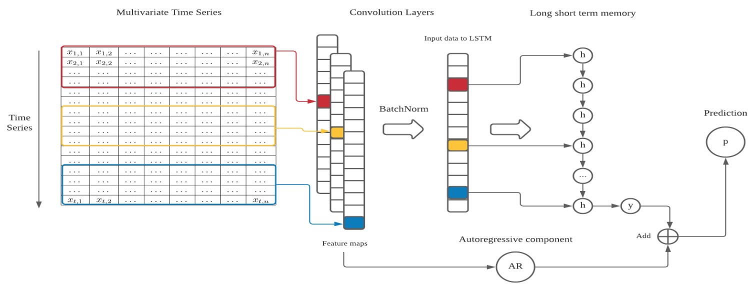

Figure 1: Overview of neural network model

2 Results

In this section, we demonstrated the result in brief for our model, the underlying asset is S&P500 used to price the European call option and compared with the hybrid EEMD GARCH option pricing model (HG) and GARCH option pricing (GARCH), and Black–Scholes model. Because of the limited page requirement in extended abstract, we will provide more detailed results in the full-length paper, including the biases in different maturity and strike price ratio (S/K).

(T-t)	S/K	HGN	HGN-BSM-bias	HG	GARCH
60	0.8	0.4026	-0.1202	0.6169	0.3707
	0.9	20.201	-0.0431	22.171	19.951
	1.0	163.63	-0.0123	162.19	164.33
	1.1	462.22	-0.0227	452.03	464.02
	1.2	782.59	0.0333	769.74	784.82
90	0.8	2.7617	-0.0183	3.41	2.5255
	0.9	43.203	-0.0528	44.78	42.407
	1.0	209.18	-0.0214	204.68	208.83
	1.1	497.22	0.0139	482.75	497.99
	1.2	806.62	0.0102	788.15	808.32

Table 1: S&P500 European call option pricing

References

[1] Jin-Chuan Duan. The garch option pricing model. *Mathematical finance*, 5(1):13–32, 1995.

[2] Qiuling Hua, Tingfeng Jiang, and Zhang Cheng. Option pricing based on hybrid garch-type models with improved ensemble empirical mode decomposition. *Quantitative Finance*, 18(9):1501–1515, 2018.

Special Sessions in memory of Prof. Peter Carr

General properties of backtestable statistics

Carlo Acerbi

Abstract: We propose a formal definition of backtestable statistic: a backtest is a null expected value involving only the statistic and its random variable, strictly monotonic in the former. We discuss the relationship with elicitability and identifiability which turn out being necessary conditions. The variance and the Expected Shortfall are not backtestable for this reason. We discuss (absolute) model validation in the context of one– or two–sided hypothesis tests, as well as (relative) model selection obtained by ranking realizations of the backtest statistic. We introduce the concept of sharpness which refers to whether a backtest is strictly monotonic with respect to the real value of the statistic and not only to its prediction. This decides whether the expected value of a backtest determines the extent of a prediction discrepancy and not only its likelihood. We show that the quantile backtest is not sharp and in fact provides no information whatsoever on the real value of the statistic. The Expectile is also not sharp; we provide bounds for its real value, which are looser for outer confidence levels. We then introduce ridge backtests, applicable to particular non–backtestable statistics, such as the variance and the Expected Shortfall, which coincide with the attained minimum of the scoring function of another elicitable auxiliary statistic. This allows to produce sharp backtest procedures in which the prediction for the auxiliary variable is also involved but with small sensitivity and known bias sign. The ridge mechanism explains why the variance has always been de–facto backtestable and allows for similar efficient ways to backtest the expected shortfall. We discuss the relevance of this result in the current debate of financial regulation (banking and insurance), where Value at Risk and Expected Shortfall are adopted as regulatory risk measures.

Generalisation of Fractional Cox-Ingersoll-Ross Process

Marc M. Mpanda[1], Safari Mukeru*, Mmboniseni Mulaudzi**

Department of Decision Sciences
University of South Africa, P. O. Box 392, Pretoria, 0003. South Africa
mpandmm@unisa.ac.za, *mukers@unisa.ac.za, **Mulaump@unisa.ac.za

Abstract

The fractional Cox-Ingersoll-Ross process of the form $X_t = Z_t^2$, where the process $(Z_t)_{t\geq 0}$ is given by $dZ_t = \left((k - \theta Z_t^2)Z_t^{-1}dt + \sigma dW_t^H\right)/2$ with $(W_t^H)_{t\geq 0}$ is a fractional Brownian motion with Hurst parameter $H \in (0, 1)$ (which is an extension of the classical Cox-Ingersoll-Ross process) was recently studied by Mishura and Yurchenko-Tytarenko [Fractional Cox-Ingerson-Ross process with non-zero "mean", *Modern Stochastics: Theory and Applications* 5(1), (2018) 99-111]. They proved that $(X_t)_{t\geq 0}$ satisfies the equation $dX_t = (k - \theta X_t)dt + \sigma\sqrt{X_t} \circ dW_t^H$ where $\circ$ refers to the Stratonovich integral. Moreover, for $H > 1/2$, $(X_t)_{t\geq 0}$ never hits zero and for $H < 1/2$, the probability of hitting zero tends to 0 provided the drift coefficient k increases to $+\infty$. In this paper, we extend these results to the general process defined by $dZ_t = \left(f(t, Z_t)Z_t^{-1}dt + \sigma dW_t^H\right)/2$ where $f(t, z)$ is a continuous function on $\mathbb{R}_+^2$. In the case where $H < 1/2$, we consider a sequence of increasing functions (f_n) and we prove that the probability of hitting zero tends to 0 as $n \to \infty$. We also analyse the generalised Cox-Ingersoll-Ross process for $H < 1/2$ defined as a square of a pointwise limit of the stochastic process $(Z_t^\epsilon)_{t\geq 0,\, \epsilon>0}$ that satisfies the stochastic differential equation $dZ_t^\epsilon = \frac{f(t,Z_t^\epsilon)}{2}\left(Z_t^\epsilon \mathbf{1}_{\{Z_t^\epsilon>0\}} + \epsilon\right)^{-1}dt + \frac{\sigma}{2}dW_t^H,\quad Z_0^\epsilon = Z_0 > 0$ as $\epsilon \to 0$. These results are illustrated with some simulations using the generalisation of the extended Cox-Ingersoll-Ross process.

Keywords: Fractional Brownian motion, Fractional Cox-Ingersoll-Ross process, Hitting times, Stratonovich integral.

[1]Corresponding author, mpandmm@unisa.ac.za

Special Sessions in memory of Prof. Tomas Bjork

Geometry of term structures

Speaker: Josef Teichmann

Abstract: Tomas Bjork was a leading researcher in introducing fine geometric reasonings to the analysis of term structure equations. We review some of his main contributions including some personal thoughts about his admired didactical approaches and we show some recent developments in the field.

Hedging Derivatives with Recalibration and Model Risk

Mark DAVIS*, Seiya GOTO† and Koichi MATSUMOTO‡

Abstract

The mathematical model plays an important role to hedge derivatives. The model usually includes some implied parameters that are determined based on the market prices. The implied parameters are assumed to be constant but keep to be revised periodically in the market, which is called recalibration. Since the recalibration is absolutely essential in the real market, we should adjust the optimal hedging strategy derived from the model, according to the movement of the implied parameters. In this study, we assume that the asset price and the implied parameter movements are represented by the mathematical model which is unknown to the investor. We study an optimal hedging strategy and the hedging error under the possible worst situation.

We consider a multi-branch model in a discrete time framework. Let $N \in \boldsymbol{N}$ be the total number of times and there are finite possible branches at each time. For $1 \leq n \leq N$, we denote the outcome of branch at time n by ω_n and the set of all outcomes at time n by Ω_n. We define $\Omega_{1,n} = \Omega_1 \times \cdots \times \Omega_n$ for $1 \leq n \leq N$ and we set $\Omega = \Omega_1 \times \cdots \times \Omega_N$. For $\omega = \omega_1\omega_2 \cdots \omega_N \in \Omega$, we use a shorter notation as $\omega_{1,n} = \omega_1\omega_2 \cdots \omega_n$ for $1 \leq n \leq N$ and we set $\omega_{1,0} = \Omega$. We assume $\mathcal{F}_0$ is trivial. We denote by $\mathcal{F}_n$ the collection of all subsets of $\Omega_{1,n}$ and we set $\mathcal{F} = \mathcal{F}_N$. A branch point of multi-branch model is called a node. Since each node corresponds to $\omega_{1,n} \in \Omega_{1,n}$ for $0 \leq n \leq N-1$, we identify a node and $\omega_{1,n}$ and we denote the set of all nodes by

$$\mathcal{N} = \{\omega_{1,n} \,|\omega \in \Omega,\ 0 \leq n \leq N-1\}.$$

We consider the market where there is a risky asset and the saving account. For the sake of simplicity, we assume the risk-free rate is zero. For $0 \leq n \leq N$, we suppose that $\mathcal{F}_n$ represents the market information. We assume that the implied parameters are given by the recalibration based on the mathematical model directly or indirectly with the market information. We denote the price of the risky asset and the implied parameter by (S_n, I_n). Suppose S_n is a positive $\mathcal{F}_n$-measurable function and I_n is a real-valued vector or matrix $\mathcal{F}_n$-measurable function. We denote by H_n the price of derivative at time n. Suppose that there exists a function h satisfying

$$H_n = h(n\triangle t, S_n, I_n), \quad 0 \leq n \leq N.$$

Here $\triangle t$ is the time interval. The pricing formula h is calculated with assumption that the implied parameter is constant rather than stochastic. The movement of implied parameter is a problem beyond of the scope of the pricing formula.

We consider an investor who hedges a derivative using the risky asset and the saving account. Let π_n be units of risky asset at time n for $1 \leq n \leq N$. Suppose that π_n is a

*Faculty of Natural Sciences, Department of Mathematics, Imperial College London
†Department of Economic Engineering, Faculty of Economics, Kyushu University
‡Department of Economic Engineering, Faculty of Economics, Kyushu University

real-valued function of the past information and we call $\pi = (\pi_1, \pi_2, \ldots, \pi_N)$ a strategy. We define the set of all strategies by

$$\mathcal{A} = \{\pi | \ \pi_n \text{ is a real-valued } \mathcal{F}_{n-1}\text{-measurable function for all } 1 \leq n \leq N \ \}.$$

Let H_0 be the initial value of the hedging portfolio. We suppose the balance between the derivative and the hedging portfolio is restored at each time, by using the saving account to make up for the hedging error. The hedging error at time n is given by

$$\triangle H_n - \pi_n \triangle S_n$$

where $\triangle H_n = H_n - H_{n-1}$ and $\triangle S_n = S_n - S_{n-1}$. We sum up the hedging errors of the whole period by the quadratic criterion. The total hedging error is given by

$$\sum_{n=1}^{N} \left(\triangle H_n - \pi_n \triangle S_n\right)^2 .$$

Let $\mathcal{P}$ be the set of all probability measures on $(\Omega, \mathcal{F})$. For $Q \in \mathcal{P}$ and $1 \leq n \leq N-1$, the conditional probability $Q[\cdot|\mathcal{F}_n]$ is denoted shortly by $Q_n[\cdot]$ and we define $Q_0[\cdot] = Q[\cdot]$. We identify a model and a probability measure. We denote by $\mathcal{Q}$ the set of all candidates for the true model. For $M \in \boldsymbol{N}$, we denote the first candidates for the true model by $P^1, P^2, \ldots, P^M \in \mathcal{Q}$. We call these probability measures elementary models. To simplify the notation, we assume that $P^m[\bar{\omega}_{1,n}] > 0$ for all nodes $\bar{\omega}_{1,n} \in \mathcal{N}$ and $1 \leq m \leq M$. We call $Q \in \mathcal{P}$ a base model if for all nodes $\bar{\omega}_{1,n} \in \mathcal{N}$, there exists $1 \leq m \leq M$ such that

$$Q_n[\omega_{n+1}](\bar{\omega}_{1,n}) = P_n^m[\omega_{n+1}](\bar{\omega}_{1,n}), \quad \omega_{n+1} \in \Omega_{n+1}.$$

A base model can be interpreted as a combination of the elementary models. We denote by M_b the total number of base models. We arrange all base models appropriately and number them as B^1, B^2, $\ldots$, B^{M_b}. Suppose the true model can be represented by a weighted average of base models. In other words, we assume $\mathcal{Q}$ satisfies

$$\mathcal{Q} = \left\{ \sum_{m=1}^{M_b} w^m B^m \middle| w = (w^1, w^2, \ldots, w^{M_b}), \quad w^1, w^2, \ldots, w^{M_b} \geq 0, \quad \sum_{m=1}^{M_b} w^m = 1 \right\}.$$

Our problem is a robust hedging problem defined by

$$\min_{\pi \in \mathcal{A}} \max_{Q \in \mathcal{Q}} E^Q [\sum_{n=1}^{N} \left(\triangle H_n - \pi_n \triangle S_n\right)^2]$$

where $E^Q[\cdot]$ is the expectation under Q. An optimal solution is called an optimal strategy and the optimal value is called the minimum hedging error.

We show that an optimal solution exists uniquely and give its concrete expression. Further we show how to calculate the optimal strategy and the minimum hedging error numerically. Finally we illustrate the effect of the model risk by the numerical analysis.

Hedging with Automatic Liquidation and Leverage Selection on Bitcoin Futures

Carol Alexander, Jun Deng, and **Bin Zou** (University of Connecticut)

Abstract

Bitcoin derivatives positions are maintained with a self-selected margin, often too low to avoid liquidation by the exchange, without notice, during periods of excessive volatility. Recently, the size and scale of such liquidations precipitated extreme discontent among traders and numerous lawsuits against exchanges. Clearly, hedgers of bitcoin should account for the possibility of automatic liquidation. That is the mathematical and operational problem that we address, deriving a semi-closed form for an optimal hedging strategy with dual objectives – to minimize both the variance of the hedged portfolio and the probability of liquidation due to insufficient collateral. An empirical analysis based on minute-level data compares the performance of major direct and inverse bitcoin hedging instruments traded on five major exchanges. The products have markedly different speculative trading scores according to new metrics introduced here. Instruments having similar hedging effectiveness can exhibit marked differences in speculative activity. Inverse perpetuals offer greater effectiveness than direct perpetuals, which also exhibit more speculation. We model hedgers with different levels of loss aversion that select their own level of leverage and collateral in the margin account. By following the optimal strategy, the hedger can reduce the liquidation probability to less than 1% and control leverage to a reasonable level, mostly below 5X.

Hermite Expansion for Multivariate Diffusions and Its Equivalence with Other Expansions

Xiangwei Wan* Nian Yang†

March 14, 2022

Abstract

In the first part of this work, we show that a small-time Hermite expansion is feasible for multivariate diffusions. By introducing *an innovative quasi-Lamperti transform*, which unitizes the diffusion matrix at the initial time, we derive explicit recursive formulas for the expansion coefficients of transition densities and European option prices for multivariate diffusions with jumps in return. These immediately available explicit formulas, particularly for the irreducbile, nonaffine, time-inhomogeneous model with different types of jump-size distribution, is new to the literature. The explicit formulas can lead to real-time derivatives pricing and hedging as well as model calibration. In the second part, we show that the derived Hermite expansion is a bridge to unify some existing methods including the expansions of Li (2013) and Yang et al. (2019). Extensive numerical experiments illustrate the accuracy and effectiveness of our approach.

KEYWORD: Hermite expansion, Irreducible diffusions, Transition densities, European option prices, Stochastic volatility models, Jumps

JEL Classification: C13, C32, G13, C63

*Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai, China. Email: xwwan@sjtu.edu.cn

†Department of Finance and Insurance, School of Business, Nanjing University, Nanjing, Jiangsu, China. Email: yangnian@nju.edu.cn

High-Frequency Liquidity in the Chinese Stock Market:
Measurements, Patterns, and Determinants

Ruixun Zhang[123], Chaoyi Zhao[1]*, Yufan Chen[1], Lintong Wu[1], Yuehao Dai[1], Ermo Chen[1], Zhiwei Yao[1], Yihao Zhou[1], and Lan Wu[123]

[1]Department of Financial Mathematics, School of Mathematical Sciences, Peking University
[2]Laboratory for Mathematical Economics and Quantitative Finance, Peking University
[3]Center for Statistical Science, Peking University

March 12th, 2022

Keywords: Liquidity; Limit order book; Price impact; Transaction costs; Bid-ask spread; Chinese stock market.

JEL Classification: C55, C58, G12, G14, G20

*zhaochaoyi@pku.edu.cn

Extended Abstract

The role of liquidity in empirical finance has grown rapidly over the past decade influencing studies in microstructure, asset pricing, market efficiency, and corporate finance. While it has long been acknowledged that liquidity matters for asset prices and risk management, a debate remains about its measurements and determinants, and the time-series and cross-sectional properties of various methods to measure liquidity (Chaieb, Errunza, and Langlois, 2021). This literature faces particularly great challenges when studying non-US markets as global intraday data is relatively expensive and growing exponentially over time (Fong, Holden, and Trzcinka, 2017).

Among the international markets, the Chinese financial market is of particular importance not only because of its growing size, but also its increasing centrality in the global financial system (Billio et al., 2022). Despite its importance, the Chinese stock market has several unique features that affect the determination and measurement of liquidity, including its high percentage of retail investors, short sales constraints, and relatively low degree of automation compared to developed markets. In addition, each stock is traded only at the exchange where it is listed, and there are no dark pools in the Chinese stock market. These factors together make research on liquidity and market microstructure in China both interesting and challenging.

In this article, we comprehensively explore, for the Chinese stock market, a broad range of high-frequency liquidity measures widely used in the literature, including measures based on trading costs (such as bid-ask spread and weighted spread), measures based on execution efficiency (such as order interarrival time, order lifetime, and order execution ratio), and measures based on price impact (such as the Kyle's lambda of Kyle (1985) and the price-impact regression of Glosten and Harris (1988)). We document their liquidity levels, conduct extensive empirical analysis for their distributional properties as well as intraday and cross-sectional patterns, and provide a model that explains the formation mechanism of liquidity based on order flows.

Our dataset is from the *SZSE Historical Tick Data*, which consists of tick-level data for all 2,081 stocks traded on the Shenzhen Stock Exchange from January 1st to December 31st, 2020, with a total of 16.7 billion events including all trades, quotes, cancellations, their timestamp and price information. Through complicated and meticulous integration of the raw event-level data, we reconstruct the sequences of all orders, trades, and detailed market snapshot at the timestamp of each event, including the price and quantities at each level of the limit order book. To the best of our knowledge, this is by far the most comprehensive study of liquidity for the Chinese stock market in the literature.

First, through survival analysis, we find that order interarrival times are better modeled by Weibull, instead of exponential, distributions, and order lifetimes follow the powerlaw, empirically confirming the intuition of Abergel et al. (2016). This implies that order interarrival time does not follow exponential distribution, hence Poisson flow is not a good model for order flow in the Chinese A-share market.

Second, we comprehensively explore the intraday patterns of liquidity measures, which follow a W-shape pattern, instead of the well-known U-shape pattern, because the Chinese stock market is closed for one and a half hours around noon. More importantly, we document novel and subtle intraday periodicities in liquidity at much higher granularity, by estimating spectral density functions for de-trended intraday time series of several liquidity measures. For example, we observe significantly larger spectral densities for the spread at whole-minute frequencies such as 1-minute, 5-minute and 10-minute than other frequencies. This effect is analogous to the price clustering effect documented in stock prices,[1] but along the time-series dimension rather than the price dimension.

Third, we document several cross-sectional patterns of liquidity measures across different stocks. Taking Kyle's lambda as an example, liquidity measures tend to be higher for companies with higher market capitalization, higher stock prices, that are more heavily held by institutional investors, and that are index constituents rather than under the special treatment (ST) status.[2] These empirical patterns highlight differences in liquidity in relation to the characteristics of the underlying asset.

Finally, to understand what determines liquidity, we develop an order-based model of the bid-ask spread. We define the *aggressive-passive imbalance* (API)—similar in spirit to the order flow imbalance of Cont, Kukanov, and Stoikov (2014)—which measures the imbalance between aggressive orders and passive orders. We show that, over short time intervals, changes in bid-ask spread are mainly driven by API. In addition, using a discrete model of bid-ask spread change, we estimate the impact of API on the direction and magnitude of spread change separately, and provide cross-sectional analysis of the degree of API's impact with respect to a few company characteristics. The significance of API as a determinant of bid-ask spread is robust across different stocks and over our entire sample period. This result highlights a universal mechanism of spread formation with respect to order flows.

Overall, our results shed light on the high-frequency liquidity measurements and patterns in the Chinese stock market, how they are affected by time-series and cross-sectional factors, and what mechanism determines their formation with respect to order flows.

[1] See, for example, the whole-dollar effect (Niederhoffer, 1965; Harris, 1991) and the stock-pinning effect on option strikes (Ni, Pearson, and Poteshman, 2005; Avellaneda, Kasyan, and Lipkin, 2012).

[2] The Shanghai and Shenzhen Stock Exchanges implement a rule that gives special treatment (ST) or delisting risk warning to certain companies with abnormal financial and nonfinancial conditions.

References

Abergel, F., M. Anane, A. Chakraborti, A. Jedidi, and I. M. Toke, 2016, *Limit Order Books* (Cambridge University Press).

Avellaneda, M., G. Kasyan, and M. D. Lipkin, 2012, Mathematical models for stock pinning near option expiration dates, *Communications on Pure and Applied Mathematics* 65, 949–974.

Billio, M., A. W. Lo, L. Pelizzon, M. Getmansky Sherman, and A. Zareei, 2022, Global realignment in financial market dynamics, SAFE Working Paper.

Chaieb, I., V. Errunza, and H. Langlois, 2021, How is liquidity priced in global markets?, *The Review of Financial Studies* 34, 4216–4268.

Cont, R., A. Kukanov, and S. Stoikov, 2014, The price impact of order book events, *Journal of Financial Econometrics* 12, 47–88.

Fong, K. Y., C. W. Holden, and C. A. Trzcinka, 2017, What are the best liquidity proxies for global research?, *Review of Finance* 21, 1355–1401.

Glosten, L. R., and L. E. Harris, 1988, Estimating the components of the bid/ask spread, *Journal of Financial Economics* 21, 123–142.

Harris, L., 1991, Stock price clustering and discreteness, *The Review of Financial Studies* 4, 389–415.

Kyle, A. S., 1985, Continuous Auctions and Insider Trading, *Econometrica* 53, 1315.

Ni, S. X., N. D. Pearson, and A. M. Poteshman, 2005, Stock price clustering on option expiration dates, *Journal of Financial Economics* 78, 49–87.

Niederhoffer, V., 1965, Clustering of stock prices, *Operations Research* 13, 258–265.

Implications of Electricity Price Regimes on Hydropower Plant Valuation

Shilei Niu and Tony S. Wirjanto

March 8, 2022

ABSTRACT

In a Markov regime switching model we investigate how the introduction of two spike regimes and the associated multiple jump sizes for electricity prices affect the hydro power plant's value with and without ramping rate restrictions. Specifically, we propose the jump size elasticity of value to measure the value of price jumps to the hydro operator. Our numerical experiments show that the plant value is positively affected by the jump size, but, to some extent, the impact depends on the restrictiveness of ramping constraints. We also provide further evidence that the hydro plant faces a larger negative impact from ramping restrictions when the expected price variation is increased due to an increase in the jump size. This makes it desirable to adjust the water release rate more substantially and frequently. The results have implications for designing ramping and competition policies in the electricity industry.

Keywords: hydro power plant; ramping rate; stochastic control; regime switching; Hamilton Jacobi Bellman Partial Differential Equation.

Shilei Niu
Xi'an Jiaotong-Liverpool University
International Business School Suzhou
111 Ren'ai Road
Suzhou, Jiangsu Province, 215123, PR China
shilei.niu@xjtlu.edu.cn

Tony S. Wirjanto
Department of Statistics & Actuarial Science, and School of Accounting & Finance
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1
twirjanto@uwaterloo.ca

IN MEMORIAM: GENERALIZATIONS OF CARR-MADAN FORMULA FOR OPTION DECOMPOSITION (BFS 2020 PROPOSAL)

SÉBASTIEN BOSSU*, PETER CARR†, AND ANDREW PAPANICOLAOU‡

Extended abstract

The Carr and Madan (1998) spanning formula is possibly the most famous result championed by Peter Carr, with far-reaching impact on the financial industry—most visibly the VIX—as well as academic research. This talk proposes to honor Peter's legacy by expounding the formula in connection with related results, and presenting two recent generalizations obtained under his broad guidance that leverage on the theoretical framework of functional analysis. In particular, a multi-asset version was derived.

Specifically, the Carr and Madan formula decomposes the payoff $f(S)$ of an arbitray European option as a combination of cash, underlying asset S, and a continuum of vanilla call and put options as

$$f(S) = \underbrace{f(S_0)}_{\text{cash}} + f'(S_0) \underbrace{(S - S_0)}_{\substack{\text{forward contr.} \\ \text{w/strike } S_0}} + \int_0^{S_0} f''(K) \underbrace{(K - S)^+}_{\substack{\text{put payoff} \\ \text{w/strike } K}} \mathrm{d}K + \int_{S_0}^{\infty} f''(K) \underbrace{(S - K)^+}_{\substack{\text{call payoff} \\ \text{w/strike } K}} \mathrm{d}K$$

where S is the underlying asset terminal price, S_0 is an arbitrary split level (typically the spot or forward price), and $t^+ := \max(0, t)$ is the positive part of the number t. By the law of one price, the non-vanilla European option price f_0 is then uniquely determined by vanilla option prices[1] for any strike in a *model-free* fashion. Splendidly enough, the log-contract $f(S) := \ln S$ introduced by Neuberger (1990) can thus be priced and hedged with vanilla options—a key result used by Peter and others to devise an effective trading strategy for the replication of *variance swaps*, and adopted by CBOE in 2002 for the calculation of the VIX[2].

The first generalization (Bossu, Carr, and Papanicolaou, 2021) uses functional analysis and integral equation theory to explore similar decompositions with "replicant" options other than vanillas, such as straddles and butterflies:

$$f(S) = c + q\, S + \int_0^{\infty} G(S, K)\phi(K)\, \mathrm{d}K,$$

where $G(S, K)$ is the known kernel of replicant option payoffs (e.g. $G(S, K) := |S - K|$ for replication with straddles) and $c, q, \phi(K)$ are unknown quantities to be determined. Formally, under mild

Date: March 15, 2022.

* WPI Business School & NYU Courant. Corresponding author and presenter.

† Died March 1st, 2022. Former chair of Finance and Risk Engineering, NYU Tandon. ‡ North Carolina State University. Andrew Papanicolaou is partially supported by NSF grant DMS-1907518.

[1] By put-call parity, the price of the forward contract is implied from vanilla option prices

[2] *The CBOE volatility index—VIX* 2009

1

conditions, $\phi(K)$ can be found by inverting the integral operator $\mathcal{G}$ induced by the replicant kernel $G(S,K)$; that is, $\phi = \mathcal{G}^{-1}f$. We show how this inversion may be done for the straddle kernel by spectral decomposition, with interesting truncation properties for fast pricing of vanilla options.

The second generalization (Bossu, Carr, and Papanicolaou, 2022, forthcoming, and Bossu, 2021, under review) explores the case of multi-asset European options with payoff $f(x_1, \cdots, x_n)$ to be decomposed with vanilla basket calls with payoff $(w_1x_1 + \cdots + w_nx_n - k)^+$ where $\mathbf{x}$ is a vector of underlying asset performances or returns, $\mathbf{w}$ is a vector of basket weights, and k is a fixed basket "strike" level. For two assets we want to find basket call quantities $\phi(w_1, w_2)$ such that

$$f(x_1, x_2) = \cdots + \iint (w_1x_1 + x_2x_2 - k)^+ \phi(w_1, w_2)\, \mathrm{d}w_1\, \mathrm{d}w_2,$$

and more generally for n assets. We show how this multidimensional integral equation can be solved for two classes of multi-asset options of interest: "standard" dispersion options with a payoff $f(|\mathbf{x}|)$ that is based on the Euclidean norm of the vector of asset performances, and the broader class of multi-asset European options with absolutely homogeneous payoff $f(\mathbf{x}; k)$. For the latter class, the solution is formally given as

$$\phi = \mathcal{R}^{-1}\left[\frac{\partial^2 f}{\partial k^2}\right]$$

where $\mathcal{R}^{-1}$ is the inverse Radon transform operator, an integral transform related to the Fourier transform which is used in medical imaging and other engineering applications. In particular, we derive first-time explicit decompositions for the dispersion call and put options, and for the two-asset best-of and worst-of call and put options. A novel generalization of the Breeden and Litzenberger (1978) formula for the multivariate implied distribution is also obtained.

References

Bossu, Sébastien (2021, under review). *Static replication of European multi-asset options with homogeneous payoff.* Working paper. NYU Courant.

Bossu, Sébastien, Peter Carr, and Andrew Papanicolaou (2021). "A functional analysis approach to the static replication of European options". In: *Quantitative Finance* 21 (4), pp. 637–655.

— (2022, forthcoming). "Static replication of European standard dispersion options". In: *Quantitative Finance.*

Breeden, Douglas T. and Robert H. Litzenberger (1978). "Prices of state-contingent claims implicit in option prices". In: *Journal of business*, pp. 621–651.

Carr, Peter and Dilip Madan (1998). "Towards a theory of volatility trading". In: *Volatility: New estimation techniques for pricing derivatives.* Ed. by Robert A. Jarrow. Vol. 29. Risk books, pp. 417–427.

Neuberger, Anthony (1990). *Volatility trading.* Working paper. London Business School.

The CBOE volatility index—VIX (2009). Tech. rep. Chicago Board Options Exchange.

INCONSPICUOUSNESS AND OPTIMAL TRANSPORT IN KYLE'S MODEL

Kerry Back
Jones Graduate School of Business and School of Social Sciences, Rice University

Reda Chhaibi
Institut de Mathématiques de Toulouse, Université Paul Sabatier -Toulouse III

Ibrahim Ekren
Department of Mathematics, Florida State University

Eunjung Noh
Department of Mathematics, Florida State University

We solve a generalized Kyle model type problem using Monge-Kantorovich duality and backward stochastic partial differential equations (BSPDEs). That is, we consider the model of asset pricing with asymmetric information, where a risk-neutral informed trader and noise traders submit order quantities to a risk-neutral market maker, who set the price to clear the market. This type of models was first studied by Kyle [1] and extended to the continuous time by Back [2], and has been widely studied in market microstructure.

First, we show that the Kyle problem can be recast into a terminal optimization problem with distributional constraints. Therefore, the theory of optimal transport between spaces of unequal dimension comes as a natural tool. Second, we analyze the structure of the Monge-Kantorovich duality, in particular, the pricing rule is established using the Kantorovich potentials. Finally, we completely characterize the optimal strategies for the informed trader by analyzing the filtering problem from the market maker's point of view. In this context, the Kushner-Zakai filtering SPDE yields an interesting BSPDE whose measure-valued terminal condition comes from the optimal coupling of measures. One main feature of our construction is to find a strategy for the informed trader which is distinct and more tractable than the classical bridge based construction.

References

[1] Albert S Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335, 1985

[2] Kerry Back. Insider Trading in Continuous Time. *The Review of Financial Studies* 5, no. 3, pages 387-–409, 1992

Invited Session: stochastic modeling of information in economics and finance

Title: Insiders and their free lunches: the role of short positions

Speaker: Delia Coculescu

Abstract: Given a stock price process, we analyse the potential of arbitrage in a context of short-selling prohibitions. We introduce the notion of minimal supermartingale measure, and we analyse its properties in connection to the minimal martingale measure. This question is more specifically analysed in the case of an investor having additional, inside information. In particular, we establish conditions when minimal martingale and supermartingale measures both fail to exist. These correspond to the case when the insider information includes some non-null events that are perceived as having null probabilities by the uninformed market investors, even as they cannot observe them. The results may have different applications, such as in problems related to the local risk-minimisation for insiders whenever strategies are implemented without short selling. Based on joint work with Aditi Dandapani.

Interpretable Counterfactual Inference via Causal Graphs

Yiyan HUANG [*] Qi WU [†] Dongdong WANG [‡] Zhixiang HUANG [§]

January 2020

Abstract

The financial arms of global e-commerce platforms typically take a non-parametric approach to access consumer credit risk. This allows machine learning tools to shine when it comes to forecasting future states, e.g., learning the map between people's digital footprints and their default outcomes. For example, the digital footprints typically consist of shopping & spending records, and the default outcomes are measures of whether they will be late on payments. However, when it comes to decision-making, i.e., adjusting people's credit lines or setting interest rates charged on unpaid balances, it is observed that changes in credit policy alter people's repayment behavior directly. These intervention effects are systematic at all levels of income, age, education, and occupation. If the decision to be made, based on states forecasted from the historical map, is going to change this map in the future, one cannot extrapolate from past outcomes to potential outcomes because the altered map due to intervention is unknown. Assuming the same map leads to erroneous decisions.

It is thus critical for credit providers to know a prior how much change there would be in a consumer's default probability had her credit line been doubled or cut in half, for example. In other words, credit risk is a function of both people's digital footprints and the platform's credit policy. The challenge, however, is that the outcomes in the historical data are only associated with past credit lines, not hypothetical credit lines one is interested in. The best online approach to resolve this problem is to conduct randomized experiments, which is, however, costly for credit business. When (historical) observational data is the only choice, one resorts to methods that are effective in removing biases in the treatment assignment process, e.g., those based on balancing scores such as the propensity score. But these methods are not reliable to remove the selection biases in unobservable variables. Further, existing models are often restricted to settings of binary treatment. For our credit problem, the treatments are credit lines and interest rates, and they take continuous values.

With these considerations, we take a causal graph approach to estimate counterfactuals in this paper. In our case, the counterfactual is the risk vs. policy function, where the risk axis is the probability that one is late on a payment, and the policy axis is the amount of uncollateralized credit line issued to her or the interest rate charged on her unpaid balance. A probabilistic graphical model is particularly useful to encode assumptions about the data-generating process if the following three properties are desirable: a) interpretable causal relations between variables; b) no need to label the potential outcomes, and c) allowing for continuous-valued treatments. In our model, the linearly connected multiple layers representing shallow causes to deep causes, allows the model to be interpretable and explicit, and the graph structure makes it easy to train with only observational factual label, instead of constructing counterfactual outcomes.

We compared our model with baselines and other state-of-art methods, e.g., k-Nearest Neighbor (kNN), Partial Linear Regression (PLR), Logistic Regression (LR), Random Forest (RF), Bayesian Additive Regression Trees (BART), Support Vector Machine (SVM) and Treatment-Agnostic Representation Network (TARNet) on semi-synthetic data by evaluating the Average Treatment Effect (ATE) and the Precision in Estimation of Heterogeneous Effect (PEHE). We also did experiments on e-commerce data to predict if one will default or not through the above methods evaluated by Area Under Curve (AUC) since the counterfactual label is unable to access, and the shape of risk vs. policy learned from counterfactual inference is significantly different from that learned by direct method, which ignores the causal relationship between treatments and covariates. Both semi-synthetic and real-world experiments show that our method performs better than many existing state-of-art models, especially when the task involves multiple treatments and inexplicable high-dimensional covariates.

Keywords: Credit Risk, Treatment Effect, Causal Inference, Graph Models, Machine Learning

[*] School of Data Science, City University of Hong Kong.
[†] School of Data Science, City University of Hong Kong. Please direct comments to qiwu55@cityu.edu.hk
[‡] Joint Lab in FinTech & Engineering, JD Digits.
[§] Joint Lab in FinTech & Engineering, JD Digits.

Irreversible reinsurance: a singular control approach

Abstract

The traditional approach for assessing reinsurance demands uses classical stochastic control framework by regarding reinsurance level as a control variable. However, reinsurance contracts are not traded, making the dynamic adjustment of the reinsurance impractical especially for reinsurance reduction. In addition, reinsurance contracts are often a long-term commitment between insurer and reinsurer. We propose a novel irreversible reinsurance framework, where the insurer enters into a long-term reinsurance contract at the best appropriate time and the contract is irreversible afterwards. When the contract is entered into, the insurer specifies the amount of reinsurance. We model the insurer's risk exposure by a mean-reverting process as inspired by the classic mortality modelling, and formulate the reinsurance decision into a two-dimensional degenerated singular control problem. The boundedness of insurer's risk retention makes our setup closely related to the finite-fuel problem. The optimal reinsurance purchase rule is triggered by a free-boundary that characterizes the optimal relationship between the risk exposure and the reinsurance cover. We conduct a numerical study to investigate the features of the reinsurance strategy and its dependency on the model parameters.

Jump-Diffusion Risk-Sensitive Benchmarked Asset Management with Traditional and Alternative Data*

Mark Davis† Sébastien Lleo‡§

December 27, 2021

This paper proposes a new method to estimate mean returns in continuous-time asset allocation models. The standard approach postulates that stochastic factors explain expected asset returns (Merton, 1971; Brennan et al., 1997). The problem is then to estimate these factors. When the factors are not observable, investors must estimate them from observed asset prices via filtering (Detemple, 1986; Gennotte, 1986; Brennan, 1998; Brennan & Xia, 2001; Nagai, 2001; Nagai & Peng, 2002). Recent advances have also combined asset prices with expert opinions to improve the estimates (see Frey et al., 2012; Fouque et al., 2017; Sass et al., 2020; Davis & Lleo, 2013, 2021, among others). However, these methods have limitations: stocks prices favor momentum strategies, and expert opinions require careful debiasing (Davis & Lleo, 2020). Therefore, we propose a jump-diffusion model in which investors estimate the factors from both traditional and alternative data.

We estimate the unobservable diffusion factors using a Kalman filter. Then we solve the benchmarked risk-sensitive stochastic control problem with the filtered estimates as an input. We prove that the value function is the unique $C^{1,2}$ solution to the Hamilton-Jacobi-Bellman Partial Differential Equation (HJB PDE). Our derivation extends the solution technique in Davis & Lleo (2011) to a benchmarked criterion. In the process, we reduce the requirements on the existence of zero-beta policies. Most importantly, we derive the optimal investment policy in quasi-closed form as the solution to a fixed point problem.

Thus, our jump-diffusion model has practical asset management implications. We find that investors construct their portfolios from a passive core and an active satellite. The passive core adds considerations for jump risk to a simple benchmark replication. The active satellite blends security selection and factor tilts with event-driven strategies unique to jump-diffusion problems. Thus, our model explains the most popular investment strategies. Furthermore, the improved expert forecast model and the introduction of alternative data provide factor tilters with new tools to sharpen their asset allocation.

*This article is dedicated to the memory of the late Mark H. A. Davis.

†Department of Mathematics, Imperial College London, London SW7 2AZ, England

‡Finance Department, NEOMA Business School, 59 rue Pierre Taittinger, 51100 Reims, France; Email: sebastien.lleo@NEOMA-bs.fr. ORCID: 0000-0002-0732-2833

§Part of this work was produced while the second author served as 2020 Bruti-Liberati Fellow at the Quantitative Finance Research Center of the University of Technology Sydney. The second author is greatly indebted to Nicholas Westray and William Ziemba for their constructive comments and helpful suggestions.

References

Brennan, M. (1998). The role of learning in dynamic portfolio decisions. *European Finance Review*, 1, 295–306.

Brennan, M. J., Schwartz, E. S., & Lagnado, R. (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control*, 21(8), 1377–1403.

Brennan, M. J. & Xia, Y. (2001). Stock price volatility and equity premium. *Journal of Monetary Economics*, 47(2), 249–283.

Davis, M. & Lleo, S. (2011). Jump-diffusion risk-sensitive asset management I: Diffusion factor model. *SIAM Journal on Financial Mathematics*, 2, 22–54.

Davis, M. & Lleo, S. (2013). Black-Litterman in continuous time: The case for filtering. *Quantitative Finance Letters*, 1(1), 30–35.

Davis, M. & Lleo, S. (2020). Debiased expert forecasts in continuous-time asset allocation. *Journal of Banking & Finance*, 113, 105759.

Davis, M. & Lleo, S. (2021). Risk-sensitive benchmarked asset management with expert forecasts. *Mathematical Finance*, 31(4), 1162–1189.

Detemple, J. (1986). Asset pricing in an economy with incomplete information. *Journal of Finance*, 41(2), 383–391.

Fouque, J., Papanicolaou, A., & Sircar, R. (2017). Perturbation analysis for investment portfolios under partial information with expert opinions. *SIAM Journal on Control and Optimization*, 55(3), 1534–1566.

Frey, R., Gabih, A., & Wunderlich, R. (2012). Portfolio optimization under partial information with expert opinions. *International Journal of Theoretical and Applied Finance*, 15(1), 1250009.

Gennotte, G. (1986). Optimal portfolio choice under incomplete information. *Journal of Finance*, 41(4), 733–746.

Merton, R. (1971). Optimal consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3, 373–413.

Nagai, H. (2001). Risk-sensitive dynamic asset management with partial information. In T. Hida, R. Karandikar, H. Kunita, B. Rajput, & S. Watanabe (Eds.), *Stochastics in Finite and Infinite Dimensions, in Honor of Gopinath Kallianpur* (pp. 321–340). Boston: Birkhauser.

Nagai, H. & Peng, S. (2002). Risk-sensitive dynamic portfolio optimization with partial information on infinite time horizon. *The Annals of Applied Probability*, 12(1), 173–195.

Sass, J., Westphal, D., & Wunderlich, R. (2020). Diffusion approximations for expert opinions in a financial market with gaussian drift. *Journal of Applied Probability*, 58(1), 197–216.

Title: Kyle-Back models with risk aversion and non-Gaussian beliefs

Abstract:

In this talk, we show that the existence of equilibrium in the Kyle-Back models can be characterized by considering a system of forward Fokker-Planck equation and a system of backward quasilinear parabolic partial differential equations coupled via an optimal transport type constraint at maturity. The methodology can be extended to a multidimensional setup as well where there are multiple stocks in the market. We also study the properties of the equilibrium obtained for small enough risk aversion parameter. In our model, the insider has exponential type utility and the belief of the market maker on the distribution of the price at final time can be non-Gaussian. This is based on joint works with Ibrahim Ekren.

Title
Kyle-Back Model with Stochastic Volatility

Abstract
We establish the existence of equilibrium in the Kyle-Back model of insider trading with stochastic liquidity of Collin-Dufresne and Fos (2016). In our model, the distribution of the fundamental price is allowed to be non-Gaussian and the increments of the volatility and the liquidity are allowed to be correlated. We characterize the equilibrium via the solution of a BSDE, a heat equation, and an optimal transport map. We obtain both the volatility dynamics and the IV curve as an output of the model. We study the impact of the distribution of the fundamental price on the dynamics of the IV curve.

Learning about latent dynamic trading demand

The price formation process in financial markets involves equating supply and demand for securities over time for arriving investors with heterogeneous trading preferences. In present day markets, large investors act on their underlying trading preferences, sometimes called *parent demands*, by splitting their trading into dynamic sequences of smaller orders, called *child orders* (see O'Hara (2015)), to minimize their price impact. Since the parent demands driving child-order trading are private information, investors use information from arriving child orders to form inferences over time about the dynamically evolving fundamental state of the market. In particular, investors learn about imbalances in the underlying aggregate parent demands and the associated pressure on future market-clearing prices and incorporate this information in their current child orders. Given the widespread prevalence of optimized order-splitting of parent orders into flows of child orders, dynamic learning about aggregate parent demands is a critical part of market dynamics.[1]

This paper is the first to provide an analytically tractable equilibrium model of dynamic learning, trading, and pricing with parent trading demands. We consider a continuous-time model with high-frequency trading at times $t \in [0, 1]$ over short time-horizons with $[0, 1]$ being a day or an hour. Trading occurs between price-sensitive optimizing traders with two different types of parent trading targets: One group has fixed individual targets, and the other group wants to track a stochastically evolving target over time. Since parent targets are initially not public, information about parent demand imbalances is partially revealed through market-clearing stock prices. Our analysis models the equilibrium dynamic learning process, stock holdings, and stock-price processes.

Our main results are:

- We construct and solve two different equilibrium models: A simpler price-friction equilibrium and a subgame perfect Nash financial-market equilibrium. In the price-friction equilibrium, price impact is due to an exogenous trading friction, but in the subgame Nash equilibrium, price impact includes both exogenous frictions and an endogenous price impact due to market clearing with constrained market asset-holding capacity. We find that these two equilibria are numerically similar.
- Intraday price drifts due to price pressure change over the trading day and are path-dependent. This leads to time-varying incentives for investors to provide liquidity to the child orders of other investors.
- A practical application of our model is that we can compute total trading costs for investors given the effects of dynamic learning and optimal trading by other investors. We show these costs are quadratic in the rebalancers' trading targets.
- Trading in our model reflects a combination of liquidity provision and speculation but not predatory trading. We conjecture that the absence of predatory trading is because our model

[1] See van Kervel and Menkveld (2019), Korajczyk and Murphy (2019), and van Kervel, Kwan, and Westerholm (2020).

replaces the exogenous price-elastic residual supply used in both Brunnermeier and Pedersen (2005) and Carlin, Lobo, and Viswanathan (2007) with endogenous demands coming from rational profit-maximizing investors.

Our paper advances several strands of research on market microstructure. First, dynamic learning and trading have been extensively studied in the context of markets with strategic investors with long-lived asymmetric information as in Kyle (1985). However, equilibrium trading, learning, and pricing with optimal dynamic order-splitting by large uninformed investors are less understood. Thus, we model price pressure to equate supply and demand rather than adverse selection. Second, Grossman and Miller (1988) model pricing and liquidity provision with impatient traders who submit single orders equal to their parent demands and with symmetric payoff information. In contrast, we model liquidity provision with optimal order-splitting of parent demands into child order flows. Third, Choi, Larsen, and Seppi (2019) construct an equilibrium with optimal dynamic trading and learning in a market with a strategic rebalancer with an end-of-day trading target and an informed investor who trades on private long-lived asset-payoff information. By filtering the order flow over time, the rebalancer learns about the underlying asset payoff, the informed investor learns about the rebalancer's trading target, and market makers learn about both when setting prices. That earlier paper provides a characterization result for equilibrium and gives numerical examples but does not have an existence proof or analytic solutions. In contrast, our model is solved analytically and gives the equilibrium in closed form. Fourth, Brunnermeier and Pedersen (2005) and Carlin, Lobo, and Viswanathan (2007) show how dynamic rebalancing by a large investor can lead to predatory trading. However, these papers abstract from the learning problem by assuming the parent trading needs are publicly observable. They also make an ad hoc assumption about the price sensitivity of a residual market-maker trading demand due to exogenous price-elastic noise traders. In contrast, our model assumes the underlying parent trading demands are private information, which leads to a learning problem. In addition, our prices are rationally set with no ad hoc residual demand. Fifth, a large body of research models optimal order-splitting strategies for a single strategic investor given an exogenous pricing rule with no learning about latent trading demands of other investors (see, e.g., Almgren and Chriss (1999, 2000), Almgren (2003), and Schied and Schöneborn (2009)). In contrast, we solve for optimal trades, learning, and pricing jointly. Van Kervel, Kwan, and Westerholm (2020) solve for optimal trading strategies for two dynamic rebalancers with learning over time about each other's latent trading demands. This leads to predictions about the effect of aggregate parent demand on individual investor child orders, which are then verified empirically. However, they assume an ad hoc linear pricing rule, and there are no existence proofs or analytic solutions. In contrast, price pressure in our Nash model is partly endogenously determined in equilibrium, and we solve our model analytically. As in van Kervel, Kwan, and Westerholm (2020), trading in our model is a combination of speculation on expected future price changes and trading-demand accommodation.

Leverage Dynamics under Managerial Discretion

Tak-Yuen Wong*

December 14, 2019

Abstract

This paper studies leverage dynamics when managers cannot commit to future financing and default policies ex-ante. Managers derive private benefits of control but debt constrains their flexibility, and they build up less excessive leverage and may even actively reduce debt over time. Thus, governance frictions weaken the leverage ratchet effect and increase the funding advantage of debt. This is in contrast with the common wisdom that manager-shareholder conflicts reduce efficiency. The paper provides three main results. First, the model delivers a novel theory of optimal maturity: shorter debt maturity disciplines the manager's leverage ratcheting incentives but increases default risks. Second, the model predicts a rich leverage dynamics. Firms with weak governance and low managerial ownership favor long-term debt, have high debt capacity, maintain low target leverage, and adjust the debt level faster towards the target. Finally, firms with high agency costs remain persistently unlevered in the no-commitment equilibrium. This result offers a potential resolution for the zero leverage puzzle.

Keywords: leverage ratchet effect, commitment, managerial agency, governance, debt maturity, capital structure, zero leverage puzzle

JEL Classification: G3, G32, G34

*School of Finance, Shanghai University of Finance and Economics; Shanghai Institute of International Finance and Economics. Email: etywong110@gmail.com. I am grateful to Patrick Bolton, Philip Bond, Dirk Hackbarth, Zhiguo He, Andrey Malenko, Nengjiu Ju, Jin Li, Alejandro Rivera, Neng Wang, Crystal Ho-Po Wong, Jinqiang Yang, Jin Yu, and seminar and conference participants at National Taipei University, National Tsing Hua University, SUFE Dynamic Finance Workshop, and Applied Theory Workshop at the Chinese University of Hong Kong for valuable discussions and comments. Jingzhou Yan provided excellent research assistance. First draft: September 28, 2019. The latest version is available at https://sites.google.com/site/etywong110/.

Log-optimal portfolio after a honest time: Existence, duality and sensitivity analysis.

Ferdoos Alharbi[1]

[1] *Department of Mathematical and Statistical Sciences, University of Alberta, Central Academic Building, Edmonton, AB T6G 2G1, Canada*

December 30, 2021

Abstract: We consider an initial market model $(S, \mathbb{F}, P)$, described by its public flow of information $\mathbb{F}$ and its tradable assets prices S that are observable through $\mathbb{F}$. To this model we add a honest time τ that might represent the default time of firm or a death time of agent and might not be an $\mathbb{F}$-stopping time. Thus, we adopt the progressive enlargement filtration of $\mathbb{F}$ with τ to incorporate the additional information contained in τ. This gives us a larger flow of information $\mathbb{G}$ that contains $\mathbb{F}$ and makes τ a stopping time. In this talk, I will address the problem of Log-optimal portfolio for the model $(S - S^\tau, \mathbb{G}, P)$ in different aspects. I am interested in answering the following questions:

1. What are the models of τ that guarantee the existence of the Log-optimal portfolio after τ?
2. How can we describe explicitly the Log-optimal portfolio after τ in terms of the $\mathbb{F}$-observable parameters of the pair (S, τ)?
3. Following the "incentive analysis" introduced in Choulli and Yansori (2021) [2] for the stopped models, I will present my quantification for the economic factors that control *the increment in maximum expected logarithmic utility from terminal wealth* for the pair of models $(S - S^\tau, \mathbb{G}, P)$ and $(S, \mathbb{F}, P)$.

In this talk, I will answer these questions fully and will highlight the main ideas behind my answers. Among these key ideas, I will present the following results that are interested in themselves:

1. The martingale decomposition of $\mathbb{G}$-martingales that live after τ. This work complements the work of Choulli/Daveloose/Vanmaele (2020) [3], and is vital for the next work.
2. Complete and explicit description of the set of all deflators for the model $(S - S^\tau, \mathbb{G}, P)$. This result is vital for our results described above on the log-optimal portfolio for the model $(S - S^\tau, \mathbb{G}, P)$.

This talk is based on the joint work [1] with Tahir Choulli (my PhD supervisor):

References

[1] Ferdoos Alharbi and T. Choulli (2021): Log-optimal portfolio after a random time: Existence, duality and sensitivity analysis, Preprint of University of Alberta.

[2] T. Choulli and Sina Yansori (2021): Log-optimal and numéraire portfolios for market models stopped at a random time, to appear in Finance and Stochastics.

[3] T. Choulli, C. Daveloose, and M. Vanmaele (2020): A martingale representation theorem and valuation of defaultable securities. Mathematical Finance, 1-38, https://doi.org/10.111/mafi.12244.

Log-optimal portfolio for market models stopped at a random time

Tahir Choulli[1]

[1]*Department of Mathematical and Statistical Sciences, University of Alberta, Central Academic Building, Edmonton, AB T6G 2G1, Canada*

December 30, 2021

Abstract: We consider a market model with two flows of information, which are mathematically represented by two filtrations. The smaller flow $\mathbb{F}$ models the public information that is available to all agents, while the larger flow $\mathbb{G}$ incorporates additional information. Herein, we focus on the case where the additional information consists of a random time τ that might model defaults times of firms or death times of agents and/or insured. As random times can not be seen before their occurrence, this boils down to suppose that $\mathbb{G}$ is the progressive enlargement of $\mathbb{F}$ with τ. In this setting, in many mathematical and economic aspects, we quantify the impact of τ on the log-optimal portfolio when passing from the model $(S, \mathbb{F}, P)$ to the model $(S^\tau, \mathbb{G}, P)$. In particular, we address the following:

1. What are the types of risks induced by τ that really affect the Log-portfolio portfolio?
2. What are the conditions on τ (preferably in terms of information theoretic concepts such as entropy) that guarantee the existence of the log-optimal portfolio of $(S^\tau, \mathbb{G}, P)$ when that of $(S, \mathbb{F}, P)$ already exists?
3. How Log-optimal portfolio can be described using the parameters of τ and those of the initial model $(S, \mathbb{F}, P)$?
4. What are the factors that fully determine *the increment in maximum expected logarithmic utility from terminal wealth* for the models $(S^\tau, \mathbb{G}, P)$ and $(S, \mathbb{F}, P)$, and how to quantify them?
5. Comparison between our progressive setting and the insider setting in which the additional information is known from the beginning of the investment time interval.

This talk is based on the following joint work with Sina Yansori (former PhD,):

T. Choulli and Sina Yansori (2021): Log-optimal and numéraire portfolios for market models stopped at a random time, to appear in Finance and Stochastics.

Special Sessions in memory of Prof. Peter Carr

MA-type Trading Strategies Maximize Utility under Gaussian Partial Information

Speaker: Roger Lee:

Abstract: Under Gaussian price dynamics with stochastic unobserved drift, including cases of mean-reversion and momentum dynamics, we solve explicitly for trading strategies which maximize expected logarithmic, exponential, and power utility. We prove that the optimal strategies depend on current price and an MA price, and in some cases current wealth – not on any other stochastic variables. We verify optimality of MA rigorously over all strategies (not just of MA type) satisfying integrability criteria, and express the optimal parameters explicitly. Joint with Xiaodong Wang

Markov Decision Processes (MDPs) may be viewed as a discrete-time stochastic control problem for sequential decision making in situations where costs are partly random and partly under the control of a decision maker. Classical MDP theory is concerned with minimizing the expected discounted total cost and, in many cases, the minimization problem is solved by establishing a Dynamic Programming Principle (DPP). Results on the vast area of MDPs may be founded in several textbooks, e.g., [5, 12, 4]. The classical expected performance criteria is, however, limited in its application and, in many cases, it is prudent to incorporate risk assessment into decision making.

One popular criterion is based on coherent risk measures [1, 9]. A naive combination of coherent risk measures and discounted total costs, however, lacks time consistency, hindering the derivation of a corresponding DPP. Roughly speaking, time consistency is about the property that smaller scores in the future epochs guarantee a smaller score at current epoch. We refer to [6] for a survey on various definitions of time consistency. There is a stem of literature (see, e.g., [11, 15, 16, 14, 7, 2]) that studies time consistency from multiple angles and/or attempts to integrate coherent risk measures and their variations into MDP.

In this paper, we focus on the framework proposed in [16] which considers deterministic costs. [16] introduces the notion of risk transition mappings and uses them to construct, in a recursive manner, a class of (discounted) dynamic risk measures. He proceeds to derive both finite and infinite (with bounded costs) time horizon DPPs for such dynamic risk measures. We also refer to [17] for the assumptions needed. [19] extends the infinite horizon DPP to unbounded costs as well as for average dynamic risk measures. The risk transition mappings involved are assumed to exhibit an analogue of a strong Feller property. [8] studies a similar infinite horizon DPP with unbounded costs with different assumptions. Recently, [3] considers unbounded latent costs and establishes the corresponding finite and infinite horizon DPP. The authors also prove sufficiency of Markovian actions against history dependent actions. They construct dynamic risk measures, for finite time horizon problems, from iterations of static risk measures that are Fatou and law invariant. The infinite horizon problems require in addition the coherent property. They also require the underlying MDP to exhibit a certain strong continuous/semi-continuous transition mechanism. Finally, it is noteworthy that the concept of risk form is introduced in [10] and is applied to handle two-stage MDP with partial information and decision-dependent observation distribution.

The main goal of this paper is to study infinite horizon risk averse MDPs in a similar framework as above, but with latent costs and randomized actions. We note that, typically, MDP theory with Polish action spaces implicitly encompasses randomized actions. In order to compare deterministic and randomized actions, however, we must characterize the randomization explicitly. In certain risk averse settings, when randomness in action is accounted for along with the random outcome, deterministic actions do not necessarily yield the optimal outcome; an insightful example is provided. To this end, we propose to study risk averse MDPs under the notion of Kusuoka-type conditional risk mappings[1], which is inspired by the Kusuoka representation for proper lower semi-continuous law invariant coherent risk measures (cf. [13] and [18, Section 6.3.5]). Kusuoka-type conditional risk mappings, in principle, covers a large class of conditional risk mappings of interest. To the best

[1] The term, conditional risk mapping, follows from [18, Section 6.5.2].

of our knowledge, the counterpart of Kusuoka representation for conditional risk mappings has not yet been established. Due to the fact that conditional average value at risk in general lacks joint measurability in the risk level and the random event, some specific technical treatment is used to develop the Kusuoka-type conditional risk mappings. The treatment may be of inerest on its own.

For simplicity, we consider bounded costs, which allows for conditional risk mappings that contain conditional essential supremum as a major ingredient – a feature that is often omitted otherwise. Moreover, the Kusuoka-type conditional risk mappings and bounded costs together allows us access to a stronger set of regularities for the related operators, and allows us to establish the DPP with mild assumptions on the remainder of the components in our setup. To be more precise, we obtain the semi-continuity of value function jointly in state and action, without needing to impose strong continuity on the transition kernel or the resulting risk transition mapping. We believe the conditional law invariant property of Kusuoka-type conditional risk mappings is essential for obtaining regularity, while the assumption on bounded costs may possibly be weaken. In static case, semi-continuity of a coherent risk measure typically requires more than weak continuity of the input; we refer to [18, Section 6.3] for detailed statements. This is possibly due to the lack of law invariance. Imposing the law invariant property, as we do here, resolves such issues. In the dynamic case, we choose to avoid the assumptions of strong continuity by implicitly leveraging a similar argument through Kusuoka-type conditional risk mappings.

The main contributions of this paper can be summarized as follows:

1. We introduce and investigate the notion of Kusuoka-type conditional risk mappings. We first study conditional average value at risk and develop an appropriate way to integrate over the quantile level. A Kusuoka-type conditional risk mapping is defined as the essential supremum of a family of integrations with random integrand and integrator. We then establish a useful representation in terms of regular conditional expectation. We then show that the Kusuoka-type conditional risk mapping is conditionally law invariant and state dependent.
2. Under mild conditions, we derive an infinite horizon DPP, for MDP with latent costs and randomized actions, subject to dynamic risk measure defined recursively via Kusuoka-type conditional risk mapping. We also derive a corresponding Q-learning version of the DPP which lends itself naturally to numerical implementation. We prove the existence of an optimal policy that is Markovian.
3. We argue that certain Markovian actions are no worse than any other history dependent actions. We also formulate a sufficient condition on the optimality of deterministic actions. Further, we provide a related heuristic discussion from the perspective of a two-player game.

REFERENCES

[1] P. Artzner, F. Delbaen, J. Eber, and D. Heath, *Coherent measures of risk*, Mathematical Finance, 9 (1999), pp. 203–228.

[2] N. Bäuerle and A. Glauner, *Minimizing spectral risk measures applied to markov decision processes*, Mathematical Methods of Operations Research, 94 (2021), pp. 35–69.

[3] N. Bäuerle and A. Glauner, *Markov decision processes with iterated coherent risk measures*, European Journal of Operational Research, 296 (2022), pp. 953–966.
[4] N. Bäuerle and U. Rieder, *Markov Decision Processes with Applications to Finance*, Springer Science & Business Media, 2011.
[5] D. P. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*, Athena Scientific, 1996.
[6] T. R. Bielecki, I. Cialenco, and M. Pitera, *A survey of time consistency of dynamic risk measures and dynamic performance measures in discrete time: Lm-measure perspective*, Probability, Uncertainty and Quantitative Risk, 2 (2017).
[7] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, *Risk-sensitive and robust decision-making: a cvar optimization approach*, Advances in Neural Processing Systems, (2015).
[8] S. Chu and Y. Zhang, *Markov decision processes with iterated coherent risk measures*, International Journal of Control, 88 (2014), pp. 2286–2293.
[9] F. Delbaen, *Coherent risk measures on general probability spaces*, Advances in Finance and Stochastics, 9 (2002), pp. 1–37.
[10] D. Dentcheva and A. Ruszczyński, *Risk forms: representation, disintegration, and application to partially observable two-stage systems*, Mathematical Programming, Series B, 181 (2020), pp. 297–317.
[11] E. R. Gianin, *Risk measures via g-expectations*, Insurance: Mathematics and Economics, 39 (2006), pp. 19–34.
[12] O. Hernandez-Lerma and J. B. Lasserre, *Discrete-Time Markov Control Processes: Basic Optimality Criteria*, Springer, New York, NY, 1996.
[13] S. Kusuoka, *On law invariant coherent risk measures*, Advances in Mathematical Economics, (2001), pp. 83–95.
[14] G. C. Pflug and A. Pichler, *Time-consistent decisions and temporal decomposition of coherent risk functionals*, Mathematics of Operations Research, 41 (2016).
[15] F. Riedel, *Dynamic coherent risk measures*, Mathematical Programming, Series B, 121 (2004), pp. 185–200.
[16] A. Ruszczyński, *Risk-averse dynamic programming for markov decision processes*, Mathematical Programming, Series B, 125 (2010), pp. 235–261.
[17] A. Ruszczyński, *Erratum to: Risk-averse dynamic programming for markov decision processes*, Mathematical Programming, Series A, 145 (2014), pp. 601–604.
[18] A. Shapiro, D. Dentcheva, and A. Ruszczynski, *Lectures on Stochastic Programming: Modeling and Theory, Third Edition*, Springer, 2021.
[19] Y. Shen, W. Stannat, and K. Obermayer, *Risk-sensitive markov control processes*, SIAM Journal on Control and Optimization, 51 (2013), pp. 3652—-3672.

Markov-Modulated Affine Processes

Kevin Kurt and Rüdiger Frey

Institute for Statistics and Mathematics
Vienna University of Economics and Business (WU Wien)

A standard affine process Y in the sense of [1] can be characterized as a strong Markov process whose generator is given in terms of so-called *admissible* parameters that form affine functions of the state y of the process, cf. [3]. Markov-modulated affine processes or MMAPs are a natural extension where the constant (non y-dependent) part of these affine functions depends on some exogenous Markov process X. A simple example is provided by a CIR process with Markov modulated mean reversion level. Formally, we introduce Markov-modulated affine processes as a class of Markov processes, which we define in terms of the generator $\mathfrak{L}$: we assume that $\mathfrak{L}$ is a linear operator of the form

$$\mathfrak{L}f(x,y) = \mathfrak{L}^X f(x,y) + \mathfrak{L}^{Y|X}(x) f(x,y), \quad f \in \mathrm{Dom}(\mathfrak{L}) \subset C_\infty(\mathcal{D}^X \times \mathcal{D}^Y).$$

Here, $\mathfrak{L}^X$ (acting on $x \mapsto f(x,y)$) is the generator of a Feller process X with state space $\mathcal{D}^X \subset \mathbb{R}^d$, and for fixed x the operator $\mathfrak{L}^{Y|X}(x)$ (acting on $y \mapsto f(x,y)$) is the generator of an affine process Y with state space $\mathcal{D}^Y \subset \mathbb{R}^n$. Thus, conditional on the path of X, the process Y can be regarded as a time-inhomogeneous affine process.

Markov-modulation adds flexibility for (financial) modelling to the class of affine processes. At the same time MMAPs largely preserve the tractability of the latter class, as the characteristic function of their marginal distributions can be computed via an extension of the Riccati equations for standard affine processes. More precisely, for an MMAP it holds that

$$\mathbb{E}_z\left[e^{\langle u, Y_t\rangle}\right] = \varphi(t,x,u) e^{\langle \psi(t,u), y\rangle}, \quad z = (x,y) \in \mathcal{D}^X \times \mathcal{D}^Y. \tag{1}$$

The function ψ is characterized by the same system of generalized Riccati equations as in the standard affine case, whereas φ is determined by means of a Cauchy problem involving ψ and the generator $\mathfrak{L}^X$ of X.

Our contribution is a substantial extension of previous work such as [4]: we make only very mild assumptions on $\mathfrak{L}^X$ and we allow for discontinuities in the coefficients of $\mathfrak{L}^{Y|X}(x)$. In particular, $\mathfrak{L}^X$ may be unbounded, so that models where X is a (jump-)diffusion or a jump process with infinite activity fall within the scope of our analysis. This is relevant for financial applications. There X often represents state variables such as market sentiment or the economic environment, and it may be more natural to model the dynamics of these state variables by continuous processes and not by processes with piecewise constant trajectories. The extension to unbounded generators $\mathfrak{L}^X$ is also interesting from a mathematical viewpoint. In that case the generator $\mathfrak{L}$ of (X,Y) does not satisfy the regularity conditions commonly imposed in the construction of Markov processes via perturbation arguments, so that classical results from semigroup theory are not readily applicable. To deal with this issue we choose a probabilistic approach and use weak convergence results to construct solutions to the martingale problem associated with $\mathfrak{L}$. If the Cauchy problem for the function φ from (1) admits a classical solution, relation (1) determines the characteristic function and hence the marginal distributions of (X,Y). Classical results from martingale problems thus guarantee the Markov property and consequently the existence of MMAPs. We go on and analyze further mathematical properties of MMAPs such as the Feller property, the existence of real exponential moments and the semimartingale characteristics of MMAPs. Finally, in order to illustrate the wide range of modelling possibilities offered by MMAPs we discuss several applications of MMAPs in finance.

Our analysis builds on the formal treatment of affine processes provided in [1]. Moreover, we make extensive use of the comprehensive treatment of Markov processes and martingale problems in [2]. We further contribute to the list of extensions to affine processes that has started to grow ever since the seminal work of [1].

On the applied side, we introduce novel credit risk and option pricing models. Two examples, which are representative for the effectiveness of Markov-modulated affine processes, are outlined in the sequel. First, we consider bond pricing with negative association between the short rate and the underlying firm's default intensity. In the standard affine setup, it is typically not feasible to introduce negative dependence between positive diffusions (i.e. processes of CIR-type). We present a simple solution to this issue by taking X as a Jacobi diffusion which has a reverse impact on the drift of the short rate and on the drift of the default intensity. Second, we propose a tractable model for the joint pricing of European options and credit-risky instruments with the same underlying. In our model, conditional on a common Markov process X the default intensity and the stochastic volatility of a firm are independent positive Markov-modulated affine processes. The underlying Markov process X controls the drift and the jump measure of the two risk factors, and creates hereby comovement among the two processes. We derive semi-explicit (explicit up to the solution of PDEs, and ODEs in special cases) pricing formulae for credit default swaps and certain European options written on the same firm.

References

[1] Duffie, D., Filipović, D., and Schachermayer, W. Affine processes and applications in finance. *Annals of Applied Probability 13*, 3 (2003), 984–1053.

[2] Ethier, S., and Kurtz, T. G. *Markov Processes: Characterization and Convergence.* Wiley, New York, 1986.

[3] Keller-Ressel, M., Schachermayer, W., and Teichmann, J. Affine processes are regular. *Probability Theory and Related Fields 151*, 3-4 (2011), 591–611.

[4] van Beek, M., Mandjes, M., Spreij, P., and Winands, E. Regime switching affine processes with applications to finance. *Finance and Stochastics 24*, 2 (2020), 309–333.

Special Sessions in memory of Prof. Tomas Bjork

Martingale transport, DeMarch-Touzi paving, and stretched Brownian motion in Rd

Speaker: Walter Schachermayer

Abstract: The paper derives eight principles that allow explaining the long-term dynamics of large stock markets, the typical distribution of the market capitalization of stocks, the risk premia for stock portfolios, the key role of optimal portfolios at the growth efficient frontier, the least expensive pricing and hedging of long-term payoffs, and other market features. By applying the concepts of entropy maximization and energy conservation in a richer modeling world than typically considered, most of these market properties follow rather directly. Furthermore, popular fundamental tools for portfolio and risk management, including the intertemporal capital asset pricing model and the preferred pricing rule, become revised.

11th World Congress of the Bachelier Finance Society
The Chinese University of Hong Kong (CUHK), June 13-17, 2022

Applicant's Details

NAME:	Camilo Hernández.
EMAIL:	c [dot] hernandez [at] imperial [dot] ac [dot] uk
PROFESSIONAL STATUS:	Chapman Fellow in Mathematics
AFFILIATION:	Mathematics Department.
	Imperial College London.
ADDRESS:	705 Weeks Hall. 16-18 Prince's Gardens. London, UK..

Talk

Title: A general approach to non-Markovian time-inconsistent stochastic control for sophisticated players
Authors: Camilo Hernandez*, Dylan Possamaï.

Based on: C. Hernández and D. Possamaï. Me, myself and I: a general theory of non–Markovian time-inconsistent stochastic control for sophisticated agents. *Accepted in Annals of Applied Probability* (*minor revisions*). *ArXiv preprint arXiv:2002.12572*, 2020.

Extended abstract

In this work we develop a theory for continuous time non-Markovian stochastic control problems which are time-inconsistent. Time-inconsistency was first mentioned in Strotz [5] where three different types of agents are described: the pre-committed agent does not revise his initially decided strategy; the naive agent revises his strategy without taking future revisions into account; the sophisticated agent revises his strategy taking possible future revisions into account, and by avoiding such makes his strategy time-consistent. In Hernández and Possamaï [4], we focused on understanding the latter type of agents. The first analytic study of sophisticated agents was carried out in Ekeland and Lazrak [2, 3] which later became the starting point of the general Markovian theory developed by Björk, Khapko, and Murgoci [1]. Nonetheless, none of these approaches was able to handle the typical non-Markovian problems that would necessarily arise in contracting problems involving a Principal and time-inconsistent Agents. This was an essential motivation for our work.

In [4] we considered a finite horizon model, $T > 0$. To illustrate our findings we present below the case of a non–exponential discounting function $f : [0,T] \longrightarrow (0,\infty)$ (with $f(0) = 1$) but our results hold for more general criteria. $\mathcal{A}$ denotes the set admissible actions used by the Agent to control the d-dimensional state process X as follows. For $\alpha \in \mathcal{A}$ and a $\mathbb{P}^\alpha$–Brownian motion W^α (depending on α), X satisfies the dynamics

$$X_t = x_0 + \int_0^t \sigma_r(X_{\cdot\wedge r},\alpha_r)\big(b_r(X_{\cdot\wedge r},\alpha_r)\mathrm{d}r + \mathrm{d}W^\alpha_r\big),\ t\in[0,T],\ \mathbb{P}^\alpha\text{–a.s.}, \tag{0.1}$$

where $X_{\cdot\wedge t}$ denotes the path up to time t of the state process X and $x \in \mathcal{X}$[1] denotes its past trajectory. Let $\mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}$ denote the augmented filtration generated by X and $\mathrm{U_A} : \mathbb{R} \longrightarrow \mathbb{R}$. The utility drawn by the Agent from an effort action α at time $t \in [0,T]$ is given by

$$\mathrm{J^A}(t,x,\alpha) := \mathbb{E}^{\mathbb{P}^\alpha}\bigg[f(T-t)\mathrm{U_A}(\xi) - \int_t^T f(r-t)c_r(X_{\cdot\wedge r},\alpha_r)\mathrm{d}r\bigg|\mathcal{F}_t\bigg],\ (t,x,\alpha)\in[0,T]\times\mathcal{X}\times\mathcal{A}. \tag{0.2}$$

Given the form of the pay-off functional (0.2), the classical problem of maximising $\alpha \longmapsto \mathrm{J^A}(0,x_0,\alpha)$ has a time-inconsistent nature. More precisely, the discount function f, in the terminal utility and the running cost, is the source of inconsistency. Of course, for exponential discounting the problem is time-consistent. Moreover, this is the only discounting structure for which this is the case. As we study this problem from a game-theoretic perspective, we revisited the notion of *equilibria* initially proposed in [2]. Our first contribution was a refinement of this concept. For any two actions $\{\alpha,\alpha'\} \subseteq \mathcal{A}$, $\alpha\otimes_t\alpha' := \alpha\mathbf{1}_{[0,t]} + \alpha\mathbf{1}_{(t,T]}$.

Definition 1 ([4]). *Let $\alpha^\star \in \mathcal{A}$. If for any $\varepsilon > 0$, there is $\ell_\varepsilon > 0$ s.t. for any $\ell \in (0,\ell_\varepsilon)$ and $(t,\alpha)\in[0,T]\times\mathcal{A}$*

$$\mathrm{J^A}(t,x,\alpha^\star) - \mathrm{J^A}(t,x,\alpha\otimes_{t+\ell}\alpha^\star) \geq -\varepsilon\ell.$$

then, we say $\alpha^\star$ is an equilibria, $\alpha^\star \in \mathcal{E}$. For $\alpha^\star \in \mathcal{E}$, we define $\mathrm{V^A}(t,x) := \mathrm{J^A}(t,x,\alpha^\star)$, $(t,x)\in[0,T]\times\mathcal{X}$.

[1] $\mathcal{X}$ denotes the space of $\mathbb{R}^d$-valued continuous functions on $[0,T]$

Given the sequential nature of the game, our definition of equilibrium can be interpreted as Player t's ε–best response for all sub-games in $[t,T]$ via coalitions. Indeed, to find her best response, Player t reasons as follows: for any $\tau \in [t, t+\ell_\varepsilon)$ knowing that Players $[\tau,T]$ are following the equilibrium $\alpha^\star$, i.e. the sub-game $[\tau,T]$ is following the equilibrium, and treating the actions of those players as fixed, the action prescribed by the equilibrium to Player t is such that he is doing at least as good as in any coalition/deviation α from the equilibrium with players $t+\ell \in [t,\tau)$, give or take $\varepsilon\ell$. The motivation for this definition was to obtain an extended dynamic programming principle satisfied by equilibria. To quote Strotz [5], "[The] problem [of a sophisticated agent] is to find the best plan among those that [he] will actually follow." Consequently, Theorem 2 below is not only the main result of [4], but also, in our opinion, a fundamental milestone for the theory of time-inconsistency in continuous time.

Theorem 2 ([4])**.** *Let $\alpha^\star \in \mathcal{E}$. For any stopping times $\sigma \le \tau$, we have that*

$$\mathrm{V}^{\mathrm{A}}_\sigma = \sup_{\alpha\in\mathcal{A}} \mathbb{E}^{\mathbb{P}^\alpha}\bigg[\mathrm{V}^{\mathrm{A}}_\tau - \int_\sigma^\tau \bigg(c_r(X_{\cdot\wedge r},\alpha) + \mathbb{E}^{\mathbb{P}^{\alpha^\star}}\bigg[f'(T-r)\mathrm{U}_{\mathrm{A}}(\xi) - \int_r^T f'(u-r)c_u(X_{\cdot\wedge r},\alpha^\star_u)\mathrm{d}u\bigg|\mathcal{F}_r\bigg]\bigg)\mathrm{d}r\bigg|\mathcal{F}_\sigma\bigg].$$

As an immediate consequence of this result, we can associate a system of BSDEs to study equilibria in the case of uncontrolled volatility, i.e. $\sigma_t(x,a) = \sigma_t(x)$. Let us introduce the system which for any $s \in [0,T]$ satisfies

$$\begin{aligned} \mathcal{Y}_t &= \mathrm{U}_{\mathrm{A}}(\xi) + \int_t^T \big(H_r(X_{\cdot\wedge r},\mathcal{Z}_r) - \partial Y^r_r\big)\mathrm{d}r - \int_t^T \mathcal{Z}_r\cdot\mathrm{d}X_r,\ t\in[0,T],\ \mathbb{P}\text{–a.s.} \\ \partial Y^s_t &= f'(T-s)\mathrm{U}_{\mathrm{A}}(\xi) + \int_t^T \nabla h_r(s, X_{\cdot\wedge r}, \partial Z^s_r, a^\star(r,X_{\cdot\wedge r},\mathcal{Z}_r))\mathrm{d}r - \int_t^T \partial Z^s_r\cdot\mathrm{d}X_r,\ t\in[0,T],\ \mathbb{P}\text{–a.s.} \end{aligned} \tag{$\mathcal{H}$}$$

where $H_t(x,z) := \sup_a\big\{\sigma_t(x)b_t(x,a)\cdot z - c_t(x,a)\big\}$, $\nabla h_t(s,x,z,a) := \sigma_t(x)b_t(x,a)\cdot z + f'(t-s)c_t(x,a)$ and $a^\star(t,x,z)$ denotes the unique, by assumption, maximiser in H. The previous system is of an infinite-dimensional nature as the second equation induces a family of BSDEs, one for every $s\in[0,T]$. Moreover, it is fully-coupled as the diagonal term of the family $(\partial Y^s)_{s\in[0,T]}$ appears in the generator of the first equation and $\mathcal{Z}$ appears, via $a^\star$, in the generator of the family of BSDEs. Be it as it may, we are able to show that: (i) $(\mathcal{H})$ is of both sufficient and necessary to characterise equilibria $\alpha^\star\in\mathcal{E}$; (ii) $\mathcal{Y}$ coincides with the value function V^{A}, and $\alpha^\star$ always arises as maximisers of the hamiltonian; and (iii) $(\mathcal{H})$ is well-defined.

We would like to mention that $(\mathcal{H})$ extends naturally to a system incorporating second-order BSDEs (2BSDEs) in the case the Agent is allowed to control the volatility. Indeed, Theorem 2 holds true in the case the Agent controls the volatility as well. For such system we are able to show (i) and (ii) still hold. Nevertheless, (iii) becomes a much more delicate matter as the existence of a solution requires the existence of an optimal measure $\mathbb{P}^\star$. As a final comment, we also address the extensions of $(\mathcal{H})$ to the case of rewards much more general than (0.2). In particular, to those with mean-variance type of criteria which are indispensable, for example, in applications in portfolio selection, or energy consumption management.

References

[1] T. Björk, M. Khapko, and A. Murgoci. On time-inconsistent stochastic control in continuous time. *Finance and Stochastics*, 21(2):331–360, 2017.

[2] I. Ekeland and A. Lazrak. Being serious about non-commitment: subgame perfect equilibrium in continuous time. *ArXiv preprint arXiv:0604264*, 2006.

[3] I. Ekeland and A. Lazrak. The golden rule when preferences are time inconsistent. *Mathematics and Financial Economics*, 4(1):29–55, 2010.

[4] C. Hernández and D. Possamaï. Me, myself and I: a general theory of non–Markovian time-inconsistent stochastic control for sophisticated agents. *Accepted in Annals of Applied Probability (minor revisions). ArXiv preprint arXiv:2002.12572*, 2020.

[5] R.H. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955.

Mean Field Game of Mutual Holding

Mao Fabrice Djete[*] Nizar Touzi[†]

March 13, 2022

Abstract

We introduce a mean field model for optimal holding of a representative agent of her peers as a natural expected scaling limit from the corresponding $N-$agent model. The induced mean field dynamics appear naturally in a form which is not covered by standard McKean-Vlasov stochastic differential equations. We study the corresponding mean field game of mutual holding in the absence of common noise. Our first main result provides an explicit equilibrium of this mean field game, defined by a bang–bang control consisting in holding those competitors with positive drift coefficient of their dynamic value. We next use this mean field game equilibrium to construct (approximate) Nash equilibria for the corresponding N–player game. We also provide some numerical illustrations of our mean field game equilibrium which highlight some unexpected effects induced by our results.

[*]Ecole Polytechnique Paris, Centre de Mathématiques Appliquées, mao-fabrice.djete@polytechnique.edu. This work benefits from the financial support of the Chairs *Financial Risk* and *Finance and Sustainable Development*.

[†]Ecole Polytechnique Paris, Centre de Mathématiques Appliquées, nizar.touzi@polytechnique.edu. This work benefits from the financial support of the Chairs *Financial Risk* and *Finance and Sustainable Development*.

Speaker: Prof. Hoi Ying Wong, The Chinese University of Hong Kong

Title: Mean-variance problems and delayed cointegration

Cointegration analysis is an econometric tool used to identify equilibrium among assets and construct a pairs trading portfolio. The discrete-time vector error correction model (VECM) for identifying cointegration includes lag difference terms as explanatory variables, thus permitting delayed adjustment of the deviations from equilibrium. The continuous-time limit of the VECM becomes a stochastic delay differential equation. This talk concerns with some dynamic mean-variance (MV) problems originated from several works by Professor Duan Li. We investigate dynamic mean-variance pairs trading strategies under such a delayed cointegration model. Both pre-commitment and equilibrium strategies are presented. We prove that the equilibrium trading of delayed cointegrated assets can lead to statistical arbitrage under certain conditions that are related to the roots of the corresponding characteristic equation. We obtain an explicit solution for a case with distributed delay. Our empirical study demonstrates the superiority of our strategy over its Markovian counterpart when the model selection result prefers a high-order VECM. In addition, we present the optimal liquidation of delay cointegrated assets considering price impact. When compared to the optimal liquidation of pairs of Markovian cointegrated assets, we show that the delay cointegration affects the time synchronization of trading among assets. This talk is based on joint works with M.C. Chiu and T. Yan.

Measure-valued processes for energy markets

Christa Cuchiero* Luca Di Persio Francesco Guida

Abstract

We introduce a framework that allows to employ (non-negative) measure-valued processes for energy market modeling, in particular for electricity and gas futures. Interpreting the process' spatial structure as time to maturity, we show how the *Heath-Jarrow-Morton* (HJM) approach can be translated to this framework, thus guaranteeing arbitrage free modeling in infinite dimensions. We derive an analog to the HJM-drift condition and then treat in a Markovian setting existence of (non-negative) measure-valued diffusions that satisfy this condition. To analyze convenient classes we build on [3] and consider measure-valued polynomial and affine diffusions, where we can precisely specify the diffusion part. Indeed, it depends on continuous functions satisfying certain admissibility conditions. For calibration purposes these functions can then be parametrized by neural networks yielding *measure-value analogs of neural SPDEs.* Exploiting the analytic tractability coming from the affine and polynomial nature as well as stochastic gradient descent methods for neural networks allows for efficient mass transport in infinite dimensions from the initial distribution, encoding the current forward curve, to the distribution of future forward curves implied from option prices.

Extended summary

The most liquidly traded products in electricity, gas and also temperature markets are future contracts as well as options written on these. These future contracts deliver the underlying commodity over a certain period rather than at one fixed instance of time. Our focus lies on these future markets as they exhibt higher liquidity than the spot energy markets. In spirit of a Heath-Jarrow-Morton (HJM) approach (see [2] in the context of energy markets), we thus model directly an analog of the forward curve, however with (non-negative) measure-valued processes rather than function-valued ones.[1]

Measure-valued processes are often used for modeling dynamical systems for which the spatial structure plays a significant role. In the current setting, time to maturity takes the role of the spatial structure. This is similar to common forward curve modeling via stochastic partial differential equation (SPDE) as for instance in [2, 1]. The potential advantage of using measure-valued processes instead of function-valued ones is that many spatial stochastic processes do not fall into the framework of SPDEs.

However, the decisive economic reason for using measure-valued processes in electricity and gas modeling lies in the very nature of the future contracts, namely as products that deliver the underlying commodity always over a certain period instead of one fixed instance in time. To formulate this mathematically, denote the price of a future with delivery over the *time interval*

*Vienna University, Department of Statistics and Operations Research, Data Science @ Uni Vienna, Kolingasse 14-16 1, A-1090 Wien, Austria, christa.cuchiero@univie.ac.at

[1] In the sequel, we shall omit "non-negative" and only use "measure" for "non-negative measure".

$[\tau_1, \tau_2]$ at time $0 \le t \le \tau_1$ by $F(t, \tau_1, \tau_2)$. Then, $F(t, \tau_1, \tau_2)$ can be written as a weighted integral of instantaneous forward prices $f(t, u)$ with delivery at *one fixed time* $\tau_1 \le u \le \tau_2$, i.e.

$$F(t, \tau_1, \tau_2) = \int_{\tau_1}^{\tau_2} w(u, \tau_1, \tau_2) f(t, u) du, \tag{0.1}$$

where $w(u, \tau_1, \tau_2)$ denotes some weight function. The crucial motivation for measure-valued process now comes from the fact that there is no trading with the instantaneous forwards $f(t, u)$ for obvious reasons. Thus, rather than using $f(t, u)du$ in the expression of the future prices, we can also use a measure, which cannot necessarily be evaluated pointwise.

To establish now a Heath-Jarrow-Morton approach we pass to the Musiela parameterization and consider instead of the maturity date rather time to maturity. More precisely, consider on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$, a measure-valued process $(\mu_t)_{t\in[0,T]}$ supported on the compact interval $E := [0, T]$ with $T > 0$ some finite time horizon, with the meaning "$\mu_t(dx) = f(t, t+x)dx$", i.e. if its Lebesgue density existed, it would correspond to instantaneous forward prices with time to maturity x.

The next step is now to specify a suitable no-arbitrage condition. In this respect the notion of *no asymptotic free lunch with vanishing risk* (NAFLVR) qualifies as an appropriate and economically meaningful condition. Mathematically it is equivalent to the existence of a so-called equivalent separating measure, which in the most relevant cases turns out to be an equivalent local martingale measure. Modulo some technical conditions, this can then be translated to the following *HJM-drift condition* on the measure-valued process $(\mu_t)_{t\in[0,T]}$: if there exists an equivalent measure $\mathbb{Q} \sim \mathbb{P}$ such that

$$\langle \phi, \mu \rangle + \int_0^{\cdot} \langle \phi'_s, \mu_s \rangle ds,$$

is $\mathbb{Q}$-martingale for all appropriate test functions ϕ, then NAFLVR holds. Here, the brackets mean $\langle \phi, \mu \rangle = \int_E \phi(x)\mu(dx)$.

The HJM-drift condition restricts of course the choice of the measure-valued process since the drift part is completely determined, but we are free to specify the martingale part as long as we do not leave the state space of (non-negative) measures. To establish existence of measure-valued diffusion processes satisfying the drift condition, we rely on the martingale problem formulation in locally compact and separable spaces which applies to our setting of measures on a compact set equipped with the weak-$*$-topology. Additionally to these requirements we look for tractable specifications coming from the class of affine and polynomial processes as introduced in [3], where we can precisely specify the diffusion part in terms of continuous functions satisfying certain admissibility conditions. For calibration and pricing purposes these functions can then be parametrized by neural networks, which allows for tractable calibration procedures via the moment formula well-known from polynomial processes and Fourier approaches as well as optimization techniques taylored to neural networks.

References

[1] F. E. Benth, N. Detering, and S. Lavagnini. Accuracy of deep learning in calibrating HJM forward curves. *Digital Finance*, pages 1–40, 2021.

[2] F. E. Benth and P. Krühner. Representation of infinite-dimensional forward price models in commodity markets. *Communications in Mathematics and Statistics*, 2(1):47–106, 2014.

[3] C. Cuchiero, L. Di Persio, F. Guida, and S. Svaluto-Ferro. Measure-valued affine and polynomial diffusions. *arXiv preprint: 2112.15129*, 2021.

MODEL-FREE BOUNDS AND DETECTION OF ARBITRAGE IN MULTI-ASSET DERIVATIVES MARKETS

ANTONIS PAPAPANTOLEON

Abstract

In this talk, we will present recent advances on improved Fréchet–Hoeffding bounds, their relation to optimal transport and their applications in mathematical finance and quantitative risk management. We will start by providing improved Fréchet–Hoeffding bounds on the copula of a random vector S that account for additional information. In particular, we will derive bounds when the values of the copula are given on a compact subset of $[0,1]^d$, the value of a functional of the copula is prescribed or different types of information are available on the lower dimensional marginals of the copula. We will apply these results to compute model-free bounds on the prices of multi-asset options or portfolio Value-at-Risk that take partial information on the dependence structure into account, such as correlations or market prices of other traded derivatives. The numerical results show that the additional information leads to a significant improvement of the option price bounds compared to the situation where only the marginal distributions are known. In order to answer certain open questions about the improved Fréchet–Hoeffding bounds, we will consider Fréchet classes of multivariate distribution functions where additional information on the joint distribution is assumed, while uncertainty in the marginals is also possible. We will derive optimal transport duality results for these Fréchet classes that extend previous results in the related literature, and show that the dual transport problem admits an explicit solution for the function $f = 1_B$, where B is a rectangular subset of $\mathbb{R}^d$. This allows us to show that the improved Fréchet–Hoeffding bounds are pointwise sharp for these classes in the presence of uncertainty in the marginals, while a counterexample yields that they are not pointwise sharp in the absence of uncertainty in the marginals, even in dimension 2. Finally, we are interested in the existence of equivalent martingale measures and the detection of arbitrage opportunities in markets where several multi-asset derivatives are traded simultaneously. Using the results on the improved Fréchet-Hoeffding bounds we can derive sufficient conditions for the absence of arbitrage and formulate an optimization problem for the detection of a possible arbitrage opportunity that can be solved numerically. The most interesting practical outcome is the following: we can construct a financial market where each multi-asset derivative is traded within its own no-arbitrage interval, and yet when considered together an arbitrage opportunity may arise.

References

[1] D. Bartl, M. Kupper, T. Lux, A. Papapantoleon, S. Eckstein: Marginal and dependence uncertainty: bounds, optimal transport, and sharpness. *SIAM J. Control and Optimization* (forthcoming).

[2] T. Lux, A. Papapantoleon: Improved Fréchet–Hoeffding bounds for d-copulas and applications in model-free finance. *Annals of Applied Probability 27*, 3633–3671, 2017.

[3] T. Lux, A. Papapantoleon: Model-free bounds on Value-at-Risk using extreme value information and statistical distances. *Insurance: Mathematics and Economics 86*, 73–83, 2019.

[4] A. Papapantoleon, P. Yanez Sarmiento: Detection of arbitrage opportunities in multi-asset derivatives markets. *Dependence Modeling* (forthcoming).

DIAM, TU Delft, Netherlands & IACM, FORTH, Greece
Email address: a.papapantoleon@tudelft.nl

1

Model-free bounds for multi-asset options using option-implied information and their exact computation

Ariel Neufeld, Antonis Papapantoleon, and Qikun Xiang

Extended Abstract

The classical paradigm in finance and theoretical economics assumes the existence of a model that provides an accurate description of the evolution of asset prices, and all subsequent computations about hedging strategies, exotic derivatives, risk measures, and so forth, are based on this model. However, academics, practitioners, and regulators have realized that all models provide only a partially accurate description of this reality, thus, either methods need to be developed in order to aggregate the results of many models, or approaches have to be devised that allow for computations in the absence of a specific model. The first approach led to the introduction of robust methods in asset pricing and no-arbitrage theory, while the second one led to model-free methods in asset pricing and no-arbitrage theory.

In this work, we consider derivatives written on multiple underlyings in a one-period financial market, and we are interested in the computation of upper and lower bounds for their arbitrage-free prices. We work in a completely realistic setting, in that we only assume the knowledge of traded prices for other single- and multi-asset derivatives, and even allow for the presence of bid–ask spread in these prices. In other words, we work in a model-free setting in the presence of option-implied information, and make no assumption about the probabilistic evolution of asset prices (*i.e.* their marginal distributions) or their dependence structure.

Our contributions are three-fold: Firstly, we provide a fundamental theorem of asset pricing for the market model described above, as well as a superhedging duality, that allows to transform the abstract maximization problem over probability measures into a more tractable problem over vectors, subject to certain constraints. Secondly, we recast this problem into a linear semi-infinite optimization problem, and provide two algorithms for its solution. These algorithms provide upper and lower bounds for the prices of multi-asset derivatives that are ε-optimal, as well as a characterization of the optimal pricing measures. These algorithms are efficient and allow the computation of bounds in high-dimensional scenarios (*e.g.* when $d = 60$), that were not possible by previous methods. Moreover, these algorithms can be used to detect arbitrage opportunities in multi-asset financial markets and to identify the corresponding arbitrage strategies. Thirdly, we perform numerical experiments using synthetic data as well as real market data to showcase the efficiency of these algorithms. These experiments allow us to understand the reduction of the no-arbitrage gap, *i.e.* the difference between the upper and lower no-arbitrage bounds, by including additional information in the form of known derivative prices. The no-arbitrage gap directly reflects the model-risk associated to a particular derivative and the information available in the market. The numerical experiments show a decrease of the model-risk by the inclusion of additional information, although this decrease is not uniform and depends on the form of information and the specific structure of the payoff functions.

Our paper has been submitted and is available at https://arxiv.org/abs/2006.14288.

Speaker: Qikun Xiang

Division of Mathematical Sciences, Nanyang Technological University, Singapore

Email: qikun001@e.ntu.edu.sg

Co-author: Ariel Neufeld

Division of Mathematical Sciences, Nanyang Technological University, Singapore

Email: ariel.neufeld@ntu.edu.sg

Co-author: Antonis Papapantoleon

Department of Mathematics, National Technical University of Athens, Greece

Email: papapan@math.ntua.gr

Special Sessions in memory of Prof. Mark H.A. Davis

Model-free portfolio theory: a rough path approach

David Promel

Abstract: Classical approaches to optimal portfolio selection problems are based on probabilistic models for the asset returns or prices. However, by now it is well observed that the performance of optimal portfolios are highly sensitive to model misspecifications. To account for various type of model risk, robust and model-free approaches have gained more and more importance in portfolio theory. Based on a rough path foundation, we develop a model-free approach to stochastic portfolio theory and Cover's universal portfolio. The use of rough path theory allows treating significantly more general portfolios in a model-free setting, compared to previous model-free approaches. Without the assumption of any underlying probabilistic model, we present pathwise Master formulae analogously to the classical ones in stochastic portfolio theory, describing the growth of wealth processes generated by pathwise portfolios relative to the wealth process of the market portfolio, and we show that the appropriately scaled asymptotic growth rate of Cover's universal portfolio is equal to the one of the best retrospectively chosen portfolio. The talk is based on joint work with Andrew Allan, Christa Cuchiero and Chong Liu.

Modeling of alternative risk-free rates beyond stochastic continuity

C. Fontana, Z. Grbac, T. Schmidt

Interbank offered rates (Ibor), such as Libor rates, have represented the key benchmark rates for interest rate derivatives. In the current reform of interest rate benchmarks, a central role is played by alternative risk-free rates (ARFRs), such as SOFR in the United States and ESTR in the Euro area. ARFRs are obtained by compounding overnight rates and, in recent years, have exhibited a spiky behavior at periodic time intervals, as a result of regulatory and liquidity constraints in the interbank market (see, e.g., [1,3,5]). This provides evidence of the presence of stochastic discontinuities (i.e., jumps occurring at pre-determined points in time) in the dynamics of ARFRs. Moreover, several Ibor fallback rates are currently under consideration to properly account for the presence of credit and funding risk in the interbank market. In this work, we propose a general modeling framework where ARFRs and forward Ibor fallback rates can have stochastic discontinuities, also allowing for jumps at fixed times in the numéraire process. By relying on the techniques recently introduced in [2], we provide a characterization of absence of arbitrage along the lines of the Heath-Jarrow-Morton approach to term structure modeling, with the additional presence of a specific no-arbitrage restriction in correspondence of the stochastic discontinuity dates. In the case where the spot Ibor fallback rate is generated by the ARFR itself, we show that the forward Ibor fallback rate solves a certain BSDE, where the structure of the driver is determined by absence of arbitrage. In general, this BSDE may admit multiple solutions and we provide sufficient conditions ensuring uniqueness of the solution. By relying on the theory of affine semimartingales (see [4]), we propose a tractable specification of our framework, leading to fully explicit valuation formulas for ARFR bonds and caplets in a Gaussian setup. Moreover, we consider the problem of hedging ARFR-based derivatives by relying on local risk-minimization in the presence of stochastic discontinuities. We exemplify this problem by considering the hedging of an ARFR-based caplet by trading in ARFR futures, illustrating the relevance of portfolio rebalancing at the stochastic discontinuity dates.

<u>References</u>:

[1] Andersen, L., Bang, D. (2020), Spike modeling in interest rate derivatives with an application to SOFR caplets, SSRN working paper.

[2] Fontana, C., Grbac, Z., Gümbel, S., Schmidt, T. (2020), Term structure modeling for multiple curves with stochastic discontinuities, *Finance and Stochastics*, 24: 465-511.

[3] Gellert, K., Schlögl, E. (2021), Short rate dynamics: a Fed Funds and SOFR perspective, SSRN working paper.

[4] Keller-Ressel, M., Schmidt, T., Wardenga, R. (2019), Affine processes beyond stochastic continuity, *Annals of Applied Probability*, 29: 3387-3437.

[5] Klinger, S., Syrstad, O. (2021), Life after LIBOR, *Journal of Financial Economics*, 141: 783-801.

Monotone mean-variance portfolio selection with convex constraints under non-Markovian regime-switching models

Kun Fan[a], **Yang Shen**[b,*], Jiaqin Wei[a], and Bin Zou[c]

[a]Key Laboratory of Advanced Theory and Application in Statistics and Data Science-MOE, School of Statistics, East China Normal University, Shanghai 200241, China

[b]**School of Risk and Actuarial Studies, University of New South Wales, Sydney, NSW 2052, Australia**

[c]Department of Mathematics, University of Connecticut, Storrs 06269-1009, USA

This Version: February 1, 2022

Abstract

We study continuous-time portfolio selection problems under monotone mean-variance (MMV) preferences, which "correct" classical mean-variance (MV) preferences on their domain of non-monotonicity. This paper distinguishes from the existing literature by considering general convex constraints on portfolio strategies, which may not be conic, and adopting non-Markovian regime-switching models for risky assets in an incomplete financial market. Due to the latter modeling feature and MMV preferences, the main problem is a stochastic max-min optimization problem with random parameters, and cannot be solved by the standard Hamilton-Jacobi-Bellman (HJB) method. Instead, we employ a combination of theories on stochastic HJB equations and backward stochastic differential equations (BSDEs) to solve this problem. We obtain the optimal investment strategy and the value function under MMV preferences in semi-closed form, subject to the unique solution to a related BSDE. Our results show that the optimal strategies under MMV preferences and MV preferences are not the same when portfolio constraints are present, but the two, to some extent, coincide when constraints are not imposed. Our findings provide both positive and negative responses to the conclusion of Strub and Li (2020) and thus offer valuable insight on the role of constraints on portfolio selection problems under MMV and MV preferences.

Our work also contributes to the literature on time-(in)consistency of portfolio selection under MV preferences. We show that the optimal investment strategy is time-consistent for constrained MMV problems, but is time-inconsistent for constrained MV problems. Upon imposing additional conditions on portfolio constraints, these two problems may become equivalent at the initial time $t_0 = 0$, but still differ at any later time $t > 0$. Although MMV preferences always dominate MV preferences (see Maccheroni et al., 2009), we show that such dominance only holds at the initial time, and fails in general dynamically over time. The main reason is that $\mathrm{d}\mathbb{Q}^*/\mathrm{d}\mathbb{P}|_{\mathcal{F}_0} = 1$ but $\mathrm{d}\mathbb{Q}^*/\mathrm{d}\mathbb{P}|_{\mathcal{F}_t} \neq 1$ for all $t > 0$, where $\mathbb{Q}^*$ is the optimal distorted probability measure in MMV preferences and $\mathbb{P}$ is the reference measure; as such, we may interpret the dynamic deviation of $\mathbb{Q}^*$ from $\mathbb{P}$ as the cost of regaining time-consistency under MMV preferences.

Keywords: Portfolio optimization; Monotone mean-variance; Non-Markovian regime-switching model; Backward stochastic differential equation; Time-inconsistency.

References

[1] Maccheroni F, Marinacci M, Rustichini A, Taboga M (2009) Portfolio selection with monotone mean-variance preferences. *Math. Finance* 19(3):487–521.

[2] Strub MS, Li D (2020) A note on monotone mean-variance preferences for continuous processes. *Oper. Res. Lett.* 48(4):397–400.

*Corresponding and presenting author. Email: y.shen@unsw.edu.au.

Multi-armed bandits with regime-switching rewards

Ningyuan Chen, Xuefeng Gao, Yi Xiong, Xiang Zhou

We study a multi-armed bandit problem where the rewards exhibit regime-switching. Specifically, the distributions of the random rewards generated from all arms depend on a common underlying state modeled as a finite-state Markov chain. The agent does not observe the underlying state and has to learn the unknown transition probability matrix as well as the reward distribution. We propose an efficient learning algorithm for this problem, building on spectral method-of-moments estimations for hidden Markov models and upper confidence bound methods for reinforcement learning. We also establish $O(T^{2/3}\sqrt{\log T})$ bound on the regret of the proposed learning algorithm where T is the unknown horizon. Finally, we conduct numerical experiments to illustrate the effectiveness of the learning algorithm.

Multi-asset optimal execution and statistical arbitrage strategies under Ornstein-Uhlenbeck dynamics

Philippe Bergault* Fayçal Drissi† Olivier Guéant‡

Extended abstract

In recent years, academics, regulators, and market practitioners have increasingly addressed liquidity issues. Amongst the numerous problems addressed, the optimal execution of large orders is probably the one that has attracted the most research works, mainly in the case of single-asset portfolios. In practice, however, optimal execution problems often involve large portfolios comprising numerous assets, and models should consequently account for risks at the portfolio level. In this paper, we address multi-asset optimal execution in a model where prices have multivariate Ornstein-Uhlenbeck dynamics and where the agent maximizes the expected (exponential) utility of her P&L.

The multivariate Ornstein-Uhlenbeck (multi-OU) model is a classical model for the multivariate dynamics of financial variables that goes beyond that of correlated Brownian motions. It is especially attractive because it is parsimonious, and yet general enough to cover a wide spectrum of multi-dimensional dynamics. Multi-OU dynamics offer indeed a large coverage since particular cases include correlated Brownian motions but also cointegrated dynamics which are heavily used in statistical arbitrage.

The main contribution of this paper is to propose a model for multi-asset portfolio execution under multi-OU price dynamics in an expected utility framework that accounts for the overall risk associated with the execution process. We focus on the problem where an agent is in charge of unwinding a large portfolio, but also illustrate the use of our results for multi-asset statistical arbitrage purposes.

We use the tools of stochastic optimal control and simplify the initial multidimensional Hamilton-Jacobi-Bellman equation into a system of ordinary differential equations (ODEs) involving a Matrix Riccati ODE for which classical existence theorems do not apply. By using *a priori* estimates obtained thanks to optimal control tools, we nevertheless prove an existence and uniqueness result for the latter ODE, and then deduce a verification theorem that provides a rigorous solution to the execution problem. Using examples based on data from the foreign exchange and stock markets, we eventually illustrate our results and discuss their implications for both optimal execution and statistical arbitrage.

*Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, philippe.bergault@etu.univ-paris1.fr

†Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, faycal.drissi@etu.univ-paris1.fr

‡Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, olivier.gueant@univ-paris1.fr *Corresponding author.*

Neural network approximation for superhedging prices

Francesca Biagini*

March 14, 2022

This talk is based on the paper

F.Biagini, L.Gonon, T. Reitsam *Neural network approximation for superhedging prices*, Preprint 2021.

1 Abstract of the talk

In this talk we present neural network approximations for the superhedging price process for a contingent claim in discrete time. Superhedging was first introduced in [11] and then thoroughly studied in various settings and market models. It is impossible to cover the complete literature here, but we name just a few references. For instance, in continuous time, for general càdlàg processes we mention [16], for robust superhedging [17], [23], for pathwise superhedging on prediction sets [1], [2], or for superhedging under proportional transaction costs [6], [10], [15], [21], [22]. Also in discrete time there are various studies in the literature, like the standard case [13], robust superhedging [8], [19], superhedging under volatility uncertainty [18], or model-free superhedging [5]. The superhedging price provides an opportunity to secure a claim, but it may be too high or reduce the probability to profit from the option. In order to solve this problem, quantile hedging was introduced in [12], where the authors propose to either fix the initial capital and maximize the probability of superhedging with this capital or fix a probability of superhedging and minimize the required capital. In this way a trader can determine the desired trade-off between costs and risk.
In certain situations it is possible to calculate explicitly or recursively the superhedging or quantile hedging price, see e.g. [7], but in general incomplete markets it may be complicated to determine superhedging prices or quantile hedging prices. Here we investigate neural network-based approximations for quantile- and superhedging prices. Neural network-based methods have been recently introduced in financial mathematics, for instance for hedging derivatives, see [4], determining stopping times, see [3], or calibration of stochastic volatility models, see [9], and many more. For an overview of applications of machine learning to hedging and option pricing we refer to [20] and the references therein.
We contribute to the literature on hedging in discrete time market models in several ways. First, we prove that the α-quantile hedging price converges to the superhedging price for α tending to

*Workgroup Financial and Insurance Mathematics, Department of Mathematics, Ludwig-Maximilians Universität, Theresienstraße 39, 80333 Munich, Germany. Emails: biagini@math.lmu.de

1. Further, we show that it is feasible to approximate the α-quantile hedging and thus also the superhedging price for $t = 0$ by neural networks. We extend our machine learning approach also to approximate the superhedging price process for $t > 0$. By the first step we obtain an approximation for the superhedging strategy on the whole interval up to maturity. By using the uniform Doob decomposition, see [13], we then only need to approximate the process of consumption B to generate the superhedging price process. We prove that B can be obtained as the the essential supremum over a set of neural networks. Finally, we present and discuss numerical results for the proposed neural network methods.

References

[1] Bartl, D., Kupper, M., and Neufeld, A. Pathwise superhedging on prediction sets. *Finance and Stochastics 24*, 1 (2020), 215–248.

[2] Bartl, D., Kupper, M., Prömel, D. J., and Tangpi, L. Duality for pathwise superhedging in continuous time. *Finance and Stochastics 23*, 3 (2019), 697–728.

[3] Becker, S., Cheridito, P., and Jentzen, A. Deep optimal stopping. *Journal of Machine Learning Research 20* (2019), 74.

[4] Buehler, H., Gonon, L., Teichmann, J., and Wood, B. Deep hedging. *Quantitative Finance 19*, 8 (2019), 1271–1291.

[5] Burzoni, M., Frittelli, M., Maggis, M., et al. Model-free superhedging duality. *Annals of Applied Probability 27*, 3 (2017), 1452–1477.

[6] Campi, L., and Schachermayer, W. A super-replication theorem in kabanov's model of transaction costs. *Finance and Stochastics 10*, 4 (2006), 579–596.

[7] Carassus, L., Gobet, E., and Temam, E. A class of financial products and models where super-replication prices are explicit. In *Stochastic Processes and Applications to Mathematical Finance*. World Scientific, 2007, pp. 67–84.

[8] Carassus, L., Obłój, J., and Wiesel, J. The robust superreplication problem: a dynamic approach. *SIAM Journal on Financial Mathematics 10*, 4 (2019), 907–941.

[9] Cuchiero, C., Khosrawi, W., and Teichmann, J. A generative adversarial network approach to calibration of local stochastic volatility models. *arXiv preprint arXiv:2005.02505* (2020).

[10] Cvitanić, J., and Karatzas, I. Hedging and portfolio optimization under transaction costs: A martingale approach 1 2. *Mathematical Finance 6*, 2 (1996), 133–165.

[11] El Karoui, N., and Quenez, M.-C. Dynamic programming and pricing of contingent claims in an incomplete market. *SIAM Journal on Control and Optimization 33*, 1 (1995), 29–66.

[12] Föllmer, H., and Leukert, P. Quantile hedging. *Finance and Stochastics 3* (1999), 251–273.

[13] Föllmer, H., and Schied, A. *Stochastic finance: an introduction in discrete time*, 4th rev. ed. ed. Walter de Gruyter, 2016.

[14] Kabanov, Y. M., and Last, G. Hedging under transaction costs in currency markets: a continuous-time model. *Mathematical Finance 12*, 1 (2002), 63–70.

[15] Kramkov, D. O. Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets. *Probability Theory and Related Fields 105*, 4 (1996), 459–479.

[16] Nutz, M. Robust superhedging with jumps and diffusion. *Stochastic Processes and their Applications 125*, 12 (2015), 4543–4555.

[17] Nutz, M., and Soner, H. M. Superhedging and dynamic risk measures under volatility uncertainty. *SIAM Journal on Control and Optimization 50*, 4 (2012), 2065–2089.

[18] Obłój, J., and Wiesel, J. Robust estimation of superhedging prices. *The Annals of Statistics 49*, 1 (2021), 508–530.

[19] Ruf, J., and Wang, W. Neural networks for option pricing and hedging: a literature review. *Journal of Computational Finance, Forthcoming* (2020).

[20] Schachermayer, W. The super-replication theorem under proportional transaction costs revisited. *Mathematics and Financial Economics 8*, 4 (2014), 383–398.

[21] Soner, H. M., Shreve, S. E., Cvitanic, J., et al. There is no nontrivial hedging portfolio for option pricing with transaction costs. *The Annals of Applied Probability 5*, 2 (1995), 327–355.

[22] Touzi, N. Martingale inequalities, optimal martingale transport, and robust superhedging. *ESAIM: Proceedings and Surveys 45* (2014), 32–47.

Non-linear Dependence and Portfolio Decisions over the Life-Cycle

Wei Jiang* Shize Li† Jialu Shen‡

March 13, 2022

Abstract

Using the Panel Study of Income Dynamics Survey, we reveal the non-linear dependence, between-squares correlation, between stock returns and earning risk exists. To understand how this non-linear dependence affects household life-cycle profile, we develop a life-cycle model that incorporates between-squares correlation and shows that this non-linear dependence can explain low participation rate and moderate risky asset shares. Empirical studies support the mode's predictions that households with higher between-squares correlations are less likely to participate in the stock market and lower their risky asset holdings conditional on participation.

JEL Classification: D31, D63, D91, E21, E32, G11.

Keywords: Countercyclical Earnings Risk, Non-linear Dependence, Life-cycle Portfolio Choice and Consumption.

*Hong Kong University of Science and Technology. E-mail: weijiang@ust.hk

†Hong Kong University of Science and Technology. E-mail: slidq@connect.ust.hk

‡Trulaske College of Business, University of Missouri. E-mail:jshen@missouri.edu

1 Introduction

Microeconomic data on household portfolios in the U.S. shows that less than 50% of households participate in the stock market and the participation rate tends to follow a life-cycle pattern (see, for example, Ameriks and Zeldes (2001), Faig and Shum (2002), Heaton and Lucas (2000), and Poterba et al. (2001)). Meanwhile, conditional on participation in asset markets, the average equity holdings is only 55

To address these two puzzles, a crucial element one needs to consider is the labor income and the risk associated with it.[1] Studies that examine how labor income affects households' portfolio decisions find mixed evidence and can be classified into two camps: 'bond-like' theory and 'stock-like' theory. 'Bond-like' theory suggests that labor income is like an extra endowment of the safe asset, implying that households should take an even higher position in the risky asset compared to models that ignore labor income.[2] On the other hand, 'stock-like' theory highlights the risk associated with labor income and shows that labor income consists of a stock component that reduces the needs of risky asset holding. For instance, Benzoni et al. (2007) shows that the young households' labor income becomes 'stock-like' because of cointegration between labor income and dividend. Storesletten et al. (2007) and Lynch and Tan (2011) show that larger variance of income shocks substantially reduces risky asset holdings.

More specifically, there are two main channels to understand labor income risk. Among the literature discussing labor income risk, several researchers use microdata to calibrate the individual labor income process (see, for example, Jagannathan et al. (1996), Davis and Willen (2000b), Campbell et al. (2001), Viceira (2001), Haliassos and Michaelides (2003), Cocco et al. (2005), Guvenen et al. (2014), Shen (2019), Catherine (2020)). Alternatively, some papers consider the dependence between labor income and stock returns. Campbell et al. (2001), Viceira (2001), Cocco et al. (2005) and Gomes and Michaelides (2005b) explore the effect of contemporaneous correlation between la-

[1] Besides the labor income, some other factors have been discussed widely contributing to household portfolio decisions, such as household preference (see, for example, Gomes (2005); Cao et al. (2005); Peijnenburg (2018); Pagel (2018)), participation costs (Vissing-Jørgensen (2002); Gomes and Michaelides (2005a)), peer effect (Hong et al. (2004)), housing (Cocco (2005); Yao and Zhang (2005)), borrowing constraints (Guiso et al. (1996); Haliassos and Michaelides (2003)).

[2] Merton (1975) introduces the risk-less labor income and suggests that households keep the share of financial wealth in stocks unchanged. Viceira (2001) shows that, as long as labor income shocks are uncorrelated with stock returns, households increase the demand for equity in the presence of risky labor income; Cocco et al. (2005) observe very low correlation and show that households almost invest all of their financial wealth in stocks.

bor income and stock market innovations. A number of studies allow for the long-run co-dependence between labor income and stock returns with the form of cointegration between these two processes (Campbell (1996), Baxter and Jermann (1997), Lucas and Zeldes (2006), Santos and Veronesi (2006), Benzoni et al. (2007), and Huggett and Kaplan (2011)). Lynch and Tan (2011) and Huggett and Kaplan (2016) consider labor income and stock returns under the context of the VAR framework.

However, the particular distributional assumption of the labor income made in the first channel still relies on the correlations between labor income and stock returns to some extent to address the low participation rates and life-cycle stock holdings conditional on participation while avoiding counter-factually high levels of risk aversion. Moreover, empirical support for correlation and cointegration is limited and mixed. Davis and Willen (2000b,a), Bonaparte et al. (2014), Catherine (2020) demonstrate relatively large correlation, suggesting a considerable part of income risk can be hedged using stocks, but Campbell et al. (2001), Vissing-Jørgensen (2002), Cocco et al. (2005) find extremely low and insignificant correlation. It is difficult to assess the magnitude for cointegration because of the data availability and thus provides little empirical support for cointegration.

In contrast to previous literature focusing on linear dependence, this paper aims to shed new light on the dependence between labor income and stock returns. In particular, this paper emphasizes the non-linear dependence and introduces the between-squares correlation, which is defined as the correlation between the squares of labor income shocks and stock returns. Between-squares correlations measure the likelihood that extreme values jointly occur in the labor income and stock market, focusing on the tail behavior. Such a specification is independent of any linear-dependence structure, such as correlation or cointegration, and thus is consistent with empirically observed low contemporaneous correlations and cointegration between market returns and labor income shocks. To the best of our knowledge, this is the first paper capturing non-linear dependence between labor income and stock returns with the form of the between-squares correlation.

The notion of non-linear dependence in the financial market is not new. Harvey and Siddique (2000) use monthly stock returns and show that co-skewness between an asset and the market portfolio can explain parts of the apparent non-systematic components in cross-sectional variation in expected returns. Cont (2001) tests for the existence of non-linear dependence by using the high-frequency stock returns. He finds the significant auto-correlation of the squared returns.[3] Our analysis differs significantly from theirs, as we introduce between-squares correlation to capture the non-linear dependence, and study the dependence structure between the labor market and financial market, instead of focusing on multiple assets within the financial market alone.

We investigate the participation decisions and optimal portfolio choices over the life-cycle for households who earn non-tradable labor income. The main innovations are twofold. First, our specification of non-linear dependence between labor income and stock returns is new. We document the existence of between-squares correlations are significant, while correlations are close to zero, using the Panel Study of Income Dynamics Survey (PSID). Second, we allow higher moments in both labor income shocks and stock returns with the mixture normal distributions. The assumption of non-linear dependence relies on the empirical evidence of higher-order moments in both labor income and stock returns. High-order moments in the labor income process are well documented and discussed in the literature (Guvenen et al. (2014), Shen (2019), and Catherine (2020)), while higher-order moments in the annual stock returns is neglected. The existence of non-normalities in daily stock returns has been recognized for at least 50 years, but higher moments of long-horizon returns are hard to measure accurately. Only a recent paper, Neuberger and Payne (2021), finds strong evidence of higher-order moments of annual stock return.

More specifically, using PSID and Ken-French Data from 1968 to 2017, we are able to document the heterogeneity in between-squares correlations. Between-squares correlations are significantly positive for both college graduates and households without college degrees and the cross-sectional distributions of the between-squares correlations are right-skewed, indicating significant fraction of households has extremely positive between-squares correlation between their income shocks and return shocks. On average, college graduates have higher between-squares correlation and the distribution of between-squares correlations are more skewed to right compared with their counterparts. We specify income shocks are between-squares correlated with stock return innovations and such a specification is consistent with empirically observed low contemporaneous correlations between market returns and changes to income flow.

We find that between-squares correlation is important in explaining low stock market

[3]Other papers that investigate the non-linear dependence include LeBaron (1988), Scheinkman and LeBaron (1989), Hsieh (1989, 1993), Edwards and Susmel (2001), Shepard (2010), Madam and Wang (2020).

participation rates and moderate equity holdings for stock market participants simultaneously. We first estimate the dynamics of labor income, annual stock returns and between-squares correlations for different groups of households conditional on their education: the observations without a college degrees and college graduates. Furthermore, we find that the stock returns and labor income shocks are highly between-squares correlated, even when the stock returns and labor income shocks are almost uncorrelated linearly (and thus most of the empirical evidence shows zero correlation between stock returns and labor income shocks). Then, we incorporate these processes into a standard life-cycle model, which not only takes into account the higher-order moments in both income process and stock returns but also the non-linear dependence between these two. The model matches the U.S. data well with a moderate coefficient of relative risk aversion $\gamma = 4.3$ and with small participation cost.

In addition, We find that households facing between-squares correlations hold less risky asset shares conditional on participation compared with those without between-squares correlations. Moreover, the wealth threshold of participation is mildly U-shaped with respect to age. The main drivers behind the U-shaped pattern of the participation threshold are the hump-shaped labor income process and the fact that young households seek to hold equity more aggressively than older households. We also observe between-squares correlations increase the wealth threshold across the ages and the increase is much more for college graduates.

Our results hold for reasonable levels of the agent's risk aversion coefficient. Qualitatively similar results obtain if we set $\gamma = 5$ and $\gamma = 3$. However, small differences in relative risk aversion can generate substantially difference in optimal risky asset holdings. This is consistent with empirical observation that asset holdings exhibit a high degree of heterogeneity. We also consider heterogeneity in elasticity of intertemporal substitution (EIS) and discount factor. EIS has limited impact on optimal risky asset holding across age groups and discount factor has moderate effect on risky asset holdings. Households with higher discount factor do not accumulate significant wealth and live from hand-to-mouth. Hence, they start holding risky assets with higher normalized cash-on-hand.

In support of our model's predictions, we examine whether a relation exists between the households' portfolio decisions (participation decisions and risky asset holdings) and between-squares correlations. We focus on households between 20 and 65 years old and control for a variety of household characteristics. First, we find that households with higher between-squares correlations are less likely to participate in the stock market. Second, households do adjust their portfolio holdings of risky assets when facing higher between-squares correlations, which is consistent with the model prediction. Quantitatively, a one-standard-deviation increase in between-squares correlations is associated with a 1.6% to 3.3% decline in equity shares.

Finally, we provide a portfolio selection perspective to understand households' risky asset holdings relative to the labor income. By doing so, we are able to observe more closely the non-linear pattern between the portfolio decisions and between-squares correlations. We consider labor income flows as a second risky asset, whose ratio to equity holdings is fixed (for example, 50%). We find that when between-squares correlations increase from zero to positive value, the portfolio skewness decreases slightly while kurtosis increases. On the other hand, when between-squares correlations drops below zero, the skewness of the portfolio decreases while kurtosis remains relatively stable. Therefore, the risk of the portfolio, measured by the sum of kurtosis and absolute value of skewness, increases whenever between-squares correlations deviate from zero. Considering that households can not change their labor income risk, in order to control the risk of the portfolio, households are less likely to enter the stock market and hold less risky assets conditional on participation when between-squares correlation is non-zero. These findings are also consistent with 'stock-like' theory. When between-squares correlations deviate from zero, labor income is more likely to serve as 'stock'.

The rest of the paper is organized as follows. Section 2 documents stylized facts regarding non-linear dependence, between-squares correlations. Section 3 introduces the model setup. Section 4 shows the calibration and Section 5 presents the quantitative analysis. Section 6 conduct empirical analysis. Section 7 provides a portfolio perspective of between-squares correlation, and section 8 concludes.

2 Stylized facts on nonlinear dependence

In this paper, we introduce the between-squares correlation (henceforth referred to as BS-Corr) to capture the nonlinear dependence between stock market returns and labor income shocks. The BS-Corr is defined as the Pearson correlation between the demeaned squares of two time series.

$$\mathrm{Corr}^{\mathrm{bs}}(X,Y) := \mathrm{Corr}\left((X-\mathbb{E}X)^2, (Y-\mathbb{E}Y)^2\right), \tag{1}$$

where $\mathrm{Corr}(X,Y)$ denotes the Pearson correlation between X and Y. BS-Corr provides a good measure of nonlinear dependence, which can be understood from two perspectives. First, BS-Corr is closely related to co-kurtosis. Specifically, the level of co-kurtosis can be represented by a linear function of BS-Corr with a multiplier depending on kurtosis.[4] BS-Corr can be considered as a normalized co-kurtosis and the advantage of BS-Corr as a measure is that it is within the range of $[-1,1]$, while co-kurtosis depends on the kurtosis of each variable. Second, two random variables with a high level of BS-Corr will tend to undergo extreme positive and negative deviations at the same time. Since BS-Corr is the correlation between squares, it will be more sensitive to large deviations from mean and thus can reveal such extreme co-movements.

Under normality assumption, BS-Corr is fully determined by the linear correlation. Specifically, we have:

$$\mathrm{Corr}^{\mathrm{bs}}(X,Y) = \mathrm{Corr}^2(X,Y), \text{where } (X,Y) \text{ is a normal vector.}$$

Therefore, two normally-distributed variables can not generate large BS-Corr, while keeps a low value of correlation. This is a common issue in many papers that discuss life-cycle models adopt normally-distributed shocks for earning shocks and stock return innovations. Opposite to these papers, this paper assumes a normal mixture for both labor income shocks and stock return innovations, which provides more freedom in dependence structures. Nikoloulopoulos (2021) studied the K-finite normal mixture copula and prove that the normal mixture is very flexible when handling various dependence patterns.

In addition, the analysis of real data also provide evidence of a strong BS-Corr between labor income shocks and stock return innovations. We use the time series of labor income shock from Panel Study of Income Dynamics Survey (PSID) and excess market returns from Ken French's data library to estimate the BS-Corr.[5] PSID was conducted annually

[4]One definition of co-kurtosis is $K(X,Y) = \dfrac{\mathbb{E}[(X-\mathbb{E}[X])^2(Y-\mathbb{E}[Y])^2]}{\sigma_X^2\sigma_Y^2}$, and the relation of it with BS-corr is $K(X,Y) = \sqrt{(\mathrm{Kurt}[X]-1)(\mathrm{Kurt}[Y]-1)}\mathrm{Corr}^{\mathrm{bs}}(X,Y)$.

[5]See `https://simba.isr.umich.edu/data/data.aspx` and `mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`; See equation (5) for the strict definition of labor income

from 1968 to 1997, but has been conducted biennially since 1997. Therefore, to make the best use of data after 1997, we consider two time-series: the annual data from 1968 to 1997 and the biennial data from 1968 to 2017. For the biennial series, data after 1997 is directly taken from PSID, and data between 1970 and 1997 is constructed from annual data. To get a more efficient estimator, when computing the correlations, we only consider households with a minimum of twenty waves of valid labor income data.[6]

Figure 1: BS-Corr distribution

This graph plots the histogram and kernel density of BS-Corr estimations between labor income shock and excess stock return. Labor income and stock return data are taken from PSID and Ken French's data library. The annual data is from 1970 to 1997 and the biennial data is from 1970 to 2017. The sample is restricted to households with a minimum of twenty waves of valid labor income data.

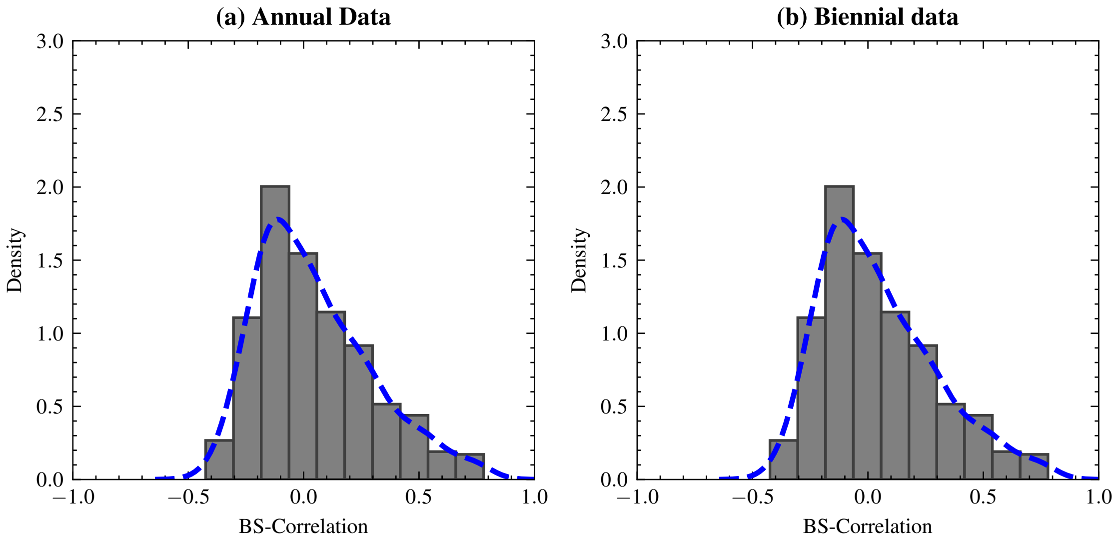

Figure 1 displays the histogram and kernel density of estimations over the full sample with different frequencies, which clearly reveals the high density of positive BS-Corr. The long right tail also shows that many households earn income which is highly exposed to stock market in the sense of BS-Corr. Although there are a few samples with negative BS-Corr, we notice that these negative estimations concentrate close to zero, which suggests that it is reasonable to assume a positive BS-Corr in the model setup. Moreover, this assumption is also consistent with the further analysis of mean and percentages as Table 1 suggests.

shock $\delta_{i,t}$. Here we estimate BS-Cor between excess stock return and the d-order difference of $\delta_{i,t}$ ($d = 1$ for annual data, $d = 2$ for biennial data). According to Ken French's data library, the excess return on the market is defined as value-weight stock return minus the Treasury bill rate.

[6]We also apply the data select principle from Nakajima and Smirnyagin (2019) to PSID data. See Appendix A for details.

Table 1: Summary Statistics of BS-corr

This table reports the mean, different percentages, and some fractions over different samples, based on the data from PSID and Ken French's data library. Panel A concludes results over the full sample. To observe BS-Corr levels of different educational groups, the sample was split in two groups: one of the households without college education, and the other of college graduates. Panel B and C show the group results. Column 'Annual data' (Biennial data) means the statistics below are based on annual (Biennial) series in $[1970, 1997]$ ($[1970, 2017]$). "Fraction $> k$" is the fraction of samples expressed in percent with BS-Corr more than k.

	Panel A: Full Sample		*Panel B: No College*		*Panel C: College*	
	Annual data	Biennial data	Annual data	Biennial data	Annual data	Biennial data
Mean	0.044	0.041	0.034	0.032	0.070	0.060
SD	0.247	0.233	0.241	0.224	0.259	0.250
p value	$< 0.2\%$	$< 0.1\%$	0.014	$< 1\%$	$< 0.3\%$	$< 0.2\%$
10th	-0.232	-0.208	-0.236	-0.214	-0.211	-0.186
25th	-0.142	-0.140	-0.152	-0.141	-0.115	-0.135
Median	0.006	0.001	0.007	0.000	0.005	0.004
75th	0.200	0.179	0.177	0.150	0.244	0.206
90th	0.400	0.378	0.357	0.366	0.433	0.406
Fraction > 0	51.26%	50.29%	51.63%	50.00%	50.39%	50.89%
Fraction > 0.1	34.71%	34.24%	34.64%	33.33%	34.88%	36.09%
Fraction > 0.2	24.83%	23.60%	22.55%	22.41%	30.23%	26.04%
No. of obs	435	517	306	348	129	169

Table 1 presents summary statistics of BS-Corr. Results from data with different frequencies both support a significant and positive level of BS-Corr. Panel A show that, over the full sample, the averages of BS-Corr are both positive (0.044 and 0.041) and significantly non-zero with p-values$< 0.2\%$. Although the medians are closed to zero, there remains substantial variability in the values of BS-Corr. Panel A also provides the inter-quartile range of BS-Corr for both data frequencies over the full sample. We find that the 75th and 90th percentages are substantially positive, with 0.2 and 0.4 respectively. Further, we calculate the fraction of estimations with a positive BS-Corr or a substantially positive BS-Corr. We find that more than half of the estimations are positive and about 1/3 feature a substantially positive value of BS-corr no less than 0.1. More remarkably, over 1/5 of estimations exceeds 0.2. Thus, we assume a positive BS-Corr and use the sample mean as our calibration target, when we set up the model.

Following Cocco et al. (2005), to control the effect of education, we split the samples into two groups: samples without a college degree and another of college graduates. Panel B and Panel C in Table 1 report the statistics of the two groups. The results are quit similar as that over the full sample. However, almost all statistics of no college group are somewhat lower than those for college group. Especially, the average of BS-Corr over college group is about twice the value for no college group. In section 4, we do separate calibrations for both groups and analysis the difference between corresponding decision rules.

3 Model

In this section, we develop a discrete time life-cycle model. Each period, a household earns labor income and decides how to allocate wealth to consumption and financial assets. The main novelty of our model is introducing BS-Corr and higher moments with mixed normally distributed shocks.

3.1 Preferences

Let t be the age of adult, generally taking $t = 0$ as age 20 and assume a maximum $T = 81$, which means adults live up to 100. At the first $K = 46$ years, adults work if alive. Let p_t denote the probability that an adult is alive at age t conditional on being alive at age $t-1$. The data of p_t is taken from the mortality tables of the National Center for Health Statistics.

Households have Epstein-Zin (1989) preferences, a recursive preference of which the elasticity of intertemporal substitution is separated from the relative risk aversion. For household i, let $X_{i,t}$ denote the total wealth at the beginning of age t, and the utility function is given by:

$$U_{i,t} = \left\{(1-\beta)C_{i,t}^{1-1/\psi} + \beta\left(E_t\left[p_{t+1}U_{i,t+1}^{1-\gamma} + b(1-p_{t+1})X_{i,t+1}^{1-\gamma}\right]\right)^{\frac{1-1/\psi}{1-\gamma}}\right\}^{\frac{1}{1-1/\psi}}, \tag{2}$$

where β is the discount factor, b determines the strength of bequest motive, γ is the coefficient of relative risk aversion and ψ is the elasticity of intertemporal substitution. For simplicity, we assume that the utility function applied to the bequest is same as that applied to the household's own consumption when alive and b controls the effect of bequest motive.

3.2 Financial Assets Return

We assume the financial market assets consists of two financial assets in which the households can invest, one risk-less and one risky. The risk-less asset has a constant gross return R_f, and the return of the risky asset is given by

$$R^S_{t+1} = R_f + \mu + \eta_{t+1}, \tag{3}$$

where η_{t+1} is the shock to returns and is independently and identically distributed as:

$$\eta_t \sim \begin{cases} N(\mu_{\eta,1}, \sigma_{\eta,1}) & \text{with prob. } p_\eta \\ N(\mu_{\eta,2}, \sigma_{\eta,2}) & \text{with prob. } 1 - p_\eta \end{cases} \tag{4}$$

Two main reasons drive us to assume a mixed normal distribution for the stock return shock. First, as argued before, mixed normal assumption allows us to deal with more complicated dependence structure, such as the BS-Corr between stock return and labor income introduced in our model. Under normality, BS-Corr is the square of Corr, while in section 2, the estimations from data show that BS-Corr is similar to Corr in magnitude. The assumption of mixed normal shocks is more flexible with BS-Corr. Second, mixed normal assumption enables us to incorporate the higher moments in the stock return innovations. Higher moments in the high-frequency returns data have been discussed and considered to be important for many years, but the long-horizon higher moments are hard to measure accurately using standard techniques. Neuberger and Payne (2021) shows that short-horizon returns can be used to estimate the higher moments of long-horizon returns, and their empirical results identify high level of negative skewness and excess kurtosis of annual U.S. equity market returns. Thus, we introduce mixed normal shocks instead of normal shocks to capture such moments.

3.3 Labor Income Process

During the working period, at age t, household i's labor income $Y_{i,t}$ is given by

$$\log Y_{i,t} = f(t, Z_{i,t}) + \delta_{i,t}, \tag{5}$$

where $f(t, Z_{i,t})$ is a deterministic function of age t and of a vector of other individual characteristics $Z_{i,t}$ and $\delta_{i,t}$ is the labor income shock. We further decompose the labor income shock $\delta_{i,t}$ into a persistent shock $\nu_{i,t}$ and a transient shock $\epsilon i, t$:

$$\delta_{i,t} = \nu_{i,t} + \epsilon_{i,t}, \tag{6}$$

We assume that $\epsilon_{i,t}$ is normally distributed as $N(0, \sigma_\epsilon^2)$ which is independent of permanent shock $\nu_{i,t}$ and stock return shock $\eta_{i,t}$. Labor income is related to stock return through the persistent shock $\nu_{i,t}$ given by:

$$\nu_{i,t} = \nu_{i,t-1} + u_{i,t}, \tag{7}$$

where $u_{i,t}$ follows a mixed normal distribution:

$$u_{i,t} = \begin{cases} u_{i,t}^{(1)} \sim N(\mu_{u,1}, \sigma_{u,1}^2) & \text{with prob. } p_u, \\ u_{i,t}^{(2)} \sim N(\mu_{u,2}, \sigma_{u,2}^2) & \text{with prob. } 1 - p_u. \end{cases} \tag{8}$$

We assume $u_{i,t}$ is correlated with the stock return shock $\eta_{i,t}$ and the dependence structure is captured by both Pearson correlation and BS-Corr. We define the correlations between components of $u_{i,t}$ and $\eta_{i,t}$ as:

$$\rho_{a,b} = \mathrm{Corr}(u_{i,t}^{(a)}, \eta_t^{(b)}), a = 1, 2, b = 1, 2. \tag{9}$$

The four correlations $\rho_{a,b}$ control the value of Corr and BS-Corr between $u_{i,t}$ and $\eta_{i,t}$.[7]

Income during retirement is assumed to be exogenous and deterministic, with all households retiring in time period K, corresponding to the retirement age 65. Income during retirement follows

$$Y_{i,t} = \lambda P_{iK}, \quad t > K, \tag{10}$$

where is a constant fraction λ of the permanent component of labor income in the last working period.

[7] Other parameters also affect Corr and BS-Corr. See Appendix B for the explicit forms.

3.4 Wealth Accumulation

We define cash-on-hand $X_{i,t}$ as the liquid wealth for consumption and savings. At each period t, households start with accumulated financial wealth and receive labor income. They consume C_t and invest the rest on a portfolio consisting of α_t share of risky assets and $1-\alpha_t$ of risk-less assets. The participation decision is represented by a dummy variable I_P, which is 1 when the household decides to invest on the risky assets and 0 otherwise. If investing in risky assets in a period, the households must pay a fixed cost F, which represents mostly the cost of acquiring information about the stock market and the transaction fee. Then, the process of cash-on-hand $X_{i,t}$ of a household satisfies:

$$X_{i,t+1} = (X_{i,t} - C_{i,t})R^p_{i,t+1} - FI_pP_{i,t} + Y_{i,t+1}, \tag{11}$$

where $R^p_{i,t+1}$ is the portfolio return given by

$$R^p_{i,t+1} = \alpha_{i,t}R^S_{t+1} + (1-\alpha_{i,t})R_f, \tag{12}$$

Finally, we restrict borrowing from risk-less assets or future labor income, and short-sales on risky assets. Specifically, households face the following constraints:

$$0 \leq \alpha_{i,t} \leq 1. \tag{13}$$

$$0 < C_{i,t} \leq X_{i,t}. \tag{14}$$

3.5 Household Optimization Problem

In each period t, households decide their consumption and risky shares $(C_{i,t}, \alpha_{i,t})$ based on cash-on-hand $X_{i,t}$ to maximize the expected utility. The optimization problem can be stated as:

$$V_{i,t} = \max_{\{\alpha_{i,u}\}_{u=t}^{T}, \{C_{i,u}\}_{u=t}^{T}} \mathbb{E}_t(U_{i,t}), \tag{15}$$

where $U_{i,t}$ is defined in equation (2) and is subject to the constraints given by equations (3) to (14).

Analytical solutions do not exist for this problem, and thus we develop a numerical solution method. We first simplify the problem based a scale-independent property. Since the permanent labor income shock $\nu_{i,t}$ is a unit-root process, we can normalize variables by the permanent labor income defined as $P_{i,t} = f(t, Z_{i,t}) + \nu_{i,t}$, and reduce the number of state variables. Specifically, let $x_{i,t} = \frac{X_{i,t}}{P_{i,t}}$ and $c_{i,t} = \frac{C_{i,t}}{P_{i,t}}$ be the normalized cash on hand and consumption. The normalized value function $v_{i,t} = \frac{V_{i,t}}{P_{i,t}}$ is

$$v_{i,t}(x_{i,t}) = \max_{\substack{\{\alpha_{i,u}\}_{u=t}^{T} \\ \{c_{i,u}\}_{u=t}^{T}}} \left\{ \begin{array}{c} (1-\beta)c_{i,t}^{1-\frac{1}{\psi}} \\ +\beta \left(\mathbb{E}_t \left(\frac{P_{i,t+1}}{P_{i,t}} \right)^{1-\gamma} \left[p_{t+1} v_{i,t+1}^{1-\gamma} + (1-p_{t+1}) b x_{i,t+1}^{1-\gamma} \right] \right)^{\frac{1-1/\psi}{1-\gamma}} \end{array} \right\}^{\frac{1}{1-1/\psi}}, \tag{16}$$

subject to

$$x_{i,t+1} = \begin{cases} (x_{i,t} - c_{i,t}) r_{i,t+1}^{p} \frac{P_{i,t}}{P_{i,t+1}} - FI_p + e^{\varepsilon_{i,t+1}} & \text{for } t \leq K \ , \\ (x_{i,t} - c_{i,t}) r_{i,t+1}^{p} \frac{P_{i,t}}{P_{i,t+1}} + \lambda & \text{for } t > K \ . \end{cases} \tag{17}$$

We use a numerical method to recursively solve the optimization problem above. In the last period, households predict a certain death and therefore the policy functions are determined by the bequest motive. We start from this terminal condition and then iterate backwards. Appendix A presents the details of the numerical solution method.

4 Calibration

The idea of our calibration is to reveal the empirical life-cycle decision rules with capturing the risk from higher order moments and nonlinear dependence. To do so, we categorize the parameters controlling the life-cycle model in two groups. First group includes the parameters shaping the dynamics of labor income and stock return. We construct panel data from PSID and Ken French's data library, and use GMM to match the empirical moments. Further, given dynamics of state variables, second group includes the parameters of preference determining the decision rules of household's optimization problem, and we do estimation through SMM targeting at the average life-cycle profiles from SCF data.

4.1 Labor Income and Stock Return

Since a more complicated dependence structure is applied, we estimate the parameters of labor income and stock return together with GMM. To get observations of labor income shock, we need to estimate the family-specific fixed effects $f(t, Z_{i,t})$ of labor income in equation (5). Following the method of Cocco et al. (2005), we use panel data of PSID from 1970 to 2017 and run a regression analysis to estimate $f(t, Z_{i,t})$.[8] The residuals from the regression can be considered as realizations of labor income shock $\delta_{i,t}$. Then we take the annual data of excess stock market return from Ken French's data library between 1970 and 2017, of which the demeaned series can be considered as realizations of stock return shock η_t. We use these data to calibrate parameters of labor income and stock return.

We construct GMM targeting[9](1) mean, variance, skewness, and kurtosis of stock return shock η_t; (2) mean and variance, skewness, kurtosis of $\delta_{i,t}$; (3) Corr and BS-Corr between $\delta_{i,t} - \delta_{i,t-1}$ and η_t. For moments in (1), Neuberger and Payne (2021) propose a new method using short-horizon stock returns to estimate long-horizon moments, and we use their results of annual log return with the standard deviation, skewness and excess kurtosis 0.216, -1.41 and 5.62 respectively. For moments in (2), following Nakajima and Smirnyagin (2019), we target the moments of different age groups. We consider four age groups indexed by $\{25, 35, 45, 55\}$ and each group contains ages ± 5 years.[10] Finally, for (3), we use the average BS-Corr in Table 1 and average Corr estimated through the same methodology.

In total, we calibrate five parameters $\{p_\eta, \mu_\eta^1, \mu_\eta^2, \sigma_\eta^1, \sigma_\eta^2\}$ that control stock return dynamics, seven parameters $\{p_u, \mu_u^1, \mu_u^2, \sigma_u^1, \sigma_u^2, k, \sigma_\epsilon\}$ that control labor income dynamics, and four parameters $\{\rho_{11}, \rho_{12}, \rho_{21}, \rho_{22}\}$ that determine the dependence structure. For simplicity, we assume that $\rho_{11} = \rho_{22}$ and $\rho_{12} = \rho_{21}$. In section 2, we divided the samples to two education groups (no college group and college group) and identifies different BS-Cor estimations. For consistency, we also do calibrations for both groups. Overall, we apply GMM with 19 moments on 14 parameters. Panel A and B in Table 2 report the calibration results.

[8]See Table 7 for variable definitions and Appendix A for data selection principle and other details.

[9]We use 3rd and 4th order moments in GMM, For simplicity, we refer to such higher order moments as skewness and kurtosis.

[10]Group $h = 25$ contains ages 22-35, group $h = 35$ contains ages 35-45, group $h = 45$ encompasses ages 45-55, and group $h = 55$ aggregates the remaining ages 55-65.

For those parameters unrelated to shocks, we set their values following Cocco et al. (2005). The risk-less rate $R_f - 1$ is set to 2% and the mean equity premium μ is 4%. Replacement rate λ is calibrated as the ratio of the average of labor income for retirees in a given education group to the average of labor income in the last working year prior to retirement, which is 0.903 for no college group and 0.945 for college group.

4.2 Preference and Bequest Motive

We calibrate the preference parameters and fixed cost rate to match the average participation rate, risky asset share and normalized wealth for different age groups using SMM. These moments summarize the life-cycle property of decision rules. The empirical moments are estimated from the triennial SCF dataset from 2007 to 2019, which provides a summary of households' portfolio decisions and has been used widely in the literature.

In addition, to capture the different behaviors between different educational households, we calibrate parameters for the two educational groups separately, except that relative risk aversion γ and fixed cost rate F are assumed to be same between groups. To generate the mixed behaviors, we use 0.3 as the weight of college group, which is taken from the PSID.

We calculate the average participation rate, risky asset share and normalized wealthy for 15 age groups in $[20, 65]$ from SCF data and obtain 45 moments in total.[11] The SMM seeks parameters that minimize

$$(\hat{m} - m)' W (\hat{m} - m), \tag{18}$$

where $\hat{m}$ are the simulated moments, m are the targets, and W is the inverse of the covariance matrix of the empirical moments, which is estimated by bootstrapping the true data.

4.3 Results

In this section, we summarize results of the two-step calibration. Table 2 reports the parameters of shocks. The two components of the stock return shock are normally dis-

[11]We construct non-overlapping age groups every 3 ages, e.g. the first group includes samples in age $[20, 22]$ and the second in age $[23, 25]$, except the last group including samples with ages in $[62, 65]$.

tributed as $N(-0.187, 0.395)$ and $N(0.038, 0.127)$, respectively. The first component pre-

Table 2: Calibrated parameters

This table reports the calibrated parameters of the life-cycle model. We use data from PSID and Ken French's data library and construct a two-step calibration. Panel A reports the parameters of the mixed normal shock controlling stock return process. Panel B reports the parameters of persistent shock $\nu_{i,t}$ (mixed normal distributed) and transient shock $\epsilon_{i,t}$ (normal distributed) controlling the labor income process for different groups. The two correlation parameters are correlations between components of stock return shock η_t and persistent labor income shock $\nu_{i,t}$. Panel C reports the preference parameters and fixed cost for different groups. We assume the two groups have same risk aversion and fixed cost.

Panel A: Stock return shock η_t			
Mixture weight	p_η	0.168	
Mean of 1-st component	$\mu_{\eta,1}$	-0.187	
Mean of 2-st component	$\mu_{\eta,2}$	0.038	
SD of 1-st component	$\sigma_{\eta,1}$	0.395	
SD of 2-st component	$\sigma_{\eta,2}$	0.127	
Panel B: Labor income shocks $\nu_{i,t}$ and $\epsilon_{i,t}$			
		No College	College
Mixture weight of $\nu_{i,t}$	p_ν	0.271	0.278
Mean of 1-st component of $\nu_{i,t}$	$\mu_{u,1}$	-0.124	-0.156
Mean of 2-st component of $\nu_{i,t}$	$\mu_{u,2}$	0.046	0.060
SD of 1-st component of $\nu_{i,t}$	$\sigma_{u,1}$	0.172	0.231
SD of 2-st component of $\nu_{i,t}$	$\sigma_{u,2}$	0.010	0.012
SD of transient shock $\epsilon_{i,t}$	σ_ϵ	0.204	0.139
Panel C: Dependence parameters			
		No College	College
Dependence parameter 1	ρ_1	0.836	0.778
Dependence parameter 2	ρ_2	-0.164	-0.214
Panel D: Preferences and fixed cost			
		No College	College
Relative risk aversion	γ	4.3	4.3
EIS	ϕ	0.9	0.3
Discount factor	β	0.90	0.98
Bequest motive	b	2.5	2.5
Fixed cost rate	F	0.008	0.008

dicts a large average loss and great volatility, which can be considered to reflect the market crash case. And the second component corresponds to the normal case. The mixture probability 0.168 is the probability of stock return shock to take the first component, and then can be understood as possibility of crashes. Such structure of stock return shock has been studied by some literatures of life-cycle model[12] and can be explained by cycle

[12] See Fagereng et al. (2017), Shen (2019) and Catherine (2020).

property or tail risk. The shock to labor income also displays a similar structure as stock return shock, and capture the higher order risk in labor income. In addition, no college group faces much more transient risk than college group, with a standard deviation 0.204 compared with 0.139 of college group. This can be explained by the fact that households with higher education are more likely to have a stable job.

The two dependence parameters determine the dependence structure, and the positive value indicates that labor income and stock market will tend to move in a same direction while the negative value implies the opposite. Although the two dependent parameters both take relatively large values (0. for positive ones and -0.2 for negative ones), the mixed dependence are not so significant. The corresponding Corr and BS-Corr are $0.046, 0.069$ for college group and 0.033, 0.038 for no college group. The results of Corr are consistent with Cocco et al. (2005), which claimed a Corr closed to zero. The values of BS-Corr get a same order of magnitude as Corr, but it is enough to generate an obvious difference in policy function, as we will show later.

Table 3 reports the fitness of calibration and shows that our model matches the data well. We obtain relative errors less than 20% for all moments except one moment, and more than half of errors are less than 5%.[13]

Panel D in Table 2 shows the calibration results of preference parameters and fixed cost rate, and Figure 2 shows the fitness. Our model matches data well with a moderate risk aversion 4.3 and low fixed cost rate 0.008 (0.8% of the household's expected annual income). Compared to the level of risk aversion used in previous literatures (over 10 in Fagereng et al. (2017), around six in Shen (2019) and Catherine (2020)), we provide a possibility to fit empirical data, especially the portfolio choice, with a lower but reasonable level of risk aversion. Taking account of BS-Corr, we actually increase the resource of risk exposure of labor income. Thus, labor income will be more stock-like and households will balance stock holdings to control the whole risk. The consideration of nonlinear dependence gives a deeper understanding of labor income risk and avoid assuming a high risk aversion, which may cause difficulties in explaining other economic behaviors.

To identify the effect of BS-Corr, we also simulate the age profiles of a model without BS-Corr for comparison. To do so, we fixed all parameters in Table 2 except the two dependence parameters. We adjust ρ_1 and ρ_2 to obtain dependence structure with the

[13]The exception is 3rd moment of age group $[25, 35]$. The relative errors are 26.47% for no college group and 25% for college group.

Table 3: Moments Fitness of Labor Income and Stock Return

This table reports relative errors of the moments targeted in the calibration of labor income and stock return. Panel A, B reports the mean, variance and higher moments of stock return shock and labor income shock. Panel C reports the error in dependence structure including Corr and BS-Corr. Moments from model are computed with calibrated parameters under mixed normal distribution. Empirical moments are from PSID and Ken French's data library.

Panel A: Moments of Stock return shock

	Model	data
SD	0.216	0.216
3rd moment	−1.42	−1.41
4th moment	5.46	5.62

Panel B: Moments of Labor income shock

		No College		College	
	Age group	Model	data	Model	data
SD	$[25,35]$	0.378	0.408	0.385	0.405
	$[35,45]$	0.428	0.432	0.412	0.395
	$[45,55]$	0.449	0.443	0.417	0.392
	$[55,65]$	0.457	0.464	0.417	0.414
3rd moment	$[25,35]$	−0.021	−0.017	−0.032	−0.026
	$[35,45]$	−0.027	−0.032	−0.034	−0.036
	$[45,55]$	−0.027	−0.032	−0.035	−0.040
	$[55,65]$	−0.029	−0.033	−0.035	−0.044
4th moment	$[25,35]$	0.066	0.081	0.076	0.086
	$[35,45]$	0.107	0.106	0.097	0.094
	$[45,55]$	0.128	0.116	0.100	0.093
	$[55,65]$	0.138	0.130	0.101	0.105

Panel C: Dependence structure

	No College		College	
	Model	data	Model	data
Corr	0.038	0.038	0.033	0.034
BS-Corr	0.046	0.046	0.069	0.070

same Corr as our main model but a zero BS-Corr. Specifically, the target of college (no college) group in the new model are 0.046 (0.038) for Corr and 0 for BS-Corr. Figure 2 shows the results. Without BS-Corr, the model generate a much more imprudent portfolio choices. Participation rate increases by 10% to 20% in different age groups, and conditional risky share increases by 20% overall. The model without BS-Corr fail to fit the data with a moderate risk aversion and low fixed cost, though we only change BS-Corr from 0.069 (0.033 for no college group) to 0.

Figure 2: Calibration results

The three graphs show the comparison of moments from main model, main model without BS-Cor and SCF data. The parameters for main model is from Table 2. For main model without BS-Cor, we only adjust the two correlation parameters to decrease BS-Cor to zero. We use triennial SCF data from 2007 to 2019 to calculate empirical moments. Moments are calculated by non-overlapping age groups from 20 to 65. Each age group contains samples of three successive ages. Figure (a), (b) and (c) show the results of participation rate, conditional risky share and normalized wealth, respectively.

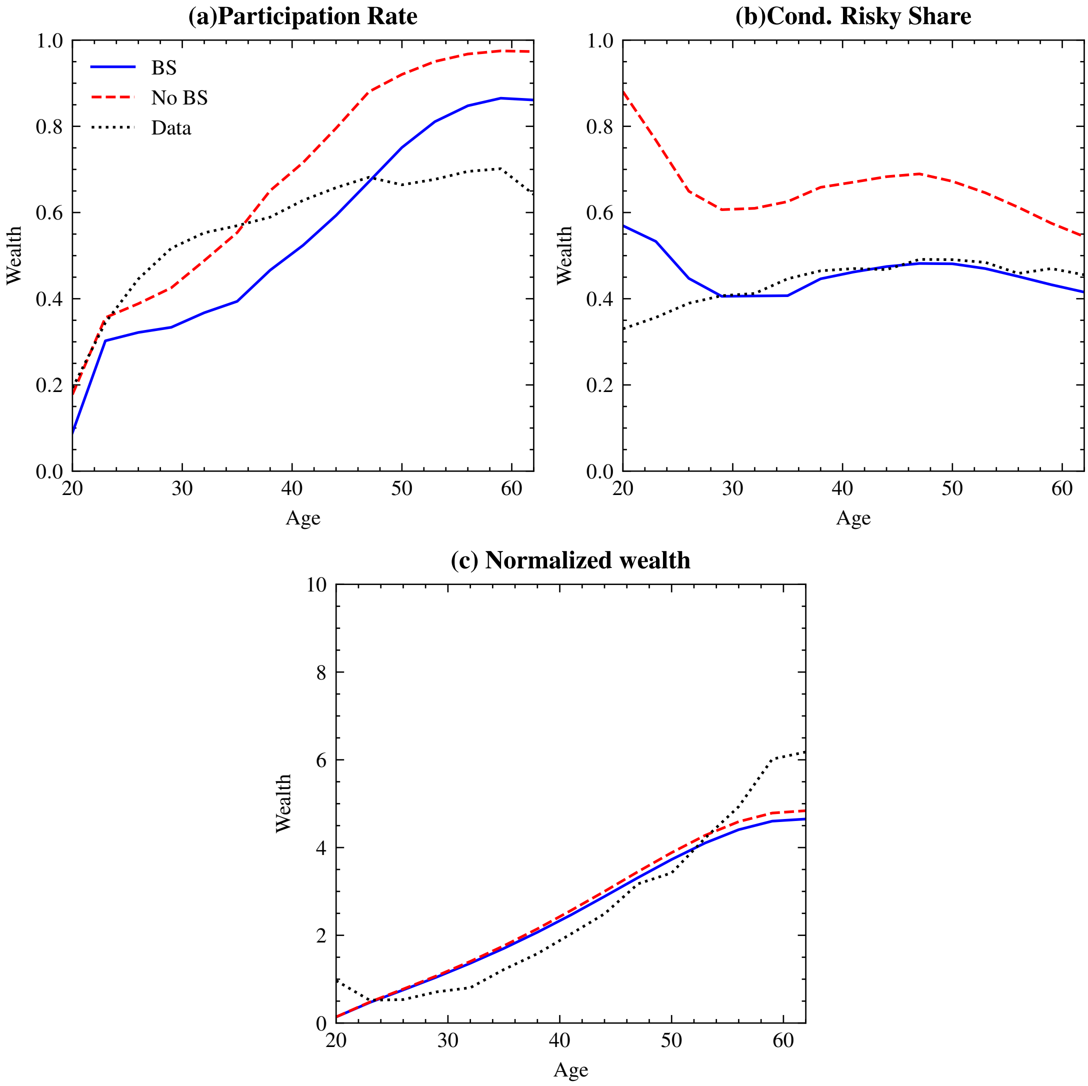

In addition, the estimations of discount rate, EIS and bequest motive also remain in the reasonable range. Further, the difference between groups provides interesting findings. No college group is more impatient than college group, with lower discount rate (0.90 vs 0.98), which means allocating more wealth to consumption. Thus, no college group accumulates less wealth and has weaker incentive to pay the fixed cost, leading to a low participation rate.[14]

5 Quantitative Analysis

As a measure of nonlinear dependence, BS-Corr provides a new resource of risk exposure for households and makes labor income more stock-like. To further understand the effect of BS-Corr on asset allocation and wealth accumulation, we discuss main quantitative results from our life-cycle model with calibrated parameters. In addition, we also compare results from models with and without BS-Corr. As stated in section 4.3, the model without BS-Corr is different from the main model only in dependence parameters, which generate a zero BS-Corr but keep Corr the same. Taking zero BS-Corr, the model is very similar to that in Gomes and Michaelides (2005b), with extra consideration of higher order moments of stock return.

5.1 Policy Function

Figure 3 plots the portfolio rules at different ages for both no college and college groups. Generally, considering a small BS-Corr, optimal risky share is lowered significantly, with the largest drop 28% from 87% conditional on participation for college group.[15] Further, the optimal portfolio choice depends on two state variables: age and cash-on-hand. For cash-on-hand, Figure 3 shows a decreasing policy function conditional on participation across all age groups, which is consistent with literatures (for example, Cocco et al. (2005)). Such pattern is from a substitution effect of labor income on risk-less assets. Households will evaluate the risk-free position from labor income and balance the portfolio choice. The effect is weaker for wealthier households, since labor income becomes trivial compared to wealth. Thus, households with less wealth will predict more position of risk-less assets through labor income and take more aggressive portfolio choice. However, labor

[14] Although no college group also get a higher EIS, discount rate will dominate in effect on consumption. See figure in section 5.1.

[15] For no college group, the largest drop is 22% from 100%.

Figure 3: Equity Share Policy Function

The six figures graphs plot the equity share policy functions for different ages and groups. The first column corresponds college group and the second corresponds no college group. The policy functions are solved with calibrated parameters.

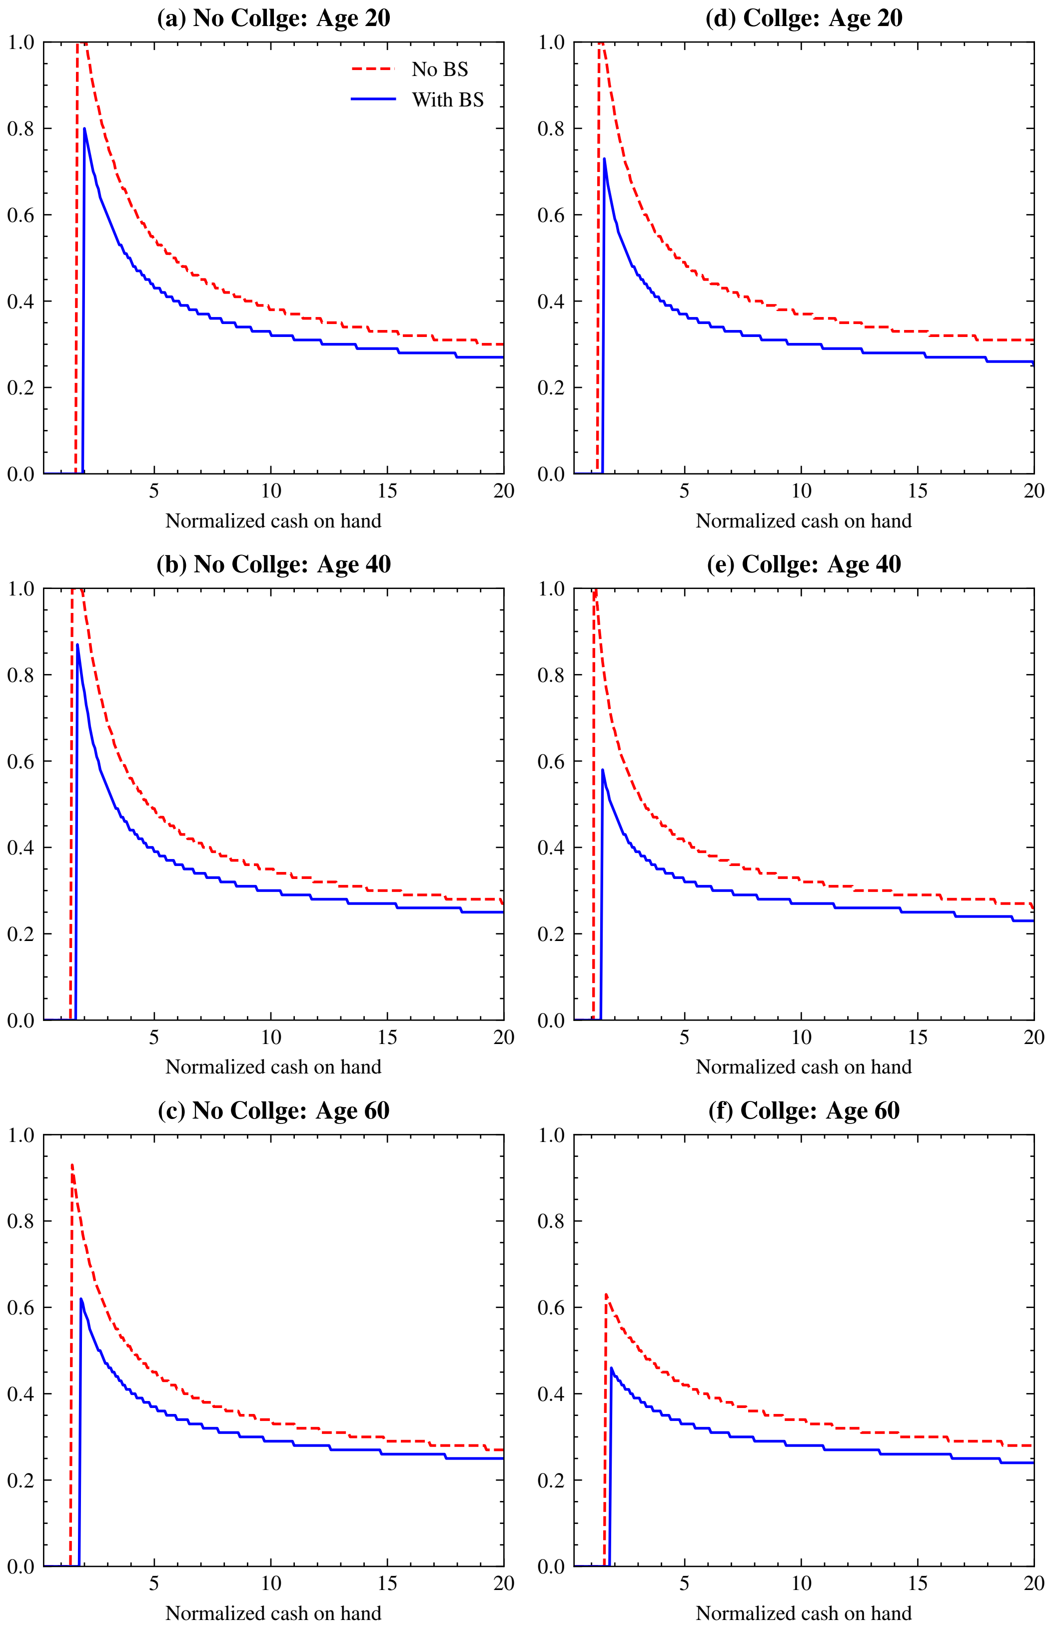

income is not a complete substitution of risk-less assets, since it has risks from different shocks. As stated in Benzoni et al. (2007), labor income also has stock-like features. When considering BS-Corr, the nonlinear dependence on stock return will make labor income more stock-like. Although the decreased property remains, the model with BS-Corr reduces the descent speed with wealth, which can be considered a heterogeneous effect on the rich and poor. To illustrate the heterogeneity, we calculate the difference between the policy functions of models with and without BS-Corr. Table 4 reports the changes in optimal risky share with BS-Corr, and shows that the change is much more obvious with less wealth. At age 20, the most significant decrease in optimal risky share (achieved at wealth ratio 2.15) is 26%, while the value is less than 7% for the wealthiest. The heterogeneity is reasonable since households with less wealth will attach more importance to the value of the labor income stream, and extra risk exposure from BS-Corr will be considered a larger position of risky assets.

Table 4: Change in optimal risky share with BS-Corr

This table reports the difference in optimal risky share α_t between models with and without BS-Corr. The data is from the policy functions discussed in section 5.1. For simplicity, we only report results for college group here. We calculate the difference for each age and wealth ratio. For instance, the second percentage −18% in row '20' is calculated as the optimal share 67% at wealth ratio = 3.0 from model with BS-Corr minus 93% from model without BS-Corr. We also report the largest difference for each age.

Age	Largest	Wealth Ratio			
		3.0	5.0	7.5	10
20	−26%	−18%	−12%	−9%	−7%
30	−27%	−15%	−10%	−7%	−6%
40	−25%	−14%	−9%	−7%	−5%
50	−21%	−13%	−9%	−6%	−5%
60	−14%	−12%	−9%	−7%	−6%

Next, we consider the participation decision rule implied by the policy function. Before deciding the risky share, households need to choose whether to participate in the stock market. The choice of not to participate is driven by two considerations. First, the existence of fixed cost will prevent the participation of households with little wealth. If fixed cost is too much relative to accumulated wealth, the optimal choice is not to pay it. Second, the risk exposure from labor income may exceed the tolerance of households. When a household predicts fairly unstable labor income, he will not participate in the stock market to avoid risk. Such case is less likely for wealthier households, since the value of labor income is not so important for them. Since BS-Corr makes labor income more stock-like, which is more risky, it will affect the participation decision through the second channel. Figure 4 presents the wealth-participation threshold over ages from policy functions discussed above. For both groups, BS-Corr implies a higher threshold for participation. The threshold increase by 13% to 27% for no college group and 12% to 38% for college group. In addition, the threshold curves keep the hump shape, which is

Figure 4: Participation Wealth Threshold

The two graphs plot the participation wealth threshold over ages for two groups. The threshold is calculated from the corresponding equity share policy function. It takes the largest value of wealth ratios with which households will not participate in the stock market.

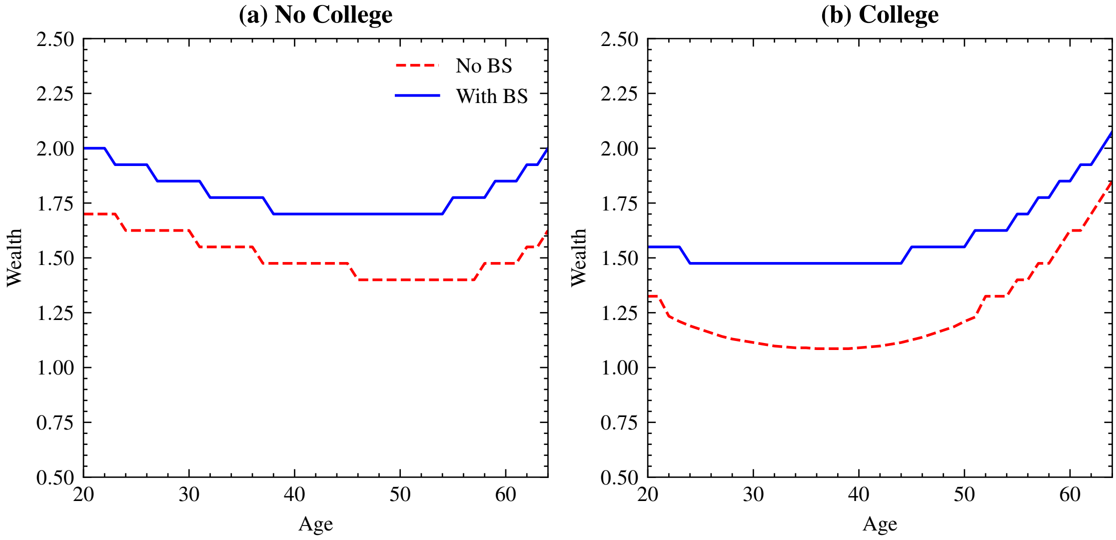

consistent with the results of Fagereng et al. (2017). At early ages, since the stock-like property becomes weaker with fewer labor income flows, households have more incentive to hold stock as age increases. When approaching retirement, the fall in the present value of future labor income gradually dominates and cause rebalancing. In addition, the heterogeneity of effect from BS-Corr also works here. The U-shape of threshold decays for both groups when considering BS-Corr, and almost become flat for college group with a larger BS-Corr.

5.2 Simulated Results

We simulate the life-cycle profile of 10, 0000 agents with policy functions discussed above. Figure 5 plots the participation rate and average conditional risky share over ages. For all models, the participation rate is increasing with ages, since wealth is gradually accumulated and households are therefore more likely to satisfy the wealth-participation threshold.

Figure 5: Life-Cycle Profile of Investment

The six figures plot the life-cycle profiles for the two groups and compare the results between models with and without BS-Corr. The profiles are calculated from a simulation of 10, 0000 households.

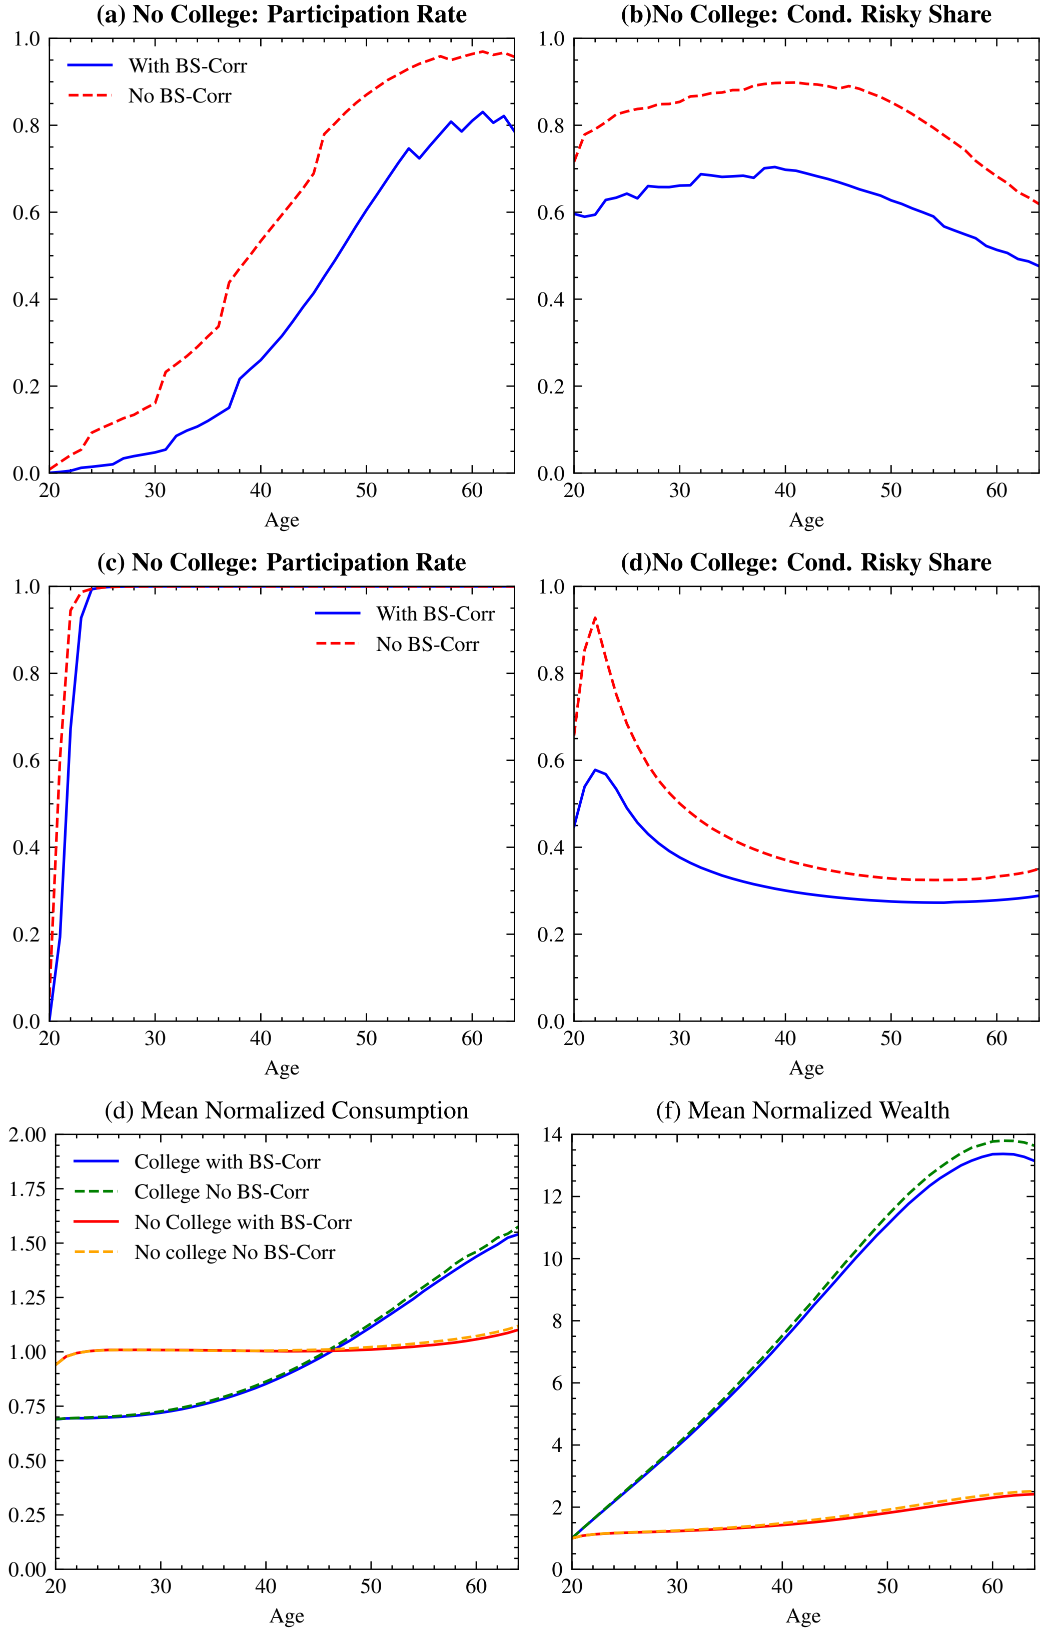

On the other hand, the average conditional risky share is roughly hump-shaped, which is consistent with the analysis of Cocco et al. (2005). In Cocco et al. (2005), the hump shape of conditional risky share is apparent only with more significant linear correlations (they used 0.2 and 0.4). Here we do not need to assume a considerable correlation, but introducing a small value of BS-Corr and risk from higher-order moments can pull the trigger. The hump shape is driven by wealth accumulation and the property of policy function. Since conditional optimal risky share is decreasing with wealth, the relatively poor investors will be very aggressive and almost invest fully in the stock market, while the relatively rich investors will invest prudently. At each period, households with wealth just passing the participation threshold will become aggressive investors. On the contrary, households already participated will decrease their risky share with wealth accumulation. At an early age, when lacking wealth accumulation, the effect of new participators will dominate and increase the mean conditional risky share. When more and more households have participated, the rebalancing effect of old investors will dominate and decrease the mean share. Thus, the mean conditional risky share is hump-shaped, and the peak location depends on the steepness of the participation rate. In our model, the college group will accumulate wealth faster and enter the stock market rapidly. Thus, the peak location of the college group is much earlier than that of no college group.

In addition, the results also support the analysis in section 5.1 regarding the effect of BS-Corr. Since the model with the BS-Corr considers a more risky labor income stream, the corresponding portfolio choice will be more prudent.

Finally, Figure 5(d) and 5(f) report the normalized mean consumption and mean normalized wealth for both groups with calibrated parameters. The results show that BS-Corr has less effect on consumption and wealth, which can be understood in twofold. First, as a measure of dependence, BS-Corr affects the prediction on labor income stream. Heuristically, it affects the evaluation of stock positions households hold through labor income. Thus, BS-Corr mainly affects the portfolio choice of households and has little effect on consumption. Seconde, the profiles of models without BS-Corr are simulated in a world with zero BS-Corr, and thus underestimate the risk in households' assets, which leads to an overestimation on wealth accumulation.

5.3 Sensitivity Analysis

Figure 6: Sensitivity Analysis

The three panels plot the mean conditional risky share over ages for different parameters. In each panel, we only change one parameter to test the robustness. The profiles are calculated from a simulation of 10,0000 households.

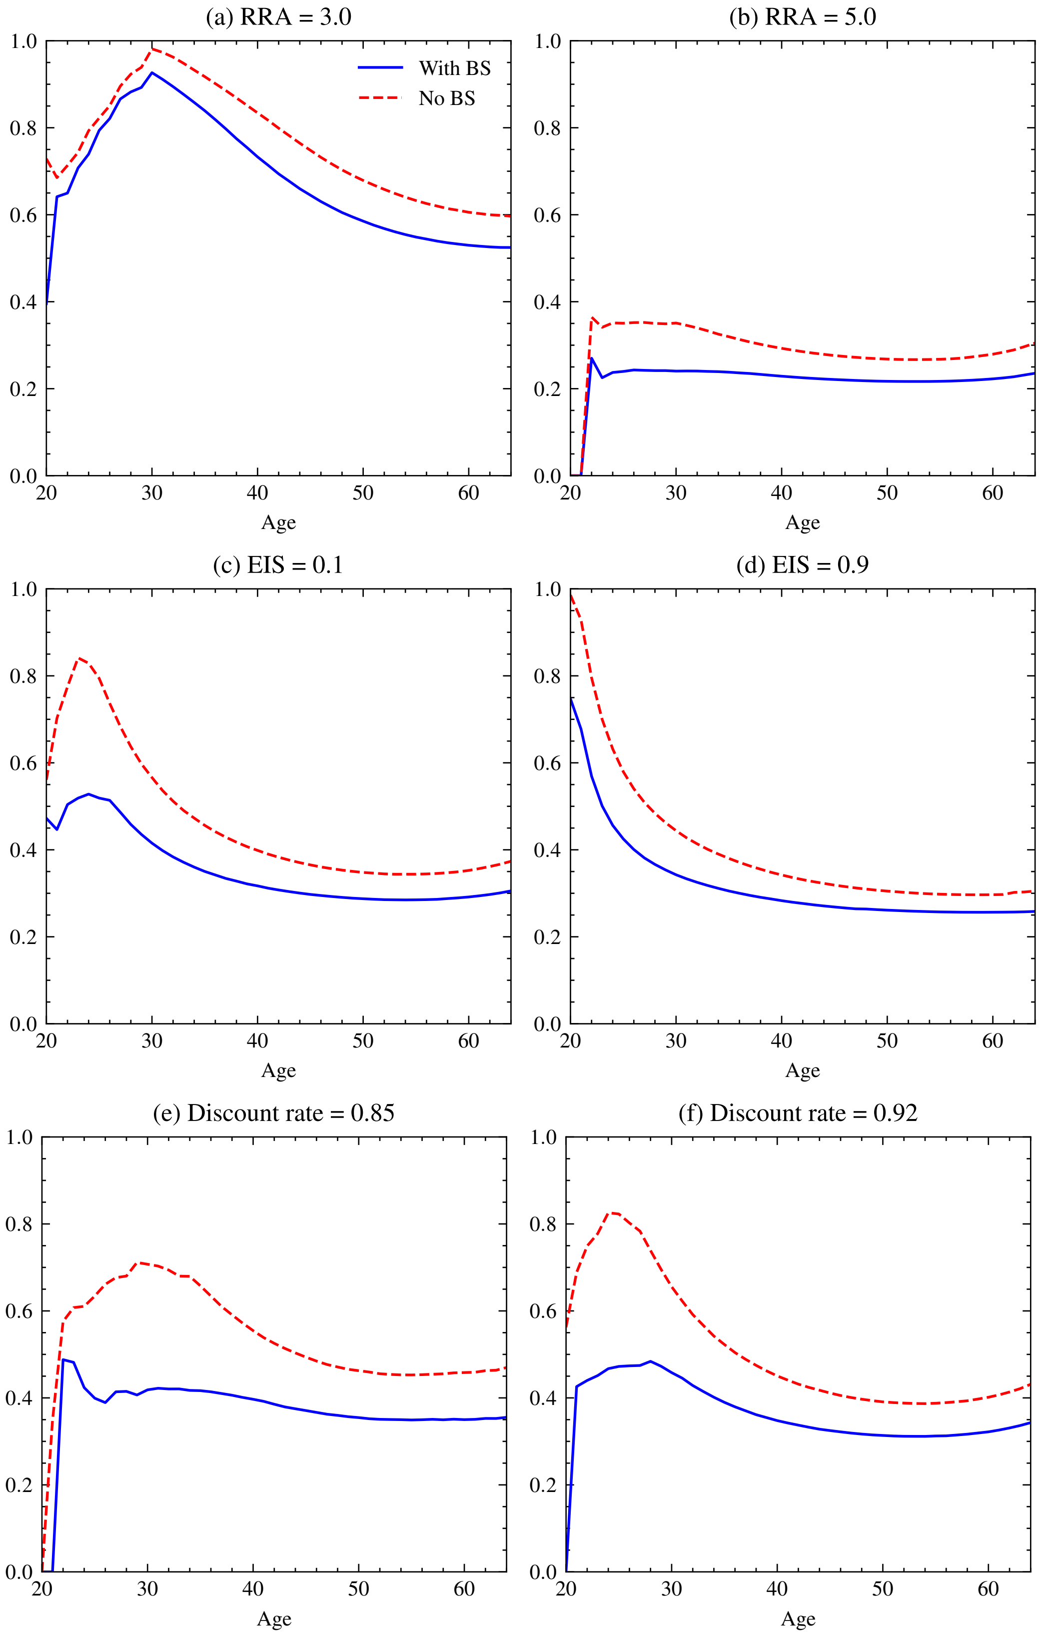

To test the robustness of the effect of BS-Corr, we perform sensitivity analysis with respect to main preference parameters: relative risk aversion, EIS, and discount rate. For simplicity, we only report the results for the college group at age 20. Results for no college group and other ages are similar. Figure 6 plots the life-cycle profile of the conditional equity share with different parameters. We find that BS-Corr has a very robust effect on portfolio choice. Although the change of parameters may reshape the risky share curve a lot, the effect of BS-Corr stays strong. The effect is relatively small with very low level of risk aversion. Figure 6(a) shows that the effect of BS-Corr is not as obvious as Figure 6(b) and Figure 6(c). It is reasonable since less risk averse households will tend to ignore the risks and invest a lot in stock market.

6 Empirical Evidence

Calibrated to the U.S. data, our model shows that the presence of BS-Corr between stock and labor income significantly affects households' portfolio decisions. In particular, households may choose not to participate in the stock market. In addition, even when households do participate in the stock market, their risky asset holdings are significantly reduced because of the BS-Corr.

To see if these predictions have empirical support, we use the U.S. variations of the degree of BS-Corr, stock investment, and stock market participation to examine the relations among the three. We use the PSID data of family level from 1997 to 2019, totalling 7321 observations. Since 1997, PSID also reports the value of stockholding.[16] Financial wealth is calculated as the sum of equity in stocks and the value in safe account, where the value in safe account is the money amount in checking and savings accounts, money market funds, certificates of deposit, government bonds, or treasury bills. We calculate the ratio of financial wealth invested in stocks, and assign value 1 if households participate in the stock market and 0 if not. To analyze the BS-Corr between stock markets and labor income, we use the annual income level and CRSP index and estimate the strength of BS-Corr for each household. Moreover, we observe a non-linear relationship between BS-Corr and households' portfolio decisions from the model.[17] Specifically, BS-Corr has an almost opposite relationship with portfolio decisions at tipping point zero. Therefore,

[16] Before that, only 1984 and 1989 waves report the data of stockholding.
[17] Please see Section 7 for details.

we suggest using an absolute value of BS-Corr, instead of BS-Corr, to better capture this non-linear relationship.

We begin our empirical analysis by estimating stock market probit participation regression. For this analysis, the dependent variable is a dummy variable as

$$I_{it} = \begin{cases} 1 & \text{if household participates, or} \\ 0 & \text{if household does not participate.} \end{cases} \tag{19}$$

and we estimate the propensity to participate in the stock market, which we denote by $p(I = 1)$. Table 5 reports the marginal effects from probit market participation

Table 5: Probit Participation Regression Estimates

This table reports marginal effects from probit regressions. The dependent variable is the dummy variable which denotes whether an individual participates the stock market. We construct probit models using different sample populations. We omit the constants and year dummies. Group 1 reports the results for no-college-degree groups and Group 2 reports the results for college graduates.

	Full Sample			Group 1			Group 2		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\lvert BSCor \rvert$	-0.500^{***}	-0.343^{***}	-0.209^{***}	-0.531^{***}	-0.407^{***}	-0.216^{***}	-0.440^{***}	-0.328^{***}	$-0.196^{\cdot}$
	(0.066)	(0.069)	(0.078)	(0.087)	(0.089)	(0.101)	(0.109)	(0.045)	(0.128)
Cor		0.050	0.053		-0.038	0.003		0.122^{*}	0.090^{*}
		(0.041)	(0.043)		(0.056)	(0.058)		(0.063)	(0.068)
$ln(y)$		0.855^{***}	0.626^{***}		0.792^{***}	0.536^{***}		0.602^{***}	0.406^{***}
		(0.035)	(0.043)		(0.049)	(0.059)		(0.055)	(0.071)
$age/10$			-0.278^{*}			-0.080			0.212
			(0.151)			(0.196)			(0.247)
$age^2/100$			0.026			0.008			0.018
			(0.018)			(0.023)			(0.029)
$Marriage$			-0.351^{***}			-0.180^{***}			-0.363^{***}
			(0.055)			(0.073)			(0.088)
$FWealth$			0.241^{***}			0.258^{***}			0.214^{***}
			(0.010)			(0.014)			(0.014)
$Std(d\delta)$			0.012			0.094			-0.088
			(0.110)			(0.145)			(0.174)
$Skew(d\delta)$			-0.034			-0.022			-0.033
			(0.023)			(0.031)			(0.036)
$Kurt(d\delta)$			-0.028^{**}			-0.041^{**}			$-0.033^{\cdot}$
			(0.014)			(0.019)			(0.022)
N	6994	6994	6994	4227	4227	4227	2767	2767	2767

$^{*}\, p < 0.1$; $^{**}\, p < 0.05$; $^{***}\, p < 0.01$

regressions. The dependent variable in these regressions is an ownership binary variable that takes the value of one if the household owns stocks. We separately examine the decision to hold stocks within college graduates and households with college degrees.

Our key independent variable is the absolute value of BS-Corr. Following the related literature, we control households' characteristics, such as the level of income and wealth, income risk, marriage status, age and education. These household-level variables would serve as a reasonable proxy for the real and perceived market participation costs.

Our results indicate that the absolute value of BS-Corr between income shocks and stock returns has a significantly negative relationship with the decision to participate in the stock market. In particular, we find that when the BS-Corr is low in the absolute value, the market participation propensity decreases. This result is strong and statistically significant in regressions with and without control variables. The coefficients of BS-Corr remains statistically significant when we look at the college graduates and no-college-degree groups respectively. The relationship is stronger for no-college-degree group -0.531 compared with college graduates -0.440. The different estimates suggest that households with college degrees are more resilient.

Table 6: Tobit Model of Risky Share

This table reports marginal effects from tobit regressions. The dependent variable is the risky share of households. We construct tobit models using different sample populations. We omit the constants and year dummies.

	Full Sample			Group 1			Group 2		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
—BSCor—	-0.297^{***}	-0.210^{***}	-0.116^{***}	-0.345^{***}	-0.267^{***}	-0.134^{***}	-0.230^{***}	-0.180^{***}	-0.101^{***}
	(0.036)	(0.034)	(0.035)	(0.053)	(0.052)	(0.053)	(0.046)	(0.045)	(0.047)
Cor		0.019	0.026		-0.026	-0.010		0.043^{*}	0.049^{*}
		(0.021)	(0.020)		(0.032)	(0.030)		(0.026)	(0.025)
ln(y)		0.422^{***}	0.338^{***}		0.442^{***}	0.300^{***}		0.253^{***}	0.200^{***}
		(0.018)	(0.019)		(0.029)	(0.030)		(0.023)	(0.026)
age/10			-0.011			-0.053			0.061
			(0.068)			(0.103)			(0.091)
$age^2/100$			0.003			-0.004			-0.005
			(0.008)			(0.012)			(0.011)
Marriage			-0.162^{***}			-0.102^{***}			-0.124^{***}
			(0.025)			(0.039)			(0.032)
$FWealth$			0.044^{***}			0.074^{***}			0.028^{***}
			(0.002)			(0.005)			(0.002)
$Std(d\delta)$			0.019			0.058			-0.022
			(0.050)			(0.076)			(0.065)
$Skew(d\delta)$			-0.017^{*}			-0.018			-0.010
			(0.010)			(0.016)			(0.013)
$Kurt(d\delta)$			-0.007			-0.009			-0.014^{*}
			(0.006)			(0.009)			(0.007)
N	6994	6994	6994	4227	4227	4227	2767	2767	2767

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Our next set of results is related to the asset allocation decisions of households. We estimate Tobit regressions for market participants. Table 6 report estimates from Tobit regressions in which the dependent variable is the fraction of wealth invested in stocks. The main independent variable is the absolute value of BS-Corr between income shocks and market returns. Like the probit market participation regression specifications, we consider a rich set of control variables. The Tobit regression estimates are consistent with the estimation results from the market participation regressions. Specifically, we find that households with low absolute value of BS-Corr allocate less wealth to risky assets. For example, in the regression 3 where we consider allocations to risky assets with a broad set of control variables, the coefficient estimate of the absolute value of BS-Corr is -0.116. Similarly, for both college graduates and no-college-degree households, the coefficient estimates are -0.101 and -0.134 respectively.

Overall, the empirical evidence supports our model's predictions. In particular, across the college graduates and no-college-graduates, as the absolute value of BS-Corr between stock returns and income shocks increases, stock investment decreases and nonparticipation in the stock market increases.

7 Portfolio Perspective of BS-Corr

Regression analysis indicates that BS-Corr has a non-linear relationship with participation decisions and conditional risky asset shares. To better understand this non-linear relationship within a simple setup, we provide a discussion under the portfolio selection framework. More specifically, we construct a portfolio with an α_{BS} share of stock and one unit of labor income flow.[18] We aim to maximize the return of this portfolio, while constrained by the portfolio risk smaller than or equal to risk level θ_0, which is defined as the sum of kurtosis and the absolute value of skewness. The optimization problem can be stated as:

$$\begin{aligned} &\max \mathbb{E}(R_t^L + \alpha_{BS} R_t^S), \\ &\text{subject to} \quad |\mathrm{Skew}(R_t^L + \alpha_{BS} R_t^S)| + \mathrm{Kurt}(R_t^L + \alpha_{BS} R_t^S) \leq \theta_0 \end{aligned} \tag{20}$$

[18]We acknowledge that labor income is non-tradable, but here for simplicity and easy explanation, we artificially construct a portfolio including labor income flows.

Fixing α_{BS} share of stock and the moments of labor income and stock returns, we change the BS-Corr alone and calculate the corresponding skewness and kurtosis of the portfolio.[19] Figure 7(a) shows the portfolio skewness and kurtosis with varying BS-Corr from -0.1 to 0.25. We find that when BS-Corr increases from negative to zero, the kurtosis increases steadily, but rise more rapidly when BS-Corr is more positive. On the other hand, skewness is hump-shaped with a saddle point at zero. When BS-Corr increases from negative to zero, skewness of the portfolio increases dramatically, while when BS-Corr increases from zero to positive, skewness decreases mildly. Considering that both skewness and kurtosis experience significant switch toward gradient around zero, the risk of the portfolio achieves minimum θ_0 when BS-Corr is zero and increases, when BS-Corr deviates from zero. Therefore, without changing the stock shares, the portfolio risk exceeds θ_0 whenever BS-Corr is not zero. In order to control the portfolio risk, households need to adjust α_{BS} accordingly, since households can not alter their labor income risk.

Figure 7: Portfolio Perspective of BS-Corr

The two figures in (a) plot the variation of portfolio higher order moments with changing BS-Corr, and the two figures in (b) plot the participation rate and conditional risky share from the corresponding policy functions. Here we use $\alpha_{BS} = 1$.

(a) Portfolio Risk (b) Policy Function

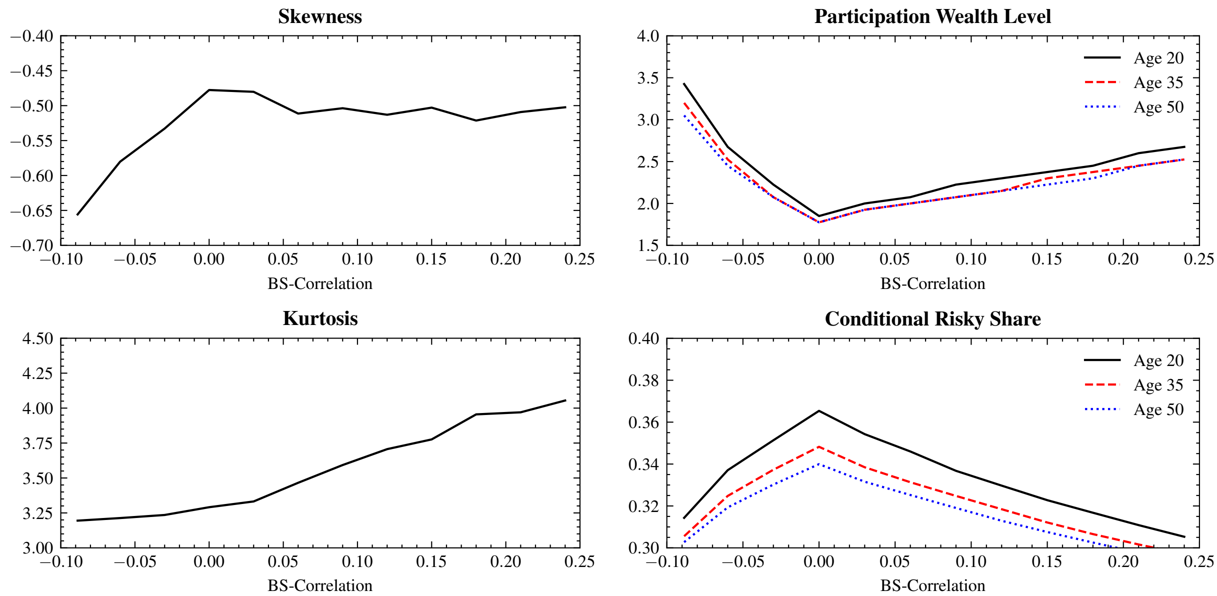

To verify such prediction, we consider policy functions for different age groups with varying BS-Corr, while fixing the shocks of labor income and stock return. We mainly

[19]To do so, we only adjust the dependence parameters of binary mixed normal distribution to change BS-Corr while keep Corr unchanged.

focus on the participation wealth threshold level and conditional risky shares. The participation wealth level is the wealth level where households begin to invest in the stock markets, and the condition risky share is the average risky share condition on participation. Figure 7(b) plots the policy functions for age 20, 35 and 50. Across all age groups, we find that households are more likely to participate in the stock market when BS-Corr is zero. If we allow α_{BS} adjust with varying BS-Corr, we find that households are willing to hold more risky asset shares conditional on participation when BS-Corr is zero. When BS-Corr differs from zero, households reduce their risky asset holdings in order to control their total portfolio risk. Overall, we can conclude that increasing BS-Corr between stocks and labor income makes labor income more 'stock-like', which suggests more risk coming from labor income flow.

8 Conclusion

In this paper, we consider the optimal portfolio decision of a household when the stock returns and income shocks are BS-correlated. We show that in the presence of BS-Corr, households are less willing to participate in the market and significant reduce their stock investment. Moreover, our paper shows that BS-Corr is independent of correlation, where literature provides mixed evidence about it. Therefore, our model complements existing studies and can potentially help explain both the limited participation puzzle and moderate risky asset holdings observed from the data. In addition, we show empirical evidence that is supportive of model's prediction and provide a portfolio selection perspective to understand households' market participation decisions and risky asset holdings relative to the labor income.

References

Ameriks, J. and Zeldes, S. P. (2001). show do household portfolio shares vary with age?, tworking paper. *Columbia University*.

Baxter, M. and Jermann, U. J. (1997). The international diversification puzzle is worse than you think. *The American Economic Review*, 87(1):170–180.

Benzoni, L., Collin-Dufresne, P., and Goldstein, R. S. (2007). Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *The Journal of Finance*, 62(5):2123–2167.

Bonaparte, Y., Korniotis, G. M., and Kumar, A. (2014). Income hedging and portfolio decisions. *Journal of Financial Economics*, 113(2):300–324.

Brunnermeier, M. K. and Nagel, S. (2008). Do wealth fluctuations generate time-varying risk aversion? micro-evidence on individuals. *American Economic Review*, 98(3):713–36.

Campbell, J. Y. (1996). Understanding risk and return. *Journal of Political economy*, 104(2):298–345.

Campbell, J. Y., Lettau, M., Malkiel, B. G., and Xu, Y. (2001). Have individual stocks become more volatile? an empirical exploration of idiosyncratic risk. *The journal of finance*, 56(1):1–43.

Cao, H. H., Wang, T., and Zhang, H. H. (2005). Model uncertainty, limited market participation, and asset prices. *The Review of Financial Studies*, 18(4):1219–1251.

Catherine, S. (2020). Countercyclical labor income risk and portfolio choices over the life cycle. *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

Cocco, J. F. (2005). Portfolio choice in the presence of housing. *The Review of Financial Studies*, 18(2):535–567.

Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, 18(2):491–533.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2):223.

Davis, S. J. and Willen, P. (2000a). Using financial assets to hedge labor income risks: estimating the benefits. Technical report, Citeseer.

Davis, S. J. and Willen, P. S. (2000b). Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice.

Edwards, S. and Susmel, R. (2001). Volatility dependence and contagion in emerging equity markets. *Journal of Development Economics*, 66(2):505–532.

Fagereng, A., Gottlieb, C., and Guiso, L. (2017). Asset market participation and portfolio choice over the life-cycle. *The Journal of Finance*, 72(2):705–750.

Faig, M. and Shum, P. (2002). Portfolio choice in the presence of personal illiquid projects. *The Journal of Finance*, 57(1):303–328.

Gomes, F. and Michaelides, A. (2005a). Optimal life-cycle asset allocation: Understanding the empirical evidence. *The Journal of Finance*, 60(2):869–904.

Gomes, F. and Michaelides, A. (2005b). Optimal life-cycle asset allocation: Understanding the empirical evidence. *The Journal of Finance*, 60(2):869–904.

Gomes, F. J. (2005). Portfolio choice and trading volume with loss-averse investors. *The Journal of Business*, 78(2):675–706.

Guiso, L., Jappelli, T., and Terlizzese, D. (1996). Income risk, borrowing constraints, and portfolio choice. *The American economic review*, pages 158–172.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Haliassos, M. and Michaelides, A. (2003). Portfolio choice and liquidity constraints. *International Economic Review*, 44(1):143–177.

Harvey, C. R. and Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of finance*, 55(3):1263–1295.

Heaton, J. and Lucas, D. (2000). Portfolio choice and asset prices: The importance of entrepreneurial risk. *The journal of finance*, 55(3):1163–1198.

Hong, H., Kubik, J. D., and Stein, J. C. (2004). Social interaction and stock-market participation. *The journal of finance*, 59(1):137–163.

Hsieh, D. A. (1989). Testing for nonlinear dependence in daily foreign exchange rates. *Journal of Business*, pages 339–368.

Hsieh, D. A. (1993). Implications of nonlinear dynamics for financial risk management. *Journal of Financial and quantitative Analysis*, 28(1):41–64.

Huggett, M. and Kaplan, G. (2011). Human capital values and returns: Bounds implied by earnings and asset returns data. *Journal of Economic Theory*, 146(3):897–919.

Huggett, M. and Kaplan, G. (2016). How large is the stock component of human capital? *Review of Economic Dynamics*, 22:21–51.

Jagannathan, R., Kocherlakota, N. R., et al. (1996). Why should older people invest less in stocks than younger people. *Federal Reserve Bank of Minneapolis Quarterly Review*, 20:11–20.

LeBaron, B. D. (1988). *Stock return nonlinearities: Some initial tests and findings*. PhD thesis, University of Chicago, Department of Economics.

Lucas, D. and Zeldes, S. P. (2006). Valuing and hedging defined benefit pension obligations-the role of stocks revisited. *Northwestern University and Columbia University, working paper, September.*

Lynch, A. W. and Tan, S. (2011). Labor income dynamics at business-cycle frequencies: Implications for portfolio choice. *Journal of Financial Economics*, 101(2):333–359.

Merton, R. C. (1975). Optimum consumption and portfolio rules in a continuous-time model. In *Stochastic optimization models in finance*, pages 621–661. Elsevier.

Nakajima, M. and Smirnyagin, V. (2019). Cyclical labor income risk. *Available at SSRN 3432213.*

Neuberger, A. and Payne, R. (2021). The skewness of the stock market over long horizons. *The Review of Financial Studies*, 34(3):1572–1616.

Nikoloulopoulos, A. K. (2021). The bivariate k-finite normal mixture ‘blanket’copula. *Journal of Statistical Computation and Simulation*, pages 1–22.

Pagel, M. (2018). A news-utility theory for inattention and delegation in portfolio choice. *Econometrica*, 86(2):491–522.

Peijnenburg, K. (2018). Life-cycle asset allocation with ambiguity aversion and learning. *Journal of Financial and Quantitative Analysis*, 53(5):1963–1994.

Poterba, J. M., Samwick, A., et al. (2001). Household portfolio allocation over the life cycle. *Aging issues in the United States and Japan*, 2.

Santos, T. and Veronesi, P. (2006). Labor income and predictable stock returns. *The Review of Financial Studies*, 19(1):1–44.

Scheinkman, J. A. and LeBaron, B. (1989). Nonlinear dynamics and stock returns. *Journal of business*, pages 311–337.

Shen, J. (2019). Countercyclical risks and optimal life-cycle profile: Theory and evidence. In *Proceedings of Paris December 2019 Finance Meeting EUROFIDAI-ESSEC*.

Storesletten, K., Telmer, C. I., and Yaron, A. (2007). Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics*, 10(4):519–548.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance*, 56(2):433–470.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of political Economy*, 110(4):825–853.

Yao, R. and Zhang, H. H. (2005). Optimal consumption and portfolio choices with risky housing and borrowing constraints. *The Review of Financial Studies*, 18(1):197–239.

A Appendix: Data and Numerical Solution

A.1 Supplementary Data

A.1.1 Panel Study of Income Dynamics

The Panel Study of Income Dynamics (PSID) is the longest running longitudinal household survey in the world, collecting data of over 18,000 individuals living in 5,000 families in the United States from 1968 to 2019. PSID interviewed on a broad range of topics, including the family information and financial situation which is required in most life cycle household models. PSID is even more appropriate for our model since we need time series of labor income to estimate its correlations with stock return. However, PSID has some well-know issues for researchers and many papers have discussed the approaches to deal with them[20]. In this paper, we roughly follow the sample selection principle of Nakajima and Smirnyagin (2019). First, years prior to 1970 are dropped. Waves of 1968 and 1969 lack some data and are not completely consistent with the waves afterwards. Second, drop SEO and Latino samples. PSID includes about 2000 low income SEO samples and 2000 Latino samples. We drop all those samples since they are colleted with unequal selection probabilities. Third, observations with missing or non-positive head/spousal labor incomes are dropped. Then, the top 1% with respect to head's and wife's labor incomes are trimmed. PSID bracket lots of variables about finance, such as labor income, with an upper boundary. This trim selection is to drop bracketed or extreme samples. Finally, households with income growth anomalies (annual log growth rate must be between 1/20 and 20) are dropped.

We construct variables following Brunnermeier and Nagel (2008). The following table summarize the variable definitions.

Table 7: Varible definitions

Variable	Definition
Labor Income Y_{it}	Includes the labor income of both reference person and spouse, labor part of farm income and business income, all transfer income of the family, and social securities.
Riskless Assets	Checking and savings accounts, money market funds, certificates of deposits, savings bonds, and treasury bills
	Bonds, bond funds, cash value in a life insurance, valuable collection for investment purposes, and rights in a trust or estate
Risky Assets	the combined value of shares of stock in publicly held corporations, mutual funds, and investment trusts
IRA Assets	Value of private annuities or Individual Retirement Accounts
Other Debts	Credit card debt, student loans, medical or legal bills, and loans from relatives
Home Equity	Value of the home minus remaining mortgage principal
Liquid Assets	Riskless Assets + Risky Assets
Financial Wealth	Liquid Assets - Other Debts + Home Equity

[20]Such as Cocco et al. (2005), Gomes and Michaelides (2005b) and Nakajima and Smirnyagin (2019).

A.1.2 SCF Data

The Survey of Consumer Finances (SCF) is normally a triennial cross-sectional survey of U.S. families. The survey data include information on families' balance sheets, pensions, income, and demographic characteristics. Information is also included from related surveys of pension providers and the earlier such surveys conducted by the Federal Reserve Board. We use data from the 2007 to 2019, 5 waves in total. Variables are constructed using the code-book and macro-variable definitions from the Federal Reserve website.

We construct variables of SCF following Gomes and Michaelides (2005b). Labor income is defined as the sum of wages and salaries (X5702), unemployment or worker's compensation (X5716) and Social Security or other pensions, annuities, or other disability or retirement programs (X5722). Then, financial wealth is constructed in a same way as variable FIN in the publicly available SCF data set, made up of LIQ (all types of transaction accounts—checking, saving, money market, and call accounts), CDS (certificates of deposit), total directly held mutual funds, stocks, bonds, total quasi-liquid financial assets (the sum of IRAs, thrift accounts, and future pensions), savings bonds, the cash value of whole life insurance, other managed assets (trusts, annuities, and managed investment accounts in which the household has equity interest), and other financial assets (includes loans from the household to someone else, future proceeds, royalties, futures, nonpublic stock, and deferred compensation). Further, data of financial assets invested in the risky asset is from variable EQUITY in the publicly available SCF data set, which consists of directly held stock, stock mutual funds, or amounts of stock in retirement accounts. We calculate the conditional risky share as (EQUITY)/(FIN) conditional on EQUITY being positive.

A.2 Numerical Solution

The model can be numerically solved using backward induction. The value function for each period depends on normalized cash on hand x_t, which is a continuous variable and thus need to be discretized. The terminal condition in the last period is determined by the bequest motive, and the value function corresponds to the bequest function. For each period, we use a grid search to optimize the value function. We compute the value associated with each grid of consumption and risky share, and choose optimal grids achieving the maximum value as the policy rules. For each period t prior to T, and for each point in the state space, this procedure is iterated backward.

To approximate the expected value at next period which is required in the backward induction, we use Gauss-Hermite for mixed normal distributed. See Appendix B.2 for detail.

B Appendix: Mixed Normal Distribution

B.1 Mixed Normal Distribution

We consider two random variables following mixed normal distribution, and they can be represented as:

$$\begin{aligned} z_1 &= I_1 x_{11} + (1 - I_1) x_{12}, \\ z_2 &= I_2 x_{21} + (1 - I_2) x_{22}. \end{aligned}$$

where $I_i \sim \mathrm{B}(1, p_i)$ and $X = (x_{11}, x_{12}, x_{21}, x_{22})'$ is normally distributed subject to:

$$\begin{aligned} \mathbb{E}(X) &= (\mu_{11}, \mu_{12}, \mu_{21}, \mu_{22})', \\ \mathrm{Cov}(X) &= \begin{bmatrix} \sigma_{11}^2 & 0 & \rho_{11}\sigma_{11}\sigma_{21} & \rho_{12}\sigma_{11}\sigma_{22} \\ 0 & \sigma_{12}^2 & \rho_{21}\sigma_{12}\sigma_{21} & \rho_{22}\sigma_{12}\sigma_{22} \\ \rho_{11}\sigma_{11}\sigma_{21} & \rho_{21}\sigma_{12}\sigma_{21} & \sigma_{21}^2 & 0 \\ \rho_{12}\sigma_{11}\sigma_{22} & \rho_{22}\sigma_{12}\sigma_{22} & 0 & \sigma_{22}^2 \end{bmatrix}. \end{aligned}$$

B.1.1 Moments

We calculate the first four moments of z_i.

$$\begin{aligned} \mu_i = \mathbb{E}[z_i] =& p_i \mu_{i1} + (1 - p_i) \mu_{i2}, \\ \sigma_i^2 = \mathrm{Var}[z_i] =& p_i \sigma_{i1}^2 + (1 - p_i) \sigma_{i2}^2 + p_i (1 - p_i)(\mu_{i1} - \mu_{i2})^2, \\ s_i = \mathrm{skew}[z_i] =& \sigma_i^{-3} \mathbb{E}[I_i \sigma_{i1} (\frac{x_{i1} - \mu_{i1}}{\sigma_{i1}}) + (1 - I_i) \sigma_{i2} (\frac{x_{i2} - \mu_{i2}}{\sigma_{i2}}) + (I_i - p_i)(\mu_{i1} - \mu_{i2})]^3 \\ =& \sigma_i^{-3} p_i (1 - p_i)(\mu_{i1} - \mu_{i2}) \left(3(\sigma_{i1}^2 - \sigma_{i2}^2) + (1 - 2p_i)(\mu_{i1} - \mu_{i2})^2 \right), \\ k_i = \mathrm{kurt}[z_i] =& \sigma_i^{-4} \big[3 p_i \sigma_{i1}^4 + 3(1 - p_i) \sigma_{i2}^4 + p_i (1 - p_i)(\mu_{i1} - \mu_{i2})^2 \big(6((1 - p_i)\sigma_{i1}^2 + p_i \sigma_{i2}^2) + \end{aligned}$$

$$(3p_i^2 - 3p_i + 1)(\mu_{i1} - \mu_{i2})^2)\big].$$

B.1.2 Pearson correlation

Denote:

$$P = \begin{bmatrix} p_1 & (1-p_1) \\ p_2 & (1-p_2) \end{bmatrix}$$

The Pearson correlation between z_1 and z_2 can be written as:

$$\begin{aligned}
\mathrm{Corr}(z_1, z_2) =& \frac{\mathrm{Cov}(I_1 x_{11} + (1-I_1)x_{12}, I_2 x_{21} + (1-I_2)x_{22})}{\sqrt{\mathrm{Var}(z_1)\mathrm{Var}(z_2)}} \\
=& \frac{1}{\sigma_1\sigma_2} \sum_{i,j=1,2} P_{1i} P_{2j} \sigma_{1i} \sigma_{2j} \rho_{ij}
\end{aligned}$$

B.1.3 Between-squares correlation

Set $y_{ij} = \dfrac{x_{ij} - \mu_{ij}}{\sigma_{ij}}, i, j = 1, 2.$ And we construct Schmidt orthogonalization:

$$\begin{aligned}
y_{11}^* =& y_{11}, \\
y_{12}^* =& y_{12}, \\
y_{21}^* =& y_{21} - \rho_{11} y_{11}^* - \rho_{21} y_{12}^*, \\
y_{22}^* =& y_{22} - \rho_{12} y_{11}^* - \rho_{22} y_{12}^* + \frac{\rho_{11}\rho_{12} + \rho_{21}\rho_{22}}{1 - \rho_{11}^2 - \rho_{21}^2} y_{21}^*.
\end{aligned}$$

And we can calculate:

$$\begin{aligned}
\mathrm{Corr}^{\mathrm{bs}}(z_1, z_2) =& \mathrm{Corr}\big((\sigma_{11} I_1 y_{11} + \sigma_{12}(1-I_1)y_{12} + (I_1 - p_1)(\mu_{11} - \mu_{12}))^2, \\
& (\sigma_{21} I_2 y_{21} + \sigma_{22}(1-I_2)y_{22} + (I_2 - p_2)(\mu_{21} - \mu_{22}))^2\big) \\
=& \frac{1}{\sigma_1^2\sigma_2^2\sqrt{(k_1-1)(k_2-1)}} \mathrm{Cov}\big((\sigma_{11} I_1 y_{11}^* + \sigma_{12}(1-I_1)y_{12}^* + (I_1 - p_1)(\mu_{11} - \mu_{12}))^2, \\
& (\sigma_{21} I_2(\rho_{11} y_{11}^* + \rho_{21} y_{12}^* + y_{21}^*) + \sigma_{22}(1-I_2)(\rho_{12} y_{11}^* + \rho_{22} y_{12}^* \\
& - \frac{\rho_{11}\rho_{12} + \rho_{21}\rho_{22}}{1 - \rho_{11}^2 - \rho_{21}^2} y_{21}^* + y_{22}^*) + (I_2 - p_2)(\mu_{21} - \mu_{22}))^2\big) \\
=& \frac{1}{\sigma_1^2\sigma_2^2\sqrt{(k_1-1)(k_2-1)}} \mathrm{Cov}\big((\sigma_{11} I_1 y_{11}^* + \sigma_{12}(1-I_1)y_{12}^* + (I_1 - p_1)(\mu_{11} - \mu_{12}))^2, \\
& ((\sigma_{21} I_2 \rho_{11} + \sigma_{22}(1-I_2)\rho_{12}) y_{11}^* + (\sigma_{21} I_2 \rho_{21} + \sigma_{22}(1-I_2)\rho_{22}) y_{12}^* + \\
& (I_2 - p_2)(\mu_{21} - \mu_{22}))^2\big)
\end{aligned}$$

$$
\begin{aligned}
=&\frac{1}{\sigma_1^2\sigma_2^2\sqrt{(k_1-1)(k_2-1)}}(2\sum_{i,j=1,2}P_{1i}P_{2j}\sigma_{1i}^2\sigma_{2j}^2\rho_{ij}^2\\
&+4(\mu_{11}-\mu_{12})(\mu_{21}-\mu_{22})p_1(1-p_1)p_2(1-p_2)\sum_{i,j=1,2}(-1)^{i+j}\sigma_{1i}\sigma_{2j}\rho_{ij}).
\end{aligned}
$$

B.2 Gauss-Hermite for Mixed Normal Distribution

During the backward induction, for each period t, we need to calculate the expectation $\mathbb{E}[v_{t+1}]$ with respect to $(\gamma_{t+1}, \epsilon_{t+1}, \eta_{t+1})$. Since mixed normal distribution consists of two normal distribution, we can use Gauss-Hermite method for normal distribution to calculate $\mathbb{E}[v_{t+1}]$. Generally, we derive a Gaussian-Hermite for $h(\gamma_{t+1}, \eta_{t+1}, \epsilon_{t+1})$, where $\gamma_{t+1} = I_1 x_{11} + (1-I_1)x_{12}$, $\eta_{t+1} = I_2 x_{21} + (1-I_2)x_{22}$ and $\epsilon_{t+1} \sim N(0, \sigma_\epsilon^2)$. We have:

$$
\mathbb{E}[h] = \sum_{i,j=0,1}\mathbb{E}[h|I_1=i, I_2=j]\mathbb{P}(I_1=i, I_2=j) = \sum_{i,j=1,2}\mathbb{E}[h(x_{1i}, x_{2j}, \epsilon_{t+1})]P_{1i}P_{2j}.
$$

Now for each component of the summation, all variables in h are normally distributed and then applicable to Gauss-Hermite method.

Extended Abstract

Although the objective functions are concave in traditional utility maximization problems, there is also a large body of literature on non-concave utility maximization in economics and finance. Examples include the S-shaped utility function in behavioral economics (e.g., Kahneman and Tversky (1979), Berkelaar, Kouwenberg and Post (2004), and Jin and Zhou (2008)), the goal-reaching problem (e.g., Browne (1999a) and Spivak and Cvitanić (1999)), delegated portfolio choices with non-concave compensation schemes (e.g., Carpenter (2000), Basak, Pavlova and Shapiro (2007), and He and Kou (2018)), and the aspiration utility maximization (e.g., Diecidue and van de Ven (2008) and Lee, Zapatero and Giga (2018)). Almost all models in the above literature rely on the concavification principle, namely, replacing a non-concave utility with its concave envelope and thus reducing the non-concave utility maximization problem to a concave one.

This paper provides a general framework for solving non-concave utility maximization with portfolio bounds where the concavification principle does not hold. By applying our methodology to the aforementioned six models in the literature of non-concave utility maximization, we find that adding portfolio bounds, which makes the concavification principle invalid, can offer distinct economic insights and implications. Theoretically, in view of the difficulties incurred by non-concave value functions, we do three things: (a) We introduce a new definition of viscosity solution. (b) Based on the new definition, we establish a novel comparison principle, which is used to prove that the value function of the non-concave utility maximization problem is the unique viscosity solution (in terms of the new definition) of the Hamilton-Jacobi-Bellman (HJB) equation. (c) We then show that a monotone, stable, and consistent finite difference scheme converges to the solution of the utility maximization problem.

Economic Insights and Implications. There are three general findings for the non-concave utility maximization with portfolio bounds. First, the concavification principle may no longer hold, because the resulting value function may not be globally concave before maturity in general. Intuitively, convex incentives would induce investors to borrow large amounts of the risk-free asset or to take large short-sale positions, which is prohibited due to portfolio bounds. Consequently, it is impossible to obtain a hypothetical value function that is the concave envelope of the original non-concave utility, yielding a non-concave value function. Indeed, we find that the non-concavity often occurs locally, depending on the time to maturity and wealth level. This implies that facing convex incentives, investors may choose to gamble or not, depending on scenarios.

Second, investors may gamble by short-selling (borrowing) a stock even with a positive (negative) risk premium as their target is yet to be reached at a time sufficiently close to maturity. For example, consider a fund manager with a convex compensation scheme who suffers a large loss near maturity and faces the no-short-selling constraint but no severe restriction on borrowing. Given

the convex incentive and the short-selling constraint, the manager would be induced to leverage on long positions, even with a negative risk premium of the stock.

Third, investors may not be myopic with respect to portfolio constraints in the sense that they may act before portfolio constraints are binding. Intuitively, due to portfolio constraints, investors may raise (reduce) their current stock investment to compensate for the potential limited borrowing (short-selling) in the future, compared to the case without portfolio constraints.

We show that these general economic insights still hold under more general market environments, such as the time-varying Gaussian mean return model, the stochastic volatility model, and the case with multiple stocks. Besides, induced by convex incentives, investors may take advantage of the stochastic investment opportunity set to get a more volatile portfolio through hedging demands.

Our study also has two empirical implications. First, Brown, Harlow, and Starks (1996) find that fund managers who are underperforming (outperforming) in the first-half-of-the-year increases (reduces) their risk level in the second-half-of-the-year. This phenomenon, known as the "risk-increasing" tournament behavior, seems to be a rational response of fund managers due to the convex incentives caused by the convex relationship of fund flow to annual relative performance. However, the results of later empirical studies on the tournament behavior are mixed and conflicting (see, e.g., Chevalier and Ellison (1997), Busse (2001), Kempf, Ruenzi and Thiele (2009), and Schwarz (2012)). Our finding provides an alternative explanation for the seemingly conflicting results. Indeed, we complement the results in Basak, Pavlova and Shapiro (2007) by showing that when a fund is significantly underperforming a benchmark, the fund manager may not alter her risk preference; only should the underperformance be within a certain range would the fund manager alter her risk preference. In other words, the tournament behavior may be local.

Second, Agarwal, Boyson, and Naik (2009) and Chen, Desai and Krishnamurthy (2013) empirically find that mutual funds that use both buying and short-selling (i.e., long-short) strategies outperform other mutual funds that do not use short-selling (i.e., long-only), in terms of portfolio returns (a risk-neutral evaluation criterion). Because mutual fund managers may have their own risk profiles and face convex incentives, it is unclear why the managers of long-short funds should do better in terms of the (risk-neutral) return criterion, which is not their objective. Our model indicates that although there may be a small utility gain from using the optimal strategy for the long-short fund, the optimal strategy is almost identical to the myopic one in terms of the return criterion. Moreover, no matter whether the optimal strategy or the myopic strategy is adopted, the long-short fund can yield a significantly higher return than the long-only fund, consistent with the empirical finding.

NUMERICAL METHODS FOR MEAN FIELD GAMES BASED ON GAUSSIAN PROCESSES AND FOURIER FEATURES

Mean field games (MFGs) [3, 4] study the behavior of a large population of rational and indistinguishable agents as the number of agents goes to infinity. The Nash equilibrium of a typical MFG is formulated by a coupled system of two partial differential equations (PDEs), a Hamilton–Jacobi–Bellman (HJB) equation and a Fokker–Plank (FP) equation. The HJB equation gives the value function of agents and the FP equation determines agents' distribution. Recently, MFG models have found widespread applications. Up to now, the well-posedness of MFGs is well understood in various settings and the first results date back to the original works of Lasry and Lions and have been given in the course of Lions at Collège de France (see [5]). However, very few MFG models admit explicit solutions. Hence, numerical computations of MFGs play an essential role in obtaining quantitative descriptions of underlying models.

Here, we propose two new algorithms to solve MFGs. The first one, called the Gaussian Process (GP) method, applies the algorithm in [2] to solve MFGs. The authors in [2] give a GP regression framework to solve nonlinear PDEs. The solution to a PDE is found by solving an optimal recovery problem, whose minimizer is viewed as a maximum *a posteriori* probability (MAP) estimator of a GP conditioned on the PDE evaluated at sample points. The main bottleneck of the GP method lies in the computation of the inverse of a square gram matrix, whose size is proportional to the product of the size of sample points times the number of linear operators in the PDE. To improve the performance, we propose the Fourier Features (FF) method, where the optimal recovery problem seeks minimizers in the space generated by sampled trigonometric functions. The FF method is based on the recent trend of approximating positive definite kernels with randomized trigonometric functions in Gaussian regressions, see [8–10]. Using the technique we propose, the dimension of the matrix needed to be inverted in the FF algorithm depends only on the number of sampled Fourier features, which is less than the size of sample points. The numerical experiments show that the FF algorithm reduces the precomputation time and the amount of storage with comparable accuracy compared to the GP method.

Meanwhile, we also prove the convergence of our algorithms. The proofs for the GP method are based on the compactness arguments in [2] and do not depend on any monotonicity condition that guarantees stability and uniqueness of the solution to a typical MFG. This feature implies the potential applications of the GP method to MFGs with non-monotone couplings. On the other hand, since the Fourier features space the FF method uses lacks compactness, the same arguments of the GP method cannot be adapted to the setting of the FF method. Instead, the Lasry–Lions monotonicity arguments provide the uniform bounds for the errors of numerical solutions and lead to the convergence of the FF algorithm. In future work, we plan to investigate the convergence of the FF method under the setting of MFGs without monotone couplings. We show the efficacy of our algorithms through experiments on a stationary MFG with a non-local coupling and on a time-dependent planning problem. We believe that the FF method can also serve as an alternative algorithm to solve general PDEs.

0.1. **Related Works.** By now, there have been various numerical methods for MFGs. Compared to the above-mentioned algorithms, our methods for MFGs admit the following features:

1. The GP and the FF algorithms are meshfree and flexible to the shape of domains, compared to algorithms in the first group mentioned above. Especially, when we solve MFGs in the whole Euclidean space, to approximate derivatives, mesh-based algorithms have to impose artificial conditions on the boundary of the domain we choose to work on. Our methods parameterize the values of derivatives

Date: March 11, 2022.
2010 *Mathematics Subject Classification.* 65L30, 65N75, 82C80, 35A01.
Key words and phrases. Mean field Games; Gaussian Process; Fourier Features.

evaluated at sample points and solve reformulated finite-dimensional minimization problems. Hence, we do not impose extra conditions to deal with derivatives near the boundary;
2. Compared to the neural network methods, our algorithms base on theories of RKHSs and Fourier series, for which the math backgrounds are well understood. On the other hand, our algorithms are equivalent to the neural network methods in the following perspective: a neural network with a single inner layer, with the activation functions being feature maps parameterized at sample points or being random sampled trigonometric functions, and with a linear output layer can be viewed as a function in RKHSs or the Fourier features space, and vice versa;
3. The choice of the kernel for the GP method and the selection of Fourier features have a profound impact on the convergence and the accuracy of our approximations. We leave the study of hyperparameter learning to future work.
4. The convergence of the GP method does not depend on the Lasry–Lions monotonicity condition of the coupling terms in MFGs, which is required by the monotone flow [1] and the Lagrangian methods [7]. Hence, we plan to study the application of the GP method to solve MFGs with displacement monotonicity [6] or non-monotone couplings in future work.
5. In general, it is less costly to compute a linear map of a Fourier feature than to calculate the same linear transformation of a kernel function. For instance, in some cases, we need to parametrize the linear operator $L = (1-\Delta)^{-1}(1-\Delta)^{-1}$ at sample points, where Δ is a Laplacian operator. In the GP method, the representer theorem gives expressions involving the computation of $LK_2(x,\cdot)$ for $x \in \mathbb{T}^2$, where K_2 is the kernel of the RKHS we choose to find an approximation for the probability density. Since $LK_2(x,\cdot)$ does not admit an explicit formula, one can use the fast Fourier transform to compute it numerically. On the other hand, for any $\omega \in \mathbb{N}^2$, we observe that $L\sin(\omega^T x) = \sin(w^T x)/(1+|\omega|^2)^2$ and $L\cos(\omega^T x) = \cos(w^T x)/(1+|\omega|^2)^2$ for $x \in \mathbb{T}^2$. Hence, if we choose trigonometric functions as features in the FF method, the action of L on a Fourier feature admits an explicit formula and is easier to compute.
6. The main bottleneck of the GP method is to compute the Cholesky decomposition of the gram matrix, whose size increases as the number of sample points grows. The FF method relieves the pressure by approximating functions in the FF space and by using the technique in Remark **??**. However, as the dimension of the problem at hand increases, we should also enlarge the FF space in the absence of information about the solution. Hence, a clever selection of features is required to apply the FF method in large dimensions. In this paper, we choose the Fourier series in the periodic settings and use random Fourier features for the non-periodic cases. We leave the study of other selections to future work.

References

[1] N. Almulla, R. Ferreira, and D. Gomes. Two numerical approaches to stationary mean-field games. *Dynamic Games and Applications*, 7(4):657–682, 2017.

[2] Y. Chen, B. Hosseini, H. Owhadi, and A. M. Stuart. Solving and learning nonlinear pdes with gaussian processes. *Journal of Computational Physics*, 2021.

[3] M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle. *Communications in Information and Systems*, 6(3):221–252, 2006.

[4] J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. I–le cas stationnaire. *Comptes Rendus Mathématique*, 343(9):619–625, 2006.

[5] P.-L. Lions. *Cours au Collège de France*, 2007-2013.

[6] A. R. Mészáros and C. Mou. Mean field games systems under displacement monotonicity. *arXiv preprint arXiv:2109.06687*, 2021.

[7] L. Nurbekyan and J. Saúde. Fourier approximation methods for first-order nonlocal mean-field games. *Portugaliae Mathematica*, 75(3):367–396, 2019.

[8] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, volume 3, page 5. Citeseer, 2007.

[9] A. Tompkins and F. Ramos. Periodic kernel approximation by index set fourier series features. In *Uncertainty in Artificial Intelligence*, pages 486–496. PMLR, 2020.

[10] F. X. Yu, A. T. Suresh, K. M. Choromanski, D. N. Holtmann-Rice, and S. Kumar. Orthogonal random features. *Advances in neural information processing systems*, 29:1975–1983, 2016.

On Asymptotically Arbitrage-free Approximations of the Implied Volatility

Masaaki Fukasawa

Abstract. Following-up Fukasawa and Gatheral (Frontiers of Mathematical Finance, 2022), we prove that the BBF formula, the SABR formula, and the rough SABR formula provide asymptotically arbitrage-free approximations of the implied volatility under, respectively, the local volatility model, the SABR model, and the rough SABR model.

1. Introduction

The implied volatility is one of the basic quantities in financial practice. The option market prices are translated to the implied volatilities to normalize in a sense their dependence on strike price and maturity. The shape of the implied volatility surface characterizes the marginal distributions of the underlying asset price process. Other than a flat surface corresponding to the Black-Scholes dynamics, no exact formula of the surface is available and so, various approximation formulae have been investigated. See [9] for a practical guide for the volatility surface.

One of the most famous and in daily use of financial practice is the SABR formula proposed by Hagan et al. [10] for the SABR model. After its original derivation by [10] based on a formal perturbation expansion, and a verification by [3] based on an asymptotic analysis of PDE, Balland [1] derived the formula (for the so-called lognormal case) by an elegant no-arbitrage argument. The no-arbitrage argument remains valid for non-Markovian models, and Fukasawa and Gatheral [8] derived an extension of the SABR formula to a rough SABR model, where the volatility process is non-Markovian.

It has been known that the implied volatility surface of an equity option market typically exhibits a power-law type term structure. Since classical local-stochastic volatility models including the SABR model are not consistent to such a term structure, the so-called rough volatility model has recently attracted attention, which is the only class of continuous price models that are consistent to the power-law; see Fukasawa [4, 6]. The rough SABR model of [8] (see also [5, 11, 7]) is a rough volatility model and the rough SABR formula derived in [8] explicitly exhibits a power-law term structure.

The aim of this paper is to verify the SABR and rough SABR formulae by the no arbitrage argument beyond the lognormal case.

References

[1] P. Balland, Forward smile, presentation at Global Derivatives, Paris, September 2006.

[2] H. Berestycki, J. Busca, and I. Florent, Asymptotics and calibration of local volatility models, *Quantitative Finance*, **2** (2002), 61–69.

[3] H. Berestycki, J. Busca, and I. Florent, Computing the implied volatility in stochastic volatility models, *Communications on Pure and Applied Mathematics*, **57** (2004), 1–22.

[4] M. Fukasawa, Short-time at-the-money skew and rough fractional volatility, *Quantitative Finance* **17** (2017), 189–198.

[5] M. Fukasawa, Hedging and Calibration for Log-normal Rough Volatility Models, presentation at the 10th World Congress of The Bachelier Finance Society, Dublin, 2018.

[6] M. Fukasawa, Volatility has to be rough, *Quantitative Finance 21* (2021), 1-8.

[7] M. Fukasawa, B. Horvath and P. Tankov, Hedging under rough volatility, preprint, arXiv:2105.04073.

[8] M. Fukasawa and J. Gatheral, A rough SABR formula, *Frontiers of Mathematical Finance*, **1** (2022), 81–97.

[9] J. Gatheral, *The Volatility Surface: A Practioner' s Guide*, 2006 (John Wiley & Sons Inc: Hoboken, NJ).

[10] P. Hagan, D. Kumar, A. Lesniewski, and D. Woodward, Managing smile risk, *Wilmott Magazine*, **1** (2002), 84–108.

[11] M. Musiela, Multivariate fractional Brownian motion and generalizations of SABR model. In: *Proceedings of the Conference "Options 45 Years After the Publication of the Black-Scholes-Merton Model"*, Jerusalem 2018. Available online at https://doi.org/10.5281/zenodo.4772004

[12] Y. Osajima, The asymptotic expansion formula of implied volatility for dynamic SABR model and FX hybrid model, Available at SSRN: https://ssrn.com/abstract=965265, 2007.

[13] D. Revuz and M. Yor, *Continuous martingales and Brownian motion*, Springer, 1999.

Graduate School of Engineering Science, Osaka University, 560-8531, Japan

Special Sessions in memory of Prof. Peter Carr

On Black’ s equation: from local volatility to local risk tolerance functions

Thaleia Zariphopolou

Abstract: In Carr et al. (1999), the authors developed an approach to construct closed form solutions of option prices for a family of local volatility models. The key ingredient is a fully non-linear equation that needs to be satisfied by the local volatility function. In a very different context, this equation also arises in the classical Merton problem for the space-time evolution of the local risk tolerance function. In this talk, I will discuss several unified properties of this equation and its role in distinct problems in mathematical finance.

Open Markets and Hybrid Jacobi Processes

David Itkin & Martin Larsson
Carnegie Mellon University

December 28, 2021

We propose a unified approach to several problems in Stochastic Portfolio Theory (SPT), which is a framework for equity markets with a large number d of stocks. Our approach combines *open markets*, where trading is confined to the top N capitalized stocks as well as the market portfolio consisting of all d assets, with a parametric family of models which we call *hybrid Jacobi processes*. We provide a detailed analysis of ergodicity, particle collisions, and boundary attainment, and use these results to study the associated financial markets. Their properties include stability of the capital distribution curve and unleveraged and explicit growth optimal strategies. The sub-class of *rank Jacobi models* are additionally shown to serve as the worst-case model for a robust asymptotic growth problem under model ambiguity and exhibit stability in the large-d limit. Our open market definition is a relaxation of existing definitions which is essential to make the analysis tractable.

SPT was introduced by Fernholz in [1] as a flexible way to model financial markets focusing on empirically observable phenomena. In the framework of SPT we takes a collection $\{S_i\}_{i=1}^d$ of stock capitalizations and the quantities of interest are the *market weights* $X_i := S_i/(S_1 + \cdots + S_d)$ which take values in the unit simplex

$$\Delta^{d-1} := \{x \in [0,\infty)^d : x_1 + \cdots + x_d = 1\}.$$

We fix an integer $N \in \{1, \ldots, d-1\}$ which will represent the size of our *open market*; that is a sub-market of the d-dimensional market introduced above. Specifically, at an fixed time $t \geq 0$ we only allow the investor to hold the assets which, at that time, occupy the top N ranks and we also allow investments in the full market portfolio. This framework differs from previous open market setups by allowing investment in the market portfolio. This is a crucial relaxation for our tractable analysis of the growth-maximization problem described below.

In a general semimartingale framework we analyze an investors growth-optimization problem in the open market of size N. We impose the standard *structure* and *finite growth* conditions shown in [2] to be necessary and sufficient for well-posedness of the growth-optimization problem. We additionally impose a condition on the interaction of the high and low capitalization stocks. To formulate this condition we introduce some notation. For $x \in \mathbb{R}^d$ define its *rank vector* $x_{()} \in \mathbb{R}^d$ to be the permutation of x satisfying $x_{(1)} \geq x_{(2)} \geq \cdots \geq x_{(d)}$. The condition then assumes that there exist processes $f_1, \ldots, f_d, g$ such that

$$d[X_{(k)}, X_{(l)}](t) = -f_l(t) f_k(t) g(t)\, dt \quad \text{for } k = 1, \ldots, N \text{ and } l = N+1, \ldots, d. \tag{1}$$

We stress that this condition does not restrict the instantaneous covariation *within* the collection of top N assets. Under these conditions we obtain the existence of a growth-optimal strategy in the open market and obtain an *explicit* expression for it in terms of the candidate growth-optimal portfolio in the unconstrained full market. Our setup is tractable in the sense that it is just as easy, or difficult, to solve as the classical growth-optimization problem in the full market.

Next we propose the use of a parametric model for the market weights, which we call *hybrid Jacobi processes*. Taking parameters $\alpha, \beta \in \mathbb{R}^d$ and $\sigma^2 > 0$ they are given by

$$dX_i(t) = \Big(\alpha_{r_i(t)} + \beta_i - X_i(t) \sum_{j=1}^d (\alpha_j + \beta_j)\Big)\, dt + \sigma \sum_{j=1}^d (\delta_{ij} - X_i(t)) \sqrt{X_j(t)}\, dW_j(t) \tag{2}$$

for $i = 1, \dots, d$ where δ_{ij} is the Kronecker delta and $r_i(t)$ is the rank occupied by name i at time t. The hybrid Jacobi processes satisfy the condition (1) since

$$d[X_{(k)}, X_{(l)}](t) = \sigma^2 X_{(k)}(t)(\delta_{kl} - X_{(l)}(t))\,dt; \qquad k,l = 1,\dots,d. \tag{3}$$

A crucial observation in SPT is the empirical stability of the *capital distribution curve* (see e.g. [1, Chapter 5]). This is captured by the hybrid Jacobi models since the ranked market weights $X_{()}$ are shown to be ergodic with explicitly given invariant density

$$q(y) \propto \sum_{\tau \in \mathcal{T}^d} \prod_{k=1}^{d} y_k^{\frac{1}{2\sigma^2}(\alpha_k + \beta_{\tau(k)}) - 1}; \qquad y \in \nabla_+^{d-1} := \{y \in \Delta^{d-1} : y_1 \geq \cdots \geq y_d > 0\}.$$

Here $\mathcal{T}^d$ are the set of permutations on $\{1,\dots,d\}$. We show that the open market induced by the market weight process (2) has a growth-optimal portfolio if and only if $\sum_{l=k}^{d} \alpha_l + \beta_{(l)} \geq \sigma^2$ for every $k = N+1,\dots,d$. In this case the optimal portfolio is explicitly given by

$$\hat{\pi}_i(t) = \begin{cases} X_i(t)\Big(1 - \dfrac{\sum_{j=1}^{d} \alpha_j + \beta_j}{2\sigma^2}\Big) + \dfrac{\alpha_{r_i(t)} + \beta_i}{2\sigma^2}, & r_i(t) \in \{1,\dots,N\}, \\ X_i(t)\Big(1 - \dfrac{\sum_{j=1}^{d} \alpha_j + \beta_j}{2\sigma^2}\Big) + \dfrac{X_i(t)}{\sum_{j=N+1}^{d} X_{(j)}(t)} \dfrac{\sum_{j=N+1}^{d} \alpha_j + \beta_{r_j(t)}}{2\sigma^2}, & r_i(t) \in \{N+1,\dots,d\}, \end{cases} \tag{4}$$

for $i = 1,\dots,d$. Here $\hat{\pi}_i(t)$ represents the proportion of wealth invested in X_i at time t.

When $\beta = 0$ we call (2) a *rank Jacobi process* and we can obtain robust properties of the model. In particular, we let $\Pi_{\geq}$ denote the set of all measures under which $X_{()}$ has invariant density q and instantaneous covariation given by (3) (as well as additional technical assumptions). The goal is to characterize the robust asymptotic growth rate

$$\lambda^N := \sup_{\pi \in \mathbb{O}^N} \inf_{\mathbb{P} \in \Pi_{\geq}} g(V^\pi; \mathbb{P}),$$

where $\mathbb{O}^N$ denotes all portfolios in the open market of size N, V^π is the investors wealth process induced by the portfolio π and

$$g(V^\pi; \mathbb{P}) := \sup\{\eta \in \mathbb{R} : \liminf_{T \to \infty} T^{-1} \log V^\pi(T) \geq \eta;\ \mathbb{P}\text{-a.s.}\}$$

is the investors asymptotic growth rate under the measure $\mathbb{P}$ when utilizing strategy π. We obtain an explicit integral representation for λ^N and show that the strategy (4) is robust growth optimal. That is $g(V^{\hat{\pi}}; \mathbb{P}) = \lambda^N$ for every $\mathbb{P} \in \Pi_{\geq}$.

Lastly, in the rank Jacobi case, under appropriate assumptions on the parameter α as d varies, we obtain weak convergence as $d \to \infty$ for the invariant measures of $X_{()}$. The limiting measure is supported on the Kingman simplex, $\nabla^\infty = \{y \in \mathbb{R}^\infty : y_1 \geq y_2 \geq \cdots > 0,\ \sum_{k=1}^{\infty} y_k = 1\}$, and is shown to be absolutely continuous to the Poisson–Dirichlet distribution with explicitly given Radon–Nikodym derivative. Financially, this corresponds to the empirical observation that the capital distributions curves have remained stable even as the number of assets in the market have grown over time.

References

[1] E. Robert Fernholz. *Stochastic portfolio theory*, volume 48 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 2002. Stochastic Modelling and Applied Probability.

[2] Ioannis Karatzas and Donghan Kim. Open markets. *Math. Finance*, pages 1–52, 2020.

Optimal brokerage contracts in Almgren-Chriss model with multiple clients

G.Alonso Alvarez, S. Nadtochiy, and K. Webster

February 2022

We consider N agents each of whom can trade a single risky asset (a constant riskless asset is also available) that follows the Almgren-Chriss model over the time interval $[0,T]$. In addition to the agents, we assume the presence of a single broker. Each agent makes a decision (once, before the trading starts) on whether he trades the asset directly or via the broker, and these decisions are represented by the vector $\theta \in \{0,1\}^N$: the i-th entry of θ is 1 if and only if the ith agent trades via the Broker. For convenience, we also denote by $\mathcal{N}(\theta) = \{n_1, \cdots, n_r\} \subset \{1, \cdots, N\}$, $0 < r < N$, the indices of the agents that trade via broker. We refer to the agents who trade via the broker as clients and to those who trade directly in the market as independent. The trading activity of an independent agent affects the price of the asset via the impact coefficients of this agent. The trading of a client is done via the broker and hence affect the price of the asset via the impact coefficients of the broker. Of course, the broker may be able to offer lower price impact to her clients, but she also charges each of them a fee for this service. Even though we refer to θ as the set of decisions of the agents, it is important to realize that these decisions are limited (or, partially controlled) by the broker who decides whether to offer a contract to a particular agent or not.

We fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and consider a standard Brownian motion B on this space. Let $\mathbb{F}^B$ be the filtration generated by the Brownian motion B. We define the set of admissible controls of a single agent as

$$\mathcal{U} := \{\nu \in L^2([0,T] \times \Omega), \quad \nu \text{ is } \mathbb{F}^B\text{-adapted}\}. \tag{1}$$

The (controlled) inventory of agent $i = 1, \dots, N$, who uses a control ν^i, is given by the process

$$X_t^i = x_0^i + \int_0^t \nu_s^i ds, \quad t \in [0,T].$$

Next, we recall the price process for the traded asset in the Almgren-Chriss model:

$$P_t = \mu t + \sigma B_t + \sum_{i=1}^N (1-\theta_i)\kappa_i \nu_t^i + \kappa_0 \sum_{i=1}^N \theta_i \nu_t^i + \sum_{i=1}^N (1-\theta_i)\lambda_i X_t^i + \lambda_0 \sum_{i=1}^N \theta_i X_t^i, \tag{2}$$

where $\{\kappa_j\}_{j=1}^N$, $\{\lambda_j\}_{j=1}^N$ are the coefficients of temporary and permanent price impacts of the agents, κ_0, λ_0 are the corresponding coefficients of the broker, $\sigma > 0$ is the volatility of the asset price, and $\mu \in \mathbb{R}$ is its drift (i.e., trading signal). For convenience, we assume that $\nu_T^i = 0$ for all i, so that the temporary impacts of the agents do not affect the terminal price.[1]

[1] This is needed to simplify the notation in (3)–(4), as the agents in our setting interact through both the permanent and the temporary impacts.

1

Let $\xi^{n_1}, \cdots, \xi^{n_r}$ be the fees that the broker charges to her clients. We assume that the fees are of the form $\xi^i = F^i(X^i, P)$ with measurable $F^i : H^1([0,T],\mathbb{R}) \times C([0,T],\mathbb{R}) \to \mathbb{R}$, where H^1 is the Sobolev space of order one, equipped with the usual norm. A client $i \in \mathcal{N}$ aims to maximize his expected profit:

$$J^{i,\theta}(\nu^i, \nu^{-i}, \xi^i) = \mathbb{E}\left(P_T X_T^i - \frac{\lambda_i}{2}(X_T^i)^2 - \int_0^T \nu_t^i P_t dt - \xi^i\right), \tag{3}$$

where $\nu^{-i} \in \mathcal{U}^{N-1}$ denotes the trading rate of the rest of agents. Similarly, an independent agent $i \notin \mathcal{N}$ maximizes his expected profit:

$$J^{i,\theta}(\nu^i, \nu^{-i}, \xi^i) = \mathbb{E}\left(P_T X_T^i - \frac{\lambda_i}{2}(X_T^i)^2 - \int_0^T \nu_t^i P_t dt\right) \tag{4}$$

We define $\mathcal{E}(\theta, \xi) \subset \mathcal{U}^{\otimes N}$ as the set of all agents' strategies that form Nash equilibria in the game defined by (3)–(4).

The objective of the broker is given by the sum of expected fees in the best equilibrium attainable with these fees:

$$J_P^\theta(\xi) := \sup_{\nu \in \mathcal{E}(\theta,\xi)} \mathbb{E} \sum_{j \in \mathcal{N}(\theta)} \xi^j. \tag{5}$$

In the above, we make the standard assumption that, given a set of admissible contracts, the agents will choose an equilibrium that is best for the principal among all attainable equilibria. To ensure that $\mathcal{E}(\theta, \xi) \neq \emptyset$ and that the agents' reservation values are met, we introduce the set of admissible fees of the broker:

$$\Sigma(\theta) := \{(\xi^j)_{j \in \mathcal{N}(\theta)} : \mathcal{E}(\theta, \xi) \neq \emptyset, \text{ and } J^j(\bar{\nu}, \xi^j) \geq R^{j,\theta}, \quad \forall j \in \mathcal{N}(\theta), \forall \bar{\nu} \in \mathcal{E}(\theta, \xi)\}, \tag{6}$$

where $R^{j,\theta}$ is the reservation value of agent $j \in \mathcal{N}(\theta)$. Thus, we obtain the following "local" maximization problem for the broker, given a set of decisions θ:

$$V_P^\theta := \sup_{\xi \in \Sigma(\theta)} J_P^\theta(\xi) = \sup_{\xi \in \Sigma(\theta)} \sup_{\nu \in \mathcal{E}(\theta,\xi)} \mathbb{E} \sum_{j \in \mathcal{N}(\theta)} \xi^j. \tag{7}$$

In our work we find the optimal contract ξ^* that attains the above supremum. Note that, in the present setting we do not assume that $R^{j,\theta}$ is given endogenously, because there is in fact a very natural endogenous definition of the reservation value of each agent. However, as the above optimal contract problem is of first-best type, with risk-neutral preferences, the value of $R^{j,\theta}$ is not important for the form of the optimal contract. Note that the "global" optimization problem of the broker is to find an optimal set $\theta^* \in \{0,1\}^N$ and the respective optimal fees ξ^*, which amounts to solving

$$V_P := \sup_{\theta \in \{0,1\}^N} V_P^\theta. \tag{8}$$

The above is a discrete optimization problem. We do not provide a complete solution to this problem herein, assuming instead that it can be solved by an exhaustive search in case of a reasonably small N or of a smaller subset of admissible $\{\theta\}$. However, even to perform such an exhaustive search, one needs to have a numerically tractable representation of the value of V_P^θ for each θ. The latter can be represented through a system of forward-backward ordinary differential equations. Finally, using numerical simulations we discuss different financially relevant questions.

Optimal Consumption with Loss Aversion and Reference to Past Spending Maximum

Xun LI Xiang YU* Qinyi ZHANG

Abstract

This paper studies an optimal consumption problem for a loss-averse agent with reference to past consumption maximum. To account for loss aversion on relative consumption, an S-shaped utility is adopted that measures the difference between the non-negative consumption rate and a fraction of the historical spending peak. We consider the concave envelope of the realization utility with respect to consumption, allowing us to focus on an auxiliary HJB variational inequality on the strength of concavification principle and dynamic programming arguments. By applying the dual transform and smooth-fit conditions, the auxiliary HJB variational inequality is solved in closed-form piecewisely and some thresholds of the wealth variable are obtained. The optimal consumption and investment control of the original problem can be derived analytically in the piecewise feedback form. The rigorous verification proofs on optimality and concavification principle are provided.

Keywords: Loss aversion, optimal relative consumption, path-dependent reference, concave envelope, piecewise feedback control, verification arguments

*Department of Applied Mathematics, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong. Email:xiang.yu@polyu.edu.hk

Optimal consumption with reference to past spending maximum

Shuoqing Deng, Xun Li, Huyên Pham and Xiang Yu

January 28, 2020

This paper studies an infinite-horizon optimal consumption problem under exponential utility, together with non-negativity constraint on consumption and the reference point to the past consumption peak. The performance is measured by the distance between the current consumption rate and a fraction of the historical consumption maximum, which renders the control problem path dependent. To apply dynamic programming arguments, the consumption running maximum process is chosen as an auxiliary state process. The associated Hamilton-Jacobi- Bellman (HJB) equation can be expressed in a piecewise manner in three different regions. By employing the dual transform of the two dimensional value function, we obtain the explicit classical solution of the dual PDE using endogenous boundary conditions and smooth-fit principle. The optimal investment and consumption strategies are provided in feedback form via a complete verification theorem. Some numerical sensitivity analysis and comparative statics are also presented.

Optimal ecological transition path of a credit portfolio distribution with banking risk assessment

Emmanuel Gobet[1] and Clara Lage[2]

[1]École Polytechnique
[2]École Polytechnique

June 22, 2021

Finance for climate change or climate finance, is new diverse field of research. It is in some instances discussed separately or often times integrated with related concepts of green finance, sustainable finance, or low-carbon finance. Particularly, climate risks are being consolidated as an important component of the low carbon transition.

According to the Basel Committee on Banking Supervision (BCBS), the impact of climate change in financial stability can be mainly divided in physical and transition risks [1]. Physical risks concern the effect of climate change in the value of assets. In this context, extreme losses need to be reevaluated including climate risks that depend on the sector and geographical location of each economic activity. Transition risks are those risks related to the transformation towards a low-carbon economy, it concerns social and political instability of policies in this period as well as technological changes.

Physical risks are recently being integrated to risk analysis in researches and by banks or insurance companies, mainly by adapting or changing models for assets losses ([2], [3]). On the other side, transition risks still difficult to model in a mathematical sense, since they depend on hardly predictable policies.

In the search to mathematically describe the transition risks in the point of view of a private actor, we consider an specific problem of a bank portfolio transition in terms of carbon score of loans and other assets and propose an optimization model adapted for a transition to a low-carbon portfolio.

The Multidate model consists in a discrete multistage transition between the actual portfolio and a target low carbon portfolio. The optimization model aims to minimize the transition cost - that can be adapted depend on the data or financial interpretation - equilibrating carbon score and risk of each portfolio on the trajectory.

In this framework, assets are characterized by an activity sector (that can include geographical location) and a carbon score. Policies are then distributions in the space of assets where capital is considered as fixed by simplicity. The model incorporates ideas of *Optimal Transport theory* considering the distance between two policies as the Kantorovich relation of the Monge problem, using a cost adapted to the financial interpretation. It also integrates a multifactor *Merton Firm model* concerning portfolio risk.

The present optimization problem is actually a bilinear saddle point problem that can be shown to have a compact and convex set of solutions in our particular case. We apply the Primal-Dual Method (PD) to solve this optimization problem and we provide a numerical analysis including convergence and a sensibility analysis of numerical solutions.

References

[1] Basel Committee on Banking Supervision. *Climate-related risk drivers and their transmission channels.* 2021.

[2] Michel De Lara and Doyen Luc. *Sustainable Management of Natural Resources Mathematical Models and Methods,* Jan. 2008. ISBN: 978-3-540-79073-0. DOI: 10.1007/978-3-540-79074-7.

[3] Xiao-Chen Yuan et al. "Risk management of extreme events under climate change". In: *Journal of Cleaner Production* 166 (2017), pp. 1169–1174. ISSN: 0959-6526. DOI: https://doi.org/10.1016/j.jclepro.2017.07.209. URL: https://www.sciencedirect.com/science/article/pii/S0959652617316645.

Optimal Execution via Continuous-Time Reinforcement Learning

Boyu Wang[†] Xuefeng Gao[‡] Lingfei Li[§]

Abstract

We solve the optimal execution problem with the mean-quadratic variation objective under the Almgren-Chriss framework by reinforcement learning (RL). We relax the control to be a distribution and consider an entropy-regularized objective to encourage exploration. We obtain a closed-form formula for the optimal value function of the relaxed stochastic control problem and show that the optimal control distribution must be Gaussian. We also prove a policy improvement theorem. Based on these results, we develop an actor-critic type RL algorithm that only involves a few parameters to learn. Our algorithm converges fast and it outperforms the TWAP strategy by a large margin on two limit order book markets that are constructed from historical data and a limit order book model, respectively.

[†]Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. E-mail: bywang@se.cuhk.edu.hk

[‡]Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. E-mail: xfgao@se.cuhk.edu.hk

[§]Corresponding author. Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. E-mail: lfli@se.cuhk.edu.hk

Optimal Execution with Stochastic Delay

Álvaro Cartea[a,b], Leandro Sánchez-Betancourt[c,d]

[a] *Mathematical Institute, University of Oxford*
[b] *Oxford-Man Institute of Quantitative Finance*
[c] *Department of Mathematics, King's College London*
[d] *LMAX Exchange, London*

Abstract

We show how traders use marketable limit orders (MLOs) to liquidate a position over a trading window when there is latency in the marketplace. MLOs are liquidity taking orders that specify a price limit and are for immediate execution only; however, if the price limit of the MLO precludes it from being filled, the exchange cancels the order. We frame our model as an impulse control problem with stochastic latency where the trader controls the times and the price limits of the MLOs sent to the exchange. We show that impatient liquidity takers submit MLOs that may walk the book (capped by the limit price) to increase the probability of filling the trades. On the other hand, patient liquidity takers use speculative MLOs that are only filled if there has been an advantageous move in prices over the latency period. Patient traders who are fast do not use their speed to hit the quotes they observe, nor to finish the execution programme early; they use speed to complete the execution programme with as many speculative MLOs as possible. We use foreign exchange data to implement the random-latency-optimal strategy and to compare it with four benchmarks. For patient traders, the random-latency-optimal strategy outperforms the benchmarks by an amount that is greater than the transaction costs paid by liquidity takers in foreign exchange markets. Around news announcements, the value of the outperformance is between two and ten times the value of the transaction costs. The superiority of the strategy is due to both the speculative MLOs that are filled and the price protection of the MLOs.

Keywords: Latency, impulse control theory, high-frequency trading, algorithmic trading.

---

*LSB acknowledges support from (a) Oriel College, Oxford, (b) the Mathematical Institute, Oxford, (c) Consejo Nacional de Ciencia y Tenconología (CONACyT), Ciudad de México, and (d) LMAX Exchange, London.

Email addresses: `alvaro.cartea@maths.ox.ac.uk` (Álvaro Cartea), `leandro.sanchez-betancourt@kcl.ac.uk` (Leandro Sánchez-Betancourt)

Preprint submitted to TBA *March 14, 2022*

We study the optimal design of a menu of funds by a fund family (such as Fidelity or Vanguard) that is required to use linear pricing and does not observe the beliefs of investors regarding one of the risky assets. The optimal menu solving this adverse selection problem involves bundling of assets and can be explicitly constructed from the solution to a calculus of variations problem that optimizes over the indirect utility that each type of investor receives. We provide a complete characterization of the optimal menu and show that the need to maintain incentive compatibility leads the manager to offer funds that are inefficiently tilted towards the asset that is not subject to the information friction. This provides a new rationale for the empirically observed phenomena of “closet indexing” and “home bias”.

Speaker: Prof. Zhongfei Li, Southern University of Science and Technology

Title: Optimal Housing, Consumption and Portfolio Choices

Abstract: This paper investigates a life-cycle housing, investment and consumption problem for a household who has stochastic income, consumes on non-housing goods, invests in financial assets, and plan to buy a house over the life cycle for dwelling. When buying a house, the household is required to pay a fraction of the housing value as the down payment and finances the rest through a mortgage loan. The household obtains utility from the consumption and renting a house before buying a house, and from the consumption and housing units after buying a house. The household aims to maximize the expected utility during her life cycle, and to find the optimal consumption and investment strategies before and after buying a house, the optimal time to buy a house and the optimal housing units. The optimization problem is solved by using the extended real option method. The derived results show that, besides including the classic Merton-type strategy, the investment strategy before buying a house includes an income adjustment term and a hedging demand related to the house price, while the one after buying a house contains the income adjustment term and the loan payment. The human capital (loan repayment) finances (reduces) the consumption and investment.

Optimal Investment under Block-Shaped Order Books

Nan Chen Min Dai Qiheng Ding Chen Yang

We study an optimal investment problem of a CARA investor trading in a block-shaped limit order book market, which synergizes three key features of market microstructure: bid-ask spread, market depth, and resilience. One salient feature of our model is its finite resilience, that is, after being consumed, the liquidity in the limit order book is replenished gradually, and a higher resilience means a faster replenishment speed and thus a higher liquidity. We derive an explicit solution under a Bachelier model for the fundamental value of the asset, in which the investor's optimal trading strategy is characterized by a buy region, a sell region, and a no-transaction region separated by two boundaries which admit an explicit form. As resilience goes to infinity, our model is reduced to the classic model in Almgren and Chriss (2001).

Specifically, we assume the fundamental price of the asset S_t follows the stochastic process

$$dS_t = \mu dt + \sigma dW_t, \quad t \geq 0. \tag{1}$$

The market is operated through an LOB. Traders submit to the LOB their supply or demand of the asset in the form of limit orders. We consider a block-shaped LOB in this paper as shown in Figure 1. S_t sits in between unfulfilled buy and sell orders. On its two sides are the bid and ask blocks of orders extending to infinite. The densities of limit orders in both the blocks are assumed to be given by two step functions, i.e.,

$$q_a(P;t) = \begin{cases} 0, & P < A_t; \\ q_a, & P \geq A_t, \end{cases} \quad q_b(P;t) = \begin{cases} 0, & P > B_t; \\ q_b, & P \leq B_t, \end{cases} \tag{2}$$

where A_t and B_t are the best ask and bid prices at time t, repsectively.

Define D_t^A (D_t^B) to be the deviation of current ask price A_t (bid price B_t) from the fundamental price S_t, i.e., $D_t^A = A_t - S_t$ ($D_t^B = S_t - B_t$), then

$$dD_t^A = -\rho_A D_t^A dt + \frac{1}{q_a} dL_t, \quad dD_t^B = -\rho_B D_t^B dt + \frac{1}{q_b} dM_t, \tag{3}$$

where dL_t and dM_t corresponds to the cumulative amounts of purchase L_t and sale M_t given by

$$L_t = \int_0^t l_s ds + \sum_{0 \leq s < t} \Delta L_s, \quad M_t = \int_0^t m_s ds + \sum_{0 \leq s < t} \Delta M_s, \tag{4}$$

with $l_t \geq 0$ for the purchase rate, $\Delta L_t \geq 0$ for a lump sum purchase, and similarly for m_t and ΔM_t. The equation (3) means that new sell limit orders come into the book at the new ask price A_t at the rate of $\rho_A D_t^A$, and new buy limit orders come into the book at the new bid price B_t at the rate of $\rho_B D_t^B$.

The objective of the trader is to maximize her utility at the terminal of the trading horizon $[0, T]$. The terminal wealth owned by the trader is $S_T X_T + Y_T$, where X_t and Y_t are the amount of asset holdings and cash respectively. Therefore the optimzation problem is

$$\max_{(l_t, \Delta L_t), (m_t, \Delta M_t)} \mathbb{E}[U(S_T X_T + Y_T) | X_0, Y_0, S_0, D_0^A, D_0^B] \tag{5}$$

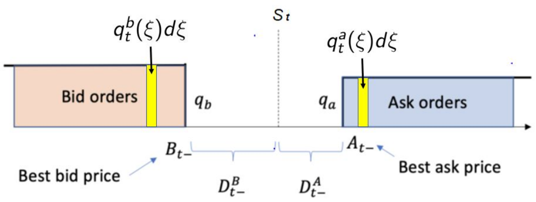

Figure 1: The block-shaped limit order book model

1

with $U(x) = 1 - \exp(-\gamma x)$. It is a singular control problem. The characterization of the state space is given in Theorem 1, and the explicit solution is given in Theorem 2.

Theorem 1. *The value function V is a viscosity solution of the following variational inequality*

$$\max\Big\{ \underbrace{\mathcal{L}V}_{\textit{No-trade}},\ \underbrace{\mathcal{B}V}_{\textit{Buy}},\ \underbrace{\mathcal{S}V}_{\textit{Sell}} \Big\} = 0, \tag{6}$$

where

$$\begin{aligned}
\mathcal{L}V &= \frac{\partial V}{\partial t} + \mu\frac{\partial V}{\partial S} + \frac{1}{2}\sigma^2\frac{\partial^2 V}{\partial S^2} - \rho_A D^A\frac{\partial V}{\partial D^A} - \rho_B D^B\frac{\partial V}{\partial D^B},\\
\mathcal{B}V &= \frac{1}{q_a}\frac{\partial V}{\partial D^A} + \frac{\partial V}{\partial X} - (S+D^A)\frac{\partial V}{\partial Y},\\
\mathcal{S}V &= \frac{1}{q_b}\frac{\partial V}{\partial D^B} - \frac{\partial V}{\partial X} + (S-D^B)\frac{\partial V}{\partial Y},
\end{aligned}$$

on $[0,T]\times\mathcal{D}$ and $\mathcal{D} := \mathbb{R}\times\mathbb{R}\times\mathbb{R}\times\mathbb{R}^+\times\mathbb{R}^+$ with $V(T,S,X,Y,D^A,D^B) = 1-\exp\{-\gamma[Y+SX]\}$.

Theorem 2. *The state space is divided as*
(i) buy region: $\frac{X_t - \frac{\mu}{\sigma^2\gamma}}{D_t^A} \le -h(t)$,
(ii) sell region: $\frac{X_t - \frac{\mu}{\sigma^2\gamma}}{D_t^B} \ge h(t)$,
(iii) no-transaction region: $\frac{\mu}{\sigma^2\gamma} - h(t)D_t^A < X_t < \frac{\mu}{\sigma^2\gamma} + h(t)D_t^B$.
where $h(t)$ is the solution to a Riccati ODE and has an explicit form.

Our result shows a significant impact of market resilience on the optimal trading strategy. In the presence of market resilience, it is optimal to target a delicate trade-off between reducing suboptimal risk exposure and reducing liquidity cost. On the one hand, the investor is incentivized to trade earlier for a better risk exposure. On the other hand, the investor is also motivated to wait for a period and trade at a lower liquidity cost due to market resilience. If the liquidity cost is high, the investor should defer trading. We find that resilience and market depth play different roles in affecting the optimal trading strategies. As market depth increases, the market provides more liquidity, and the investor trades more aggressively, in the sense that both the current and future trading size are greater. In contrast, as market resilience increases, since the investor knows that the best trading price will recover quickly and would like to benefit from the future smaller spread, the current trading size decreases, and after a short time the investor will trade faster to enjoy the benefit of resilience.

We also study an extension of the above model that incorporates return-predicting signals. In this case, we derive an asymptotic expansion for the strategies under small signal changes despite the lack of an explicit form. The return signal μ now follows the Vasicek model

$$d\mu_t = \varepsilon[(a - b\mu_t)dt + \beta d\widetilde{W}_t], \tag{7}$$

where ε is the asymptotic parameter. The existing literature shows that the investor exploits the signal by trading towards a moving target which is a weighted average of both the current and future expected portfolios. By combining market resilience and return signal, our extended model unveils the additional trade-off between the signal and resilience. Specifically, while it is still optimal to rebalance the portfolio towards the moving target

$$Aim = (1-p(t))\frac{\mu_t}{\sigma^2\gamma} + p(t)\frac{a/b}{\sigma^2\gamma}, \tag{8}$$

the investor should only do so when the portfolio is sufficiently away from the target. Furthermore, for a lower resilience, the investor will put more weight on the expected future portfolio, since the investor concerns more about the future as the degree of current liquidity is lower compared to future.

Due to the increased fragmentation of financial markets, exchanges have needed to find new ways to attract liquidity. One such way is through maker-taker fee systems, wherein a large market maker is offered a fee rebate for executed limit orders. The problem of finding the optimal incentive contract for the exchange to offer market makers was solved by El Euch et. al. in [1]. They used a principal-agent approach, in which the agent is the market maker, and the principal is the exchange who cannot directly control the liquidity offered by the agent (in the form of the bid-ask spreads). The contract is allowed to depend only on the order flow and price process, but not the bid-ask spreads directly. Since the principal cannot directly contract the bid-ask spreads set by the agent, he must anticipate how the agent will respond to a given compensation. In El Euch et. al., the principal makes the *subgame perfect* assumption that the agent will respond optimally to any contract offered to her.

Here, we are interested in the case that the market maker tries to obtain a more favorable contract by responding, or at least threatening to respond, sub-optimally to non-preferred contracts. This could, for instance, occur if a market maker is active on more than one exchange; if the contract offered on one exchange is not preferred, the market maker may act to encourage submission of orders to another exchange on which the market maker's contract is more favorable. Allowing the market maker to have this bargaining power leads to a greater variety of contracts that can be offered in equilibrium, and enables her to demand a contract that gives her greater utility than the one found in [1]. In this work we introduce a notion of equilibrium that captures this bargaining power.

In the model of El Euch et. al., we characterize the contracts that can be offered in equilibrium. We compute the equilibrium contract which is most favorable to the market maker in a quasi-explicit from. We find that the addition of bargaining power to the market maker does not affect the equilibrium bid-ask spreads, and therefore provides the same liquidity improvement and reduction in trading costs as the situation without bargaining power. The negotiation between the market maker and exchange only affects the size of a fixed payment between them.

References

[1] O. El Euch, T. Mastrolia, M. Rosenbaum, N. Touzi. Optimal make-take fees for market making regulation. *Mathematical Finance* 31, pp109-148 (2020).

Invited Session: stochastic modeling of information in economics and finance

Title: Optimal reinsurance via BSDEs in a partially observable contagion model

Speaker: Claudia Ceci

Abstract: We study an optimal reinsurance problem under the criterion of maximizing the expected utility of terminal wealth when the loss process exhibits jump clustering features and the insurance company has restricted information about the claims arrival intensity. By solving the associated filtering problem we reduce the original problem to a stochastic control problem under full information. Since the classical Hamilton-Jacobi-Bellman approach does not apply, due to the infinite dimensionality of the filter, we choose an alternative approach based on Backward Stochastic Differential Equations (BSDEs). Precisely, we characterize the value process and the optimal reinsurance strategy in terms of a BSDE driven by a marked point process. Based on joint work with M. Brachetta, G. Callegaro and C. Sgarrra.

Optimal risk transfer in insurance networks and its implications for systemic risk

Walter Farkas
University of Zurich, Department of Banking and Finance
ETH Zürich, Department of Mathematics
walter.farkas@bf.uzh.ch

Andreas Kull
University of Basel, Department of Mathematics and Computer Science
andreas.kull@unibas.ch

Alexander Smirnow
University of Zurich, Department of Banking and Finance
alexander.smirnow@bf.uzh.ch

Abstract

This paper aims to shed new light on risk transfer as a key risk and capital management tool. Insurance companies can transfer their risk to other entities, like reinsurance companies, in order to mitigate loss potential and reduce the cost of capital. However, risk sharing will potentially increase the interconnectedness between insurers in a network and thus contribute to higher systemic risk in that network.

For a single insurance company optimising risk transfer is straightforward as it means to compare the price for risk transfer to its benefits in terms of mitigated losses and reduced cost of risk capital. For a network of insurers, however, the optimisation can take different perspectives. In the context of a general network of insurance companies with passive and active reinsurance, N. Ettlin, A. Kull, and A. Smirnow have shown in a recent paper [1] that the optimisation of profitability across this network can be achieved through risk transfer. More precisely, risk based capital of each entity in the network is calculated as the Conditional Expected Shortfall. It is conditioned on the total net underwriting profit generated by all insurers and then the sum over all risk based capital holdings is minimised.
In this setup, uniqueness is not possible without further assumptions. However, using concepts from cooperative game theory, they showed that a unique risk transfer scheme exists and results in the Aumann-Shapley allocation, which in the sense of M. Denault [2] is considered fair and optimal.

While the profitability of the network is optimised and risk is distributed across the network, improving the resilience at the level of individual insurers, it remains to assess the systemic risk in this network. Risk transfer may lead to a reduction of overall risk capital in the network and thus, contribute to a higher systemic risk exposure. In this paper, we adjust the above

mentioned proposed network model to the network model of L. Eisenberg and T. H. Noe [3] so that we incorporate the possibility of stand-alone and contagious defaults. We use the fair optimal risk allocation of N. Ettlin, A. Kull, and A. Smirnow from [1] and consider several systemic risk measures. We show that while the cost of capital is reduced by the optimisation, the systemic risk in the network might rise as a consequence of the increased interconnectedness between the insurance companies.

References

[1] Ettlin, Nicolas and Kull, Andreas and Smirnow, Alexander (2019), "Optimal risk-sharing across a network of insurance companies." *Preprint*, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3417016`, *Submitted.*

[2] Denault, Michel (1999), "Coherent allocation of risk capital", *Journal of Risk*, vol. **4**, pp. 1 - 34.

[3] Eisenberg, Larry and Noe, Thomas H. (2001), "Systemic Risk in Financial Systems", *Management Science*, vol. **47**(2), pp. 236-249.

Abstract

Problem Definition: To develop supply chain, manufacturers need to contract with stockists to distribute goods, manage inventory, service retailers, and even make market. In this project, we investigate (i) how a firm selects stockist and designs incentive contract under asymmetric information and learning cost; (ii) how a stockist operates under career concern and switching cost. We study these questions in an emerging market, where the level of mistrust is high, learning and contract switching are costly, and competition is intense. **Methodology:** We derive a unified two-period model with one firm and two stockists, incorporating bandit selection, Bayesian learning, contract theory, and structural estimation. We solve the model explicitly and provide characterizations for the optimal contract design and optimal bandit selection. Using contract and sales data gathered from an Indian potato chips manufacture, we calibrate the model through simulated method of moments and perform counterfactual analysis. **Results**: Our model shows the optimal contract consists of three components: competition effect, career concerns, and compensation. Competition and career concern have quite opposite effects on both firm and stockists' decision-making. Competition induces stockists to put more effort to win the contracts, and therefore allows the firm to offer a lower level of performance-based payoff. In contrast, career concern motivates stockists to make less effort to win the future contract since all surplus is distributed to the firm, and thus firm needs to offer a contract with higher performance-based payoff. We show by using our structural estimation that there exist two types of stockists in the India supply chain: one has higher mean ability but also higher variance compared to the other. Furthermore, the low-productive stockist faces higher switching cost and production cost, indicating that switching to another firm is difficult, and exerting effort is costly. From the quantitative analysis, we show firm's stockist selection faces trade-offs between stockist's mean ability and associated risks, such as idiosyncratic risk and learning risk. **Implications**: Our analysis is helpful for manufactures to select stockists and design compensation contracts in supply chain. We also provide an analysis of data from India, and offer managerial insights of developing efficient supply chain in emerging markets.

Optimal stopping under model ambiguity: A time-consistent equilibrium approach

Xiang Yu
Department of Applied Mathematics
The Hong Kong Polytechnic University

Abstract

An unconventional approach for optimal stopping under model ambiguity is introduced. Besides ambiguity itself, we take into account how *ambiguity-averse* an agent is. This inclusion of ambiguity attitude, via an α-maxmin nonlinear expectation, renders the stopping problem time-inconsistent. We look for subgame perfect equilibrium stopping policies, formulated as fixed points of an operator. For a one-dimensional diffusion with drift and volatility uncertainty, we show that any initial stopping policy will converge to an equilibrium through a fixed-point iteration. This allows us to capture much more diverse behavior, depending on an agent's ambiguity attitude, beyond the standard worst-case (or best-case) analysis. In a concrete example of real options valuation under model ambiguity, all equilibrium stopping policies, as well as the *best* one among them, are fully characterized under appropriate conditions. It demonstrates explicitly the effect of ambiguity attitude on decision making: the more ambiguity-averse, the more eager to stop—so as to withdraw from the uncertain environment. The main result hinges on a delicate analysis of continuous sample paths in the canonical space and the capacity theory. To resolve measurability issues, a generalized measurable projection theorem, new to the literature, is also established. Joint work with Yu-Jui Huang.

1

Optimal Stopping via Randomized Neural Networks

Calypso Herrera Florian Krach Pierre Ruyssen Josef Teichmann

In this paper, we propose two neural network based algorithms to solve the optimal stopping problem for Markovian settings: a backward induction and a reinforcement learning approach. The idea is inspired by randomized neural networks (Cao et al., 2018; Huang et al., 2006). Instead of learning the parameters of all layers of the neural network, those of the hidden layers are randomly chosen and fixed and only the parameters of the last layer are learned. Hence, the non convex optimization problem is reduced to a convex problem that can be solved with linear regression. The hidden layers form random feature maps, which can be interpreted as random basis functions. In particular, in this paper we show that there is actually no need for complicated or a large number of basis functions. Our algorithms are based on the methods proposed by Longstaff and Schwartz (2001) (backward-induction approach) and Tsitsiklis and Van Roy (2001) (reinforcement learning approach). The difference is that we use a randomized neural network instead of a linear combination of basis functions. However, a randomized neural network can also be interpreted as a linear combination of random basis functions. On the other hand, our algorithms can also be interpreted as the neural network extensions of these methods, where not the entire neural network but only the last layer is trained. Moreover, we provide a randomized recurrent neural network approach for non-Markovian settings. We compare our algorithms to the most relevant baselines in terms of accuracy and computational speed in different option pricing problems. With only a fraction of trainable parameters compared to existing methods, we achieve high quality results considerably faster.

Methods. We approximate American options with payoff function g by Bermudan options and assume for simplicity that we have N equidistant exercise dates. Moreover, we assume to have access to a procedure to sample discrete paths of X under a fixed martingale measure $\mathbb{Q}$. Therefore, we can sample m realizations of the stock price paths, where the i-th realization is denoted by $x_0, x_1^i, x_2^i, \dots, x_N^i$, with the fixed initial value x_0. For each realization, the cash flow realized by the holder when following the optimal stopping strategy is given by the backward recursion

$$p_N^i := g(x_N^i), \qquad p_n^i := \begin{cases} g(x_n^i), & \text{if } g(x_n^i) \geq c_n(x_n^i), \\ \alpha p_{n+1}^i, & \text{otherwise.} \end{cases}$$

As p_1^i are samples of $\alpha^{\tau_1 - 1} g(X_{\tau_1})$, we have by the strong law of large numbers that the price of the Bermudan option is almost surely given by

$$U_0 = \max\left(g(X_0), \lim_{m\to\infty} \frac{1}{m} \sum_{i=1}^{m} \alpha p_1^i \right). \tag{1}$$

In our *randomized least squares Monte Carlo (RLSM)* algorithm, we approximate the continuation value $c_n(x_n^i) = \mathbb{E}[\alpha U_{n+1} | X_n = x_n^i]$ by a randomized neural network $c_{\theta_n}(x) := A_n^\top \sigma(Ax + b) + b_n$, i.e. where (A, b) are randomly sampled and fixed weights, σ is an activation function and $\theta_n = (A_n, b_n)$ are learned with ordinary least squares in a backward recursion.

Instead of approximating the continuation value at each single date with a different NN, we can directly learn the continuation function which also takes the time as argument. This we do in our *randomized fitted Q-iteration (RFQI)* algorithm, where we define $c_\theta(n, x) = A_\theta \sigma(A\tilde{x}_n + b) + b_\theta$, where $\theta = (A_\theta, b_\theta)$ are trainable and (A, b) are randomly sampled and fixed. Instead of using backward recursion (with N steps), we iteratively improve the approximation c_θ. More precisely, we start with some (random) initial weight θ_0 and then iteratively improve it by minimizing the difference between the continuation function c_{θ_ℓ} and the prices p computed with the previous weight $\theta_{\ell-1}$. Moreover, differently than for RLSM, we use the continuation value for the decision whether to continue *and* for the approximation of the discounted future price, as in (Tsitsiklis and Van Roy, 2001).

For non-Markovian processes, the continuation value $c_n := \mathbb{E}[\alpha g(X_{n+1}) \,|\, \mathcal{F}_n]$, where $\mathcal{F}_n$ denotes the information available up to time n, depends on the entire history. In the *randomized recurrent LSM (RRLSM)* algorithm we therefore replace the randomized feed-forward neural network of LSM by a randomized recurrent neural network (randomized RNN), which can utilize the entire information of the path up to the current time. In particular, we define $h_{-1} = 0$ and

$$\begin{cases} h_n &= \sigma(A_x x_n + A_h h_{n-1} + b), \\ c_{\theta_n}(h_n) &= A_n^\top h_n + b_n = \theta_n^\top \phi(x_n, h_{n-1}), \end{cases} \tag{2}$$

where (A_x, A_h, b) are randomly sampled and fixed and $\theta_n = (A_n, b_n)$ are learned with OLS in a backward recursion.

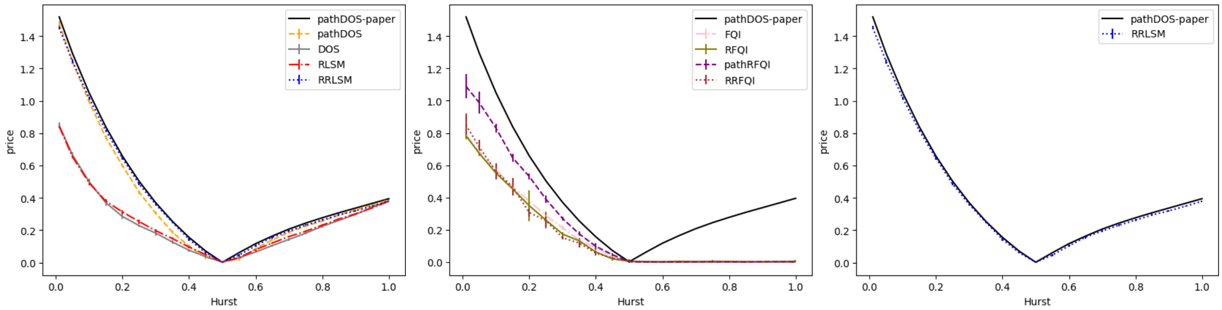

Figure 1: Left: algorithms processing path information outperform. Middle: reinforcement learning algorithms do not work well in non-Markovian cases. Right: RRLSM achieves similar results as reported in (Becker et al., 2019), while using only 20K paths instead of 4M for training which took only $1s$ instead of the reported $430s$.

Theoretical results. For all three proposed methods we can prove guarantees that the price computed with them converges to the correct price of the discretized American option.

Experiments. To test the methods in a Markovian case, we compare RLSM and RFQI to three backward induction algorithms (least squares Monte Carlo (LSM) (Longstaff and Schwartz, 2001), a neural network version of LSM (NLSM) (Lapeyre and Lelong, 2021) and deep optimal stopping (DOS) (Becker et al., 2019)) and one reinforcement learning approach (fitted Q-iteration (FQI) (Tsitsiklis and Van Roy, 1997)) on different underlying stock models with different parameters, different payoffs and for varying spot prices and number of underlying stocks d. For a fair comparison, we use comparable sizes of the underlying NN architectures and similar training procedures. We compare the mean and standard deviation of the achieved prices as well as the median of the computation times for each of the algorithms. Here, we only show the results of one of the numerous experiments, but we note that the qualitative results are similar in all of them. For a min-put option on a Heston model (Heston, 1993)

$$\begin{aligned} dX_t &= (r-\delta)X_t dt + \sqrt{v_t} X_t dW_t, \\ dv_t &= -\kappa(v_t - v_\infty)dt + \sigma\sqrt{v_t}dB_t \end{aligned} \tag{3}$$

with $X_0 = 100$ and $v_0 = 0.01$, where $(W_t)_{t\geq 0}$ and $(B_t)_{t\geq 0}$ are two d-dimensional Brownian motions correlated with coefficient $\rho = -30\%$, and $r = 2\%$, $\delta = 0\%$, $\sigma = 20\%$, $v_\infty = 0.01$, $\kappa = 2$ and $m = 20'000$ sampled paths with $N = 10$ equidistant exercise dates, the results are given in Table 1. We observe that RLSM is about 7 times faster than NLSM and more than 50 times faster than DOS for high dimensions. Furthermore RFQI is about 1.5 times as fast as RLSM. For $d \leq 50$ all algorithms yield very similar prices and for larger d the highest prices are always achieved by RFQI , whereby the prices computed with RFQI never deviate more than 0.2% from those computed with FQI. Moreover, the prices computed with RLSM are never more than 0.5% smaller than those computed with LSM. In addition, RLSM achieves the second highest prices for high dimensions.

	price						duration					
d	LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI
5	12.34 (0.05)	12.31 (0.05)	12.16 (0.11)	12.29 (0.06)	12.35 (0.09)	12.37 (0.09)	30s	6s	3s	0s	8s	0s
10	16.48 (0.07)	16.52 (0.08)	16.09 (0.13)	16.55 (0.06)	16.64 (0.07)	16.61 (0.08)	1m31s	6s	3s	0s	28s	0s
50	22.86 (0.05)	25.56 (0.04)	24.03 (0.42)	24.85 (0.08)	25.72 (0.03)	25.71 (0.07)	39m57s	9s	4s	0s	1h21m59s	1s
100	-	29.13 (0.04)	27.30 (0.46)	28.50 (0.06)	-	29.33 (0.07)	-	16s	6s	0s	-	1s
500	-	36.26 (0.05)	34.74 (0.31)	36.28 (0.04)	-	36.95 (0.05)	-	1m21s	24s	3s	-	2s
1000	-	38.62 (0.08)	38.19 (0.20)	39.32 (0.03)	-	39.93 (0.05)	-	3m18s	45s	6s	-	4s
2000	-	39.22 (0.13)	41.05 (0.21)	42.22 (0.04)	-	42.81 (0.04)	-	12m51s	1m37s	13s	-	8s

Table 1: Min put option on Heston (with variance) for different number of stocks d.

To test the methods in a non-Markovian case, we consider the problem of optimally stopping a fractional Brownian motion $(W_t^H)_{t\geq 0}$ as in (Becker et al., 2019). Results of the different tested methods and a comparison to the results reported in (Becker et al., 2019) are given in Figure 1.

Conclusion We presented new methods based on randomized NNs, which are easily scalable to high dimensional optimal stopping problems. In our empirical study we saw that our methods are considerably faster than existing algorithms for high dimensional problems, while still achieving high quality results. In our Markovian experiments, reinforcement learning methods surpass backward induction methods. In our non-Markovian experiment, our randomized recurrent neural network algorithm RRLSM achieves similar results as the path-version of DOS. This example also brought up the limitations of reinforcement learning based approaches, which do not work well in non-Markovian experiments.

References

Weipeng Cao, Xizhao Wang, Zhong Ming, and Jinzhu Gao. A review on neural networks with random weights. *Neurocomputing*, 275:278–287, 2018.

Guang-Bin Huang, Lei Chen, Chee Kheong Siew, et al. Universal approximation using incremental constructive feedforward networks with random hidden nodes. *IEEE Transactions on Neural Networks*, 17(4):879–892, 2006.

Francis A. Longstaff and Eduardo S. Schwartz. Valuing american options by simulation: a simple least-squares approach. *Review of Financial Studies*, 2001.

John Tsitsiklis and Benjamin Van Roy. Regression methods for pricing complex american-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703, 2001.

Bernard Lapeyre and Jérôme Lelong. Neural network regression for bermudan option pricing. *Monte Carlo Methods and Applications*, 2021.

Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Deep optimal stopping. *Journal of Machine Learning Research*, 20:74, 2019.

John Tsitsiklis and Benjamin Van Roy. Optimal stopping of markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control*, 44:1840–1851, 1997.

Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2):327–343, 1993.

Optimal Tax-Timing Strategy with Transaction Costs

Min Dai * Yaoting Lei † Hong Liu ‡

January 21, 2020

Abstract

We develop a dynamic tax-timing model that takes into account both proportional transaction costs and capital gains taxes. Different from extant literature but in line with real tax rules, capital gains taxes realized are assumed to be paid at the end of each calendar year. We consider both the full rebate case in which there is no limitation for tax rebate and the carry over case in which capital losses can be carried forward to offset futures capital gains. Our model reveals that consistent with empirical findings, even with tiny transaction costs, e.g. 0.5%, investors may defer large capital losses. Moreover, our model shows that investors have an incentive to realize capital gains in early months for both the full rebate and carry-over cases, and the earlier the month the stronger the incentive; investors have an incentive to realize capital losses in late months for the full rebate case, and the late the month the stronger the incentive.

Morover, we find that in contrast to all existing portfolio choice models with proportional transaction costs in which a lump-sum sale or purchase never occurs except at initial time, our model may result in frequent lump-sum sale or purchase that is not induced by wash sale. For example, investors are inclined to realize a lump-sum capital gain at the beginning of a calendar year, as most of the no-trading region at the end of a calendar year is contained in the sell region at the beginning of a calendar provided there is a capital gain.

Existing discrete time portfolio choice models with capital gains tax often assume that investors realize capital gains in January and capital losses in December. In contrast, the optimal

* Department of Mathematics, Risk Management Institute, and Suzhou Research Institute, National University of Singapore. Tel: (65) 6516 2754. Email: matdm@nus.edu.sg.

† School of Economics and Management, Nanchang University, 999 Xuefu Road, Nanchang, China 330031,Email: leiyaoting@ncu.edu.cn.

‡ Olin Business School, Washington University in St. Louis, One Brookings Drive, St. Louis, MO 63130. Tel: (314) 935 5883. Email: liuh@wustl.edu.

investment policy derived from our continuous time model with full rebate suggests that investors may also realize gains or losses in other months, and the amounts of gain (loss) realizations decrease (increase) from January to December; moreover, we find that the realized gains in January and the realized losses in December are significantly higher than those in other months, and there are also more trading chances to realize gains in January and losses in December. In the carry over case, we can observe similar behavior of gain realizations as in the full rebate, but the loss realizations are independent of time.

Keywords: portfolio selection, transaction costs, capital gains tax, year-end tax.

Optimal Tax-Timing with Inflation Risk and Indexed Capital Gains

Yaoting Lei, Jing Xu *

Abstract

The current U.S. tax code specifies that capital gains taxes are levied on nominal capital gains. A direct consequence of this tax rule is that inflation can increase investors' real tax burdens. This is because inflation can reduce the size of real gains (in terms of purchasing power), while the associated tax liabilities are not adjusted. Motivated by this fact, the concept of capital gains indexation (indexation hereafter) was proposed in the United States as an alternative tax deduction method. The core idea of indexation is to adjust the tax basis for inflation so that taxes are levied on real capital gains.

Although indexation has been debated for decades, its potential implications for investors' portfolio choices and welfare have not been seriously examined. In this paper, we attempt to fill this void by studying a dynamic portfolio model that incorporates capital gains tax and inflation risk. Unlike the existing tax-timing literature which exclusively focuses on a fixed cost basis, we allow the investor's cost basis to be dynamically adjusted for inflation. This unique feature allows our model to capture the core idea of indexation.

Our paper offers several novel insights. First, we show that inflation risk can substantially strengthen the investor's incentive to defer gains realizations, leading to a first-order increase in the value of her tax-deferral option. Second, we show that indexation could make the investor realize gains sooner (even later) if she can make full (limited) use of losses. Third, we demonstrate that comparing with the current practice of taxing nominal capital gains at lower rates (such as the long-term capital gains tax rates), taxing indexed capital gains at higher rates (such as the marginal income tax rates) could both improve investors' welfare and increase capital gains tax revenue. Fourth, we argue that although low income investors

*Yaoting Lei, School of Economics and Management, Nanchang University, Nanchang 330031, China, Email: leiyaoting@ncu.edu.cn. Jing Xu (Corresponding author), School of Finance, Renmin University of China, Beijing 100872, China, Email: jing.xu@ruc.edu.cn.

are subject to lower tax rates on capital gains, they could benefit from indexation more than wealthy investors do by taking advantage of losses credit.

Keywords: Portfolio choice, Inflation risk, Capital gains tax, Capital gains indexation.

JEL Classification: G11, H24, K34.

OPTIMAL TRANSPORT AND RISK AVERSION IN KYLE'S MODEL OF INFORMED TRADING

KERRY BACK
Jones Graduate School of Business and School of Social Sciences, Rice University

FRANCOIS COCQUEMAS
College of Business, Florida State University

IBRAHIM EKREN
Department of Mathematics, Florida State University

ABRAHAM LIOUI
EDHEC Business School

This version: August 11, 2021

We establish connections between optimal transport theory and the dynamic version of the Kyle model, including new characterizations of informed trading profits via conjugate duality and Monge-Kantorovich duality. We use these connections to extend the model to multiple assets, general distributions, and risk-averse market makers. With risk-averse market makers, liquidity is lower, assets exhibit short-term reversals, and risk premia depend on market maker inventories, which are mean reverting. We illustrate the model by showing that implied volatilities predict stock returns when there is informed trading in stocks and options and market makers are risk averse.

Kerry Back: kerry.e.back@rice.edu
Francois Cocquemas: fcocquemas@fsu.edu
Ibrahim Ekren: iekren@fsu.edu
Abraham Lioui: abraham.lioui@edhec.edu
I. Ekren gratefully acknowledges financial support from the NSF Grant DMS-2007826.

Optimization of conditional convex risk measures

Tiexin Guo[a,*], Yachao Wang[a], George Yuan[b]

[a] *School of Mathematics and Statistics, Central South University, ChangSha 410083, China*
[b] *Centre for Financial Engineering, Soochow University, Suzhou 215006, China.*

Abstract

Optimization of conditional convex risk measure is a central theme in dynamic portfolio selection theory, which has not yet systematically studied in the previous literature perhaps since conditional convex risk measures are neither random strictly convex nor random coercive. The purpose of this paper is to give some basic results on the existence and uniqueness on this theme, in particular our results for conditional monotone mean–variance and conditional entropic risk measures are complete and deep. As the basis for the work of this paper, this paper first begins with a brief introduction to random functional analysis, including the historical backgrounds for its birth and some important advances. This paper then further surveys some recent progress in random convex analysis and its applications to conditional convex risk measures. Finally, based on these, we establish a concise sufficient and necessary condition for a return to be a solution to the optimization problem of conditional monotone mean–variance. We also make use of the recently developed theory of L^0–convex compactness to establish the existence of the optimization problem of conditional entropic risk measure when the conditional mean of returns is given and the returns fall within a random closed ball. Besides, the related uniqueness problems are also solved.

Keywords: complete random normed module, L^0–convex compactness, conditional convex risk measure; conditional optimization, portfolio selection, existence and uniqueness.

*Corresponding author

Email addresses: tiexinguo@csu.edu.cn (Tiexin Guo), wyachao@csu.edu.cn (Yachao Wang), george_yuan99@yahoo.com (George Yuan)

OPTIMUM THRESHOLDING USING MEAN SQUARE ERROR FOR TEMPERED-STABLE-LIKE LÉVY MODELS

RUOTING GONG

We consider a semimartingale model X with jumps for the log-return of an asset price. Separately identifying the contribution of the Brownian part (through the integrated variance) and the one of the jumps to the asset price variations when we can observe prices discretely is crucial in many respects, for instance, for model assessing and for improving volatility forecasting. The correct identification of a model has a significant impact on option pricing and on risk management and thus on assets allocation. With discrete (non-noisy) observations, one of the most efficient estimators for non-parametrically disentangling the jumps from integrated variance is the so-called Truncated (or Threshold) Realized Variance (TRV), especially in the presence of infinite activity jump component. However, the choice of the truncation level (threshold) has an impact on the estimation performance with finite samples.

When X is a Lévy process, it was shown in [2] that there exists a unique threshold ε_h^* that minimizes the mean-square error of the TRV estimator of the continuous variance σ^2, where $h > 0$ is the observation step. The asymptotic behavior of ε_h^*, as $h \to 0$, is derived in [2] for the cases when X has finite activity jumps or X is symmetric and strictly stable. In this talk, we generalize the above work to a large class of Lévy models, the so-called tempered-stable-like (TSL) models proposed in [1]. We show that the asymptotic behavior of the optimal threshold, obtained in [2], holds true for TSL models. In order to devise a completely data-driven threshold method for volatility estimation we revisit the problem of estimating the index of jump activity based on high-frequency data. To this end we adapt and generalize the recent results of [3] and propose a novel efficient estimation method that simultaneously estimate the volatility and the index of jump activity. Our numerical study demonstrates that the finite-sample performance of the estimator is superior to any other method proposed in the literature.

This is the joint work with José E. Figueroa-López and Yuchen Han.

References

[1] J. E. Figueroa-López, R. Gong, and C. Houdré. High-Order Short-Time Expansions for ATM Option Prices of Exponential Lévy Models. *Math. Finance*, 26(3):516−557, 2016.

[2] J. E. Figueroa-López and C. Mancini. Optimum Thresholding Using Mean and Conditional Mean Squared Error. *J. Econom.*, 208(1), 179−210, 2019.

[3] F. Mies. Rate-Optimal Estimation of the Blumenthal-Getoor Index of a Lévy Process. *Preprint*, 2019. arXiv:1906.08062

Department of Applied Mathematics, Illinois Institute of Technology, Chicago, IL 60616, U.S.A. (`rgong2@iit.edu`).

Overcoming the curse of dimensionality in the approximation of semilinear Black-Scholes PDEs

Philippe von Wurstemberger, ETH Zurich

Parabolic partial differential equations (PDEs) are key mathematical tools to model natural phenomena and man made complex systems. In particular, Black-Scholes-type PDEs are widely used in the financial industry to determine fair prices of financial derivatives. Amongst the manifold variations and extensions of the original Black-Scholes equation, some models suggest that the price process of a financial derivative satisfies certain semilinear Black-Scholes PDEs. This includes, for example, models which take into account the possibility of a defaulting counterparty (cf., e.g, [3, 4]). A major issue in the scientific literature is that most approximation methods for nonlinear PDEs suffer under the so-called curse of dimensionality and nearly all approximation methods for nonlinear PDEs in the scientific literature have not been shown not to suffer under the curse of dimensionality. The PDEs appearing in financial engineering applications are often high dimensional since the dimension typically corresponds to the number of considered underlying financial assets. As a result, in most cases no feasible numerical algorithm to approximate solutions of such nonlinear PDEs is known.

In this talk, we present a novel approximation method for solutions of high-dimensional semilinear Black-Scholes PDEs (see [1]) and reveal that it overcomes the course of dimensionality under the assumption that the nonlinearity is Lipschitz continuous. This result proves, for the first time, that the numerical approximation of solutions of semilinear Black-Scholes equations is a polynomially tractable approximation problem, in the sense that the computational effort of the proposed method grows at most polynomially in both the dimension $d \in \mathbb{N}$ and the reciprocal of the prescribed approximation accuracy $\varepsilon > 0$. Our algorithm is based on a Feynman-Kac formula for semilinear parabolic PDEs and carefully combines Picard iterations with a multilevel Monte Carlo approach to approximate expectations. In view of this our scheme is called a *multilevel Picard* (MLP) algorithm. It is related to and builds on previous work on MLP algorithms (cf., e.g, [2]).

References

[1] Hutzenthaler, M., Jentzen, A., von Wurstemberger, P., Overcoming the curse of dimensionality in the approximative pricing of financial derivatives with default risks. arXiv:1903.05985, (2019).

[2] Hutzenthaler, M., Jentzen, A., Kruse, T., Nguyen, T. A., von Wurstemberger, P., *Overcoming the curse of dimensionality in the numerical approximation of semilinear parabolic partial differential equations.* arXiv:1807.01212, (2018).

[3] Burgard, C., and Kjaer, M., Partial differential equation representations of derivatives with bilateral counterparty risk and funding costs. *The Journal of Credit Risk 7*, 3 (2011), 1-19.

[4] Henry-Labordere, P., Counterparty Risk Valuation: A Marked Branching Diffusion Approach. arXiv:1203.2369, (2012).

Philippe von Wurstemberger, ETH Zurich, Rämistrasse 101, 8092 Zurich, Switzerland
philippe.vonwurstemberger@math.ethz.ch

Permanent Income Hypothesis with Income Disaster and Asset Pricing*

Seyoung Park†

Abstract

We study the impact of income disaster–which can be triggered by various reasons such as pandemics and technological disruption–on Friedman's permanent income hypothesis. Although the hypothesis is still retained, our general equilibrium model offers two insights. First, with income disaster the equilibrium interest rate can be a decreasing function of the risk aversion while the equity premium is an increasing function, helping to disentangle the risk-free rate puzzle from the equity premium puzzle. Second, the model can better match empirical marginal propensities to consume numbers.

Keywords: Income Disaster, Permanent Income Hypothesis, Asset Pricing, Marginal Propensities to Consume

JEL Codes: D15, D58, G11, G12

*The author is grateful for the helpful discussions with Darrell Duffie, Paul Glasserman, Steven Kou, Jussi Keppo, Karl Schmedders, Alan Morrison, David Bell, Hato Schmeiser, Abhay Abhyankar, Alain Bensoussan, Alistair Milne, Chiaki Hara, Eckhard Platen, Paul Embrechts, Hyeng Keun Koo, Phillip Yam, Huainan Zhao, Jiro Akahori, and the seminar participants at 55th Annual Meeting of the Eastern Finance Association (EFA), 4th World Risk and Insurance Economics Congress (WRIEC), 4th International Conference on Econometrics and Statistics (EcoSta), IFABS 2021 Oxford Conference, ETH Zurich, Seoul National University, Yonsei University, Korea University, Korean Advance Institute of Science and Technology (KAIST), Pohang University of Science and Technology (POSTECH), Ulsan National Institute of Science and Technology (UNIST), Pusan National University, Ajou University, Kyung Hee University, Sookmyung Women's University for helpful comments. All errors are the author's own responsibility.

†Nottingham University Business School, University of Nottingham, E-mail: `seyoung.park@nottingham.ac.uk`

Perpetual American standard and lookback options with event risk and asymmetric information

Pavel V. Gapeev ((Corresponding author) London School of Economics, Department of Mathematics, Houghton Street, London WC2A 2AE, United Kingdom; e-mail: p.v.gapeev@lse.ac.uk)
and Libo Li (University of New South Wales, School of Mathematics and Statistics, Sydney NSW 2052, Australia; e-mail: libo.li@unsw.edu.au)

We derive closed-form solutions to the perpetual American standard and floating-strike lookback put and call options in an extension of the Black-Merton-Scholes model with event risk and asymmetric information. It is assumed that the contracts are terminated by their writers with linear or fractional recoveries at the last hitting times for the underlying asset price process of its ultimate maximum or minimum over the infinite time interval which are not stopping times with respect to the reference filtration. We show that the optimal exercise times for the holders are the first times at which the asset price reaches some lower or upper stochastic boundaries depending on the current values of its running maximum or minimum. The proof is based on the reduction of the original optimal stopping problems to the associated free-boundary problems and the solution of the latter problems by means of the smooth-fit and normal-reflection conditions. The optimal exercise boundaries are proven to be the maximal or minimal solutions of some first-order nonlinear ordinary differential equations.

It is well established that American-styled financial contracts can be exercised by their holders (investors) at some random times which are stopping times with respect to the filtrations reflecting the information available in the appropriate financial markets. In this paper, motivated by the study of optimal buyback time for short sellers facing recall risk, we study a model where such contingent claims are terminated or cancelled prematurely by their writers (counterparties) which have competing interests with the holders as well as an access to certain additional insider information which is not available to the average investors trading in the markets. The studies of financial models with insider information have been an object of interest of many works, where the authors typically assumed that the informed investors know the terminal values of the underlying risky assets and hence have access to the initial enlargement of the market filtration with the terminal values. In contrast to the aforementioned works, we suppose that the writers can take advantage of knowing certain additional information that create global peaks or bottoms in the charts of the market prices of the underlying risky assets. These assumptions correspond to the observations of Associate Professor Daniel Taylor from the Wharton School in an NBC article on the insider trading that: *"One of the most well-accepted facts from decades of research on insider trading is that corporate insiders buy near bottoms and sell near peaks"*. We stress that we do not claim that the writers or some third parties communicating with the writers know the future dynamics of the risky asset prices in advance but rather they are insiders who can either influence or determine the timing at which the price processes of the underlying risky assets reach their global peaks or bottoms. In order to model the advantages which are available to the writers of the contingent claims, we assume that the counterparties can terminate the contracts at the last times at which the prices reach their ultimate maxima or minima and some linear or fractional recovery amounts are paid to

https://www.nbcnews.com/business/business-news/peloton-insiders-sold-nearly-500-million-stock-big-drop-rcna12741 the holders. In other words, we assume that the insiders have an access to the progressive enlargements of the market filtrations with the last passage times of the underlying asset prices on the levels of their ultimate maxima or minima.

For a precise formulation of the problems, we consider a probability space $(\Omega, \mathcal{F}, P)$ with a standard Brownian motion $B = (B_t)_{t \geq 0}$. Assume that the process $X = (X_t)_{t \geq 0}$ describing the price of a risky asset in a financial market is given by:

$$X_t = x \, \exp\Big(\big(r - \delta - \sigma^2/2 \big) \, t + \sigma \, B_t \Big) \tag{0.1}$$

so that it solves the stochastic differential equation:

$$dX_t = (r - \delta) \, X_t \, dt + \sigma \, X_t \, dB_t \quad (X_0 = x) \tag{0.2}$$

where $x > 0$ is fixed, and $r > 0$, $\delta > 0$, and $\sigma > 0$ are some given constants. Here, r is the riskless interest rate, δ is the dividend rate paid to the asset holders, and σ is the volatility rate. Let the processes $S = (S_t)_{t \geq 0}$ and $Q = (Q_t)_{t \geq 0}$ be the *running maximum and minimum* of X defined by:

$$S_t = s \vee \Big(\max_{0 \leq u \leq t} X_u \Big) \quad \text{and} \quad Q_t = q \wedge \Big(\min_{0 \leq u \leq t} X_u \Big) \tag{0.3}$$

for some arbitrary $0 < q \leq x \leq s$. To model the event horizon, we also introduce the random times $\theta = \sup\{ t \geq 0 \, | \, X_t = S_t \}$ and $\eta = \sup\{ t \geq 0 \, | \, X_t = Q_t \}$, which are not stopping times with respect to the natural filtration $(\mathcal{F}_t)_{t \geq 0}$ of the process X, but they are honest times.

The main aim of this paper is to compute closed-form expressions for the values of the discounted optimal stopping problems:

$$\overline{V}_i = \sup_{\tau} E\big[e^{-r\tau} \, G_{i,1}(X_\tau, S_\tau) \, I(\tau < \theta) + e^{-r\theta} \, (\varphi_i + \psi_i \, X_\theta) \, I(\theta \leq \tau) \big] \tag{0.4}$$

and

$$\overline{U}_i = \sup_{\zeta} E\big[e^{-r\zeta} \, G_{i,2}(X_\zeta, Q_\zeta) \, I(\zeta < \eta) + e^{-r\eta} \, (\xi_i + \chi_i \, X_\eta) \, I(\eta \leq \zeta) \big] \tag{0.5}$$

with $G_{1,1}(x,s) = L_1 - x$, $G_{2,1}(x,s) = s - L_2 x$, for some $L_1 > 0$ and $L_2 \geq 1$, while $G_{1,2}(x,q) = x - K_1$, $G_{2,2}(x,q) = K_2 x - q$, for some $K_1 > 0$ and $K_2 \geq 1$, where $\varphi_i, \xi_i \in \mathbb{R}$, and $\psi_i, \chi_i \in (-1,1)$, for $i = 1,2$, are fixed, and $I(\cdot)$ denotes the indicator function. Suppose that the suprema in (0.4) and (0.5) are taken over all stopping times τ and ζ with respect to the filtration $(\mathcal{F}_t)_{t \geq 0}$, and the expectations there are taken with respect to the risk-neutral probability measure P. In this view, the values $\overline{V}_i$ and $\overline{U}_i$, for $i = 1,2$, in (0.4) and (0.5) provide the rational (no-arbitrage) prices of the perpetual American defaultable standard and floating-strike lookback options in an extension of the Black-Merton-Scholes model with event risk and asymmetric information, when we formally set $s = x$ and $q = x$ in (0.3).

In particular, we show that the perpetual American defaultable floating-strike lookback put and call options may be exercised when the processes (X,S) or (X,Q) start in certain subsets of the edges of their state spaces, under specific relations on the parameters of the model. These properties represent new features of the optimal stopping problems for the running maximum and minimum processes.

The case of zero recoveries $\varphi_i = \xi_i = 0$ and $\psi_i = \chi_i = 0$, for $i = 1,2$, was recently studied in Gapeev and Li (2021), Optimal stopping problems for maxima and minima in models with asymmetric information, published online first in *Stochastics* (28 pp).

Policy Evaluation and Temporal–Difference Learning in Continuous Time and Space: A Martingale Approach

We propose a unified framework to study policy evaluation (PE) and the associated temporal difference (TD) methods for reinforcement learning in continuous time and space. We show that PE is equivalent to maintaining the martingale condition of a process. From this perspective, we find that the mean–square TD error approximates the quadratic variation of the martingale and thus is not a suitable objective for PE. We present two methods to use the martingale characterization for designing PE algorithms. The first one minimizes a "martingale loss function", whose solution is proved to be the best approximation of the true value function in the mean–square sense. This method interprets the classical gradient Monte-Carlo algorithm. The second method is based on a system of equations called the "martingale orthogonality conditions" with test functions. Solving these equations in different ways recovers various classical TD algorithms, such as TD(λ), LSTD, and GTD. Different choices of test functions determine in what sense the resulting solutions approximate the true value function. Moreover, we prove that any convergent time-discretized algorithm converges to its continuous-time counterpart as the mesh size goes to zero, and we provide the convergence rate. We demonstrate the theoretical results and corresponding algorithms with numerical experiments and applications.

Policy Gradient and Actor–Critic Learning in Continuous Time and Space: Theory and Algorithms

Yanwei Jia YJ2650@COLUMBIA.EDU
Department of Industrial Engineering and Operations Research
Columbia University
New York, NY 10027, USA

Xun Yu Zhou XZ2574@COLUMBIA.EDU
Department of Industrial Engineering and Operations Research &
The Data Science Institute
Columbia University
New York, NY 10027, USA

Abstract

We study policy gradient (PG) for reinforcement learning in continuous time and space under the regularized exploratory formulation developed by **?**. We represent the gradient of the value function with respect to a given parameterized stochastic policy as the expected integration of an auxiliary running reward function that can be evaluated using samples and the current value function. This effectively turns PG into a policy evaluation (PE) problem, enabling us to apply the martingale approach recently developed by **?** for PE to solve our PG problem. Based on this analysis, we propose two types of the actor–critic algorithms for RL, where we learn and update value functions and policies simultaneously and alternatingly. The first type is based directly on the aforementioned representation which involves future trajectories and hence is offline. The second type, designed for online learning, employs the first-order condition of the policy gradient and turns it into martingale orthogonality conditions. These conditions are then incorporated using stochastic approximation when updating policies. Finally, we demonstrate the algorithms by simulations in two concrete examples.

Portfolio Liquidation Games with Self-Exciting Order Flow

Guanxing Fu* Ulrich Horst† Xiaonyu Xia‡

Model

We analyze novel portfolio liquidation games with self-exciting order flow. Both the N-player game and the mean-field game are considered. We assume that players' trading activities have an impact on the dynamics of future market order arrivals thereby generating an additional transient price impact. To be precise, we study the following liquidation problem

$$\inf \mathbb{E}\left[\int_0^T \eta_s \xi_s^2 + \xi_s Y_s + \lambda_s X_s^2 \, ds\right] \tag{0.1}$$

subject to the state dynamics

$$\begin{aligned} dX_t &= -\xi_t \, dt \\ X_0 &= \mathcal{X}, \quad X_T = 0 \\ dY_t &= \{-\rho_t Y_t + \gamma_t(\xi_t - (\beta-\alpha)C_t + \alpha(\mathcal{X} - \mathbb{E}[X_t]))\} \, dt \\ Y_0 &= 0 \\ dC_t &= -(\beta-\alpha)C_t + \alpha(\mathcal{X} - \mathbb{E}[X_t]) \, dt \\ C_0 &= 0. \end{aligned} \tag{0.2}$$

Here, all market parameters are random. The quantity X_t denotes the number of shares the investor needs to sell at time $t \in [0,T]$, while ξ_t denotes the rate at which the stock is traded at that time. The process Y describes the persistent price impact. It can be viewed as a shift in the mid quote price caused by past trades where the impact is measured by impact factor. Our key conceptual contribution is to allow for an additional feedback effect into the above model that accounts for the possibility of an additional order flow C ("child orders") triggered by the large investor's trading activity. The dynamics of C can be obtained by Hawkes process theory with exponential kernel.

The novelty of our model in mathematics is two-fold. First, because of the dynamics of child order flow, there is a mean-field term. Thus, the investor faces a mean-field control problem. Second, the additional order flow make the optimization non-convex; it makes the verification non-standard.

Solve the non-convex optimization

By a maximum principle argument, the optimal control of the optimization problem (0.1)-(0.2) can be characterized by an *singular* FBSDE. The singularity comes from the liquidation constraint $X_T = 0$ in

*The Hong Kong Polytechnic University, Department of Applied Mathematics; email: guanxing.fu@polyu.edu.hk

†Humboldt University Berlin, Department of Mathematics and School of Business and Economics; email: horst@math.hu-berlin.de

‡Wenzhou University, Department of Mathematics; email: xiaonyu.xia@wzu.edu.cn

(0.2). Making a standard affine ansatz the system with unknown terminal condition can be replaced by an FBSDE with known initial and terminal condition, yet singular driver. Proving the existence of a small time solution to this FBSDE is not hard. The challenge is to prove the existence of a global solution on the whole time interval. Extending the continuation method for singular FBSDEs established in [1] to our higher dimensional system we prove that the FBSDE system does indeed have a unique solution in a certain space under a weak interaction condition. Weak interaction conditions have been extensively used in the game theory literature before; see, e.g. [2] and references therein.

Subsequently, we establish a novel verification argument for the N-player game1 from which we deduce that the solution to the FBSDE system does indeed give the desired optimal trading rate. Our maximum principle does not require convexity of the cost function as it is usually the case. In fact, unlike in existing results, in our model the players' optimization problems are not convex and hence standard verification arguments do not apply. Instead, we establish a novel maximum principle that strongly relies on the liquidation constraint. Our idea is to decompose trading costs into a sum of optimal costs plus round-trip costs and then to show that deviations from the optimal strategy are costly. The decomposition result provides a sufficient condition for our impact model to be viable. In particular, based on our novel verification argument, we can prove rigorously there is no beneficial round-trip.

Extension

The game-theoretic extension of (0.1)-(0.2) including N-player game and mean-field game, can also be studied. We characterize open-loop Nash equilibria in both games in terms of a novel (multi-dimensional or mean-field) FBSDE system with unknown terminal condition. The novel verification argument can be applied to this multi-player setting. The power of our verification argument is that we can prove round trips always exist in the game-theoretic liquidation problem.

References

[1] G. Fu, P. Graewe, U. Horst, and A. Popier. A mean field game of optimal portfolio liquidation. *Mathematics of Operations Research*, 46(4):1235–1657, 2022.

[2] U. Horst. Stationary equilibria in discounted stochastic games with weakly interacting players. *Games and Economic Behavior*, 51(1):83–108, 2005.

Portfolio liquidation under transient price impact - theoretical solution and implementation with 100 NASDAQ stocks*

Ying Chen †, Ulrich Horst‡ and Hoang Hai Tran §

December 22, 2019

Abstract

We derive an explicit solution for deterministic market impact parameters in the Graewe and Horst (2017) portfolio liquidation model. The model allows the combination of various forms of market impact, namely, instantaneous, permanent, and temporary. We show that the solutions to the two benchmark models of Almgren and Chriss (2001) and of Obizhaeva and Wang (2013) are obtained as special cases. We relate the different forms of market impact to the microstructure of limit order book markets and show how the impact parameters can be estimated from public market data. We investigate the numerical performance of the derived optimal trading strategy based on high-frequency limit order books of 100 NASDAQ stocks that represent a range of market impact profiles. This shows that the strategy achieves significant cost savings compared with the benchmark models of Almgren and Chriss (2001) and Obizhaeva and Wang (2013).

AMS Subject Classification: 62M10, 62P20
Keywords: Liquidity Risk, Optimal Trading Strategy, Portfolio Liquidation, Hawkes process

References

Almgren, R. and Chriss, N. (2001), 'Optimal execution of portfolio transactions', *Journal of Risk* **3**, 5–40.

Graewe, P. and Horst, U. (2017), 'Optimal trade execution with Instantaneous price impact and stochastic resilience', *SIAM Journal on Control and Optimization* **55**(6), 3707–3725.

Obizhaeva, A. and Wang, J. (2013), 'Optimal trading strategy and supply/demand dynamics', *Journal of Financial Markets* **16**(1), 1–31.

*This research has been supported through the profile partnership program between Humboldt-Universität zu Berlin and National University of Singapore.The authors gratefully acknowledge the financial support of the Singapore Ministry of Education Academic Research Fund Tier 1 at National University of Singapore.

†Department of Mathematics and Risk Management Institute, National University of Singapore, Lower Kent Ridge Road 10, 119076 Singapore, Singapore; email: matcheny@nus.edu.sg

‡Department of Mathematics, and School of Business and Economics, Humboldt-Universität zu Berlin Unter den Linden 6, 10099 Berlin, Germany; email: horst@math.hu-berlin.de

§Department of Statistics & Applied Probability, National University of Singapore, 6 Science Drive 2, Singapore 117546, Singapore; email: e0045275@u.nus.edu

Title: Portfolio Rebalancing with Realization Utility

Authors: Min Dai, Cong Qin and Neng Wang

Abstract:
We develop a model where a realization-utility investor (Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013) optimally targets her liquid-illiquid wealth ratio at a constant $w*$. By saving in the risk-free asset ($w* > 0$), she makes smaller bets in the illiquid asset and realizes gains/losses more frequently. By leveraging ($w* < 0$), she makes bets larger than her equity and realizes gains/losses less frequently. For a discontinuous/jump-diffusion price process, the solution features four regions: loss-realization, gain-realization, and two disconnected (deep-loss and normal) holding regions. We generate a quantitatively significant non-monotonic propensity to realize losses consistent with evidence. This work is joint with Cong Qin and Neng Wang

Portfolio Selection with Deep Learning

Zhang Chao (zufezc@126.com)
School of Finance, Zhejiang University of Finance and Economics
Hangzhou 310018, PR China
Rongda Chen (rongdachen@zufe.edu.cn)
School of Finance, Zhejiang University of Finance and Economics
Hangzhou 310018, PR China
Yongzeng Lai (ylai@wlu.ca)
Department of Mathematics, Wilfrid Laurier University
Waterloo, Ontario, N2L 3C5, Canada
Gaoxiang Wang (gx_mattew_wang@126.com)
School of Finance, Zhejiang University of Finance and Economics
Hangzhou 310018, PR China

Abstract

In this talk, we propose a two-stage framework to construct portfolios based on deep learning algorithms. Both NASDAQ100 and CSI300 are selected as representatives of developed markets and emerging markets, respectively. At stage 1, once a stock index is selected, for each component stock in the index, we use principal component analysis (PCA), auto-encoder(AE), and restricted Boltzmann machine (RBM) as data representation methods to reconstruct the stock prices, and select outstanding stocks to enter the portfolio according to the characteristics of data reconstruction. At stage 2, taking the selected stock index as the target, we train the artificial neural networks to construct portfolios and to test investment strategies by validation. Our results show that (1) there is no significant difference in the performance of different data representation methods; (2) the contribution of communal information to the optimal portfolio decreases with the number of selected stocks; (3) the characteristics of different types of markets obtained by deep learning are different; (4) this approach achieves good results for different trading frequency data.

Portfolio selection with exploration of new investment opportunities

Luca De Gennaro Aquino* Didier Sornette† Moris S. Strub‡

Abstract

We introduce a model for portfolio selection with an extendable investment universe where an agent with mean-variance preferences faces a trade-off between exploiting existing and exploring for new investment opportunities. If the option to explore for new investment opportunities is exercised, the agent chooses to devote a part of her wealth for exploration. This amount represents the costs associated with exploration, for example for hiring a team of analysts or acquiring information. Exploration results in the discovery of a new asset, whose distributional properties depend on the amount devoted to exploration in a way that the asset becomes more attractive to the agent when a larger amount is devoted to exploration. After discovery of the new asset, the agent then distributes the remainder of her wealth among the assets of the extended investment universe in order to optimize her preferences over the resulting terminal wealth. We first show that the problem is well-posed when the Sharpe ratio of the newly discovered asset has reasonable asymptotic elasticity, and characterize the optimal amount devoted to exploration. We determine that incremental exploration does not pay off: one must put a significant amount at risk in order to harvest the potential benefits of exploring for new investment opportunities. Our model also shows that it is increasingly worthwhile to explore in worse market environments, and that investment performance measured by the Sharpe ratio is increasing in the initial wealth of the agent indicating that richer agents can make better use of new investment opportunities.

*Department of Information Systems and Management Engineering, College of Business, Southern University of Science and Technology (SUSTech), Shenzhen, China. Email: degennaroaquino@sustech.edu.cn

†Department of Management, Technology, and Economics, ETH Zürich, Zürich, Switzerland; Institute of Risk Analysis, Prediction, and Management, Southern University of Science and Technology (SUSTech), Shenzhen, China; and Tokyo Tech World Research Hub Initiative (WRHI), Institute of Innovative Research, Tokyo Institute of Technology, Tokyo, Japan. Email: dsornette@ethz.ch

‡Department of Information Systems and Management Engineering, College of Business, Southern University of Science and Technology (SUSTech), Shenzhen, China. Email: strub@sustech.edu.cn.

In order to obtain a parsimonious model, our base setting, and the analysis thereof, rests on some simplifying assumptions. Namely, we suppose that there is zero correlation between the existing investment universe and the newly discovered asset, and that there is a single asset to be explored. We further suppose that the agent knows ex ante the distributional characteristics of the newly discovered asset as a function of the amount devoted for exploration: there is no uncertainty in our base setting. In the final part of the paper, we relax these assumptions and consider possible extensions of our model. Allowing for correlation between the newly discovered asset and the assets of the existing universe is a technical extension and the results conform with those of the base model. If the newly discovered asset is more strongly correlated with some assets of the existing universe than with others, the integration of the newly discovered asset into the investment universe will lead to a rebalancing of the portfolio that would be optimal without an option for exploration as one could expect. Our qualitative results also remain robust when the agent can explore for more than one new asset. One interesting new observation is that, when economies of scale for the cost of information acquisition occur, it is possible that there are two local optima for the number of newly discovered assets: a small local optimum if the agent focuses on reaping low hanging fruits and a larger number reflecting the attempt to benefit from economies of scale. The case where there is uncertainty about the distributional properties of the new asset is conceptually interesting due to the dynamic inconsistency of the variance, but our main findings remain again robust and this extension even completely reduces to the base model when there is a concurrent decision about how much to devote for exploration and how to invest the remaining wealth in available assets. We finally examine how the model could be extended to a dynamic setting and investigate a simple two-period model in greater detail. Our initial results on this point show that devoting resources to discover new investment opportunities is mostly beneficial at the beginning of the investment horizon and, subsequently, when (and if) wealth decreases below a certain value. On the contrary, when the agent is already close to reach her objective at maturity, investing additional funds for the discovery of new assets is not optimal.

Portfolios from additional information

TOMOYUKI ICHIBA
Department of Statistics and Applied Probability
University of California Santa Barbara

In this talk we discuss a family of portfolios generated by functions in equity market from Stochastic Portfolio Theory and determine relative arbitrage opportunities over market portfolio. If there is an available additional information about distribution of future values in the market, an investor may update portfolio and reallocate wealth, in order to optimize relative arbitrage opportunity. We formulate it as a portfolio selection problem with Knightian uncertainty, evaluate the updated portfolio wealth and discuss the value of additional information.

Post-trade netting and contagion

Luitgard Anna Maria Veraart [*]

London School of Economics and Political Science

Yuliang Zhang [†]

London School of Economics and Political Science

We analyse how post-trade netting in over-the-counter derivatives markets affects systemic risk. In particular, we focus on portfolio rebalancing and portfolio compression, which are two post-trade services using multilateral netting techniques. First, we provide a mathematical characterisation of portfolio rebalancing. We show how it can be used as a general representation for post-trade netting mechanisms and relate it to portfolio compression. Then, we analyse the effects of portfolio rebalancing on the financial system from a network perspective by considering contagion arising from only partial repayments in networks of variation margin payments. We provide sufficient conditions for portfolio rebalancing to reduce systemic risk. We show using examples that portfolio rebalancing can be harmful. Finally, we investigate the implications of post-trade netting when financial institutions strategically react to liquidity stress by delaying their payments. In this setting, we show that netting that preserves counterparty relationships always reduces systemic risk, whereas netting that does not preserve counterparty relationships can be harmful to the system.

[*]London School of Economics and Political Science, Department of Mathematics, Houghton Street, London WC2A 2AE, UK, Email: l.veraart@lse.ac.uk.

[†]London School of Economics and Political Science, Department of Mathematics, Houghton Street, London WC2A 2AE, UK, Email: y.zhang300@lse.ac.uk.

Precommitted Strategies with Initial-time and Intermediate-time VaR Constraints

Chufang Wu * Jia-Wen Gu † Wai-Ki Ching ‡

Abstract

Expected utility maximization (EUM) is a fundamental task in mathematical finance and has been widely studied in the literature. One important research problem is EUM with downside risk constraints. Most of the downside risk constrained EUM problems have been confined to one terminal risk constraint settings. Basak and Shapiro (2001) considered a continuous-time EUM problem with a risk constraint on the terminal wealth. The risk is quantified by different risk measures including VaR and limited-expected-losses of the terminal wealth. Gundel and Weber (2007, 2008) solved the portfolio selection problem with a utility-based shortfall risk limit in a semimartingale market setting and even in an incomplete market setting, in which the existence of optimal solution is proved. Cuoco et al. (2008) extended the static model to a dynamic version and formulate a dynamic consistent EUM model where VaR limit is reevaluated dynamically. However, their model setting does not fully reflect the optimization problem of financial institutions in practice. Generally, VaR constraints are imposed over a yearly, monthly or weekly time horizon instead of continuously. Besides, it is infeasible to recalculate the risk constraint continuously.

Multiple risk constraints framework arises naturally under Solvency II and Basel III type regulation. Cuoco and Liu (2006) solved the discrete-time optimal investment problem subject to capital requirements based on self-reported VaR estimates. Shi and Werker (2012) studied the effects of imposing repeated VaR short-horizon regulatory constraints on long-term investors. Chen et al. (2018) considered a utility maximization problem with two VaR constraints. In addition to the terminal wealth VaR constraint, there is an extra VaR constraint imposed on the intermediate-time wealth. Their problem was transformed to two decoupled one-VaR constrained optimizations at initial time and intermediate time, hence equilibrium strategies were then obtained.

In this paper, we study a problem that a fund manager aims at maximizing the expected utility of terminal wealth, subject to two terminal wealth VaR constraints imposed at the initial time and intermediate time, respectively. The motivation to consider such a problem is intuitive: there can be both initial time and midterm risk evaluation or regulation requirement for fund managers and financial institutions. Bearing in mind of these risk constraints, fund managers make optimal investment strategies at initial time, i.e., precommitted strategies are derived. In Table 1, we show the comparison of features covered in this paper and some existing papers.

The contribution of our paper is three-fold. Firstly, by introducing an intermediate-time terminal wealth risk constraint, we develop a more realistic EUM model compared to the one terminal risk constraint setting. Different from the previous works in the literature, we assume that there is a conditional VaR constraint during the investment period. Thus,

*Department of Mathematics, The University of Hong Kong, Pokfulam Road, Hong Kong, China; Department of Mathematics, Southern University of Science and Technology, Shenzhen, China. E-mail: wucf@connect.hku.hk.

†Corresponding author. Department of Mathematics, Southern University of Science and Technology, Shenzhen, China. E-mail: kaman.jwgu@gmail.com.

‡Department of Mathematics, The University of Hong Kong, Pokfulam Road, Hong Kong, China. E-mail: wching@hku.hk.

Table 1: Mainstream of papers in constrained EUM literature

	Initial-time Risk Constraints	Intermediate-time Risk Constraints	Intermediate Wealth Constraints	Precommitted Strategy
Basak and Shapiro (2001)	Yes	No	No	Yes
Kraft and Steffensen (2013)	Yes	No	Yes	No
Chen et al. (2018)	No	Yes	Yes	No
Our Paper	Yes	Yes	No	Yes

the investment strategies must satisfy the specified terminal risk limit at the intermediate time, rather than only at the initial time. Secondly, the derived precommitted strategy is dominant over all other feasible strategies at initial time. We find that it is also optimal at the intermediate time for "bad" market states and performs significantly better than the one terminal-wealth risk constraint solutions under the relative loss ratio measure. Thirdly, we construct a contingent claim on Merton's portfolio to replicate the optimal constrained portfolio.

References

Basak, S. and Shapiro, A. (2001). Value-at-Risk-Based Risk Management: Optimal Policies and Asset Prices. *The Review of Financial Studies*, 14(2):371–405.

Chen, A., Nguyen, T., and Stadje, M. (2018). Risk Management with Multiple VaR Constraints. *Mathematical Methods of Operations Research*, 88(2):297–337.

Cuoco, D., He, H., and Isaenko, S. (2008). Optimal Dynamic Trading Strategies with Risk Limits. *Operations Research*, 56(2):358–368.

Cuoco, D. and Liu, H. (2006). An Analysis of VaR-Based Capital Requirements. *Journal of Financial Intermediation*, 15(3):362–394.

Gundel, A. and Weber, S. (2007). Robust Utility Maximization with Limited Downside Risk in Incomplete Markets. *Stochastic Processes and Their Applications*, 117(11):1663–1688.

Gundel, A. and Weber, S. (2008). Utility Maximization under A Shortfall Risk Constraint. *Journal of Mathematical Economics*, 44(11):1126–1151.

Kraft, H. and Steffensen, M. (2013). A Dynamic Programming Approach to Constrained Portfolios. *European Journal of Operational Research*, 229(2):453–461.

Shi, Z. and Werker, B. J. (2012). Short-Horizon Regulation for Long-Term Investors. *Journal of Banking and Finance*, 36(12):3227–3238.

Predictable Forward Performance Processes: Infrequent Evaluation and Robo-Advising Applications*

Gechun Liang† Moris S. Strub‡ Yuwei Wang§

Abstract

We study discrete-time predictable forward processes when trading times do not coincide with performance evaluation times in the binomial tree model for the financial market. The key step in the construction of these processes is to solve a linear functional equation of higher order associated with the inverse problem driving the evolution of the predictable forward process. We provide sufficient conditions for the existence and uniqueness and an explicit construction of the predictable forward process under these conditions. Furthermore, we show that these processes are time-monotone in the evaluation period. Finally, we argue that predictable forward preferences are a viable framework to model preferences for robo-advising applications and determine an optimal interaction schedule between client and robo-advisor that balances a tradeoff between increasing uncertainty about the client's beliefs on the financial market and an interaction cost.

*Author order is alphabetical. All authors are co-first authors of this paper. Helpful comments and suggestions from Martin Herdegen is gratefully acknowledged. Moris Strub gratefully acknowledges funding through the National Natural Science Foundation of China under Grant No. 72050410356.

†Department of Statistics, University of Warwick. Email: g.liang@warwick.ac.uk

‡Department of Information Systems and Management Engineering, Southern University of Science and Technology. Email: strub@sustech.edu.cn

§Corresponding author. Department of Statistics, University of Warwick and Department of Information Systems and Management Engineering, Southern University of Science and Technology. Email: yuwei.wang.2@warwick.ac.uk

Price Impact of Order Flow Imbalances: Multi-level, Cross-asset and Forecasting

Rama Cont[1], Mihai Cucuringu[1,2,3], and Chao Zhang[2]

[1]Mathematical Institute, University of Oxford, Oxford, UK
[2]Department of Statistics, University of Oxford, Oxford, UK
[3]The Alan Turing Institute, London, UK

Abstract

We investigate the impact of order flow imbalance (OFI) on price movements in equity markets in a multi-asset setting. First, we show that taking into account multiple levels of the order book, when defining the order book imbalance, leads to higher explanatory power for the contemporaneous **price impact** of OFI. By performing a principal component analysis of OFI across order book levels, we define a notion of *integrated OFI* which shows superior explanatory power for market impact, both in-sample and out-of-sample.

Second, we examine the notion of **cross-impact** and show that, once the information from multiple levels is incorporated in the OFI, multi-asset models with cross-impact do not provide additional explanatory power for contemporaneous impact compared to a sparse model without the cross-impact terms. However, we find evidence that cross-impact terms do provide additional information for intraday forecasting of future returns.

Introduction

Price dynamics in most major stock exchanges are driven by the intraday flow of buy and sell orders through a centralized limit order book. Empirical studies [3] on intraday data have identified the Order Flow Imbalance (OFI), defined as the net order flow resulting from the imbalance of buy and sell orders, as the main driver of price changes. In the present paper, we systematically examine the impact of order flow imbalances on *contemporaneous price-impact*, *contemporaneous cross-impact* and *forecasting of future returns*.

Price impact on contemporaneous returns. The relationship between contemporaneous order flow imbalance and price movement, i.e. price-impact, has drawn substantial attention in recent decades, due to its importance for studying price dynamics, and estimating trading costs. Cont et al. [3] use a linear model to describe the contemporaneous impact of the order flow imbalance on price dynamics. However, Cont et al. [3] only investigate the in-sample R^2 and ignore the generalization error of their regression models.

Contemporaneous cross-impact. Broadly speaking, contemporaneous cross-impact assesses whether one stock's prices/returns are influenced by the order flow imbalances of other stocks. The existence of cross-impact naturally relates to problems such as optimal execution of portfolios, statistical arbitrage, etc. Capponi and Cont [1] decompose the order flow imbalances into a common factor and idiosyncratic components, in order to verify that, as long as the common factor is involved in the model, adding cross-impact terms improves the explained proportion of the variance only by 0.5%. In addition to examining the cross-impact of best-level OFIs as in [1], we also consider the cross-impact from multi-level OFIs, in order to gauge a comprehensive understanding of the relations between multi-level OFIs of different assets and individual returns.

Forecasting future returns. We extend the above studies to future returns, and probe into the forward-looking price-impact and cross-impact models. Chinco et al. [2] observe that cross-asset returns can help forecast future returns. Nonetheless, to the best of our knowledge, multi-asset order flow imbalances have not been considered as predictors for forecasting future returns in the literature, which is a direction we explore in this work.

Main Results

First, we show that taking into account multiple levels of the order book when defining order book imbalance leads to higher explanatory power for the contemporaneous price impact of OFI. Our empirical evidence reveals that, as OFIs at deeper order book levels are included as features, the explained percentage of variance steadily increases in the in-sample tests. Using a principal component analysis (PCA) of OFI across order book levels, we observe that the first principal component of multi-level OFIs can explain more than 80% of the total variance. Based on these results, we define a notion of *integrated OFI* which shows superior explanatory power for price impact over best-level OFIs and multi-level OFIs in both in-sample and out-of-sample. It is worth emphasizing that we use both in-sample and out-of-sample tests to assess the model's performance, while previous works [3, 1] only focus on in-sample tests.

We revisit the cross-impact with an approach which takes into account model sparsity, using the Least Absolute Shrinkage and Selection Operator (LASSO) to estimate a linear model of contemporaneous impact with self- and cross-impact terms. By comparing with the price impact model using best-level OFIs, we find that the cross-impact model using the best-level OFIs of multiple assets as candidate features can provide small but significant additional explanatory power for price movements. Moreover, our results show that, once the information from multiple levels is included in OFI, multi-asset models with cross-impact do not provide additional explanatory power for contemporaneous impact compared to a sparse model without cross-impact terms. These findings suggest that consolidating multi-level OFIs, such as our integrated OFIs, is a more effective method to model price dynamics than introducing cross-impact terms.

Finally, we investigate the use of OFI to forecast future returns. We study a multi-period price-impact model using lagged OFIs to predict returns, and find evidence that cross-impact terms provide additional information for intraday forecasting of future returns. In particular, the results reveal that involving multi-asset OFIs can increase both in-sample and out-of-sample R^2. We subsequently demonstrate that this increase in out-of-sample R^2 leads to additional profits, when incorporated in common trading strategies.

References

[1] Francesco Capponi and Rama Cont. Multi-asset market impact and order flow commonality. *Available at SSRN*, 2020.

[2] Alex Chinco, Adam D Clark-Joseph, and Mao Ye. Sparse signals in the cross-section of returns. *The Journal of Finance*, 74(1):449–492, 2019.

[3] Rama Cont, Arseniy Kukanov, and Sasha Stoikov. The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88, 2014.

Pricing Interest Rate Derivatives under Volatility Uncertainty

Julian Hölzermann*

The present paper deals with the pricing of interest rate derivatives under volatility uncertainty in the sense of Knightian uncertainty or model uncertainty, also referred to as ambiguous volatility. Due to the assumption of a single, known probability measure, traditional models in finance are subject to *model uncertainty*—that is, the uncertainty about using the correct probability measure—since it is not always possible to specify the probabilistic law of the underlying. Therefore, a new stream of research, called *robust finance*, emerged in the literature, examining financial markets in the presence of a family of probability measures (or none at all) to obtain a robust model. The most frequently studied type of model uncertainty is volatility uncertainty: the volatility determines the probabilistic law of the underlying, but there are many ways to model the volatility of an underlying and it is unknown which describes the future evolution of the volatility best. The literature on robust finance has led to pricing rules that are robust with respect to the volatility. The aim of this paper is to develop robust pricing rules for contracts traded in fixed income markets.

The initial setting is an arbitrage-free bond market under volatility uncertainty. The uncertainty about the volatility is represented by a family of probability measures, called *set of beliefs*, consisting of all beliefs about the volatility. This framework naturally leads to a sublinear expectation and a G-Brownian motion. A G-Brownian motion is basically a standard Brownian motion with an ambiguous volatility—the volatility is completely uncertain but bounded by two extremes. We model the bond market in the spirit of the Heath-Jarrow-Morton (HJM) methodology; that is, we model the instantaneous forward rate as a diffusion process, which is driven by a G-Brownian motion. The remaining quantities on the bond market are defined in terms of the forward rate in accordance with the HJM methodology. We model the forward rate in such a way that it satisfies a suitable drift condition, ensuring the absence of arbitrage on the bond market. Additionally, we assume that the diffusion coefficient of the forward rate is deterministic, which enables us to derive pricing methods for typical derivatives and corresponds to an HJM model in which the foward rate is normally distributed.

In the presence of volatility uncertainty, we obtain a sublinear pricing measure for additional contracts we add to the bond market, which yields either a single price or a range of prices. Within the framework described above, we consider additional contracts, which we want to price without admitting arbitrage. The pricing of contracts under volatility uncertainty is different from the classical approach, since the expectation—which corresponds to the pricing measure in the classical case without volatility uncertainty—is sublinear in this setting. In contrast to the classical case, we use the sublinear expectation to determine the price of a contract or its bounds; hence, we refer to it as the *risk-neutral sublinear expectation*. To show that this approach indeed yields arbitrage-free prices, we define trading strategies and arbitrage on the bond market extended by the additional contract. Then we show that the extended bond market is arbitrage-free, meaning that we can use this approach to find no-arbitrage prices for contracts.

To simplify the pricing of single cashflows, we introduce a counterpart of the forward measure, called *forward sublinear expectation*. The forward measure is used for pricing discounted cashflows in classical models without volatility uncertainty. We define the forward sublinear

*Center for Mathematical Economics, Bielefeld University, Bielefeld, Germany

expectation by a G-backward stochastic differential equation and show that it corresponds to the expectation under the forward measure. Similar to the forward measure, the forward sublinear expectation has the advantage that computing the sublinear expectation of discounted cashflows reduces to computing the forward sublinear expectation of cashflows, discounted with the bond price. Under the forward sublinear expectation, we obtain several results needed for pricing cashflows of typical fixed income products. As a by-product, we obtain a robust version of the expectations hypothesis under the forward sublinear expectation. Moreover, we provide pricing methods for options on forward prices. The prices of such options are characterized by nonlinear partial differential equations (PDEs) or, in some cases, by the prices from the corresponding HJM model without volatility uncertainty.

In addition, we develop pricing methods for contracts consisting of several cashflows. In traditional models without volatility uncertainty, there is no distinction between pricing single cashflows and pricing a stream of cashflows, since the pricing measure is linear. However, when there is uncertainty about the volatility, the nonlinearity of the pricing measure implies that we cannot generally price a stream of cashflows by pricing each cashflow separately. Therefore, we provide different schemes for pricing a family of cashflows. If the cashflows of a contract are sufficiently simple, we can price the contract as in the classical case. In general, we use a backward induction procedure to find the price of a contract. When the contract consists of a family of options on forward prices, the price of the contract is characterized by a system of nonlinear PDEs or, in some cases, by the price from the corresponding HJM model without volatility uncertainty.

With the tools mentioned above, we derive robust pricing formulas for all major interest rate derivatives. We consider typical linear contracts, such as fixed coupon bonds, floating rate notes, and interest rate swaps, and nonlinear contracts, such as swaptions, caps and floors, and in-arrears contracts. Due to the linearity of the payoff, we obtain a single price for fixed coupon bonds, floating rate notes, and interest rate swaps; the pricing formula is the same as the one from classical models without volatility uncertainty. Due to the nonlinearity of the payoff, we obtain a range of prices for swaptions, caps and floors, and in-arrears contracts; the range is bounded from above, respectively below, by the price from the corresponding HJM model without volatility uncertainty with the highest, respectively lowest, possible volatility. Therefore, the pricing of common interest rate derivatives under volatility uncertainty reduces to computing prices in models without volatility uncertainty. For other (less common) contracts the pricing procedure requires (novel) numerical methods.

The pricing formulas show that volatility uncertainty is able to naturally explain empirical findings that many traditional term structure models fail to reproduce. According to empirical evidence, volatility risk in fixed income markets cannot be hedged by trading solely bonds, which is termed *unspanned stochastic volatility* and inconsistent with traditional term structure models. Since the presence of volatility uncertainty naturally leads to market incompleteness, the pricing formulas derived in this paper show that it is no longer possible to hedge volatility risk in fixed income markets with a portfolio consisting solely of bonds when there is uncertainty about the volatility. Moreover, the pricing formulas are in line with the empirical findings mentioned above.

Apart from giving a natural explanation for empirical findings, the theoretical results can be used in practice for different purposes. One can use the pricing procedure for stress testing by pricing contracts in the presence of different levels of volatility uncertainty and investigating how the pricing bounds behave compared to the price from the corresponding HJM model without volatility. One can also fit the pricing bounds to bid-ask spreads of quoted prices to obtain the bounds for the volatility and use them to price other contracts. Alternatively, the bounds for the volatility can be inferred from historical data on the volatility in the form of confidence intervals to generally price contracts.

PRICING OPTIONS ON FLOW FORWARDS BY NEURAL NETWORKS IN HILBERT SPACE

LUCA GALIMBERTI

Abstract

In commodity markets, options are typically written on forward and futures contracts. In some markets, like for example electricity and gas, as well as freight and weather markets on temperature and wind, the forwards deliver the underlying commodity or service over a contracted delivery period, and not at a specified delivery time in the future. Such forwards are sometimes referred to as *flow forwards*.

There is a large literature on neural networks and financial derivatives, mostly focusing on approximating option prices. As is well-known in mathematical finance, option prices can be re-cast as solutions of partial differential equations based on Feynman-Kac formulae for diffusion processes. This connection has been utilised in, say, Beck *et al.* [1, 2], Han, Jentzen and E [9] and Hutzenthaler *et al.* [10] in studying deep neural network approximations. In all these papers, the main argument for introducing deep neural networks is to overcome the curse of dimensionality. Thus, deep neural networks can be applied to price options on a high-dimensional underlying, like basket options, say.

In this talk we bring this to the "ultimate" high-dimensional case, considering deep neural networks approximating option prices on *infinite-dimensional* underlyings. Indeed, option prices on flow forwards are in general functions of functions, as the underlying will be a curve (i.e., the term structure) rather than a vector of points (i.e., prices of the underlying assets).

We propose to approximate this non-linear option price functional by a neural network in Hilbert space. More specifically, motivated by the approach of Beck *et al.* [1], we generalise the option pricing problem to deep neural networks in Hilbert space, appealing to the general neural nets in Fréchet space and their universal approximation of continuous mappings developed in Benth, Detering and Galimberti [3].

To be more specific about the problem we are dealing with and the motivation for our approach, we recall briefly options on power. In electricity markets, such as the European Energy Exchange (EEX) and Nord Pool one can trade in call and put options written on forward contracts delivering power over a contracted period of time. I.e., for a contracted delivery period $[T_1, T_2]$, where $0 \leq T_1 < T_2$, we denote the price at time $t \leq T$ by $\hat{F}(t, T_1, T_2)$. For a strike price $\mathcal{K}$ and exercise time $\tau \leq T_1$, the price of a call option at time $t \leq \tau$ is defined as (working directly under the martingale measure $\mathbb{Q}$)

$$V(t, \tau) = \mathbb{E}[\max(\hat{F}(\tau, T_1, T_2) - \mathcal{K}, 0) \,|\, \mathcal{F}_t]. \tag{1}$$

One would like to have a model for the forward prices across different delivery periods which avoids calendar arbitrage, that is, arbitrage from investing in forwards with different delivery periods. The most convenient way to do this (compare [5]) is via fixed-delivery forwards. Therefore, assuming that $F(t, T)$ is the price of an "artificial" forward at time t delivering the underlying commodity at time $T \geq t$, the following relationship holds (see Benth, Šaltytė Benth and Koekebakker [7, Prop. 4.1.]):

$$\hat{F}(t, T_1, T_2) = \frac{1}{T_2 - T_1} \int_{T_1}^{T_2} F(t, T) dT. \tag{2}$$

If one now has a dynamic model for the forward curve $F(t, \cdot)$, then prices for related forwards and options can be obtained via this last relationship and (1). Dynamic models for the forward curve have been proposed and analysed in, say, Clewlow and Strickland [8], Benth and Koekebakker [5] and Benth and Krühner [6]. Such a dynamic is usually defined

in terms of some stochastic partial differential equation (SPDE). However, while models for the forward curve lead to coherent arbitrage free prices across options with different delivery periods, they pose computational challenges because of their infinite dimensional nature.

The paper [4] on which this talk is based tries to overcome some of these challenges by proposing a numerical method for pricing options on flow forwards based on neural networks. We first derive properties of the pricing function which ensure that one can actually recast the pricing problem (1) into an optimization problem over a space of continuous functions defined on a Hilbert space of functions. We then show by a density argument that it is actually sufficient to optimize only over a restricted set of continuous functions, namely Hilbert space neural networks. These neural networks have been proposed in Benth, Detering and Galimberti [3] for approximating functionals defined on a Fréchet space $\mathfrak{X}$. We show how one can efficiently implement these neural networks by standard machine learning packages as TensorFlow/Keras and use the stochastic gradient descent algorithm for the optimization task. Our method delivers option prices automatically for a wide range of initial market conditions. This has the tremendous advantage that expensive simulations are performed only once for the training step of the neural network, and they do not have to be repeated if market conditions change.

We test our methodology in some numerical case studies, and find that it works very well with high dimensional noise which is in line with the general perception that numerical methods based on neural networks can often overcome the course of dimensionality. In the case studies, we also compare our approach with the direct classical neural network approximation where the term structure curves are simply sampled and turned into high-dimensional input objects. Numerical evidence talks strongly in favour of our infinite-dimensional network in Hilbert space as being superior.

This is a joint work with Fred Espen Benth and Nils Detering.

References

[1] C. Beck, S. Becker, P. Grohs, N. Jaafari, and A. Jentzen. Solving the Kolmogorov PDE by means of deep learning. *Journal of Scientific Computing*, 88(3), Jul 2021.

[2] C. Beck, M. Hutzenthaler, A. Jentzen, and B. Kuckuck. An overview on deep learning-based approximation methods for partial differential equations, 2021.

[3] F. E. Benth, N. Detering, and L. Galimberti. Neural networks in Fréchet spaces, 2021.

[4] F. E. Benth, N. Detering, and L. Galimberti. Pricing options on flow forward by neural networks in hilbert spaces, 2022.

[5] F. E. Benth and S. Koekebakker. Stochastic modeling of financial electricity contracts. *Energy Economics*, 30(3):1116–1157, 2008.

[6] F. E. Benth and P. Krühner. Representation of infinite dimensional forward price models in commodity markets. *Communications in Mathematics and Statistics*, 2(1):47–106, 2014.

[7] F. E. Benth, J. Šaltytė Benth, and S. Koekebakker. *Stochastic Modelling of Electricity and Related Markets*. World Scientific, Singapore, 2008.

[8] L. Clewlow and C. Strickland. *Energy Derivatives – Pricing and Risk Management*. Lacima Publications, London, 2000.

[9] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.

[10] M. Hutzenthaler, A. Jentzen, T. Kruse, and T. A. Nguyen. A proof that rectified deep neural networks overcome the curse of dimensionality in the numerical approximation of semilinear heat equations. *SN Partial Differential Equations and Applications*, 1:10, 2020.

Principal-agent mean-field games in renewable energy certificate markets

Arvind Shrivats*, Dena Firoozi, Sebastian Jaimungal*****
*Department of Operations Research & Financial Engineering, Princeton University, Princeton, NJ, USA (shrivats@princeton.edu)
**Department of Decision Sciences, HEC Montréal, Montreal, QC, Canada (dena.firoozi@hec.ca)
***Department of Statistical Sciences, University of Toronto, Toronto, ON, Canada (sebastian.jaimungal@utoronto.ca)

Principal agent games are a growing area of research which focuses on the optimal behaviour of a principal and an agent, with the former contracting work from the latter, in return for providing a monetary award. While this field canonically considers a single agent, the situation where multiple agents, or even an infinite number of agents are contracted by a principal are growing in prominence and pose interesting and realistic problems. Here, agents form a Nash equilibrium among themselves, and a Stackelberg equilibrium between themselves as a collective and the principal.
We apply this framework to the problem of implementing renewable energy certificate (REC) markets. We do so while incorporating market clearing as well as agent heterogeneity, and distinguish ourselves from extant literature by incorporating the probabilistic approach to MFGs as opposed to the analytic approach, with the former lending itself more naturally for our problem. For a given market design, we find the Nash equilibrium among agents using techniques from mean field games. We then provide results for the optimal market design from the perspective of the regulator, who aims to maximize revenue and overall environmental benefit.

More specifically, in this work, we build off [1], flipping the point of view to now incorporate the potential goals of a regulator in a REC market. We focus on the principal agent problem where the principal aims to introduce a REC market to induce investment into some renewable energy generation. Their goals are to maximize revenue and said investment into renewable energy.

Meanwhile, the agents each aim to navigate the market at minimum cost by modulating their planned REC generation and trading activities. By taking the infinite-player limit of the model, we can apply results from MFG theory to solve the agents' cost minimization problem. Given a solution to the agents' cost minimization problem we then consider how the principal may optimize the design of the REC market they institute for their own goals. We impose a constraint of the principal having to ensure the average agent achieves a cost no worse than some exogenous reservation cost, to ensure that the market they choose is `fair' in some sense to the agents.

We discuss some results for the optimal penalty function for the principal to impose under these conditions, utilizing tools from the nascent field of PA-MFGs. In particular, we restrict ourselves to considering penalty functions for which the MFG the agents experience have a solution. The methods used in this process marry the ideas of [2], [3], and [4], with the McKean-Vlasov control techniques espoused by [5]. Ultimately, we are able to find an optimal penalty function.

In particular, we find that the optimal penalty function is linear in the terminal RECs of the agents (plus some constant). This implies a constant marginal benefit to REC acquisition, which is an interesting and unintuitive result that would have fascinating implications for the design of REC markets generally, as it implies that the optimal REC market (under these conditions and assumptions) is more akin to a tax than a market.

Nonetheless, in providing the mathematical framework contained in this paper, we have produced a valid extension of a coherent structure under which REC markets can be studied holistically, including from the perspective of a regulator. In particular, this perspective while incorporating the mean field aspect of agent behaviour represented a gap that was not yet filled in this area. We reflect that this is potentially of great use to regulatory bodies in these systems, and will work to make the improvements we have outlined above.

[1] Shrivats, A., Firoozi, D., & Jaimungal, S. (2022) A Mean-Field Game Approach to Equilibrium Pricing, Optimal Generation, and Trading in Solar Renewable Energy Certificate (SREC) Markets, *Mathematical Finance*, doi: 10.1111/mafi.12345.

[2] Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, *75*(3), 957-984.

[3] Cvitanić, J., Possamaï, D., & Touzi, N. (2018). Dynamic programming approach to principal– agent problems. *Finance and Stochastics*, *22*(1), 1-37.

[4] Elie, R., Mastrolia, T., & Possamaï, D. (2019). A tale of a principal and many, many agents. *Mathematics of Operations Research*, *44*(2), 440-467.

[5] Carmona, R., & Delarue, F. (2015). Forward–backward stochastic differential equations and controlled McKean–Vlasov dynamics. *The Annals of Probability*, *43*(5), 2647-2700.

Special Sessions in memory of Prof. Tomas Bjork

Principles for modeling long-term market dynamics

Speaker: Eckhard Platen

Abstract: The paper derives eight principles that allow explaining the long-term dynamics of large stock markets, the typical distribution of the market capitalization of stocks, the risk premia for stock portfolios, the key role of optimal portfolios at the growth efficient frontier, the least expensive pricing and hedging of long-term payoffs, and other market features. By applying the concepts of entropy maximization and energy conservation in a richer modeling world than typically considered, most of these market properties follow rather directly. Furthermore, popular fundamental tools for portfolio and risk management, including the intertemporal capital asset pricing model and the preferred pricing rule, become revised.

Invited Session for stochastic modeling of information in economics and finance
Title: Progressive enlargement of filtrations and control problems for step processes

Speaker: Paolo Di Tella

Abstract

In the present paper we address stochastic optimal control problems for a step process (X, F) under a progressive enlargement of the filtration. The global information is obtained adding to the reference filtration F the point process $H = 1_{[\tau,+\infty)}$. Here τ is a random time that can be regarded as the occurrence time of an external shock event. We study two classes of control problems, over $[0, T]$ and over the random horizon $[0, T \wedge \tau]$. We solve these control problems following a dynamical approach based on a class of BSDEs driven by the jump measure μ^Z of the semimartingale $Z = (X, H)$, which is a step process with respect to the enlarged filtration G. The BSDEs that we consider can be solved in G thanks to a martingale representation theorem which we also establish here. To solve the BSDEs and the control problems we ensure that Z is quasi-left continuous in the enlarged filtration G. Therefore, in addition to the F-quasi left continuity of X, we assume some further conditions on τ. This is a joint work with Elena Bandini (Università di Bologna) and Fulvia Confortola (Politecnico di Milano).

RANDOM FEATURE NEURAL NETWORKS LEARN BLACK-SCHOLES TYPE PDES WITHOUT CURSE OF DIMENSIONALITY

LUKAS GONON

This work investigates the use of random feature neural networks for learning option prices and Kolmogorov partial (integro-)differential equations associated to Black-Scholes and more general Lévy models. Random feature neural networks Huang et al. (2006), Rahimi and Recht (2008), Rahimi and Recht (2009) are single-hidden-layer feedforward neural networks in which only the output weights are trainable. To be more specific, a random neural network takes the form

$$H_W^{A,B}(x) = \sum_{i=1}^{N} W_i \varrho(A_i \cdot x + B_i), \quad x \in \mathbb{R}^d,$$

where $\varrho\colon \mathbb{R} \to \mathbb{R}$ is a fixed activation function, $A = (A_1, \dots, A_N)$ and $B = (B_1, \dots, B_N)$ are randomly generated weights and W is a vector of output weights that is trained based on data. The non-convex optimization problem that needs to be solved in order to train a standard neural network reduces to a convex optimization problem here. This simplifies both training in practice and theoretical analysis. On the other hand, allowing only parts of the parameters to be trained reduces the approximation capabilities and so, at least a priori, it is not clear if random neural networks still have any of the powerful approximation properties of general deep neural networks. Indeed, deep neural networks have been proved to be able to approximate option prices and solutions to various classes of PDEs without the curse of dimensionality, see for instance Grohs et al. (2018), Gonon and Schwab (2021a), Hutzenthaler et al. (2020), Reisinger and Zhang (2019).
In this work we show that expressivity is not reduced in the context of learning option prices and partial integro-differential equations (PIDEs) associated to non-degenerate exponential Lévy models from data: if a random neural network is used as a learning technique, then this network is capable of learning option prices / PIDEs on a hypercube $[-M, M]^d$ without the curse of dimensionality. Interestingly – in contrast to works addressing approximation by deep neural networks – in the case of random neural networks all error components (approximation, generalization and optimization) can be addressed. We derive bounds for the prediction error and show that the derived bounds do not suffer from the curse of dimensionality, i.e., for each of these error components we obtain polynomial convergence rates which do not depend on the dimension d of the underlying PDE and constants which grow at most polynomially in d.
Finally, we also provide examples and apply the results in the context of pricing max-call and basket options and validate the bounds numerically.

References

L. Gonon and C. Schwab. Deep ReLU network expression rates for option prices in high-dimensional, exponential Lévy models. *Preprint, arXiv 2101.11897, to appear in Finance Stoch.*, 2021a.

L. Gonon and C. Schwab. Deep relu neural network approximation for stochastic differential equations with jumps. *Preprint, arXiv 2102.11707*, 2021b.

P. Grohs, F. Hornung, A. Jentzen, and P. von Wurstemberger. A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of Black-Scholes partial differential equations. *To appear in Memoirs of the American Mathematical Society; arXiv:1809.02362*, page 124 pages, 2018.

G.-B. Huang, L. Chen, and C.-K. Siew. Universal approximation using incremental constructive feedforward networks with random hidden nodes. *Trans. Neur. Netw.*, 17(4):879–892, July 2006. ISSN 1045-9227. doi: 10.1109/TNN.2006.875977. URL `https://doi.org/10.1109/TNN.2006.875977`.

M. Hutzenthaler, A. Jentzen, T. Kruse, T. A. Nguyen, and P. von Wurstemberger. Overcoming the curse of dimensionality in the numerical approximation of semilinear parabolic partial differential equations. *Proc. A.*, 476(2244):630–654, 2020. ISSN 1364-5021. doi: 10.1098/rspa.2019.0630. URL `https://doi.org/10.1098/rspa.2019.0630`.

A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2008.

A. Rahimi and B. Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21, pages 1313–1320. Curran Associates, Inc., 2009. URL `https://proceedings.neurips.cc/paper/2008/file/0efe32849d230d7f53049ddc4a4b0c60-Paper.pdf`.

C. Reisinger and Y. Zhang. Rectified deep neural networks overcome the curse of dimensionality for nonsmooth value functions in zero-sum games of nonlinear stiff systems. *arXiv:1903.06652*, page 34 pages, 2019.

Title: Rational Expectations Equilibrium With Heterogeneous Information Flows

Co-authors: Jerome Detemple (Boston University), Sergio Pulido (ENSIIE)

Abstract: In this talk, we consider equilibria in the presence of dynamically evolving asymmetric information. "Insiders" obtain private signals about the terminal value of an asset at intermediate times, and at each signal time, the "uninformed" agent receives a noisy version of the respective insider's signal. Both the filtration and asset price process jump at each signal time, and markets are in general incomplete. After establishing existence of a dynamic noisy rational expectations equilibrium for a finite set of signal times, we pass to the continuous time limit where the fundamental filtration is enlarged by a continuous signal flow process. Results are valid for constant absolute risk aversion investors, and where the factor process follows a multi-dimensional OU process. After presenting OU results, we will discuss extending to general factor processes.

Realization utility, market regimes, and the disposition effect

Xue Dong He* Shengcheng Shao* Moris S. Strub†

Abstract

We study a model of realization utility in a financial market where the drift of the stock switches between a bull and a bear market driven by a hidden Markov chain. In this model, an investor trading between a risky, regime switching stock and a risk-free bond seeks to determine an optimal sequence of purchasing and selling times for the stock to balance utility bursts experienced when realizing a gain or loss of the stock and utility derived from terminal wealth. We characterize the optimal value functions when holding the bond and stock respectively as the unique solutions to a system of coupled variational equalities and accordingly obtain optimal hodling and selling regions for the bond and stock. The presence of market regimes and preferences incorporating realization utility have opposite effects on trading behavior: Models of realization utility without market regimes predict extreme versions of the disposition effect while a model incorporating market regimes but not realization utility predicts trend following strategies, i.e., the opposite of the disposition effect. Including both features generates a rich set of trading behavior including voluntarily realizing gains and losses with optimal liquidation points depending on the assessment of the market state, selling of the stock with or without immediate repurchase, and reasonable levels of the disposition effect and other statistics conforming with recorded behavior of individual investors.

Keywords: Realization utility; regime switching model; disposition effect

*Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. Email: {xdhe@se.cuhk.edu.hk, david.shao.91@gmail.com}

†Department of Information Systems and Management Engineering, Southern University of Science and Technology. Email: strub@sustech.edu.cn

Special Sessions in memory of Prof. Tomas Bjork

Recalling some of Tomas Bjork’ s contributions to the mathematical theory of interest rates

Speaker: Wolfgang Runggaldier

Abstract: Tomas Bjork may well be considered as one of the main actors in the development of the mathematical theory of the term structure of interest rates. I have had the privilege of working with him on this topic from very early onwards. In this talk I try to recall some of the main issues in this area that Tomas had been working on, partly jointly with me, and what were some of his contributions there.

Recover utility of rational inattentive agent and applications on robo-advising

Hao Xing* Zeyu Zhu†

Abstract

We consider a rational inattentive agent who acquires costly signal to make decisions. By observing agent's actions, we formulate an inverse reinforcement learning problem to recover agent's utility. We propose an efficient numeric algorithm and prove its convergence. The framework is applied to robo-advising problems to recover investors' utilities by observing their investment strategies in both a mean-variance and a target date investment setting.

More specifically, following Sims (1998, 2003) and Matejka and Mckay (2015), we consider a decision problem where it is costly for the decision maker to acquire information. The decision maker needs to determine the optimal signal structure to balance the utility of choices and the cost of information acquisition. For a finite state space X, a finite action space A, and a utility function $u : X \times A \to \mathbb{R}$, the decision maker determines a choice rule $p(a|x)$, which is the conditional probability of choosing the action a for given the state x, and the default rule $q(a)$, which is the unconditional probability of choosing the action a.

Given the state-action demonstrations $\{(x_i, a_i); i \geq 1\}$ from data, we introduce an inverse problem for the aforementioned rational inattention problem, in order to recover decision maker's utility from the observed actions:

$$\min_u \left\{ \psi(u) + \max_{p,q} \left\{ H(p,q) + \sum_{x,a} u(x,a)\mu(x) \big(p(a \mid x) - p_E(a \mid x) \big) \right\} \right\},$$

where μ is the prior distribution of states and

$$H(p,q) = \sum_x \mu(x) \sum_a p(a \mid x) \ln \frac{p(a \mid x)}{q(a)}$$

is the relative entropy between p and q. Our formulation is motivated by entropy inverse reinforcement learning problem (Finn et. al. (2016) and Ho and Ermon (2016)), but it has important difference comparing to the inverse reinforcement learning problems.

*Boston University, haoxing@bu.edu.
†Boston University, zz918@bu.edu.

1

Starting from a static setting, we characterize the recovered utility and identify its uniqueness in an equivalence class. We propose an efficient numeric algorithm and prove its convergence. We also introduce a dynamic formulation and its associated numeric algorithm.

We apply our methodology to two robo-advising problems, where investors have mean-variance preference or utility for consumption at a future investment target date. Our methodology perform well in both settings.

2

Special Sessions in memory of Prof. Tomas Bjork

Reduced form framework under model uncertainty

Speaker: Francesca Biagini

Abstract: In this talk we present a reduced form framework under model uncertainty. In particular we introduce a sublinear conditional operator with respect to a family of possibly nondominated probability measures in presence of a single or multiple ordered default times, and use it for the valuation of credit portfolio derivatives under model uncertainty. In addition, we extend this framework to include mortality intensities following an affine process under parameter uncertainty.
This talk is based on joint work with several co-authors: A. Mazzon, K.Oberpriller, and Y. Zhang.

Regret and Asset Pricing*

Jorgo T.G. Goossens†

Most recent version here

December 23, 2021

Abstract

I investigate the consequences of regret aversion for asset prices in an otherwise standard model of financial markets. This paper shows that accounting for investors' regret aversion can help explain the risk-free rate puzzle, excess volatility, the downward sloping term structure of equity risk premiums, and the predictability of stock returns both in the time series and in the cross section. The model also evaluates bond behavior and predicts a downward sloping real yield curve. I provide an empirical measure of regret which confirms empirically the main model's predictions. This paper is the first to document the linkage between regret aversion and many stylized facts concerning asset prices.

Keywords: equity, bonds, asset pricing puzzles, stylized facts, regret aversion

JEL Codes: G12, G41

*I am particularly grateful to Bas Werker and Ralph Koijen. I thank Paul Bokern, Fabio Braggion, Stefano Cassella, Aditya Chaudry, Alex Clyde, Joost Driessen, Rik Frehen, Rob van den Goorbergh, Marike Knoef, Stefan Nagel, Antoon Pelsser, Eduard Ponds, Nikolaus Schweizer, Martijn de Vries, Ole Wilms and Marcel Zeelenberg, and seminar participants at the European Winter Meeting of the Econometric Society (2021), EEA (2021), the RCEA Money-Macro-Finance Conference (2021), the Netspar Pension Day (2021), the ENTER Jamboree (2021), the KVS New Paper Sessions (2021), the Netspar International Pension Workshop (2021) and Tilburg University (2021, 2020) for useful comments.

†Tilburg University, department of Econometrics and Operations Research. Corresponding author: j.t.g.goossens@uvt.nl

REINFORCEMENT LEARNING WITH DYNAMIC CONVEX RISK MEASURES

Anthony Coache
Department of Statistical Sciences
University of Toronto
anthony.coache@mail.utoronto.ca
`https://anthonycoache.ca/`

Sebastian Jaimungal
Department of Statistical Sciences
University of Toronto
sebastian.jaimungal@utoronto.ca
`http://sebastian.statistics.utoronto.ca/`

Reinforcement learning (RL) provides a (model-free) framework for learning-based control. RL problems aim at learning dynamics in the underlying data and finding optimal behaviors while collecting data via an interactive process. It differs from *supervised learning* that attempts to learn classification functions from labeled data, and *unsupervised learning* that seeks hidden patterns in unlabeled data. The agent makes a sequence of decisions while interacting with the data-generating process and observing feedback in the form of costs. The agent aims to discover the best possible actions by interacting with the environment and consistently updating their actions based on their experience, while often, also taking random actions to assist in exploring the state space – the classic exploration-exploitation trade-off (Sutton and Barto, 2018).

In RL, *uncertainty* in the data-generating process can have substantial effects on performance. Indeed, the environmental randomness may result in algorithms optimized for "on-average" performance to have large variance across scenarios. For example, consider a portfolio selection problem: a risk-neutral optimal strategy (where one optimizes the expected terminal return) focuses on assets with the highest returns while ignoring the risks associated with them. As a second example, consider an autonomous car which should account for environmental uncertainties such as the weather and road conditions when learning the optimal strategy. Such an agent may wish to account for variability in the environment and the results of its actions to avoid large potential "losses". For an overview and outlook on RL in financial mathematics, see e.g. Jaimungal (2022).

In the extant literature, there are numerous proposals for accounting for risk sensitivity, where most of them replace the expectation in the optimization problem by risk measures. *Risk-awareness* or *risk-sensitivity* in RL offers a remedy to the data-generating process uncertainty by quantifying low-probability but high-cost outcomes, provides strategies that are more robust to the environment, and allows more flexibility than traditional approaches as it is tuned to the agent's risk preference. The specific choice of risk measure is a decision the agent makes considering their goals and risk tolerances.

Several authors address risk evaluation for sequential decision making problems by applying risk measures recursively to a sequence of cost random variables, and optimizing the risk in a *dynamic framework* as additional information becomes available. For instance, Ruszczyński (2010) evaluates the risk at each period using dynamic Markov coherent risk measures, while Chu and Zhang (2014) and Bäuerle and Glauner (2021) propose iterated coherent risk measures, where they both derive risk-aware dynamic programming (DP) equations and provide policy iteration algorithms. While they focus on coherent risk measures, various classes of risk measures have already been extended to a dynamic framework, such as distribution-invariant risk measures (Weber, 2006), coherent risk measures (Riedel, 2004), convex risk measures (Frittelli and Gianin, 2004), and dynamic assessment indices (Bielecki et al., 2016), among others – however, these works do not look at how to perform model free optimization, i.e., they do not look at RL.

An interesting approach to risk-aware learning stems from Tamar et al. (2015; 2016), where they provide policy search algorithms to solve risk-aware RL problems in a dynamic framework. Both studies investigate *stationary policies*, restrict themselves to *coherent* risk measures, and apply their methodology to simple financial engineering applications. More specifically, they evaluate their algorithms when learning policies for (perpetual) American options and trading in static portfolio optimization problems. Several real-world applications require a level of flexibility that these limitations preclude.

Here, we develop a model-free approach to solve a wide class of (non-stationary) risk-aware RL problems in a time-consistent manner. Our contributions may be summarized as follows: (i) we extend Tamar et al. (2015; 2016); Ahmadi et al. (2020); Kose and Ruszczynski (2021) by focusing on the broad class of *dynamic convex* risk measures and consider finite-horizon problems with *non-stationary policies*; (ii) we devise an *actor-critic* algorithm to solve this class of RL problems using neural networks to allow continuous state-action spaces; (iii) we derive a recursive formula for efficiently computing the policy gradients; and (iv) we demonstrate the performance and flexibility of our proposed approach on three important applications: optimal trading for statistical arbitrage, obstacle avoidance in robot control, and hedging financial options. We demonstrate that our approach appropriately accounts for uncertainty and leads to strategies that mitigate risk.

References

Mohamadreza Ahmadi, Ugo Rosolia, Michel D Ingham, Richard M Murray, and Aaron D Ames. Constrained risk-averse markov decision processes. *arXiv preprint arXiv:2012.02423*, 2020.

Nicole Bäuerle and Alexander Glauner. Markov decision processes with recursive risk measures. *European Journal of Operational Research*, 2021.

Tomasz R Bielecki, Igor Cialenco, Samuel Drapeau, and Martin Karliczek. Dynamic assessment indices. *Stochastics*, 88(1):1–44, 2016.

Shanyun Chu and Yi Zhang. Markov decision processes with iterated coherent risk measures. *International Journal of Control*, 87(11):2286–2293, 2014.

Marco Frittelli and E Rosazza Gianin. Dynamic convex risk measures. *Risk measures for the 21st century*, pages 227–248, 2004.

Sebastian Jaimungal. Reinforcement learning and stochastic optimisation. *Finance and Stochastics*, 26(1):103–129, 2022.

Umit Kose and Andrzej Ruszczynski. Risk-averse learning by temporal difference methods with markov risk measures. *J. Mach. Learn. Res.*, 22:38–1, 2021.

Frank Riedel. Dynamic coherent risk measures. *Stochastic processes and their applications*, 112(2):185–200, 2004.

Andrzej Ruszczyński. Risk-averse dynamic programming for markov decision processes. *Mathematical programming*, 125(2):235–261, 2010.

Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.

Aviv Tamar, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. Policy gradient for coherent risk measures. *Advances in Neural Information Processing Systems*, 28:1468–1476, 2015.

Aviv Tamar, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. Sequential decision making with coherent risk. *IEEE Transactions on Automatic Control*, 62(7):3323–3338, 2016.

Stefan Weber. Distribution-invariant risk measures, information, and dynamic consistency. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 16(2):419–441, 2006.

RETHINKING HETEROGENEOUS MODELS AND MEAN FIELD GAMES IN DISCRETE TIME

Andrew Lyasoff
alyasoff@bu.edu; http://www.andrewlyasoff.tech

Keywords: Heterogeneous Models, Mean Field Games, Incomplete Markets, General Equilibrium

JEL: D52, C68, D58, E21, MSC: 91A16, 91A50, 91B50.

This paper is prompted by a recent discovery – see Lyasoff (2019) – showing that the well known solution procedure for the classical Bewley-Aiyagari heterogeneous model, as described in, for example, Ch. 18 of Ljungqvist and Sargent (2018), fails to produce equilibrium in the very example (borrowed from ibid.) that is meant to illustrate it, if the approximation grid is refined sufficiently.[1] The Bewley-Aiyagari heterogeneous model is in effect the workhorse in macroeconomics, and is concerned with a very large number of agents who live forever, constantly face idiosyncratic employment shocks (modeled by statistically identical but independent Markov chains), and invest in a risk-free private lending instrument which is in aggregate supply of 0. The classical solution method comes down to these steps: (a) start by fixing the value of the interest rate r; (b) compute (with the chosen r) the optimal policy (demand for the bond) of a *generic* (stand alone) agent, as a function of her state (amount of wealth and employment level), by solving the associated Bellman equation; (c) calculate the steady state distribution of the optimal state of a generic agent, and *reinterpret* that distribution as equal to the cross-sectional distribution of the entire population of agents over the range of wealth and employment categories; (d) calculate the aggregate demand, $a(r)$, for the bond, by integrating the optimal policy function of a generic agent against the stipulated cross-sectional distribution of agents; (e) if $a(r)$ is close to 0, then stop and declare r to be the equilibrium rate, if $a(r) > 0$ decrease r and go back to (b), if $a(r) < 0$ increase r and go back to (b). The outcome from these steps, applied to the example borrowed from Ljungqvist and Sargent (2018) is shown in Figure 1 below. The reason for the failure of the right plot in Figure 1 to locate the equilibrium is illustrated by the left panel in Figure 2: while the policy function depends on r continuously, there is an apparent bifurcation in the dependence on r of the long-run probability distribution of the optimal state[2] (interest rates that are within 10^{-6} of each other generate vastly different distributions).

Cardaliaguet et al. (2019) note that the main equations in theory of MFG were developed as "the continuous-time analogues of equations that appear in the analysis of dynamic stochastic general equilibria in heterogeneous agent models (Aiyagari (1994), Bewley (1986), Huggett (1993), Krusell & Smith (1998))." Indeed, the dynamics of a mean field game are usually described in terms of a McKean-Vlasov equation (Carmona and Delarue 2018, 3.1.2-3). Figuratively speaking, these are the dynamics of an "agent" who is aware (when deciding how to play) not only of her state, but also of the distribution of that state. To put it another way, the optimal state of each player is "exactly distributed according to the state of the population" (Cardaliaguet et al. 2019, 1.1.3). Alternatively, such games can be written as coupled systems of a backward HJB equation, in which the running cost depends on the distribution, and a forward Kolmogorov equation that governs that distribution.

As Figure 1 makes clear, contrary to the common belief, the discrete-time analog of the theory of MFG cannot be applied to certain situations that are of paramount importance in macroeconomics.[3] The present paper develops a discrete time alternative to MFG, and *without the stipulation that the optimal state of each player is "exactly distributed according to the state of the population."* The strategy on which this new approach is based is analogous to that of Dumas and Lyasoff (2012)[4] and is reminiscent to the technology of FBSDEs and decoupling fields (discrete time version). The key innovation is in the way these technologies are coordinated among a large population of agents and across time. The main steps in the application to heterogeneous models are outlined in Lyasoff (2019) and the resolution of

[1] Somehow, the classical method appears to almost work on a coarse grid (see the left panel on Figure 1), which is the likely reason for the failure of the method to remained unnoticed.

[2] A "state" is a pair of an employment category (one of 7 possible in this example) and a level of wealth. The cross-sectional distribution of agents, or the distribution of the optimal state of a generic agent, can be expressed as a list of 7 distributions over the range of wealth, which get selected at random with the steady-state probabilities attached to the respective employment categories.

[3] Although analogous, MFG in continuous time, or, more generally, stochastic differential games, are not exactly equivalent to their discrete time counterparts. Just as an example, assuming a constant diffusion coefficient is a step that has no analog in situations where the diffusion process is replaced by a generic Markov chain – as in the example illustrated in Figure 1.

[4] The present paper may be viewed as an extension of Dumas and Lyasoff (2012) for heterogeneous incomplete market models with infinitely many agents. Just as in MFG, passing to the limit as the number of agents increase to ∞ becomes a highly nontrivial task.

the example from Figure 1 is illustrated by the right panel in Figure 2: the equilibrium rate of $\approx 3.7\%$, computed with the new model, is substantially larger than the rate one might infer from the left panel in Figure 1, and the equilibrium cross-sectional distribution of agents is vastly different from the long-run stand-alone (derived independently from everything else) distribution of the individual optimal state. This demonstrates the significance of removing the stipulation noted above in italics and illustrates the difference between the approach developed in the present paper and the classical MFG theory and methods. Contrary to what one sees in Figure 1, the cross-sectional distribution from the right panel in Figure 2 clears the market within ≈-1.76078736e-6, with maximal error in all variables of ≈9.04753562e-5. The classical Bewley-Aiyagari heterogeneous model can now be solved.

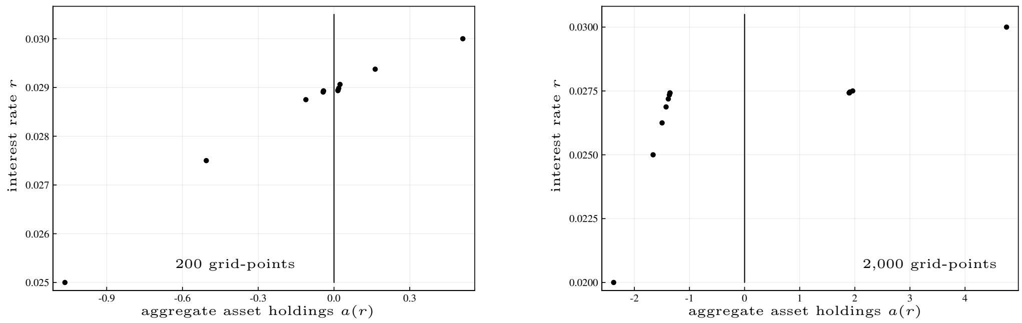

Figure 1: The classical solution to the pure credit example borrowed from Ljungqvist and Sargent (2018).

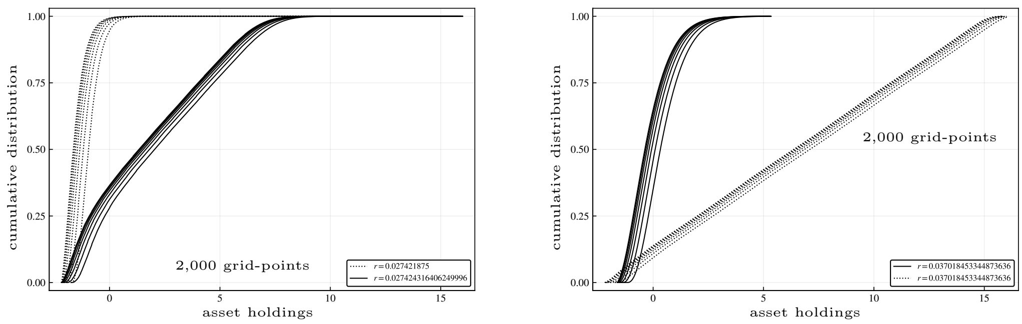

Figure 2: Left panel: the bifurcation in the dependence on r of the long-run distribution of the optimal state. Right panel: the actual cross-sectional distribution of agents (the solid lines) in equilibrium, computed with the new method, vs. the stand-alone (independent from anything else) long-run distribution of the optimal individual state, computed with the same *equilibrium* rate of $\approx 3.7\%$ by way of conventional dynamic programming.

REFERENCES

Cardaliaguet, P., F. Delarue, J.-M. Lasry, and P.-L. Lions. (2019). *The Master Equation and the Convergence Problem in Mean Field Games*. Princeton, NJ: Princeton University Press.

Carmona, R., and F. Delarue. (2018a). *Probabilistic Theory of Mean Field Games with Applications I: mean field FBSDEs, control, and games*. Cham, CH: Springer.

Dumas, B., and A. Lyasoff. (2012). Incomplete-market equilibria solved recursively on an even tree. *Journal of Finance* **67** 1897-1941.

Ljungqvist, L., and T. J. Sargent. (2018). *Recursive Macroeconomic Theory (4th Ed.)*. Cambridge, MA: MIT Press.

Lyasoff, A. (2019). Bewley's incomplete market model revisited. *Working Paper*.

Risk contributions of lambda quantiles

A. INCE*,[⁎], I. PERI*,† and S. PESENTI‡

*Department of Economics, Mathematics and Statistics, Birkbeck, University of London, Malet Street, Bloomsbury, London WC1E 7HX, UK
‡Department of Statistical Sciences, University of Toronto, 700 University Avenue, Toronto, ON M5G 1Z5, Canada

(*December 30, 2021*)

Risk contributions of portfolios form an indispensable part of risk adjusted performance measurement. The risk contribution of a portfolio, e.g., in the Euler or Aumann-Shapley framework, is given by the partial derivatives of a risk measure applied to the portfolio profit and loss in direction of the asset units. For risk measures that are not positively homogeneous of degree 1, however, known capital allocation principles do not apply. We study the class of lambda quantile risk measures that includes the well-known Value-at-Risk as a special case but for which no known allocation rule is applicable. We prove differentiability and derive explicit formulae of the derivatives of lambda quantiles with respect to their portfolio composition, that is their risk contribution. For this purpose, we define lambda quantiles on the space of portfolio compositions and consider generic (also non-linear) portfolio operators.

We further derive the Euler decomposition of lambda quantiles for generic portfolios and show that lambda quantiles are homogeneous in the space of portfolio compositions, with a homogeneity degree that depends on the portfolio composition and the lambda function. This result is in stark contrast to the positive homogeneity properties of risk measures defined on the space of random variables which admit a constant homogeneity degree. We introduce a generalised version of Euler contributions and Euler allocation rule, which are compatible with risk measures of any homogeneity degree and non-linear but homogeneous portfolios. These concepts are illustrated by a non-linear portfolio using financial market data.

Keywords: Lambda Quantiles; Capital Allocation; Risk Contribution; Lambda Value-at-Risk; Euler Allocation

History: Earlier versions have been presented at Birkbeck, University of London PhD Jamboree, the *Mathematical and Statistical Methods for Actuarial Sciences and Finance 2020 Conference (eMAF)* (virtual), and the *10th General Advanced Mathematical Methods for Finance (AMaMeF) 2021 Conference* (virtual).
[⁎]Corresponding author. Email: aince02@student.bbk.ac.uk
†Email: i.peri@bbk.ac.uk
‡Email: silvana.pesenti@utoronto.ca

Igor Cialenco, Illinois Institute of Technology

Risk-Averse Control of Systems with Model Uncertainty

Abstract. We consider a Markov decision process subject to model uncertainty in a Bayesian framework, where we assume that the state process is observed but its law is unknown to the observer. In addition, while the state process and the controls are observed at time t

, the actual cost that may depend on the unknown parameter is not known at time t. The controller optimizes these running costs by using a family of special risk measures, that we call risk filters and that are appropriately defined to take into account the model uncertainty of the controlled system. These key features lead to non-standard and non-trivial risk-averse control problems, for which we derive the Bellman principle of optimality. We illustrate the general theory on several practical examples. This is a joint work with Andrzej Ruszczynski and Tomasz R. Bielecki.

Risk-Sensitive Credit Portfolio Optimization under Partial Information and Contagion Risk

Lijun Bo * Huafu Liao † Xiang Yu ‡

Abstract

This paper studies the finite time risk-sensitive portfolio optimization in a regime-switching credit market with physical and information-induced default contagion. The Markovian regime-switching process is assumed to be unobservable, which has countable states that affect default intensities of surviving assets. The stochastic control problem is formulated under partial observations of asset prices and default events. By proving an innovative martingale representation theorem based on incomplete and phasing out filtration, we characterize the value function in an equivalent but simplified form. This allows us to connect the previous control problem to a quadratic BSDE with jumps that is new to the literature, in which the driver term has non-standard structures and carries the conditional filter as an infinite-dimensional parameter. By proposing some novel truncation techniques, we obtain the existence of solution to this new BSDE using the delicate convergence of solutions associated to some truncated BSDEs. The verification theorem and the characterization of the optimal trading strategy can be concluded with the aid of our newly established BSDE results.

Keywords: Risk-sensitive control; default contagion; partial observations; countable regime states; quadratic BSDE with jumps; martingale representation theorem; verification theorem.

1 Introduction

This paper aims to employ the risk-sensitive criteria for the dynamic optimal portfolio among multiple credit risky assets, and we further recast the problem into a more practical setting when the underlying regime-switching process is not observable by the investor.

2 The Model and Main Results

The hidden regime-switching process I is described by a continuous time Markov chain with transition rate matrix $Q = (q_{ij})_{1\leq i,j\leq m}$, where $2 \leq m \leq +\infty$. We denote $H = (H_i(t);\ i = 1, \ldots, n)_{t\geq 0}$

*Email: lijunbo@ustc.edu.cn, School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui Province, 230026, China, and Wu Wen Tsun Key Laboratory of Mathematics, Chinese Academy of Science, Hefei, Anhui Province 230026, China.

†Email: lhflhf@mail.ustc.edu.cn, School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui Province, 230026, China.

‡E-mail: xiang.yu@polyu.edu.hk, Department of Applied Mathematics, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong.

as the default indicator process of the n risky assets with the state space $S_H = \{0,1\}^n$.

The observed market information is described by filtration $\mathbb{F}^{\mathrm{M}} := (\mathcal{F}_t^{\mathrm{M}})_{t\geq 0}$, which is generated by default indicator process H and the pre-default price processes of defaultable assets. This assumption clearly matches with the real life situation that the investor can no longer perceive any information from the asset once it exits the market. Furthermore, such incomplete information filtration in this paper possesses a phasing out feature due to sequential defaults of multiple assets.

We start with the martingale representation theorem under $\mathbb{F}^{\mathrm{M}}$. Based on the foresaid martingale representation theorem, we investigate the quadratic BSDE with jumps

$$\begin{cases} dY(u) = f(p^{\mathrm{M}}(u), H(u), Z(u), V(u))du + Z(u)^\top dW^{o,\tau}(u) + V(u)^\top d\Upsilon^*(u), \quad u \in [t,T]; \\ \\ Y(T) = 0, \end{cases} \tag{2.1}$$

where the function $f(p,z,\xi,v)$ maps $(0,1)^m \times S_H \times \mathbb{R}^n \times \mathbb{R}^n$ to $\mathbb{R}$ and grows quadratically in ξ.

There is a vast literature on quadratic BSDE with jumps. The main results in Morlais (2009), Morlais (2010) and Kazi-Tani et al. (2015) rely on a quadratic-exponential structure of the drivers in their BSDEs, which BSDE (2.1) doesn't possess. On the other hand, the conventional progressive enlargement argument in Ankirchner et al. (2010) and Kharrobi and Lim (2014) only works for the filtration generated by Brownian motions and default processes, while in BSDE (2.1) $\mathbb{F}^{\mathrm{M}}$ is generated by stopped Brownian motions $W^{o,\tau}$ and compensated default processes Υ^*.

In order to prove the existence of solution to BSDE (2.1), we adopt and modify some approximation arguments in Kobylanski (2000) to fit into our setting with jumps. Using the solution just announced, we characterize the value function and give the optimal trading strategy.

References

Ankirchner, S., C. Blanchet-Scalliet, and A. Eyeraud-Loisel (2010): Credit risk premia and quadratic BSDEs with a single jump. *Inter. J. Theor. Appl. Finan.* 13, 1103-1129.

Kharroubi, I., and T. Lim (2014): Progressive enlargement of filtrations and backward stochastic differential equations with jump. *J. Theor. Probab.* 27, 683-724.

Kobylanski, M. (2000): Backward stochastic differential equations and partial differential equations with quadratic growth. *Ann. Probab.* 28, 558-602.

Morlais, M.A. (2009): Utility maximization in a jump market model. *Stochastics.* 81, 1-27.

Morlais, M.A. (2010): A new existence result for quadratic BSDEs with jumps with application to the utility maximization problem. *Stoch. Process. Appl.* 120, 1966-1995.

Kazi-Tani, N., D. Possamai, and C. Zhou (2015): Quadratic BSDEs with jumps: a fixed point approach. *Electron. J. Probab.* 66, 1-28.

Special Sessions in memory of Prof. Mark H.A. Davis

Risk-sensitive investment management

Sebastien Lleo

Abstract: I discuss risk-sensitive stochastic control and its applications to investment management, focusing on Mark's considerable influence and contributions through a series of 11 articles, four chapters, and a book. Risk-sensitive stochastic control directly includes the agent's risk aversion in the control criterion. Thus, it is ideal for tackling a wide range of portfolio management problems. The talk starts with a presentation of the seminal works by Bielecki, Pliska, Kuroda, and Nagai. Then I introduce control criteria for benchmarked asset management and ALM. Next, I sketch solution techniques for jump-diffusion problems. Finally, I outline a partial observation model with expert opinions à la Black Litterman.

Robust deep hedging

Eva Lütkebohmert[1], Thorsten Schmidt[2], Julian Sester[3]

February 11, 2022

[1] *Department of Quantitative Finance,
Institute for Economic Research, University of Freiburg,
Rempartstr. 16, 79098 Freiburg, Germany.*

[2] *Department of Mathematical Stochastics, Mathematical Institute,
University of Freiburg, Ernst-Zermelo-Str. 1, 79104 Freiburg, Germany.*

[3] *NTU Singapore, Division of Mathematical Sciences,
21 Nanyang Link, Singapore 637371*

1 Extended abstract

Uncertainty, as coined by Frank Knight, refers to the case where a number of models (technically: probability measures) are available and one is not able to distinguish between them. This applies for example to the prediction of the evolution of a stock in the future. Even if we have a reliable and rich source of historic information, predicting the future evolution, the future variance or even the whole future distribution is highly complicated. On the one side, this is due to the classical estimation problem: estimated parameters allow for confidence intervals which need to be taken into account for the prediction. On the other side, in particular in financial markets, changes in the underlying dynamics are rather the rule than the exception and additional uncertainty and model risk come into effect, resulting in a widening of confidence intervals. For pricing, one can efficiently rely on the calibration to option surfaces with all its difficulties. For hedging, when one wants to incorporate the performance under the objective measure and not under the risk-neutral one, this becomes much more challenging.

Our paper addresses exactly this setting and suggests a deep learning approach for hedging under parameter uncertainty. The basis for our work is the recently developed class of affine processes under parameter uncertainty. We extend this approach to those Markovian processes which satisfy

$$dX_t = (b_0 + b_1 X_t)dt + (a_0 + a_1 X_t)^\gamma dW_t, \tag{1.1}$$

where we allow for parameter uncertainty in all the parameters b_0, b_1, a_0, a_1, and γ. We develop the theory for this class of processes which we call nonlinear generalized affine (NGA) processes. The associated pricing problem under parameter uncertainty is solved by utilizing a general dynamic programming principle and establishing the nonlinear Kolmogorov equation, which opens the door for fast (and well-known) numerical approaches. In order to solve the hedging problem under parameter uncertainty, we propose a deep hedging approach which efficiently solves the hedging problem under parameter uncertainty. To the best of our knowledge, this is the first attempt of this kind.

We numerically evaluate this method first on simulated data and show that the robust deep hedging outperforms existing hedging approaches when parameter uncertainty is present.

For a realistic data application, we consider the COVID-19 period. In this period, stock markets experienced unexpectedly high volatility and variation in the price paths, which poses a huge challenge to classical hedging approaches. When applying robust methods, the first challenge is to find reliable estimates for the parameter intervals specifying the uncertainty in the considered model class. To that end, we propose a sliding-window maximum-likelihood estimation approach whose maximal and minimal parameter estimates lead to the targeted intervals. With this uncertainty specification at hand, we are able to show that in the mentioned data example showing high volatility, the robust deep hedging approach leads to a remarkably smaller hedging error in comparison to classical hedging strategies.

Abstract

In this paper, we consider a general dynamic mean-variance framework and propose a novel definition of the robust equilibrium strategy. Under our definition, a classical solution to the corresponding PDE system must imply a robust equilibrium strategy. We then explicitly solve the PDE system for some special stochastic dynamics arising from portfolio selection problems. It is found that in all special models, the worst-case scenario for the investor is independent of her/his wealth level and time to maturity.

ROBUST ESTIMATION OF SUPERHEDGING PRICES

JOHANNES OBLOJ AND JOHANNES WIESEL

Abstract. We consider statistical estimation of superhedging prices using historical stock returns in a frictionless market with d traded assets. We introduce a plugin estimator based on empirical measures and show it is consistent but lacks suitable robustness. To address this we propose novel estimators which use a larger set of martingale measures defined through a tradeoff between the radius of Wasserstein balls around the empirical measure and the allowed norm of martingale densities. We then extend our study, in part, to estimation of risk measures, to the case of markets with traded options, to a multiperiod setting and to settings with model uncertainty. We also study convergence rates of estimators and convergence of superhedging strategies.

Given stock returns $r_1, \dots, r_N$, this article's main focus is on the non-parametric estimation of the $\mathbb{P}$-a.s. superhedging price of an exotic option g. As a natural first step, we examine the simple plugin approach

$$\inf\{x \in \mathbb{R} \,|\exists H \in \mathbb{R}^d \text{ s.t. } x + H(r-1) \geq g(r) \ \forall r \in \{r_1, \dots, r_N\}\}$$

and establish its asymptotic consistency, but show that it lacks robustness in the sense of Tukey-Hampel-Huber. In consequence we propose several novel estimators, which overcome the shortcomings of the plugin approach. These new estimators are based on the dual formulation of the superhedging price

$$\inf\{x \in \mathbb{R} \,|\exists H \in \mathbb{R}^d \text{ s.t. } x + H(r-1) \geq g(r) \ P\text{-a.s}\} = \sup_{Q \sim P, \mathbb{E}_Q[r]=1} \mathbb{E}_Q[g]. \tag{1}$$

In order to achieve statistical robustness and to obtain better control over the point estimates it is necessary to consider a larger class of martingale measures. A natural choice is

$$\pi_{\mathcal{Q}_N}(g) = \sup_{Q \in \mathcal{Q}_N} \mathbb{E}_Q[g],$$

where $\mathcal{Q}_N$ is a subset of all martingale measures $\mathcal{M}$. The plugin estimator corresponds to taking $\mathcal{Q}_N = \{Q \in \mathcal{M} : Q \sim \hat{P}_N\}$ and we suggest to replace it with

$$\mathcal{Q}_N = \{\mathbb{Q} \in \mathcal{M} : \exists \tilde{P} \in B_N(\hat{P}_N) \text{ s.t. } \mathbb{Q} \sim \tilde{P}\},$$

where $B_N(\hat{P}_N)$ is some "ball" in the space of probability measures around the empirical measure $\hat{P}_N$. We show that this can lead to a consistent estimator if a sufficiently strong metric is used, e.g., the Wasserstein infinity metric $\mathcal{W}^\infty$. In general however such $\mathcal{Q}_N$ is too large. Instead, the main insight is to consider a tradeoff between the size of the balls and the behaviour of martingale densities:

$$\hat{\mathcal{Q}}_N := \{\mathbb{Q} \in \mathcal{M} \mid \|dQ/d\tilde{P}\|_\infty \leq k_N \text{ for some } \tilde{P} \in B^p_{\varepsilon_N}(\hat{P}_N)\},$$

where $B^p_{\varepsilon_N}(\hat{P}_N)$ denotes the p-Wasserstein ball of radius ε_N around $\hat{P}_N$ and $\varepsilon_N \to 0$ as well as $k_N \to \infty$. With a suitable choice of ε_N and k_N inspired by confidence

Date: December 31, 2021.

bounds in Wasserstein distance between the true and empirical measure of Fournier and Guillin [2015], we establish consistency of $\pi_{\hat{\mathcal{Q}}_N}(g)$ for a regular g, and also robustness in a Hausdorff-Wasserstein metric. This furthermore allows to study the cases when the estimator naturally extends to the setting of superhedging under model uncertainty about $\mathbb{P}$. We also extend the analysis to the case when risk is assessed not through superhedging capital but rather via a generic risk measure ρ admitting a Kusuoka representation, see Kusuoka [2001].

References

N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probab. Theory Relat. Fields*, 162(3-4):707–738, 2015.

S. Kusuoka. On law invariant coherent risk measures. In *Adv. Math. Econ.*, pages 83–95. Springer, 2001.

Johannes Obloj
Mathematical Institute and St. John's College
University of Oxford
AWB, ROQ, Woodstock Rd, Oxford, OX2 6GG, UK
E-mail address: jan.obloj@maths.ox.ac.uk

Johannes Wiesel
Columbia University, Department of Statistics
1255 Amsterdam Avenue
New York, NY 10027, USA
E-mail address: johannes.wiesel@columbia.edu

Special Sessions in memory of Prof. Mark H.A. Davis

Robust portfolio choice with sticky wages

Sara Biagini

Abstract: We present a robust version of the life-cycle optimal portfolio choice problem in the presence of labor income, as introduced in Biffis, Gozzi and Prosdocimi and Dybvig and Liu. In particular, in Biffis, Gozzi and Prosdocimi the influence of past wages on the future ones is modelled linearly in the evolution equation of labor income, through a given weight function. The optimisation relies on the resolution of an infinite dimensional HJB equation. We improve the state of art in three ways. First, we allow the weight to be a Radon measure. This accommodates for more realistic weighting of the sticky wages, like, e.g., on a discrete temporal grid according to some periodic income. Second, there is a general correlation structure between labor income and stocks market. This naturally affects the optimal hedging demand, which may increase or decrease according to the correlation sign. Third, we allow the weight to change with time, possibly lacking perfect identification. The uncertainty is specified by a given set of Radon measures K, in which the weight process takes values. This renders the inevitable uncertainty on how the past affects the future, and includes the standard case of error bounds on a specific estimate for the weight. Under uncertainty averse preferences, the decision maker takes a maxmin approach to the problem. Our analysis confirms the intuition: in the infinite dimensional setting, the optimal policy remains the best investment strategy under the worst case weight.

Robusts model in band ambiguity

Corina Birghila
University of Waterloo (Canada)

A major topic in finance and actuarial science is the identification, among all admissible models, of the most influential model for risk quantification. The set of these models is often called the ambiguity set, and it captures uncertainty concerning the data. In this regard we represent the set using the likelihood ratio with respect to a reference model $\hat{P}$, and use this ratio to construct a band ambiguity set, introduced in Shapiro and Ahmed [2004]. The band set is parametrized by two constants that control the size of the set, or equivalently, the amount of uncertainty. We propose to find the optimal likelihood ratio that minimizes the f-divergence in the sense of Csiszar [1967], such that the general reward-risk ratio of a random return is exceeding a prescribed acceptance level β. The solution of the robust optimization problem is derived in both a continuous and a discrete state space. In the latter case, due to the connection between any f-divergence and the Kullback-Leibler divergence (KL), we derive bounds for the optimal value in terms of the KL divergence. Via simulation experiments, we analyse the sensitivity of the optimal solution to the change in the band ambiguity parameters, as well as the change in acceptance level.

Key words: robust model, reward-risk ratio, f-divergence.

References

Imre Csiszar. Information-type measures of difference of probability distributions and indirect observation. *studia scientiarum Mathematicarum Hungarica*, 2:229–318, 1967.

Alexander Shapiro and Shabbir Ahmed. On a class of minimax stochastic programs. *SIAM Journal on Optimization*, 14(4):1237–1249, 2004.

Special Sessions in memory of Prof. Mark H.A. Davis

Rogue Traders

Paolo Guasoni

Abstract: Investing on behalf of a firm, a trader can feign personal skill by committing fraud that with high probability remains undetected and generates small gains, but that with low probability bankrupts the firm, offsetting ostensible gains. Honesty requires enough skin in the game: if two traders with isoelastic preferences operate in continuous-time and one of them is honest, the other is honest as long as the respective fraction of capital is above an endogenous fraud threshold that depends on the trader’ s preferences and skill. If both traders can cheat, they reach a Nash equilibrium in which the fraud threshold of each of them is lower than if the other one were honest. More skill, higher risk aversion, longer horizons, and greater volatility all lead to honesty on a wider range of capital allocations between the traders.

ROUGH MULTIFACTOR VOLATILITY MODELS FOR SPX AND VIX

ANTOINE JACQUIER, AITOR MUGURUZA, AND ALEXANDRE PANNIER (SPEAKER)

Abstract. We provide explicit small-time formulae for the at-the-money implied volatility of a large class of underlying assets, as well as its first and second derivatives with respect to the strike. Chiefly interested in VIX options, we draw insights from several past results arguing in favour of multi-factor stochastic (rough) volatility models, which are conducive to capture the smile behaviour. Therefore our asymptotic analysis includes underlyings measurable with respect to a multidimensional Brownian motion. Our proofs build on the framework of Alòs, Garcá-Lorite and Muguruza: they are based on Malliavin calculus techniques—in particular an anticipative Itô's formula, and our formulae are expressed in terms of the Malliavin derivatives of the underlying.

We apply our general results to VIX options in a two-factor rough Bergomi model, with correlated Brownian motions, and derive closed-form expressions for the short-time at-the-money implied volatility level, skew, and curvature. These formulae only depend on the parameters of the model. Hence, they give analytical insights on the interplay between the different parameters, allowing to understand the behaviour of at-the-money VIX options prior to performing numerical tests. Nevertheless, to support our theoretical results, we also provide numerical examples and calibration with real data.

Finally, we present a similar analysis of the smile of the S&P index and, thanks to the additional degrees of freedom offered by the correlation parameters, we show how the two-factor rough Bergomi model succeeds in jointly calibrating the SPX and the VIX. Despite the extensive academic literature, options on VIX and S&P index still display different volatility surfaces, betraying the lack of a proper modelling framework. This issue is well-known as a the SPX-VIX joint calibration problem and has motivated a number of creative modelling innovations in the past 15 years. Recently, the new paradigm of rough volatility has shown new promise in resolving this challenging problem.

Rough volatility: fact or artefact?
Extended abstract

Rama CONT Purba DAS

Date: March 23, 2022

1 Abstract

We investigate the statistical evidence for the use of 'rough' fractional processes with Hurst exponent $H < 0.5$ for the modeling of volatility of financial assets, using a model-free approach. We introduce a non-parametric method for estimating the roughness of a function based on discrete sample, using the concept of normalized p-th variation along a sequence of partitions, and discuss the consistency of the estimator in a pathwise setting. We investigate the finite sample performance of our estimator for measuring the roughness of sample paths of stochastic processes using detailed numerical experiments based on sample paths of fractional Brownian motion and other fractional processes. We then apply this method to estimate the roughness of realized volatility signals based on high-frequency observations. Detailed numerical experiments based on stochastic volatility models show that, even when the instantaneous volatility has diffusive dynamics with the same roughness as Brownian motion, the realized volatility exhibits rough behaviour corresponding to a Hurst exponent significantly smaller than 0.5. Comparison of roughness estimates for realized and instantaneous volatility in fractional volatility models with different values of Hurst exponent shows that, irrespective of the roughness of the spot volatility process, realized volatility always exhibits 'rough' behaviour with an apparent Hurst index $\hat{H} < 0.5$. These results suggest that the origin of the roughness observed in realized volatility time-series lies in the microstructure noise rather than the volatility process itself.

2 Contribution

In the paper, we address these questions in detail by re-examining the statistical evidence from high-frequency financial data, in an attempt to clarify whether the assertion that 'volatility is rough' (i.e. rougher than typical paths of Brownian motion) is supported by empirical evidence. We investigate the statistical evidence for the use of 'rough' fractional processes with Hurst exponent $H < 0.5$ for the modelling of volatility of financial assets, using a non-parametric, model-free approach.

We introduce a non-parametric method for estimating the roughness of a function/path based on a (high-frequency) discrete sample, using the concept of normalized p-th variation along a sequence of partitions, and discuss the consistency of our estimator in a pathwise setting. We investigate the finite sample performance of our estimator for measuring the roughness of sample paths of stochastic processes using detailed numerical experiments based on sample paths of fractional Brownian motion and other fractional processes. We then apply this method to estimate the roughness of realized volatility signals based on high-frequency observations. Through a detailed numerical experiment based on a stochastic volatility model, we show that even when the instantaneous (spot) volatility has diffusive dynamics with the same roughness as Brownian motion, the realized volatility exhibits rough behaviour corresponding to a Hurst exponent significantly smaller than 0.5. Similar behavior is observed in financial data as well, which suggests that the origin of the roughness observed in realized volatility time-series lie in the estimation error rather than the volatility process itself. Comparison of roughness estimates for realized and instantaneous volatility in fractional volatility models for different values of Hurst parameter H shows that whatever the value of H for the (spot) volatility process, realized volatility always exhibits 'rough' behaviour.

Our results are broadly consistent with the points raised by Rogers [3], but we pinpoint more precisely the origin of the apparent 'rough' behaviour of volatility as being the microstructure noise inherent in the estimation of realized volatility. In particular, our results question whether the empirical evidence presented from high-frequency volatility estimates supports the 'rough volatility' hypothesis.

3 Rough volatility ... or microstructure noise?

Given the large literature on 'rough volatility' in quantitative finance, it is somewhat surprising that the initial claim [2] that one needs to model the spot volatility process using a 'rough' fractional noise with Hurst exponent $H < 1/2$ has not been examined more closely, especially given that several follow-up studies point to the fact the observations in [2] may well be compatible with a Brownian diffusion model for volatility [1, 3].

Our detailed examples illustrate that, for stochastic-volatility diffusion models driven by Brownian motion, the realized volatility has a roughness index ≈ 0.3 so seems to exhibit a significantly 'rougher' behaviour than the instantaneous volatility, both in terms of normalized p-th variation statistics described in the paper and also in terms of the linear- regression method used by Gatheral et al. [2]. These results suggest that one cannot hastily conclude that the roughness observed in realized volatility is an indicator of similar behaviour in spot volatility, as implicitly assumed in the 'rough volatility' literature; the observations in high-frequency financial data are in fact compatible with a stochastic volatility model drive by Brownian motion and the origin of this apparent roughness may very well lie in microstructure noise rather than the noise process driving spot volatility.

A further example in the paper shows that, the rough behaviour of realized volatility does not lead us to reject the hypothesis that the underlying spot volatility may be modelled with a Brownian diffusion model. This observation, together with "Occam's razor", pleads for the use of diffusion-based stochastic volatility models which seem compatible with the empirical evidence but are far more tractable.

We are thus drawn to concur with Rogers [3] that "the notion that volatility is rough, that is, governed by a fractional Brownian motion (with $H < 1/2$), is not an incontrovertible established fact; simpler models explain the observations just as well."

References

[1] M. Fukasawa, T. Takabatake, and R. Westphal, *Consistent estimation for fractional stochastic volatility model under high-frequency asymptotics*, Mathematical Finance, (To appear).

[2] J. Gatheral, T. Jaisson, and M. Rosenbaum, *Volatility is rough*, Quantitative Finance, 18 (2018), pp. 933–949.

[3] L. Rogers, *Things we think we know*, (2019). `https://www.skokholm.co.uk/wp-content/uploads/2019/11/TWTWKpaper.pdf`.

Sensitivity to large losses and ρ-arbitrage for convex risk measures

Martin Herdegen Nazem Khan

1 Extended Abstract

The aim of a risk measure is to quantify the risk of a financial position by a single number. This number can be interpreted in different ways in different contexts: In banking, it represents the capital requirement to regulate a risk; in insurance, it calculates the premium for an insurance contract; and in economics it ranks the preference of a risk for a market participant.

In this work, we interpret the risk measure ρ as a regulatory constraint imposed by the regulator on a financial agent seeking to optimise a portfolio in a one-period model. Because our focus is on the effectiveness of the risk constraint, we ignore any idiosyncratic risk aversion of the agent. Denoting by X_π the excess return of a portfolio $\pi \in \mathbb{R}^d$, we consider the following two mean-ρ portfolio selection problems:

(1) Given a minimal desired expected excess return $\nu^* \geq 0$, minimise the risk $\rho(X_\pi)$ among all portfolios $\pi \in \mathbb{R}^d$ that satisfy $\mathbb{E}[X_\pi] \geq \nu^*$;

(2) Given a maximal risk threshold $\rho^* \geq 0$, maximise the return $\mathbb{E}[X_\pi]$ among all portfolios $\pi \in \mathbb{R}^d$ that satisfy $\rho(X_\pi) \leq \rho^*$.

There is a large literature on mean-ρ portfolio selection, yet it is still open for debate what properties the risk measure ρ should possess. As of today, the most popular risk measures are monotone, cash-invariant and positively homogeneous with Value at Risk (VaR) and Expected Shortfall (ES) being the most famous examples. Moreover, ES is the current industry standard. However, mean-ρ portfolio selection for positively homogeneous (monotone and cash-invariant) risk measures may be ill-posed in the sense that (1) or (2) have no solutions. This phenomena is referred to as ρ-arbitrage and has been studied extensively in Herdegen and Khan [3] from a theoretical perspective, and the practical relevance has been discussed by Armstrong and Brigo in [1]. In particular, [3, Theorem 3.23] implies that regulators cannot exclude (strong) ρ-arbitrage a priori when imposing a positively homogeneous cash-invariant risk measure – unless ρ is as conservative as the worst-case risk measure. Since a worst-case approach to risk is infeasible in practise, this indicates that one should move beyond the class of positively homogeneous risk measures for effective risk constraints in the context of portfolio selection.

Our main objective is to study mean-ρ portfolio selection where ρ is monotone, normalised and star-shaped, i.e., $\rho(\lambda X) \geq \lambda\rho(X)$ for $\lambda \geq 1$. For some of our results, we assume in addition that ρ is convex, cash-invariant and satisfies the Fatou property. Assuming that ρ lives on some Riesz space $L^1 \subset L \subset L^\infty$ and is $(-\infty, \infty]$-valued, we seek to answer the following questions:

- **Absence of ρ-arbitrage.** Given a one-period market, what are necessary and sufficient conditions to ensure that the market does not admit ρ-arbitrage?
- **Suitable for portfolio selection.** When is ρ suitable for portfolio selection, i.e., for every one-period arbitrage-free market, and for every $\nu^* \geq 0$ and $\rho^* \geq 0$, the mean-ρ problems (1) and (2) admit at least one solution with finite risk?

The crucial ingredient for the characterisation of ρ-arbitrage is the new axiom of strong sensitivity to large losses, that is, for all X with $\mathbb{P}[X < 0] > 0$, there exists $\lambda > 0$ such that $\rho(\lambda X) > 0$. Its economic meaning is simple and intuitive: Apart from the riskless portfolio, any portfolio has a positive risk if it is scaled by a sufficiently large amount.

We then provide a dual characterisation for the absence of ρ-arbitrage. To this end, a key methodological tool is to consider ρ^{∞}, the smallest positively homogeneous risk functional that dominates ρ. A key observation is that ρ satisfies strong sensitivity to large losses if and only if ρ^{∞} does. If ρ has a dual representation, then so does ρ^{∞}, and we can lift the results from [3] on the dual characterisation of ρ-arbitrage for coherent risk measures to the case of convex risk measures.

From a regulator's perspective, the notion of being suitable for portfolio selection is highly desirable. Here again, the key methodological tool is to consider ρ^{∞} and to use that ρ satisfies strong sensitivity to large losses if and only if ρ^{∞} does. Assuming that ρ is a convex risk measure that satisfies the Fatou property, we show that ρ is suitable for portfolio selection if and only if it is real-valued and ρ^{∞} is the worst-case risk measure.

While the above fully answers the second question from a theoretical perspective, it leaves open the question how large the subclass of risk measures suitable for portfolio selection is or how concrete examples look like. Perhaps surprisingly, we can describe all such risk measures in a dual way if L is an Orlicz heart, which includes all L^p-spaces for $p \in [1, \infty)$. In particular, Optimised Certainty Equivalent (OCE) risk measures are suitable for portfolio selection if and only if the associated loss function satisfies $\lim_{x\to-\infty} l(x)/x = 0$ and $\lim_{x\to\infty} l(x)/x = \infty$.

Of course, of special interest is the case $L = L^1$. We first show that an important subset of risk measures that are suitable for portfolio selection on L^1 are given by a subclass of so-called Adjusted Expected Shortfall risk measures, recently studied by Burzoni et al. [2]. In particular, we introduce the new risk measure of Loss Sensitive Expected Shortfall, which is not more complicated to compute than Expected Shortfall (ES), but unlike ES, is suitable for portfolio selection on L^1. We believe that this new risk measure could become of great relevance to the regulator because it keeps many attractive features of ES, while being strongly sensitive to large losses – which ES is not.

References

[1] J. Armstrong and D. Brigo, *Risk managing tail-risk seekers: Var and expected shortfall vs s-shaped utility*, J. Bank. Finance **101** (2019), 122–135.

[2] M. Burzoni, C. Munari, and R. Wang, *Adjusted expected shortfall*, J. Bank. Finance (2021), forthcoming.

[3] M. Herdegen and N. Khan, *Mean-ρ portfolio selection and ρ-arbitrage for coherent risk measures*, Math. Finance (2021), forthcoming.

Signature Method for Option Pricing Problem With Path-Dependent Features

Abstract, Qi Feng

The classical models for asset processes in math finance are SDEs driven by Brownian motion of the following type $X_t = x + \int_0^t b(s, X_s)ds + \int_0^t \sigma(s, X_s)dB_s$. Then $u(t, X_t) = \mathbb{E}[g(X_T)|\mathcal{F}_t^X]$ is a deterministic function of X_t and $u(t, x)$ solves a parabolic PDE. In this talk, I will talk about two types of path-dependent option pricing problems. In the first scenario, the option function depends on the whole path of the Markov process X, the option pricing problem is indeed to compute $\mathbb{E}^{\mathbb{P}}[g(X_{[0,T]})]$; In the second scenario, we consider the option pricing problem for rough volatility models, where the volatility of the asset process follows a Volterra SDE. The option function depends on the terminal value of a non-Markovian asset process. In both cases, the function $u(t, \cdot)$ solves the so-called Path-Dependent PDEs. Due to the path-dependent feature, the standard numerical algorithms are not efficient for both cases. We introduce the "Signature" idea from the Rough Path theory to our numerical algorithms to improve the efficiency. Our first algorithm is based on "deep signature" and deep learning methods for BSDEs; Our second algorithm is based on cubature formula for "Volterra signature", which is motived from the "cubature formula" for the signature of Brownian motion. In the end, I will present some future works based on the Volterra signature. The talk is based on two joint works with Man Luo, Zhaoyu Zhang, and Jianfeng Zhang.

1

Signature Methods in Stochastic Portfolio Theory

Extended Abstract

Christa Cuchiero, Janka Möller

March 15, 2022

Abstract

We introduce a novel class of functionally generated portfolios, which we call signature portfolios. In stochastic portfolio theory, one usually considers a portfolio generating function $\mathbf{S}(\mu)$ and constructs portfolios via $S^i(\mu) := D_i(\log(\mathbf{S}(\mu)))$. We, however, do not require the $S^i(\mu)$ to be the log-gradient of a function and therefore, we refer to the $S^i(\mu)$ directly as portfolio generating functions. Signature portfolios are *universal* in the sense that every continuous[1] (path-dependent) portfolio generating function $S^i(\mu)$ can be uniformly approximated by a signature portfolio generating function. We study the optimization task of maximizing the log-relative-value of a signature portfolio and find that this leads to a *convex quadratic optimization problem.* We use these theoretical results to implement the learning task of maximizing the expected log-relative-value, in two numerical examples. Indeed, the trained signature portfolios are remarkably close to the theoretical growth-optimal portfolios on the test-data.

1 Introduction

The concept of the signature of a path is particularly important in the field of rough path theory [6]. However, the signature is also interesting from a machine learning perspective, due to some of its remarkable properties. Two of which are:

- A time-extended path is uniquely determined by its signature.
- Linear functions on the signature uniformly approximate continuous (with respect to certain variation metrics) path functionals on compact sets.

Therefore, signature methods have recently entered the field of mathematical finance, see for example [2, 5, 1]. We use signature methods in the context of stochastic portfolio theory (SPT) by introducing *signature portfolios.* Due to the aforementioned properties of the signature, the signature portfolios are more general than the previously considered functionally generated portfolios in [4, 3].

2 The Signature

We briefly introduce the signature of an $\mathbb{R}^d$-valued continuous semimartingale X. The signature of X, denoted by $\mathbb{X}$ is an element of the extended tensor algebra $T((\mathbb{R}^d))$ defined as

$$T((\mathbb{R}^d)) := \{(a_0, a_1, ...)|a_n \in (\mathbb{R}^d)^{\otimes n}, n \geq 0\}$$

with $(\mathbb{R}^d)^{\otimes 0} := \mathbb{R}$. Let $I = \{i_1, ..., i_m\}$ be a multi-index and denote its length by $|I| = m$. For $e_I := e_{i_1} \otimes ... \otimes e_{i_m}$ a basis element of $(\mathbb{R}^d)^{\otimes m}$, the corresponding element of the signature of X at time t is given by

$$\langle e_I, \mathbb{X}\rangle_t = \int_{0\leq t_1 \leq ... \leq t_m \leq t} \circ dX_{t_1}^{i_1} \cdots \circ dX_{t_m}^{i_m},$$

where $\circ$ means that the integrals are to be understood in the Stratonovic sense. The signature *truncated* at order n is defined as

$$\mathbb{X}_t^{(n)} = 1 + \sum_{1\leq |I| \leq n} \langle e_I, \mathbb{X}\rangle_t e_I.$$

We use the convention that $|I| = 0$ means $I = \emptyset$. Hence, in this notation $\langle e_\emptyset, \mathbb{X}\rangle_t = 1$.

3 Signature Portfolios

We assume the market to consist of d stocks, with market weights given by a vector $\mu_t = (\mu_t^1, ..., \mu_t^d)$ of d continuous semimartingales.

We introduce portfolios generated by linear functions on the (truncated) signature of the time-extended market weights $\hat{\mu}_t = (t, \mu_t)$. We consider two types of such portfolios, denoted by π and η,

$$\pi_t^i = \mu_t^i (F^i(\hat{\mu}_{[0,t]}) + 1 - \sum_{j=1}^d \mu_t^j F^j(\hat{\mu}_{[0,t]})) \tag{1}$$

$$\eta_t^i = F^i(\hat{\mu}_{[0,t]}) + \mu_t^i (1 - \sum_{j=1}^d F^j(\hat{\mu}_{[0,t]})) \tag{2}$$

for portfolio generating functions

$$F^i(\hat{\mu}_{[0,t]}) = \sum_{0\leq |I| \leq n} l_I^{(i)} \langle e_I, \hat{\mathbb{\mu}}\rangle_t, \tag{3}$$

with (optimization) parameters $l_I^{(i)} \in \mathbb{R}$. Here, we use the notation $F^i(\hat{\mu}_{[0,t]})$ to make explicit that the functions F^i can depend on the entire path of $\hat{\mu}$. We call such portfolios *signature portfolios.* When we refer to *portfolio generating functions*, we mean the functions F^i themselves and do not require them to be the log-gradient of a portfolio generating function in the classical sense of SPT.

[1] with respect to certain variation metrics

3.1 Optimization Task

For a portfolio[2] π, we denote its value at time t by V_t^π, assume $V_0^\pi = 1$ and call the process $\log \frac{V_t^\pi}{V_t^\mu}$ the *log-relative-value process* of a portfolio. For a signature portfolio this process is given by

$$\log \frac{V_t^\pi}{V_t^\mu} = \sum_{i=1}^{d} \sum_{0 \le |I| \le n} l_I^{(i)} \int_0^t \langle e_I, \hat{\mathbb{\mu}} \rangle_s d\mu_s^i \tag{4}$$
$$- \frac{1}{2} \sum_{i,j=1}^{d} \sum_{\substack{0 \le |I| \le n \\ 0 \le |J| \le n}} l_I^{(i)} l_J^{(j)} \int_0^t \langle e_I ⧢ e_J, \hat{\mathbb{\mu}} \rangle_s d[\mu_i, \mu_j]_s,$$

where $[\mu_i, \mu_j]$ denotes the covariation process of μ^i, μ^j and $⧢$ is the shuffle-product. Clearly, (4) is quadratic in the (optimization) parameters $\{l_I^{(i)}\}_{1 \le i \le d, 0 \le |I| \le n}$.

Therefore, maximizing the log-relative-value of a signature portfolio is a quadratic optimization task and can be brought to the form

$$\min_{\mathbf{l}} \frac{1}{2} \mathbf{l}^T Q \mathbf{l} - \mathbf{c}^T \mathbf{l}, \tag{5}$$

with $\mathbf{l}$ being a vector of the optimization parameters $\{l_I^{(i)}\}$. We can show that (5) is a *convex quadratic optimization problem*, which makes it *computationally tractable.*

Moreover, we can consider several modifications of the problem, which again lead to convex quadratic optimization tasks. One of which is the *Monte-Carlo type* maximization of the *expected* log-relative-value over M samples, given by

$$\min_{\mathbf{l}} \frac{1}{2} \mathbf{l}^T \left(\sum_{m=1}^{M} Q_m \right) \mathbf{l} - \left(\sum_{m=1}^{M} \mathbf{c_m} \right)^T \mathbf{l}. \tag{6}$$

Remark. Since our only assumption on the market weights is that they are continuous semimartingales, all of the above also holds for the corresponding *ranked* market weights.

3.2 Numerical Examples

We consider two different examples of a Monte-Carlo type optimization of the expected log-relative-value (for M=99'968):

(i) Signature portfolio η (see (2)) for $n = 2$ in the market model:
$$\frac{dS_t}{S_t} = \mathbf{a}_t dt + \Sigma \, dW_t \tag{7}$$
where $(\mathbf{a}_t)_i = \sum_{0 \le |I| \le 2} \alpha_I^{(i)} \langle e_I, \hat{\mathbb{\mu}} \rangle_t$, $\alpha_I^{(i)} \in \mathbb{R}$, W_t is a three-dimensional Brownian motion and Σ is a constant 3×3 matrix.

(ii) Signature portfolio π (see (1)) for $n = 3$ in a correlated Black-Scholes model.

We present the theoretical growth-optimal weights and the learnt signature weights in Fig. 1. The learnt signature weights are remarkably close to the theoretical weights in both examples. The mean log-relative value, evaluated on 100'000 test samples, of the signature and the growth-optimal portfolio differ by only 8-20bp.

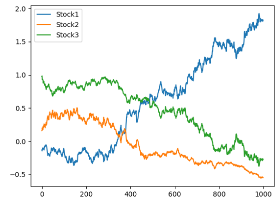

(i): Growth-Optimal Weights

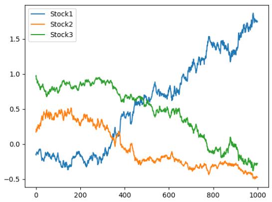

(i): Signature Weights

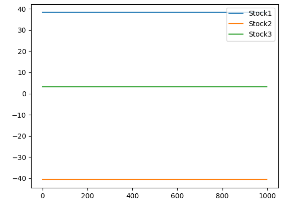

(ii): Growth-Optimal Weights

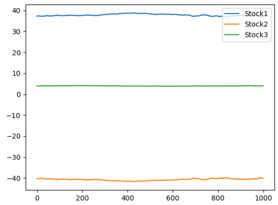

(ii): Signature Weights

Figure 1: The theoretical growth-optimal weights (left) and the signature portfolio's weights (right) for markets (i) (top) and (ii) (bottom) respectively, evaluated at one test sample.

Note that although the growth-optimal weights in the Black-Scholes market are constant, the signature portfolio generating function $F^i(\mu_{[0,t]})$ had to learn to approximate $\frac{\beta_i}{\mu_i}$ for some $\beta_i \in \mathbb{R}$.

We conclude that signature portfolios are indeed versatile enough to learn optimal strategies, while being computationally efficient. We want to highlight that signature portfolios are purely trained on the path trajectories of the market weights.

References

[1] I. P. Arribas, C. Salvi, and L. Szpruch. *Sig-SDEs Model for Quantitative Finance.* In: Proceedings of the First ACM International Conference on AI in Finance. 2020.

[2] H. Bühler et al. *A Data-Driven Market Simulator for Small Data Environments.* Available at SSRN 3632431. 2020.

[3] S. Campbell and T.-K. L. Wong. *Functional portfolio optimization in stochastic portfolio theory.* 2021. arXiv: `2103.10925 [q-fin.PM]`.

[4] R. Fernholz. *Stochastic Portfolio Theory.* Springer, 2002.

[5] T. Lyons, S. Nejad, and I. P. Arribas. *Non-parametric Pricing and Hedging of Exotic Derivatives.* Applied Mathematical Finance. 2019.

[6] T. J. Lyons. *Differential equations driven by rough signals.* Revista Matemática Iberoamericana. 1998.

[2]For the case of the portfolios η, the derivations are similar.

Signature-based models: theory and calibration

Christa Cuchiero* Guido Gazzani† Sara Svaluto-Ferro‡

February 24, 2022

Abstract

Signature methods represent a non-parametric way for extracting characteristic features from time series data which is essential in machine learning tasks. This explains why these techniques become more and more popular in econometrics and mathematical finance. Indeed, signature based approaches allow for data-driven and thus more robust model selection mechanisms, while first principles like no arbitrage can still be easily guaranteed. Here we focus on signature-based models whose characteristics are linear functions of a primary underlying process, which can range from a (market-inferred) Brownian motion to a general multidimensional tractable stochastic process. The framework is universal in the sense that any classical model can be approximated arbitrarily well and that the model characteristics can be learned from all sources of available data by simple methods. In view of option pricing and calibration, key quantities that need to be computed in these models are the expected value or Fourier Laplace transform of the signature of the primary underlying process. Surprisingly this can be achieved via techniques from affine and polynomial processes. These formulas can then be used in the calibration procedure to option prices, while calibration to time series data just reduces to a simple regression.

Keywords: signature methods, calibration of financial models, affine and polynomial processes
MSC (2010) Classification: 91B70, 62P05, 65C20.

1 Introduction and the model

In the past few years data driven models have successfully entered the area of stochastic modeling and mathematical finance. The paradigm of calibrating a few well interpretable parameters has changed to learning the model's characteristics as a whole, thereby exploiting all available sources of data. Thus highly parametric and overparametrized models methods have gained more and more importance. We consider here so-called signature based models, meaning that the model itself or its characteristics are parametrized as linear functions of the signature of a primary underlying process. This underlying process can either be a classical driving signal, e.g. a Brownian motion or a Lévy process, but also a more general tractable stochastic model describing well observable quantities. The notion of the signature of a path goes back to Chen [1957] and plays a particular important role in the context of rough path theory initiated by Lyons [1998]. In particular, the Stone-Weierstrass theorem yields a Universal Approximation Theorem, showing that the signature serves as linear regression basis for continuous path functionals, since

- it is point-separating, as the signature of a path extended by time uniquely determines the path.
- linear functions on the signature form an algebra that contains 1. More precisely every polynomial on signature may be realized as a linear function via the so-called shuffle product $\sqcup\!\sqcup$.

These properties motivate why these techniques become more and more popular in econometrics and mathematical finance, see Buehler et al. [2020], Perez Arribas et al. [2020], Bayer et al. [2021] and references therein.
We consider here signature-based methods with the goal to provide a data-driven, universal, tractable and easy to calibrate model for a set of traded assets $S = (S^1, \dots, S^m)$. To achieve this the main ingredient is a primary underlying process

*Vienna University, Department of Statistics and Operations Research, Data Science Uni Vienna, Kolingasse 14-16 1, A-1090 Wien, Austria, christa.cuchiero@univie.ac.at
†Vienna University, Department of Statistics and Operations Research, Kolingasse 14-16 1, A-1090 Wien, Austria, guido.gazzani@univie.ac.at
‡University of Verona, Department of Economics, Via Cantarane 24, 37129 Verona, Italy, sara.svalutoferro@univr.it
The authors gratefully acknowledge financial support through grant Y 1235 of the FWF START-program.

$(\widehat{X}_t)_{t\geq 0} = (t, X_t^1, \dots, X_t^d)_{t\geq 0}$ with $d \leq m$, where X is a continuous semimartingale . We suppose here that time-series data of $\widehat{X}$ is available and that its signature denoted by $\widehat{\mathbb{X}}$ serves a linear regression basis for S.
We shall now introduce the most essential concepts in order to rigorously define signature in the current context.
For any $I = (i_1, \dots, i_n)$ multi-index with entries in $\{1, \dots, d\}$ denote with $e_I = e_{i_1} \otimes \cdots \otimes e_{i_n}$ the basis element of $(\mathbb{R}^d)^{\otimes n}$. Additionally, given $\mathbf{a} \in \bigoplus_{n\geq 0} (\mathbb{R}^d)^{\otimes n}$ with $(\mathbb{R}^d)^{\otimes 0} := \mathbb{R}$, we write $\langle e_I, \mathbf{a}\rangle$ to extract the I^{th} component from a_n. Moreover, the coordinate signature indexed by a multi-index $I = (i_1, \dots, i_n)$ of an $\mathbb{R}^d$-valued semimartingale $\widehat{X}$ is defined via iterated Stratonovich integrals (denoted by $\circ$):

$$\langle e_I, \widehat{\mathbb{X}}_T\rangle := \int_{0<t_1<\cdots<t_n<T} \circ\, \mathrm{d}\widehat{X}_{t_1}^{i_1} \cdots \circ\, \mathrm{d}\widehat{X}_{t_n}^{i_n}.$$

Hence $\widehat{\mathbb{X}}_T = 1 + \sum_{n=1}^{\infty} \sum_{|I|=n} \langle e_I, \widehat{\mathbb{X}}_T\rangle \in \bigoplus_{n\geq 0} (\mathbb{R}^d)^{\otimes n}$.
Let us now describe the precise modeling framework. The traded assets $S = (S^1, \dots, S^m)$ are modeled via $S(\ell)$: for $j = 1, \dots, m$ and $t > 0$

$$S_n^j(\ell^j)_t := \ell_0^j + \sum_{0<|I|\leq n} \ell_I^j \langle e_I, \widehat{\mathbb{X}}_t\rangle, \tag{Sig-model}$$

where $\widehat{\mathbb{X}}$ is the signature of $\widehat{X}$, $n \in \mathbb{N}$ is the degree of truncation and $\ell_0^j, \ell_I^j \in \mathbb{R}$ are coefficients to be found from data.

For notational simplicity we shall in the sequel set $m = 1$. The attractiveness of this model class arises from the following features:
No arbitrage: The model can also be expressed in terms of stochastic integrals with respect to local martingales, from which conditions for no-arbitrage can be easily deduced. This translates into a change of basis $\{\tilde{e}_I : |I| \leq n\}$.
Calibration to time-series data: This task reduces to a simple linear regression. Indeed we aim to match N market prices $(S_{t_1}^M, \dots, S_{t_N}^M)$ and we are given $(\widehat{X}_{t_1}, \dots, \widehat{X}_{t_N})$, e.g. market-inferred Brownian motions, thus we aim to solve

$$\text{argmin}_\ell \sum_{i=1}^{N} \left(\ell_0 + \sum_{0<|I|\leq n} \ell_I \langle \tilde{e}_I, \widehat{\mathbb{X}}_{t_i}\rangle - S_{t_i}^M\right)^2.$$

Since the dimension of ℓ is typically high, introducing a regularization (Lasso, Ridge) is necessary. In our findings we show how we can approximate trajectories of classical stochastic volatility models.
Calibration to options: When calibrating to option data the goal is to match N option prices $(\pi_1, \dots, \pi^N)$ corresponding to European payoffs $F_i(S_{T_i})$, typically Call and Put options with different strikes. Given $M > 0$ samples of $\widehat{\mathbb{X}}$ at maturities, we aim to solve

$$\text{argmin}_\ell \sum_{i=1}^{N} \omega_i \left(\frac{1}{M} \sum_{j=1}^{M} F_i(S_n(\ell)_{T_i}) - \pi_i\right)^2,$$

where ω_i are e.g. normalized Vega-weights known to match implied volatility well. The advantages of the Sig-model is that all Monte-Carlo samples of $(\widehat{\mathbb{X}}_{T_i})_{i=1,\dots,N}$ can be easily precomputed and re-used so that, by linearity of the model, the calibration reduces to a simple optimization task without any Monte-Carlo simulation within an optimization step. This procedure is shown to fit perfectly the implied volatility smile of the S&P 500 on a given trading day.

References

C. Bayer, P. Hager, S. Riedel, and J. Schoenmakers. Optimal stopping with signatures. *ArXiv e-prints*, 2021. URL `https://arxiv.org/abs/2105.00778`.

H. Buehler, B. Horvath, T. Lyons, I. Perez Arribas, and B. Wood. A Data-Driven Market Simulator for Small Data Environments. *Available at SSRN 3632431*, 2020.

K.T. Chen. Integration of paths, geometric invariants and a generalized Baker-Hausdorff formula. *Annals of Mathematics*, page 163–178, 1957.

T. J Lyons. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310, 1998.

I. Perez Arribas, C. Salvi, and L. Szpruch. Sig-SDEs model for quantitative finance. *ArXiv e-prints*, 2020. URL `https://arxiv.org/abs/2006.00218`.

Signatures, Functional Itô Calculus and Claim Decomposition

European option payoffs can be generated by combinations of hockey stick payoffs or monomials. Interestingly, path dependent options can be generated by combinations of signatures, which are the building blocks of path dependence.

We focus on the case of one asset together with time, typically the evolution of the price x as a function of time t. The signature of a path for a given word with a letter in the alphabet {t,x} is an iterated integral with respect to the letters of the word and it plays the role of a monomial in a Taylor expansion.

We first review and establish properties of the signature in that context, and then study the Taylor expansion of a functional. It is a sum over the words of the functional derivative from the functional Itô calculus with respect to this word at the origin times the signature associated to the word. We explore the implications of this expansion in terms of pricing and hedging of exotic options. This approach provides a fast alternative to Deep Hedging.

We compare various decompositions/approximations of functionals such as Taylor expansion, Wiener chaos and Volterra series.

SIMPLIFIED STOCHASTIC CALCULUS WITH APPLICATIONS IN ECONOMICS AND FINANCE

ALEŠ ČERNÝ AND JOHANNES RUF

ABSTRACT. The paper introduces a simple way of recording and manipulating general stochastic processes without explicit reference to a probability measure. In the new calculus, operations traditionally presented in a measure-specific way are instead captured by tracing the behaviour of jumps (also when no jumps are physically present). The calculus is fail-safe in that, under minimal assumptions, all informal calculations yield mathematically well-defined stochastic processes. The calculus is also intuitive as it allows the user to pretend all jumps are of compound Poisson type. The new calculus is very effective when it comes to computing drifts and expected values that possibly involve a change of measure. Such drift calculations yield, for example, partial integro–differential equations, Hamilton–Jacobi–Bellman equations, Feynman–Kac formulae, or exponential moments needed in numerous applications. We provide several illustrations of the new technique, among them a novel result on the Margrabe option to exchange one defaultable asset for another.

Keywords. drift; Émery formula; Girsanov's theorem; simplified stochastic calculus

1. INTRODUCTION

Anyone who has attempted stochastic modelling with jumps will be aware of the sudden increase in mathematical complexity in models that are not of compound Poisson type. The difficulty is such that experienced researchers readily forgo generality in order to reduce the technical burden placed on their readers; see, for example, Feng and Linetsky (2008), Cai and Kou (2012), Hong and Jin (2018), and Aït-Sahalia and Matthys (2019).

In this paper we introduce an intuitive calculus that works for general processes but retains the simplicity of compound Poisson calculations. To achieve this, a change of paradigm is required. Classical Itô calculus is based on decomposing the increments of every process into signal (drift, expected change) and noise (Brownian motion, zero-mean shock). This is at once convenient and mathematically expedient. The convenience of knowing the drift is immediate. Many tasks where stochastic processes are concerned involve computation of the drift of some quantity. Hamilton–Jacobi–Bellman equations in optimal control, for example, express the fact that the optimal value function plus the integrated historical cost is a martingale and therefore has zero drift. Similarly, Feynman–Kac formulae reflect zero drift of an integral of costs discounted at a specified stochastic killing rate. Closer to home, the Black–Scholes partial differential equation can be obtained by setting the drift of the discounted option price process to zero under the risk-neutral measure.

The expediency of the signal–noise decomposition comes from the early construction of the Itô integral where the drift is integrated path-by-path but the Brownian motion integral is performed, loosely speaking, by summing up uncorrelated square-integrable random variables with zero mean. The paradigm shift is applied here: we separate how a process is recorded from

BAYES BUSINESS SCHOOL, CITY, UNIVERSITY OF LONDON, 106 BUNHILL ROW, LONDON, EC1Y 8TZ, UK.
DEPARTMENT OF MATHEMATICS, LSE, HOUGHTON STREET, LONDON, WC2A 2AE, UK.
E-mail addresses: ales.cerny.1@city.ac.uk, j.ruf@lse.ac.uk.
MSC [2010]. (Primary) 60H05, 60H10; 60G44, 60G48; (Secondary) 91B02, 91B25; 91G10.
Previous versions of this paper were circulated under the title "Finance without Brownian motions: An introduction to simplified stochastic calculus." We thank Jan Kallsen, Jan-Frederik Mai, Lola Martinez Miranda, Johannes Muhle-Karbe, and two anonymous referees for helpful comments.

the drift calculation. In other words, we do not carry the drift with us at all times but only evaluate it when the drift is really needed. This feels a little uncomfortable at first but there are ample rewards for the small intellectual effort required.

The consequences of the subtle change in perspective are far-reaching. By recording processes in a measure-invariant manner the technicalities of stochastic integration fall away, the importance of Brownian motions and Poisson processes recedes, and one begins to see deeper into the fundamental relationships among the modelled variables, which now take center stage. Measure change, too, becomes an easy application of the simplified calculus, for as long as the new measure is directly driven by the variables being studied, which is overwhelmingly the case in practice.

We have prepared the paper with two audiences in mind. First and foremost, the paper is intended for the research community whose members do not consider themselves experts in mathematics in general, or stochastic analysis in particular, but who nevertheless use stochastic calculus as a modelling tool. To this readership we want to demonstrate that the new calculus is easy to understand and apply in practice. Second, but no less important, we address colleagues specializing in stochastic analysis whom we wish to convince that all our arguments are mathematically rigorous.

The stated goal is not without its challenges. Where practical, we contrast the new approach with the more involved classical notation. In order to perform such comparison, one has to introduce some advanced concepts, such as the Poisson random measure, which are needed in classical stochastic calculus. Plainly, the lay readers will not be acquainted with some of the advanced concepts, nor do they need to be. Familiarity with Brownian motion, compound Poisson processes, some version of Itô's formula, and perhaps Girsanov's theorem ought to be enough to sufficiently appreciate the backdrop against which the new calculus is constructed. The new calculus itself only needs a grasp of drift, volatility, and jump arrival intensity, plus three basic rules that are self-evident on an informal level.

The paper is organized as follows. In the rest of this introduction, we trace how the novel concept of this paper, the semimartingale representation (1.18), arises from classical Itô calculus. Section 2 provides a thorough introduction to the simplified stochastic calculus. It also explains how the proposed approach facilitates computation of drifts and expected values; in particular, it tackles the introductory example in the presence of jumps. Section 3 demonstrates the strength of the proposed approach on three additional examples. Section 4 amplifies this point by showcasing calculations that also require a change of measure. In particular, Example 4.3 contains a new result that makes use of a non-equivalent change of measure. Section 5 highlights the robustness of the proposed approach whereby, for a given task, the same representation applies in both discrete and continuous models. Such unification is unattainable in standard calculus.

The examples in the paper are inspired by applications in Economics and Financial Mathematics but the broader lessons are clearly applicable to Science at large. We explore wider repercussions of the proposed methodology and briefly mention applications to Statistics and Engineering in the concluding Section 6.

Smart Stochastic Discount Factors

SOFONIAS ALEMU KORSAYE, ALBERTO QUAINI and FABIO TROJANI*

January 22, 2020

ABSTRACT

We introduce model-free Smart Stochastic Discount Factors (S–SDFs) minimizing various notions of SDF dispersion, under general convex constraints on non zero pricing errors. S–SDFs can be naturally motivated by market frictions, asymptotic no-arbitrage conditions in an APT framework or a need for SDF regularization. More broadly, we show that they are always supported by a suitable viable economy with transaction costs. Minimum dispersion S–SDFs give rise to new nonparametric bounds for asset pricing models, under weaker assumptions on a model's ability to price cross-sections of assets. They arise from a simple transformation of the optimal payoff in a dual penalized portfolio problem. We clarify the deep properties of S–SDFs induced by various economically motivated pricing error structures and develop a systematic tractable approach for their empirical analysis. For various APT settings, we demonstrate the improved out-of-sample pricing performance of minimum dispersion S–SDFs, which directly corresponds to highly profitable dual portfolio strategies.

Keywords: SDF, Pricing Errors, Fundamental Theorem of Asset Pricing, Market Frictions, Model-free SDF, Minimum Dispersion SDF, Arbitrage Pricing Theory, Factor Models, Tests of Asset Pricing Models, SDF Bounds, Transaction Costs, Portfolio Regularization, Lasso, Ridge, Elastic Net.

*Sofonias Alemu Korsaye (email: Sofonias.Korsaye@unige.ch) and Fabio Trojani (email: Fabio.Trojani@unige.ch) are with University of Geneva, Geneva Finance Research Institute & Swiss Finance Institute. Alberto Quaini (email: Alberto.Quaini@unige.ch) is with University of Geneva, Geneva Finance Research Institute. We thank Caio Almeida, Federico Bandi, Tony Berrada, Federico Carlini, Ines Chaieb, George Constantinides, Jerome Detemple, Patrick Gagliardini, Eric Ghysels, Lars Hansen, Oliver Linton, Stefan Nagel, Olivier Scaillet, Paul Schneider, Michael Weber, Dacheng Xiu, conference participants at the 2019 SoFiE annual Meeting in Shanghai, the 2019 European Econometric Society Meeting in Manchester, the 2019 Workshop on Big Data and Economic Forecasting in Ispra, the 2019 Conference on Quantitative Finance and Financial Econometrics in Marseille, the 2019 Financial Econometrics Conference in Toulouse, the 2019 Vienna Congress on Mathematical Finance, the 2019 Swiss Finance Institute Research days, the 2019 ESSEC Workshop on Monte Carlo Methods and Approximate Dynamic Programming with Applications in Finance, the 2019 Paris December Finance Meeting and seminar participants at BI Norwegian Business School, Bocconi University, University of Zurich, University of Lugano, University of Geneva, University of Lund, JRC in Ispra and Luiss Business School. All errors are ours.

Social Contagion and the Survival of Diverse Investment Styles*

David Hirshleifer[†], Andrew W. Lo[‡], and Ruixun Zhang[§]

Extended Abstract

The Efficient Markets Hypothesis is based upon the premise that investors with correct beliefs will grow richer at the expense of agents with incorrect beliefs. However, there is evidence of social contagion of investment behavior in financial markets that is not always explained by rational information processing.

In this article, we examine the contagion of investment ideas in a multiperiod setting in which investors are more likely to transmit their ideas to other investors after experiencing higher payoffs in one of two investment styles with different return distributions. We show that heterogeneous investment styles are able to coexist in the long run, implying a greater diversity than traditional theory predicts. We characterize the survival and popularity of styles in relation to the return characteristics of the underlying securities, including their mean returns, betas, and idiosyncratic volatilities.

*We thank seminar participants at New York University, Peking University, and the conference on Evolution and Financial Markets for very helpful comments. Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged.

†Distinguished Professor of Finance and Merage Chair in Business Growth, Merage School of Business, University of California, Irvine.

‡Charles E. and Susan T. Harris Professor, MIT Sloan School of Management; Director, MIT Laboratory for Financial Engineering; Principal Investigator, MIT Computer Science and Artificial Intelligence Laboratory; External Faculty, Santa Fe Institute. alo-admin@mit.edu (email).

§Assistant Professor, Peking University; Research Affiliate, MIT Laboratory for Financial Engineering. zhangruixun@pku.edu.cn (email).

In addition, our framework incorporates important psychological forces that affect an investor's receptiveness toward other investors' choices of style. We demonstrate that conformist preference—the phenomenon that investors view others as being well-informed and therefore follow their choices—can lead to oscillations and bubbles in certain financial environments. Attention to novelty—the phenomenon that investors are more likely to pay attention to a novel investment choice if it is *very* different from the most popular ones—leads to an even higher degree of diversity in the long run.

Our model offers empirically testable predictions, leading to explanations for certain puzzles about returns that are difficult to reconcile within traditional asset pricing models, including the "beta puzzle" and the "idiosyncratic volatility puzzle." These results also provide new insights into the persistence of the wide range of investment strategies used by individual investors, hedge funds, and other professional portfolio managers.

Given recent developments in information technology and the growth of social networks, to understand economic and financial behavior it is important to incorporate the impact of contagion via social interactions. Our model, and more generally, the evolutionary finance approach, offers a possible framework for modeling how social interactions affect financial behaviors and market prices.

Keywords: Contagion; Investment Styles; Investor Behavior; Investor Psychology; Adaptive Markets

JEL Classification: G40, G11, G12, G23

State-dependent temperature control for Langevin diffusions

We study the temperature control problem for Langevin diffusions in the context of non-convex optimization. The classical optimal control of such a problem is of the bang-bang type, which is overly sensitive to errors. A remedy is to allow the diffusions to explore other temperature values and hence smooth out the bang-bang control. We accomplish this by a stochastic relaxed control formulation incorporating randomization of the temperature control and regularizing its entropy. We derive a state-dependent, truncated exponential distribution, which can be used to sample temperatures in a Langevin algorithm, in terms of the solution to an HJB partial differential equation. We carry out a numerical experiment on a one-dimensional baseline example, in which the HJB equation can be easily solved, to compare the performance of the algorithm with three other available algorithms in search of a global optimum.

Option Exercise Games and the q Theory of Investment

Min Dai* Zhaoli Jiang† Neng Wang ‡

The real-options approach towards corporate investment is regularly taught in business school classrooms and widely used by practitioners. A key insight underpinning the real-options approach is that the opportunity of increasing a firm's capital stock is analogous to exercising an expansion option on the firm's business opportunity generating incremental stochastic cash flows. Almost all existing literature on real-options do not model the strategic interactions across firms. In those models, investment (real option exercise) strategies are formulated in isolation, without taking into account competitors' investment strategies.

In reality, however, competition is a first-order consideration for many firms. The value of a firm crucially depends not only on its own business investment strategy but also on its competitors' strategies. Grenadier (2002) is among the first to develop an elegant, tractable model of oligopoly real-option exercise game by integrating a classic real options (monopolist) model based on Abel and Eberly (1996) into a Cournot setting where firms choose investment and output facing imperfect competition.

Back and Paulsen (2009) point out that the equilibrium solution concept in Grenadier (2002) is in open-loop form. Open-loop strategies only depend on the history of the exogenous demand shock process, regardless of competitors' investment strategies. However, these commitments may not be credible. In practice, as a firm observes its competitor's investment strategies and capital stock dynamics, it has incentives to adjust its own investment strategy in response. This reasoning suggests that closed-loop equilibrium, which explicitly take into account the feedback between investment strategies and the state variables, is a plausible solution concept.

Back and Paulsen (2009) emphasize the importance of closed-loop equilibria but note that "there are difficulties in even defining the game in closed-loop form." The closed-loop equilibrium identified in their paper is the one that features the perfectly competitive outcome in which all firms make zero profits. In this equilibrium, the option value is thus completely

*The Hong Kong Polytechnic University. E-mail: mindai@polyu.edu.hk
†National University of Singapore. Email: rmijz@nus.edu.sg
‡Columbia University and NBER. E-mail: neng.wang@columbia.edu

eroded to zero. However, this prediction is inconsistent with the observation that oligopolies make positive profits in reality.

In this paper, we take up the challenges posed in Back and Paulsen (2009) and make the following contributions. First, we propose a definition of real option exercise game in closed-loop form where a firm can respond optimally to its competitor's investment strategy over time. Second, we provide and prove a verification theorem for closed-loop equilibria. Additionally, we provide economically intuitive interpretations for the conditions under which our verification theorem holds. Third, we solve for optimal closed-loop investment strategies and corresponding value functions in closed form for a set of closed-loop equilibria in an economy where the demand function is linear in industry output. We then identify the closed-loop equilibrium with the highest option value and show that this equilibrium is unique. Finally, we show that a firm is able to preserve a significant portion of its option value even in closed-loop equilibria, because the competitors' off-the-equilibria overinvestment punishment strategies deter the firm from excessive production, thus supporting significant option values and profits for all firms in equilibrium. Our result complements to but differs from the results that competition significantly erodes the option value as emphasized in Grenadier (2002) and Back and Paulsen (2009).

Finally, our model also generates predictions on marginal q and Tobin's average q under imperfect competition, thus generalizes the q theory of investment under perfect competition and monopoly settings (e.g., Lucas Jr and Prescott (1971); Hayashi (1982); Abel and Eberly (1996)) to Cournot settings with imperfect competition. We show that imperfection competition substantially enriches the predictions of the neoclassical q theory of investment.

References

Abel, A. B. and Eberly, J. C. (1996). Optimal investment with costly reversibility, *The Review of Economic Studies* **63**(4): 581–593.

Back, K. and Paulsen, D. (2009). Open-loop equilibria and perfect competition in option exercise games, *The Review of Financial Studies* **22**(11): 4531–4552.

Grenadier, S. R. (2002). Option exercise games: An application to the equilibrium investment strategies of firms, *The Review of Financial Studies* **15**(3): 691–721.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation, *Econometrica: Journal of the Econometric Society* pp. 213–224.

Lucas Jr, R. E. and Prescott, E. C. (1971). Investment under uncertainty, *Econometrica: Journal of the Econometric Society* pp. 659–681.

Strategies of naive and sophisticated gamblers under weighted-average risk preferences *

December 17, 2019

Abstract

In the recent two papers by Ebert and Strack (2015, 2018), the authors show that in continuous time, naive agent never stops gambling and sophisticated gambler never starts to do so when their risk preferences are characterized by cumulative prospect theory. In this paper, we investigate the strategies of the above-mentioned two types of agents, whose risk preference is induced by a weighted average of expected utility theory and cumulative prospect theory. Furthermore, the definition of weak equilibrium in control theory is employed to study the strategies of sophisticated agent. Our results then generalize the ones in Ebert and Strack (2015, 2018). We show that under weighted-average preferences, naive gambler still never stops for a wide range of specifications, but sophisticated gambler does not always stop immediately under certain circumstances.

Key words: naive agents; sophisticated agents; weak equilibrium; cumulative prospect theory

*Acknowlegement

Stripping the Discount Curve—a Robust Machine Learning Approach*

Damir Filipović† Markus Pelger‡ Ye Ye§

March 15, 2022

Abstract

We introduce a robust, flexible and easy-to-implement method for estimating the yield curve from Treasury securities. This method is non-parametric and optimally learns basis functions in reproducing Hilbert spaces with an economically motivated smoothness reward. We provide a closed-form solution of our machine learning estimator as a simple kernel ridge regression, which is straightforward and fast to implement. We show in an extensive empirical study on U.S. Treasury securities, that our method strongly dominates all parametric and non-parametric benchmarks. Our method achieves substantially smaller out-of-sample yield and pricing errors, while being robust to outliers and data selection choices. We attribute the superior performance to the optimal trade-off between flexibility and smoothness, which positions our method as the new standard for yield curve estimation.

Keywords: yield curve estimation, U.S. Treasury securities, term structure of interest rates, nonparametric method, machine learning in finance, reproducing kernel Hilbert space
JEL classification: C14, C38, C55, E43, G12

*We thank Nicolas Camenzind, Rüdiger Fahlenbrach, Kay Giesecke, Thomas Viehmann, and Sander Willems and seminar and conference participants at Stanford, ETH Zurich, Columbia University, Swiss Re, Ettore Majorana Foundation and Centre For Scientific Culture, Vienna University of Economics and Business, the World Online Seminars on Machine Learning in Finance, INFORMS and the SoFiE Financial Econometrics Summer School for helpful comments. We thank the China Merchants Bank for generous research support.

†École Polytechnique Fédérale de Lausanne and Swiss Finance Institute, Email: damir.filipovic@epfl.ch
‡Stanford University, Department of Management Science & Engineering, Email: mpelger@stanford.edu.
§Stanford University, Department of Management Science & Engineering, Email: yeye@stanford.edu.

Speaker: Prof. Xun Li, The Hong Kong Polytechnic University

Title: Survey on Multi-period Mean-Variance Portfolio Selection Model

Abstract: Due to the non-separability of the variance term, the dynamic mean-variance (MV) portfolio optimization problem is inherently difficult to solve by dynamic programming. Li and Ng (2000) and Zhou and Li (2000) develop the pre-committed optimal policy for such a problem using the embedding method. Following this line of research, researchers have extensively studied the MV portfolio selection model through the inclusion of more practical investment constraints, realistic market assumptions, and the various financial applications. As the principle of optimality no longer holds, the pre-committed policy suffers from the time-inconsistent issue, i.e., the optimal policy computed at the intermediate time \$t\$ is not consistent with the optimal policy calculated at any time before time \$t\$. The time inconsistency of the dynamic MV model has become an important yet challenging research topic. This paper mainly focuses on the multi-period mean-variance (MMV) portfolio optimization problem, reviews the essential extensions, and highlights the critical development of time-consistent policies.

Abstract

We study the system of heterogeneous interbank lending and borrowing based on the relative average of log-capitalization given by the linear combination of the average within groups and the ensemble average and describe the evolution of log-capitalization by a system of coupled diffusions. The model incorporates a game feature with homogeneity within groups and heterogeneity between groups where banks search for the optimal lending or borrowing strategies through minimizing the heterogeneous linear quadratic costs in order to avoid to approach the default barrier. In order to analyze the heterogeneous behavior influenced by the external factor given by the market or economic environment, the factor model is proposed. Due to the complicity of the lending and borrowing system, the closed-loop Nash equilibria and the open-loop Nash equilibria are both driven by the coupled partial differential equations (PDEs). The existence of the equilibria in the two-group case where the number of banks goes to infinity in each group is guaranteed by the solvability for the coupled PDEs. The equilibria are consisted of the mean-reverting term identical to the one group game and the group average owing to heterogeneity. In addition, the corresponding heterogeneous mean filed game with the arbitrary number of groups is also discussed. The existence of the ϵ-Nash equilibrium in the general d heterogeneous groups is also verified. Finally, in the financial implication, we observe the Nash equilibria governed by the mean-reverting term and the linear combination of the ensemble averages of individual groups and study the influence of the relative parameters on the corresponding liquidity rate through the numerical analysis.

Systemic Risk in Markets with Multiple Central Counterparties

Luitgard A. M. Veraart [*] Iñaki Aldasoro[†]

Abstract

We provide a framework for modelling risk and quantifying payment shortfalls in cleared markets with multiple central counterparties (CCPs). Building on the stylised fact that clearing membership is shared among CCPs, we show that stress in this shared membership can transmit across markets through multiple CCPs. We provide stylised examples to lay out how such stress transmission can take place, as well as empirical evidence based on publicly available data to illustrate that the mechanisms we study could be relevant in practice. Furthermore, we show how stress mitigation mechanisms such as variation margin gains haircutting by one CCP can have spillover effects to other CCPs. Finally, we discuss how the framework can be used to enhance CCP stress-testing. The current "Cover 2" standard requires CCPs to be able to withstand the default of their two largest clearing members. We show that who these two clearing members are can be significantly affected by higher-order effects arising from interconnectedness through shared membership.

[*]London School of Economics and Political Science, Department of Mathematics, Houghton Street, London WC2A 2AE, UK, l.veraart@lse.ac.uk

[†]Bank for International Settlements, Centralbahnplatz 2, 4051 Basel, Switzerland. inaki.aldasoro@bis.org

Systemic Risk Quantification via Shock Amplification in Financial Networks

Dohyun Ahn[*], Nan Chen[†]
The Chinese University of Hong Kong

Kyoung-Kuk Kim[‡]
Korea Advanced Institute of Science and Technology

Abstract

Under the Eisenberg-Noe model for financial networks, we are interested in the probability that at least one bank in a specific group fails when partial network information is available. The motivation of this problem comes from the protection of a financial system by regulating the so-called systemically important financial institutions (SIFIs), e.g., higher loss absorbency requirements, and from the fact that the information of interbank transactions is not fully known in practice. This makes it very essential for regulators to estimate and understand the default probability of one or more SIFIs under incomplete information. However, little research has been carried out in this direction.

Based on the characterization of shock amplification caused by the network structure and bankrupcy costs, we propose a robust optimization approach to assess the said probability. In particular, we investigate the worst-case default probability by formulating a mixed-integer linear program that characterizes the worst-case default event that takes into account all possible network configurations under limited information. Our new approach facilitates the unbiased estimation of the worst-case default probability and is applicable to a variety of partial information that is practically observable. The existing robust optimization methods either provide more conservative solutions than ours or require computationally intractable procedures. In addition, we compute the worst-case capital amount required to keep all SIFIs solvent with a certain probability in order to protect the financial system from falling into a crisis, which is directly associated with the aforementioned motivation of this work, and provide practical insights into network information collection from regulators' perspective through extensive numerical experiments with real-world data.

[*]Assistant Professor, Department of Systems Engineering & Engineering Management, `dohyun.ahn@cuhk.edu.hk`
[†]Professor, Department of Systems Engineering & Engineering Management, `nchen@se.cuhk.edu.hk`
[‡]Associate Professor, College of Business, `kkim@kaist.ac.kr`

Tail-GAN: Nonparametric Scenario Generation for Tail Risk Estimation

Rama Cont[1], Mihai Cucuringu[1,2,3], Renyuan Xu[4], and Chao Zhang[2]

[1]Mathematical Institute, University of Oxford, Oxford, UK
[2]Department of Statistics, University of Oxford, Oxford, UK
[3]The Alan Turing Institute, London, UK
[4]Department of Industrial and Systems Engineering, University of Southern California, Los Angeles, US

Abstract

The estimation of loss distributions for dynamic portfolios requires the simulation of scenarios representing realistic joint dynamics of their components, with particular importance devoted to the simulation of *tail risk* scenarios. Commonly used parametric models have been successful in applications involving a small number of assets, but may not be scalable to large or heterogeneous portfolios involving multiple asset classes.

We propose a novel data-driven approach for the simulation of realistic multi-asset scenarios with a particular focus on the accurate estimation of tail risk for a given class of static and dynamic portfolios selected by the user. By exploiting the joint elicitability property of Value-at-Risk (VaR) and Expected Shortfall (ES), we design a Generative Adversarial Network (GAN) architecture capable of learning to simulate price scenarios that preserve tail risk features for these benchmark trading strategies, leading to consistent estimators for their Value-at-Risk and Expected Shortfall. We demonstrate the accuracy and scalability of our method via extensive simulation experiments using synthetic and market data. Our results show that, in contrast to other data-driven scenario generators, our proposed scenario simulation method correctly captures tail risk for both static and dynamic portfolios.

Introduction

Scenario simulation is extensively used in finance for estimating the loss distribution of portfolios and trading strategies, often with a focus on the estimation of risk measures such as Value-at-Risk and Expected Shortfall [2]. The estimation of such risk measures for static and dynamic portfolios involves the simulation of scenarios representing realistic joint dynamics of their components. This requires both a realistic representation of the temporal dynamics of individual assets (*temporal dependence*), as well as an adequate representation of their co-movements (*cross-asset dependence*).

A common approach in scenario simulation is to use parametric models, either time series models such as GARCH models, diffusion and jump processes, stochastic volatility models, and copulas. While these approaches have been successful in low-dimensional applications involving a small number of assets, very few of them are scalable to large or heterogeneous portfolios involving multiple asset classes. In most high-dimensional applications, Gaussian models with constant coefficients (dubbed as the "multivariate Black-Scholes model") have been often used as the default approach, even when there is ample econometric evidence for their inadequacy, given the stylized features –heavy tails, volatility clustering, tail dependence, etc.– observed in many financial time series [1].

Generative Adversarial Networks (GANs) [4] have emerged in the recent years as an efficient alternative to parametric models for the simulation of patterns whose features are extracted from complex and high-dimensional data sets. GANs are generative models between two competing neural networks: a *generator* network G and a *discriminator* network D. The generator network G attempts to fool the discriminator network by converting random noise into sample data, while the discriminator network D tries to identify whether the input sample is fake or real. The goal for the generator G is to output (generated) samples that resemble as best as possible the true samples under a certain criterion.

Main contributions

To simulate high-dimensional financial scenarios across assets with accurate estimation of tail risk measures for a given class of benchmark strategies, we introduce a novel approach for scenario simulation which specifically addresses the above mentioned issue of lacking problem-dependent statistics as evaluation criterion. In particular, we develop a new framework denoted as *Tail Generative Adversarial Network* (Tail-GAN) to simulate multivariate financial time series data that preserve the same specified *statistics* for benchmark strategies as the training data. Compared to the GANs studied in the literature, the proposed Tail-GAN framework has the following novel ingredients which has not been investigated before:

(1) **Tail risks as evaluation criterion:** Tail risk refers to the risk of large portfolio losses. Value at Risk (VaR) and Expected Shortfall (ES) are commonly used statistics for measuring the tail risk of a variety of trading strategies. To correctly measure the tail risks of trading strategies applied to the input price scenarios and the generated price scenarios, we incorporate VaR and ES of a given set of trading strategies (introduced below) in the loss function by utilizing the joint *elicitability* property of VaR and ES [3]. Namely, $(\mathrm{VaR}, \mathrm{ES})$ can be written as a minimizer of some score function therefore it can be learned by solving some optimization problem.

(2) **Discrimination via dynamic trading strategies:** In the loss function of the Tail-GAN architecture, we introduce the risk measures of multiple static portfolios and dynamic trading strategies applied to the price series, with static portfolios capturing cross-asset dependence information and dynamic strategies capturing temporal dependence information in the financial scenarios. Dynamic trading strategies are *nonlinear* functions and possibly path-dependent functions of the underlying price scenarios. On one hand, the computational complexity of the training process increases when including these strategies into the discriminator's loss function. On the other hand, this exploration in the space of nonlinear functions of the underlying price scenarios *significantly improves* the simulation accuracy of GAN models. Previous works on using GAN for simulating financial scenarios only focus on *matching return distributions* between the simulated data and the input data [5]. This corresponds to considering a naive buy-and-hold strategy as the evaluation criterion and resulting in no guarantee for consistent behaviors of other (nonlinear) strategies applied to the simulated data and the input data.

References

[1] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2):223, 2001.

[2] Paul Glasserman. *Monte Carlo methods in financial engineering*. Springer, 2003.

[3] Tilmann Gneiting. Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106(494):746–762, 2011.

[4] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27:2672–2680, 2014.

[5] Magnus Wiese, Robert Knobloch, Ralf Korn, and Peter Kretschmer. Quant GANs: Deep generation of financial time series. *Quantitative Finance*, pages 1–22, 2020.

Speaker：Prof. Duy-Minh Dang, University of Queensland

Title: TAB

Abstract: TAB

Term structure modelling using cylindrical measure-valued forward rates

Philipp Harms *

1 Term structure models with measure-valued forward rates

Discontinuous term structures. In their seminal work [15] Heath, Jarrow, and Morton (HJM) proposed a model for the bond market which invokes the trading of a "continuum of default-free discount bonds." Their key idea was to parameterize bond prices via forward rates. A recent extension is to allow the term structure of bond prices to be discontinuous in the time to maturity [1, 2, 3]. These discontinuities reflect jumps in the underlying at predictable stopping times. The importance of such jumps is widely acknowledged in the financial literature; see [5, 12, 16, 18]. Prominent examples are elections, in particular the Brexit vote, quarterly reports, and dividend payments. Similarly, on energy markets, scheduled power-plant outtakes, abnormal loads, or temperature extremes drive the spike risk at electricity markets and create discontinuities in the futures' term structure [20, 6, 7, 8].

Measure-valued forward rates. A modeling implication of the above discussion is that the forward rates become measure-valued. For instance, a simplified version of [2, 1] reads as follows. Let τ denote the default time, i.e., the stopping time at which all zero-recovery bonds loose their value. Then the defaultable bond prices are given by

$$P(t,T) = \mathbb{1}_{\{\tau>t\}} \exp\left(-\int_{(t,T]} X_t(du)\right), \qquad 0 \le t \le T, \tag{1}$$

where the forward rate X_t is a random measure on $\mathbb{R}_+$, given in a certain finite-dimensional parametric form. While these measure-valued forward rates X_t formally satisfy an HJM equation, the development of genuinely infinite-dimensional measure-valued HJM equations meets severe technical difficulties. The next section points out some of these difficulties and shows that they can be circumvented using a new approach based on cylindrical stochastic analysis. This approach has applications to a variety of large financial markets, including for interest rate, energy, and insurance products.

*This presentation is based on joint research with C. Fontana, G. Szulda, and T. Schmidt.

1

2 Measure-valued Heath–Jarrow–Morton equations

Challenges. It is difficult to give rigorous meaning to HJM equations of the general form

$$\mu_t = S_t\mu_0 + \int_0^t S_{t-s}b_s(\mu_s)ds + \int_0^t S_{t-s}\sigma_s(\mu_s)dW_s, \qquad t \in \mathbb{R}_+, \tag{2}$$

where μ, b, and σ are measure-valued, S is the shift semigroup, and W is Brownian motion. Weighted Sobolev spaces, which are commonly used in term structure modeling [13, 17], are either too small to contain all shifted Dirac measures or too large to allow for evaluations against intervals. Therefore, there does not appear to be an easy alternative to using signed measures as a state space for such HJM equations. Unfortunately, the shift semigroup S, which appears in the HJM equation, is not strongly measurable and, even worse, not Rademacher-bounded on the space of signed measures. Consequently, measure-valued stochastic convolutions are ill-defined. For instance, if δ_0 is the Dirac measure at zero and W is standard Brownian motion, then the stochastic convolution $\int_0^t S_{t-s}\delta_0 dW_s$ defines a random distribution with almost surely unbounded total variation, which for this reason cannot be a random measure. However, as explained next, it is a cylindrical random measure, and this can be used to make sense of (2).

Well-posedness in the cylindrical sense. Cylindrical measures, random variables, and processes are defined by the requirement that all linear projections to Euclidean spaces are bona fide measures, random variables, and processes, respectively. They came up in the 1960s and 1970s in probability theory on topological vector spaces [21, 22, 23, 9], and there is by now an extensive literature on many aspects of the topic; see e.g. [4, 11, 19, 10, 14]. However, the theory has not yet been applied to large financial markets and, more specifically, measure-valued term structure equations.

For the present purpose, cylindrical measures on $\mathbb{R}_+$ shall be understood as continuous linear maps $\mu\colon C_0(\mathbb{R}_+) \to L^2(\Omega)$. An example is the above stochastic convolution $\int_0^t S_{t-s}\delta_0 dW_s$. Thanks to the reflexivity of $L^2(\Omega)$, cylindrical measures have Riesz representations as countably additive $L^2(\Omega)$-valued vector measures. In particular, evaluations against intervals are well-defined. Any existing scalar stochastic integral can be turned into a cylindrical stochastic integral without further ado. Note that this is a major simplification compared to stochastic analysis on Banach spaces. Consequently, HJM equations (2) have cylindrical meaning for a wide variety of driving noises and are well-posed under the usual Lipschitz conditions.

Summary and outlook. Cylindrical stochastic analysis is simpler than standard stochastic analysis and provides additional generality, which is needed to formulate measure-valued SPDEs with financial applications. The results open the door to the study of cylindrical semimartingales and cylindrical affine processes with applications to e.g. energy, volatility, or insurance markets.

References

[1] C. Fontana and T. Schmidt. "General dynamic term structures under default risk". In: *Stochastic Processes and their Applications* 128.10 (2018), pp. 3353–3386.

[2] F. Gehmlich and T. Schmidt. "Dynamic defaultable term structure modelling beyond the intensity paradigm". In: *Mathematical Finance* 28.1 (2018), pp. 211–239.

[3] M. Keller-Ressel, T. Schmidt, and R. Wardenga. "Affine processes beyond stochastic continuity". In: *Annals of Applied Probability* (2019), to appear.

[4] D. Applebaum and M. Riedle. "Cylindrical Lévy processes in Banach spaces". In: *Proc. Lond. Math. Soc. (3)* 101.3 (2010), pp. 697–726.

[5] A. Bélanger, S. E. Shreve, and D. Wong. "A general framework for pricing credit risk". In: *Mathematical Finance* 14.3 (2004), pp. 317–350.

[6] F. E. Benth and P. Krühner. "Derivatives pricing in energy markets: an infinite-dimensional approach". In: *SIAM Journal on Financial Mathematics* 6.1 (2015), pp. 825–869.

[7] F. E. Benth and P. Krühner. "Representation of infinite-dimensional forward price models in commodity markets". In: *Communications in Mathematics and Statistics* 2.1 (2014), pp. 47–106.

[8] F. E. Benth and F. Paraschiv. "A space-time random field model for electricity forward prices". In: *Journal of Banking & Finance* (2017).

[9] V. I. Bogachev. *Gaussian measures*. 62. American Mathematical Soc., 1998.

[10] D. Criens. "Cylindrical martingale problems associated with Lévy generators". In: *Journal of Theoretical Probability* (2018), pp. 1–54.

[11] G. Da Prato and J. Zabczyk. *Stochastic equations in infinite dimensions*. Cambridge University Press, 2014.

[12] E. F. Fama. "Efficient capital markets: a review of theory and empirical work". In: *The journal of Finance* 25.2 (1970), pp. 383–417.

[13] D. Filipović. *Consistency problems for Heath–Jarrow–Morton interest rate models*. Vol. 1760. Lecture Notes in Mathematics. Springer Verlag. Berlin Heidelberg New York, 2001.

[14] G. Grafendorfer. "Infinite-dimensional affine processes". PhD thesis. ETH Zurich, 2016.

[15] D. Heath, R. A. Jarrow, and A. J. Morton. "Bond pricing and the term structure of interest rates". In: *Econometrica* 60 (1992), pp. 77–105.

[16] M. H. Miller and K. Rock. "Dividend policy under asymmetric information". In: *The Journal of Finance* 40.4 (1985), pp. 1031–1051.

[17] S. Peszat and J. Zabczyk. "Heath-Jarrow-Morton-Musiela equation of bond market". In: *Preprint IMPAN 677, Warsaw* (2007).

[18] M. Piazzesi. "Bond yields and the Federal Reserve". In: *Journal of Political Economy* 113.2 (2005), pp. 311–344.

[19] M. Riedle. "Ornstein-Uhlenbeck processes driven by cylindrical Lévy processes". In: *Potential Analysis* 42.4 (2015), pp. 809–838.

[20] T. Schmidt. "Modelling energy markets with extreme spikes". In: *Mathematical Control Theory and Finance*. Ed. by R. Grossinho, M. Guerra, A. Sarychev, and A. Shiryaev. Springer, 2008.

[21] L. Schwartz. *Radon measures on arbitrary topological spaces and cylindrical measures*. Oxford University Press, 1973.

[22] L. Schwartz and P. R. Chernoff. "Geometry and probability in Banach spaces". In: *Springer Lecture Notes in Mathematics* (1981).

[23] N. Vakhania, V. Tarieladze, and S. Chobanyan. *Probability distributions on Banach spaces*. Vol. 14. Springer Science & Business Media, 1987.

Risk neutral variance term structures are characterized by their time elasticities. They are synthesized by space scaling and time changing self decomposable laws at unit time. Monotone time elasticities are modeled using exponential functions while gamma functions permit humps. Results for both cases as time changes are followed by simultaneously space scaling and time changing. Space scaling contributes towards front end options while time changing works on the back end. Splitting the space scaling and time changing for the positive and negative moves delivers models with rising absolute skewness and kurtosis. The space scaled and time changed densities are those of additive processes made up by aggregating over time the space scaled innovations of the sum of two underlying time changed Lévy processes taken at log time. The space scaled process is also a space scaled solution to a time varying OU equation driven by a time changed Lévy process taken at log time, where the OU mean reversion rates are the time elasticities. The two processes are termed the space scaled and time change components and their relative contributions, space to time are determined to be twice the ratio of their variance elasticities.

The Hurst roughness exponent and its model-free estimation

Xiyue Han* and Alexander Schied

March 2, 2022

Abstract

We say that a continuous real-valued function x admits the Hurst roughness exponent H if the p^{th} variation of x converges to zero if $p > 1/H$ and to infinity if $p < 1/H$. For the sample paths of many stochastic processes, such as fractional Brownian motion, the Hurst roughness exponent exists and equals the standard Hurst parameter. In our main result, we provide a mild condition on the Faber–Schauder coefficients of x under which the Hurst roughness exponent exists and is given as the limit of the classical Gladyshev estimates $\widehat{H}_n(x)$. This result can be viewed as a strong consistency result for the Gladyshev estimators in an entirely model-free setting, because no assumption whatsoever is made on the possible dynamics of the function x. Nonetheless, our proof is probabilistic and relies on a martingale that is hidden in the Faber–Schauder expansion of x. Since the Gladyshev estimators are not scale-invariant, we construct several scale-invariant estimators that are derived from the sequence $(\widehat{H}_n)_{n\in\mathbb{N}}$. We also discuss how a dynamic change in the Hurst roughness parameter of a time series can be detected. Our results are illustrated by means of high-frequency financial time series.

The impact of stochastic volatility on initial margin and MVA in a post LIBOR world

J.H. Hoencamp[†] [a], J.P. de Kort[b], B.D. Kandhai[a,b]

[a]Computational Science Lab, University of Amsterdam, Science Park 904, 1098XH Amsterdam, Netherlands;
[b]Quantitative Analytics, ING Bank, Foppingadreef 7, 1102 BD Amsterdam, The Netherlands
[†]Corresponding author. Email: j.h.hoencamp@uva.nl

In this research we investigate the impact of stochastic volatility on future initial margin (IM) and margin valuation adjustment (MVA) calculations for interest rate derivatives. We incorporate the latest risk-free rate benchmarks (RFR), which are selected to be the new standard after the LIBOR transition. An analysis is performed under different market conditions, namely during the peak of the Covid-19 crisis when the markets were stressed and during Q4 of 2020 when volatilities were low. The Cheyette short-rate model is extended by adding a stochastic volatility component, which is calibrated to fit the EUR swaption volatility surfaces. We extend modern Fourier pricing techniques to accommodate the new RFR benchmark and derive closed-form sensitivity expressions, which are used to model IM profiles in a Monte Carlo simulation framework. The various results are compared to the deterministic volatility case. The results reveal that the replacement of LIBOR by RFR intrinsically alters the IM valuation of swaps and caps/floors and qualitatively impacts their future margin profiles. It is further observed that the inclusion of a stochastic volatility component can have a considerable impact on non-linear derivatives, especially for far out-of-the-money swaptions. The effect is particularly pronounced if the market exhibits a substantial skew or smile in the implied volatility curve. This can have severe consequences for funding cost valuation and risk management.

The learning, timing, and pricing of the option to invest with guaranteed debt and asymmetric information☆

Pengfei Luo[a], Huamao Wang[b], Zhaojun Yang[c,*]

[a]*School of Finance and Statistics, Hunan University, Changsha, China.*
[b]*Kent Business School, University of Kent, Canterbury, Kent, CT2 7FS, UK.*
[c]*Department of Finance, Southern University of Science and Technology, Shenzhen, 518055, China.*

Abstract

We consider an entrepreneur (borrower) who must borrow money from a bank (lender) to start an irreversible but delayable project. The debt is secured by an insurer who takes the project and pays the lender all the outstanding principal and interest in the event of default. Thus the lender is exposed to no risk of loss. In return for the guarantee, the borrower grants the insurer a fraction of the money borrowed (fee-for-guarantee swap, FGS) or of the project's equity (equity-for-guarantee swap, EGS) as guarantee cost. We assume there is information asymmetry on the project profitability, which is uniquely determined by the expected growth rate of the project's cash flow, between the borrower and insurer: The borrower knows the expected growth rate of the project but the insurer only knows that it follows a two-point distribution, i.e. the project is either high-type or low-type.

We develop a signaling game model with learning to determine the guarantee cost and the timing and pricing of the option to invest. In our study,

☆We are grateful for the helpful comments of Sang Hu from the Chinese University of Hong Kong (Shenzhen) and other participants on the Ninth Annual Conference of FEFRM Branch of OR Society of China. The research reported in this paper was started in 2015 and supported by the National Natural Science Foundation of China (Project No. 71371068) and by Special Innovation Projects of Universities in Guangdong (project No. 2018WTSCX131). The order of the authors is alphabetical in the surnames and the contributions of the authors are the same.

*Corresponding author. Tel: +86 755 8801 8603.

Email addresses: `pfluo@hnu.edu.cn` (Pengfei Luo), `h.wang@kent.ac.uk` (Huamao Wang), `yangzj@sustech.edu.cn` (Zhaojun Yang)

Preprint submitted to the 2020 Bachelier Hong Kong *January 25, 2020*

the borrower signals to the insurer the project profitability not only through timing investment decisions but also through proposing the amount of money borrowed and guarantee cost. We address whether the signal is sufficient to identify the project type. If yes, we derive a separating equilibrium and if not, we get a *dynamic* pooling equilibrium, where the insurer demands a mean guarantee cost, according to her/his updated belief on the expected growth rate of the project's cash flow (project type). In contrast to the literature, the pooling equilibrium that is reached by our model is a *dynamic* one where the insurer can learn about the project type from the evolution of the project's cash flow and dynamically update her/his belief on the project type.

We discover that asymmetric information makes a higher profitable borrower pays more guarantee cost than that required in a perfect information case. The learning can reduce the guarantee cost and the cost of adverse selection through updating the insurer's belief on the high-type borrower, which allows the high-type to accelerate investment. A sufficiently high initial belief on a high-type borrower or a sufficiently large funding gap makes the dynamic pooling equilibrium with Bayesian learning Pareto dominate a separating equilibrium since a high initial belief or a large amount of debt facilitates the insurer's learning. These findings are verified by both the FGS and EGS swap. Finally, we show that FGSs are superior to EGSs since the former provides borrowers with a larger firm value and a lower adverse selection cost than the latter.

Keywords: Real options, Secured debt, Asymmetric information, Bayesian learning, Signaling game.
JEL: G13, G31, D82, C73

2

Annika Kemper and Maren Diane Schmeck

The Market Price of Risk for Delivery Periods: Pricing Swaps and Options in Electricity Markets

KEYWORDS: Electricity Swaps · Seasonal Stochastic Volatility · Option Pricing · Delivery Period Risk

Abstract

In electricity markets futures deliver the underlying over a period and thus function as a swap contract. In this paper we introduce a market price of risk for delivery periods of electricity swaps. In particular, we suggest a weighted *geometric average* of an artificial geometric electricity futures price over the corresponding delivery period. This leads to a geometric electricity swap price dynamics without any approximation requirements. Our framework allows to include typical features as the Samuelson effect, seasonalities as well as a stochastic volatility in the absence of arbitrage. We show that our suggested model is suitable for pricing options on electricity swaps using the Heston method. Especially, we illustrate the related pricing procedure for electricity swaps and options in the setting of Arismendi et al. (2016), Schneider and Tavin (2018) and Fanelli and Schmeck (2019).

References

Arismendi, J. C., Back, J., Prokopczuk, M., Paschke, R., and Rudolf, M. (2016): "Seasonal Stochastic Volatility: Implications for the Pricing of Commodity Options". In: *Journal of Banking and Finance* 66, pp. 53–65.

Fanelli, V. and Schmeck, M. D. (2019): "On the Seasonality in the Implied Volatility of Electricity Options". In: *Quantitative Finance* 19, pp. 1321–1337.

Heston, S. L. (1993): "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options". In: *The Review of Financial Studies* 6, pp. 327–343.

Schneider, L. and Tavin, B. (2018): "From the Samuelson Volatility Effect to a Samuelson Correlation Effect: An Analysis of Crude Oil Calendar Spread Options". In: *Journal of Banking & Finance* 95, pp. 185–202.

Theory of Cryptocurrency Interest Rates

Dorje C. Brody[1], Lane P. Hughston[2] and Bernhard K. Meister[3]

[1] *Department of Mathematics, University of Surrey, Guildford GU2 7XH, UK*
[2] *Department of Computing, Goldsmiths College, University of London, New Cross, London SE14 6NW, UK*
[3] *School of Management, Swansea University, Swansea SA1 8EN, UK*

(Dated: December 19, 2019)

A term structure model in which the short rate is zero is developed as a candidate for a theory of cryptocurrency interest rates. The price processes of crypto discount bonds are worked out, along with expressions for the instantaneous forward rates and the prices of interest-rate derivatives. The model admits functional degrees of freedom that can be calibrated to the initial yield curve and other market data. Our analysis suggests that strict local martingales can be used for modelling the pricing kernels associated with virtual currencies based on distributed ledger technologies.

The paper has been accepted for publication by SIAM Journal on Financial Mathematics, and can be found at arXiv:1904.05472.

Keywords: cryptocurrencies, distributed ledger technologies, blockchains, interest rate models, pricing kernels, foreign exchange, derivatives.

1. Introduction

Over a decade has passed since the release of the white paper establishing bitcoin — the progenitor of the burgeoning open-source cryptocurrency movement [1]. In the intervening years the ability of digital currencies to flourish with market capitalizations in the billions of dollars without the backing of sovereign states has been widely publicized. One of the key factors that has to this date hindered cryptocurrencies from reaching the mainstream and expanding beyond fringe applications to link with the real economy in significant ways, despite numerous preliminary attempts, has been the inability of regulated trading consortiums to develop broad platforms for cryptocurrency derivatives and structured products, and more generally to put in place the robust infrastructure needed to support financial markets analogous to those we take for granted in connection with sovereign currencies. But there are signs that such markets are beginning to develop. Cash-settled bitcoin futures contracts have been listed since December 2017 by Chicago Mercantile Exchange Group (CME), a major derivatives exchange, and were available at Chicago Board Options Exchange (CBOE) from December 2017 to June 2019. Such faltering progress can be typical in the emergence of new markets (one saw this in the early days of the credit derivative markets in the 1990s), so one will have to wait and see how it plays out. Liquidity in the CME futures contracts has so far been concentrated mainly in the first two monthly contracts, with daily trading volumes in the tens to hundreds of millions of dollars. Options on bitcoin futures are scheduled to be launched by CME in January 2020. In addition, an assortment of cryptocurrency-focused exchanges in various parts of the world and covered by a range of regulatory frameworks have emerged, with a variety of different types of contracts being traded. One rather unusual example is a so-called perpetual contract traded on Bitmex, which has no expiry or settlement, but has periodic payments exchanged between longs and shorts every eight hours, hence mimicing a margin-based spot market. Some of these new exchanges also provide platforms for short-term borrowing and lending between market participants for up to a month or so in a number of cryptocurrencies and fiat currencies.

Currently, a large number of start-ups are vying to establish products and infrastructures for cryptocurrency financial markets. This activity ranges from cryptocurrency currency loans collateralized by holdings in liquid fiat currencies, or other assets, to more ambitious projects such as UMAproject, dYdX Trading, and many others, that aim to provide protocols for so-called smart contracts. These are contracts that implement financial derivatives on the Ethereum blockchain, and other decentralized ledgers, without human intervention.

In economies based on well-established sovereign currencies, interest rate derivatives are central to the functioning of financial markets. A statistical bulletin published in December 2018 by the Bank of International Settlement illustrates how the notional amount for fixed-income derivatives trumps every other category by a multiple of five or more [2]. For an interest-rate derivatives market to function, products need to be priced, payoffs need to be replicated, and positions need to be hedged — and for pricing, replicating, and hedging, it is essential that financial institutions, regulators, and other market participants should have at their disposal a diverse assortment of serviceable interest-rate models that are well-adapted to characteristics of the currencies in which the instruments being traded are based. When it comes to cryptocurrencies, this requirement immediately poses a challenge, since cryptocurrencies by their nature offer no short-term interest, and it seems impossible *prima facie* that one should be able to build an interest rate model for which the short rate is identically zero. In the case of bitcoin, the regularly-updated blockchain represents the distributed ledger and therefore holdings in the cryptocurrency. The technical details of the updating process are highly involved, but do not impinge upon the general statement that holdings of bitcoin as recorded in the ledger do not earn interest. No central bank exists. New coins are not issued to existing holders of the coins, but instead are awarded to "miners" successfully solving numerical puzzles necessary to update the blockchain. Other prominent cryptocurrencies provide variations on this general theme of not paying by accommodating additional functionality. For example, ethereum, the second largest cryptocurrency, incorporates distributed computing.

What are the implications for interest rate modelling? To begin, it may be helpful if we recall what is generally meant by a "conventional" model for interest rates. This will allow us to identify some of the differences between the conventional theory and the version required for cryptocurrencies. Such conventional interest rate models exist in abundance, and include, for instance, most of the models constructed in the papers appearing in the collections [3] and [4]. One typically assumes the existence of a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with an associated filtration $\{\mathcal{F}_t\}_{t\geq 0}$ satisfying the usual conditions, upon which an adapted short-rate process $\{r_t\}_{t\geq 0}$ is defined. Here, time 0 denotes the present. By a unit discount bond we mean a financial instrument that delivers a single cash flow of one unit of currency at time T and derives its value entirely from that cash flow. The price P_{tT} of a T-maturity unit discount bond at time $t < T$ is then given by a conditional expectation of the form

$$P_{tT} = \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s \,\mathrm{d}s \right) \Big| \, \mathcal{F}_t \right], \tag{1}$$

taken under a suitably-specified risk-neutral measure $\mathbb{Q}$ that is equivalent (in the sense of agreeing on null sets) to the physical measure $\mathbb{P}$. Glancing at equation (1) one might conclude that it is not possible to have a nontrivial discount-bond system if $r_t = 0$ for all $t \geq 0$, and indeed this is true in "conventional" models. For example, in the well-known Heath-Jarrow-Morton (HJM) model [5], the instantaneous forward rates, which determine the discount bond prices via the relation

$$P_{tT} = \exp\left(-\int_t^T f_{tu} \,\mathrm{d}u \right), \tag{2}$$

for $t < T$, are given by

$$f_{tT} = \frac{1}{P_{tT}} \mathbb{E}^{\mathbb{Q}}\left[r_T \exp\left(-\int_t^T r_s \,\mathrm{d}s \right) \Big| \, \mathcal{F}_t \right]. \tag{3}$$

Thus f_{tT} is the forward price per unit notional, made at time t, for purchase of the rights to a cash flow of the amount r_T per unit notional at time T. It follows that if the short rate vanishes then so do the instantaneous forward rates, in which case the discount bond system trivializes and takes the form $P_{tT} = 1$ for all $t < T$. The same conclusion follows directly from (1). Hence, if the short rate vanishes in a conventional interest rate model, it follows immediately that so do all the term rates.

Nevertheless, just because the short rate vanishes we are not necessarily forced to conclude that the discount bond system is trivial. Beginning with the work of Constantinides [6], finance theorists have learned to think about interest-rate modelling in a more general way, in terms of so-called pricing kernels. The pricing kernel method avoids some of the technical issues that arise with the introduction of the risk neutral measure and the selection of a preferred numeraire asset in the form of a money market account, and at the same time it leads to interesting new classes of interest rate models (see, for example, [7–10] and references cited therein). By a pricing kernel, we mean an $\{\mathcal{F}_t\}$-adapted càdlàg semimartingale $\{\pi_t\}_{t\geq 0}$ satisfying (a) $\pi_t > 0$ for $t \geq 0$, (b) $\mathbb{E}[\,\pi_t\,] < \infty$ for $t \geq 0$, and (c) $\liminf_{t\to\infty} \mathbb{E}[\pi_t] = 0$, with the property that if an asset with value process $\{S_t\}_{t\geq 0}$ delivers a bounded $\mathcal{F}_T$-measurable cash flow H_T at time T and derives its value entirely from that cash flow, then

$$S_t = \mathbb{1}_{\{t<T\}} \frac{1}{\pi_t} \mathbb{E}[\pi_T H_T | \mathcal{F}_t]. \tag{4}$$

Here we write $\mathbb{E}$ for expectation under $\mathbb{P}$. We adopt the convention that the value of such an asset drops to zero at the instant the cash flow occurs. Thus $\lim_{t\to T} S_t = H_T$, whereas $S_t = 0$ for $t \geq T$. This is in keeping with the usual analysis of stock prices when a stock goes ex-dividend, and respects the requirement that the price process should be càdlàg. Our assumptions imply that $\pi_T H_T$ is integrable and that $\{\pi_t S_t + \mathbb{1}_{\{t\geq T\}}\, \pi_T H_T\}_{t\geq 0}$ is a uniformly integrable martingale. The existence of an established pricing kernel is equivalent, in a broad sense, to what we mean by market equilibrium and the absence of arbitrage. In fact, it can be shown under rather general conditions [11] that with a few reasonable assumptions any pricing formula for contingent claims takes the form (4). Then if $H_T = 1$, we obtain an expression for the price at time t of a bond that pays one unit of currency at T, given by

$$P_{tT} = \mathbb{1}_{\{t<T\}} \frac{1}{\pi_t} \mathbb{E}[\pi_T | \mathcal{F}_t]. \tag{5}$$

To be sure, models of the type represented by formula (1) can be obtained as instances of models of the type represented by formula (5), but it is not the case that all interest rate models are of type (1). As we shall demonstrate, models of type (5) can be constructed for which the unit of conventional currency is replaced by a unit of cryptocurrency, and in such a way that we arrive at a nontrivial interest rate model for which the cryptocurrency condition $r_t = 0$ is satisfied for all $t \geq 0$ and yet for which term rates are non-vanishing.

The present paper explores a class of such models achieved by allowing the crypto pricing kernel to be a strict local martingale. The reasoning behind this proposal is as follows. The pricing kernel methodology requires that the price $\{S_t\}$ of a non-dividend-paying asset should have the property that the product $\{\pi_t S_t\}$ should be a martingale. Now, suppose the market admits a unit-initialized money market account with value process $\{B_t\}$ of the form

$$B_t = \exp\left(\int_0^t r_s \, \mathrm{d}s\right). \tag{6}$$

Then the process $\{\Lambda_t\}$ defined by $\Lambda_t = \pi_t B_t$ is a martingale, and the pricing kernel as a consequence is given by

$$\pi_t = \exp\left(-\int_0^t r_s \, \mathrm{d}s\right) \Lambda_t. \tag{7}$$

In that case, if we introduce the Radon-Nikodym derivative defined by

$$\left.\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right|_{\mathcal{F}_t} = \Lambda_t, \tag{8}$$

we can make a change the measure in (4), and we are led to the well-known risk-neutral valuation formula

$$S_t = \mathbb{1}_{\{t<T\}}\, \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\int_t^T r_s \, \mathrm{d}s \right) H_T \middle| \mathcal{F}_t \right]. \tag{9}$$

Then if we set $H_T = 1$ we recover the class of "conventional" interest rate models defined by a bond price of the form (1) along with the money market account (6).

In the class of models we consider for crypto interest rates, which will be introduced in Section 2, we exclude the existence of an instantaneous money-market asset from the model. This can be achieved by choosing the pricing kernel to be a strict local martingale, which implies that the short rate vanishes and hence that the process $\{B_t\}$ defined by (6) is constant. But if $\{B_t\}$ is constant, then $\{\pi_t B_t\}$ is not a martingale, and thus there is no money-market asset. We consider a model in which the pricing kernel is given by the reciprocal of a Bessel process of order three. This reciprocal process, introduced in [12], is a well-known example of a strict local martingale, and has the advantage of being highly tractable (chapter VI §3 of [13], [14–16]). The idea that this process can be used as a pricing kernel appears in §8.3.4 of [17], where it is recognized that the resulting interest rate model does not admit a representation for the bond price in the form (1). Here we develop a model of this type in detail in the context of cryptocurrency bonds. In Section 3 we derive explicit expressions for the discount bond system and the various associated rates. The results are applied in Section 4 to obtain pricing formulae for digital options on discount bonds and caplets on simple crypto rates. In Section 5, we introduce a class of related models based on Bessel(n) processes with $n \geq 4$, and the fourth order model is worked out in detail. In Section 6, we introduce a class of models based on a complexification of the Bessel(3) process. We conclude in Section 7 with some remarks about options on cryptocurrency exchange rates.

2. A model of no interest

The pricing kernel formalism allows us to identify where the theory of cryptocurrencies deviates from the conventional one: namely, there is no money market account. But how is it possible to construct an interest rate model without a money market account? This apparent impossibility in the context of a conventional interest rate model is nonetheless possible in a pricing kernel framework. The argument is as follows. First, we observe, by virtue of (5), that a necessary and sufficient condition for the bond price to be a decreasing function of T for any fixed t such that $t < T$ is that $\{\pi_t\}$ should be a supermartingale. This implies that the interest rate system associated with $\{\pi_t\}$ is non-negative. Now, cryptocurrencies are by their nature storable assets, with negligible storage costs, and hence by a standard arbitrage argument cannot be borrowed at a negative rate of interest. Thus it is reasonable to assume that the crypto pricing kernel is a supermartingale. We recall that a semimartingale is a local martingale if for any increasing sequence of stopping times $\{\tau_n\}_{n\in\mathbb{N}}$ with $\lim_{n\to\infty} \tau_n = \infty$, the stopped process is a martingale for each value of n. A strict local martingale is a local martingale that is not a true martingale. Now, it is well known that a positive local martingale is necessarily a supermartingale. Thus, if we admit the possibility that the pricing kernel is a strict local martingale, then the positivity of the pricing kernel implies that it is a supermartingale with vanishing drift. This suggests that the cryptocurrency interest-rate term structure can be modelled by letting the pricing kernel be a strict local martingale.

As an illustration, we examine an interest rate model based on a pricing kernel given by the reciprocal of the Bessel process of order three. Specifically, the pricing kernel is constructed as follows. Let $\{W_t^{(1)}, W_t^{(2)}, W_t^{(3)}\}_{t\geq 0}$ be three independent standard Brownian motions on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\{\mathcal{F}_t\}_{t\geq 0}$, and define

$$X_t = \int_0^t \sigma_s \, \mathrm{d}W_s^{(1)}, \quad Y_t = \int_0^t \sigma_s \, \mathrm{d}W_s^{(2)}, \quad Z_t = \int_0^t \sigma_s \, \mathrm{d}W_s^{(3)}, \tag{10}$$

where $\{\sigma_t\}_{t\geq 0}$ is a strictly positive, left-continuous, deterministic function such that $0 < \inf_t(\sigma_t) \leq \sup_t(\sigma_t) < \infty$. We then define a model for the pricing kernel by setting

$$\pi_t = \frac{1}{\sqrt{(X_t - a)^2 + (Y_t - b)^2 + (Z_t - c)^2}}\,, \tag{11}$$

where a, b, c are constants, not all equal to zero. The initial condition $\pi_0 = 1$ requires that $a^2 + b^2 + c^2 = 1$, and rotational symmetry implies that the vector (a, b, c) can lie on any point on the unit sphere in $\mathbb{R}^3$. We note that the function $u : \mathbb{R}^3 \to \mathbb{R}^+ \cup \infty$ defined by

$$u(x, y, z) = \frac{1}{\sqrt{(x - a)^2 + (y - b)^2 + (z - c)^2}} \tag{12}$$

is a solution of Laplace's equation

$$\frac{\partial^2}{\partial x^2} u(x, y, z) + \frac{\partial^2}{\partial y^2} u(x, y, z) + \frac{\partial^2}{\partial z^2} u(x, y, z) = 0 \tag{13}$$

on $\mathbb{R}^3 \backslash (a, b, c)$. One can think of $\{u(x, y, z)\}$ as the Coulomb potential generated by a unit charge situated at the point (a, b, c). Then $\pi_t = u(X_t, Y_t, Z_t)$, and we find that the pricing kernel satisfies the stochastic differential equation

$$\mathrm{d}\pi_t = -\sigma_t \, \pi_t^2 \, \mathrm{d}W_t\,, \tag{14}$$

with vanishing drift, where the process $\{W_t\}_{t\geq 0}$ defined by the relation

$$\mathrm{d}W_t = \pi_t \left[(X_t - a)\mathrm{d}W_t^{(1)} + (Y_t - b)\mathrm{d}W_t^{(2)} + (Z_t - c)\mathrm{d}W_t^{(3)} \right] \tag{15}$$

is a standard Brownian motion. It follows by a standard argument that we can define a localizing sequence $\{\tau_n\}$ by setting

$$\tau_n = \inf \left\{ t : \int_0^t \sigma_s^2 \, \pi_s^4 \, \mathrm{d}s \geq n \right\} \tag{16}$$

with the property that for each n the process $\{\pi_t^{(n)}\}_{t\geq 0}$ defined by $\pi_t^{(n)} = \pi_{t\wedge\tau_n}$ is a martingale, which establishes that $\{\pi_t\}$ is a local martingale. But a calculation shows that $\mathbb{E}\,[\pi_t] < \pi_0$ for all $t > 0$, which means that $\{\pi_t\}$ is a strict local martingale.

The interpretation of (14) is as follows. Suppose we consider a generic market model in which the pricing kernel is a strictly positive Ito process driven by a Brownian motion $\{W_t\}$ and satisfies a stochastic differential equation of the form

$$\mathrm{d}\pi_t = \alpha_t \, \pi_t \, \mathrm{d}t + \beta_t \, \pi_t \, \mathrm{d}W_t\,. \tag{17}$$

Let $\{S_t\}$ be the price of a non-dividend-paying asset driven by the same Brownian motion. Then $\pi_t S_t = M_t$ for some positive martingale $\{M_t\}$ driven by $\{W_t\}$. Let the dynamics of $\{M_t\}$ be given by

$$\mathrm{d}M_t = \nu_t M_t \, \mathrm{d}W_t \,. \tag{18}$$

A calculation using Ito's formula then shows that

$$\mathrm{d}S_t = [-\alpha_t - \beta_t(\nu_t - \beta_t)] \, S_t \, \mathrm{d}t + (\nu_t - \beta_t) \, S_t \, \mathrm{d}W_t \,. \tag{19}$$

Thus if one writes $r_t = -\alpha_t$, $\lambda_t = -\beta_t$ and $\gamma_t = \nu_t - \beta_t$, it follows that the stochastic differential equation of the asset takes the familiar form

$$\mathrm{d}S_t = (r_t + \lambda_t \gamma_t) \, S_t \, \mathrm{d}t + \gamma_t \, S_t \, \mathrm{d}W_t \,. \tag{20}$$

We see that r_t is the short rate of interest, λ_t is the market price of risk, γ_t is the volatility, and that the pricing kernel satisfies

$$\mathrm{d}\pi_t = -r_t \pi_t \, \mathrm{d}t - \lambda_t \pi_t \, \mathrm{d}W_t \,. \tag{21}$$

Combining (14) and (21), we deduce that $r_t = 0$ for all $t \geq 0$ in the Bessel(3) model for the pricing kernel, and the market price of risk is given by $\lambda_t = \sigma_t \pi_t$. The vanishing of the short rate does not, however, imply the vanishing of rates of finite tenor, such as Libor rates and swap rates, as we shall see in what follows.

3. Discount bonds and yields

To work out the bond price process we shall be using the pricing formula (5). Thus, we have

$$P_{tT} = \frac{1}{\pi_t} \mathbb{E}_t \left[\frac{1}{\sqrt{(X_T - a)^2 + (Y_T - b)^2 + (Z_T - c)^2}} \right], \tag{22}$$

where $\mathbb{E}_t[\,\cdot\,] := \mathbb{E}[\,\cdot\,|\,\mathcal{F}_t]$. Writing $X_T - a = (X_T - X_t) + (X_t - a)$, and similarly for $Y_T - b$ and $Z_T - c$, we observe that the increments $X_T - X_t$, $Y_T - Y_t$, and $Z_T - Z_t$ are independent of $\mathcal{F}_t$, whereas $X_t - a$, $Y_t - b$, and $Z_t - c$ are $\mathcal{F}_t$-measurable. Thus, if we define $X = X_T - a$, $Y = Y_T - b$, and $Z = Z_T - c$, then conditionally on $\mathcal{F}_t$ we have $X \sim N(X_t - a, \Sigma_{tT})$, $Y \sim N(Y_t - b, \Sigma_{tT})$, and $Z \sim N(Z_t - c, \Sigma_{tT})$, where

$$\Sigma_{tT} = \int_t^T \sigma_s^2 \, \mathrm{d}s \,. \tag{23}$$

If we define the vectors $\boldsymbol{R} = (x, y, z)$ and $\boldsymbol{\xi}_t = (X_t - a, Y_t - b, Z_t - c)$, and their squared norms $R^2 = \boldsymbol{R} \cdot \boldsymbol{R}$ and $\xi_t^2 = \boldsymbol{\xi}_t \cdot \boldsymbol{\xi}_t$, one can work out the conditional density for the distribution of X_T, Y_T, Z_T, and we find that the bond price is given by a Gaussian integral of the form

$$P_{tT} = \frac{1}{\pi_t} \frac{1}{(\sqrt{2\pi \Sigma_{tT}}\,)^3} \int_{\mathbb{R}^3} \frac{1}{R} \, \mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1} |\boldsymbol{R} - \boldsymbol{\xi}_t|^2} \mathrm{d}^3 \boldsymbol{R} \,. \tag{24}$$

where $R = \sqrt{x^2 + y^2 + z^2}$. It should be emphasized that (24) is obtained by using the specific structure of Σ_{tT}. Then, introducing the spherical representation

$$\mathrm{d}^3 \boldsymbol{R} = R^2 \sin\theta \, \mathrm{d}R \, \mathrm{d}\theta \, \mathrm{d}\phi \tag{25}$$

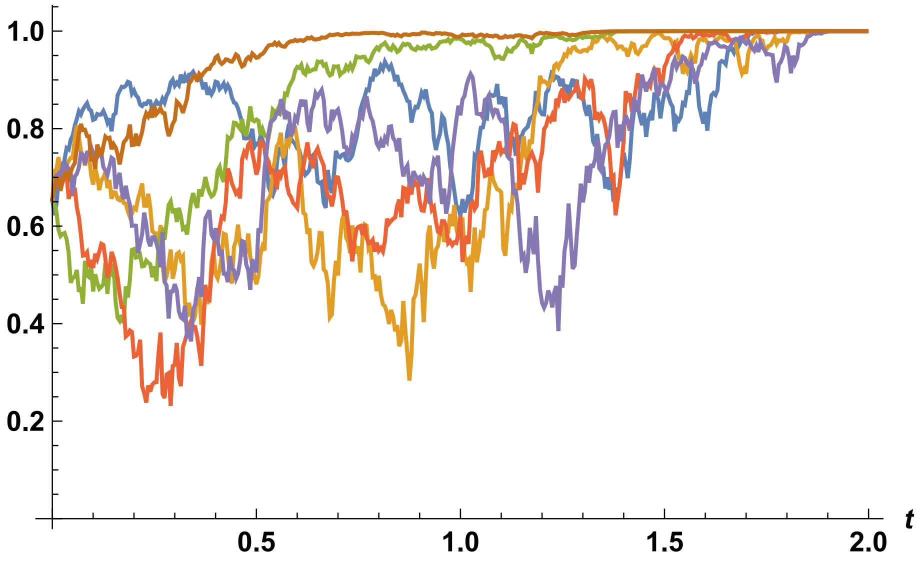

FIG. 1: *The discount bond price process for bond maturity* $T = 2$ *and volatility* $\sigma_t = 0.75$. Six sample paths are displayed, illustrating the qualitative behaviour of the bond price process in the Bessel(3) model.

for the volume element in $\mathbb{R}^3$ we deduce that

$$\begin{aligned} P_{tT} &= \frac{1}{\pi_t}\frac{1}{\sqrt{2\pi}\Sigma_{tT}^{3/2}}\int_0^\infty R^2\,\frac{1}{R}\int_0^\pi \sin\theta\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,(R^2-2R\xi_t\cos\theta+\xi_t^2)}\mathrm{d}\theta\,\mathrm{d}R \\ &= \frac{1}{\pi_t}\frac{1}{\sqrt{2\pi}\Sigma_{tT}^{3/2}}\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,\xi_t^2}\int_0^\infty R\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,R^2}\int_0^\pi \sin\theta\,\mathrm{e}^{R\xi_t\cos\theta/\Sigma_{tT}}\mathrm{d}\theta\,\mathrm{d}R \\ &= \frac{1}{\pi_t}\frac{1}{\sqrt{2\pi}\Sigma_{tT}^{3/2}}\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,\xi_t^2}\int_0^\infty R\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,R^2}\left[\left.\frac{-\mathrm{e}^{R\xi_t\cos\theta/\Sigma_{tT}}}{R\xi_t/\Sigma_{tT}}\right|_0^\pi\right]\mathrm{d}R \\ &= \frac{1}{\sqrt{2\pi\Sigma_{tT}}}\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,\xi_t^2}\int_0^\infty \mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,R^2}\left(\mathrm{e}^{R\xi_t/\Sigma_{tT}}-\mathrm{e}^{-R\xi_t/\Sigma_{tT}}\right)\mathrm{d}R\,, \end{aligned} \tag{26}$$

where

$$\xi_t = \sqrt{(X_t-a)^2+(Y_t-b)^2+(Z_t-c)^2}\,. \tag{27}$$

Completing the squares in the exponents in (26), we obtain

$$\begin{aligned} P_{tT} &= \frac{1}{\sqrt{2\pi\Sigma_{tT}}}\int_0^\infty\left(\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,(R-\xi_t)^2}-\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,(R+\xi_t)^2}\right)\mathrm{d}R \\ &= \frac{1}{\sqrt{\pi}}\int_{-\xi_t/\sqrt{2\Sigma_{tT}}}^{\xi_t/\sqrt{2\Sigma_{tT}}}\mathrm{e}^{-u^2}\mathrm{d}u\,. \end{aligned} \tag{28}$$

Let us now define the error function as usual by setting

$$\mathrm{erf}(z) = \frac{1}{\sqrt{\pi}}\int_{-z}^{z}\mathrm{e}^{-u^2}\mathrm{d}u = \frac{2}{\sqrt{\pi}}\sum_{n=0}^{\infty}\frac{(-1)^n}{n!(2n+1)}\,z^{2n+1}, \tag{29}$$

which is an entire function on the complex plane. In particular, we have

$$\mathrm{erf}(z) = N(\sqrt{2}\,z) - N(-\sqrt{2}\,z)\,, \tag{30}$$

where the normal distribution function is given by

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} \mathrm{e}^{-\frac{1}{2}z^2}\,\mathrm{d}z\,. \tag{31}$$

. Then we obtain the following expression for the bond price:

$$P_{tT} = \mathrm{erf}\left(\sqrt{\frac{(X_t-a)^2+(Y_t-b)^2+(Z_t-c)^2}{2\,\Sigma_{tT}}}\right). \tag{32}$$

Equivalently, we have

$$P_{tT} = \mathrm{erf}\left(\frac{\xi_t}{\sqrt{2\,\Sigma_{tT}}}\right), \tag{33}$$

where $\{\xi_t\}$ is the Bessel(3) process. For illustration, we have shown in Figure 1 some sample paths for the bond price. As the price process of an asset, the discount bond has the property that $\{\pi_t P_{tT}\}$ is a martingale, which follows at once from the tower property:

$$\mathbb{E}_s[\pi_t P_{tT}] = \mathbb{E}_s[\mathbb{E}_t[\pi_T]] = \mathbb{E}_s[\pi_T] = \pi_s P_{sT}\,. \tag{34}$$

Alternatively, the martingale condition can be checked directly from expression (32) for the bond price, if we make use of the identity

$$\frac{1}{\sqrt{\pi\alpha}}\int_0^\infty \left(\mathrm{e}^{-\frac{1}{\alpha}(\xi-x)^2} - \mathrm{e}^{-\frac{1}{\alpha}(\xi+x)^2}\right)\mathrm{erf}\left(\frac{\xi}{\sqrt{\beta}}\right)\mathrm{d}\xi = \mathrm{erf}\left(\frac{x}{\sqrt{\alpha+\beta}}\right). \tag{35}$$

We have seen that the model entails no interest since the drift of the pricing kernel is identically zero. Yet, bond prices give rise to discounting; in other words, P_{tT} is a strictly decreasing function of T for each $t < T$, since Σ_{tT} is a strictly increasing function of T for each $t < T$. One might therefore wonder whether the short rate vanishes if one employs the alternative definition of the short rate given by

$$r_t = -\lim_{t\to T}\frac{\partial P_{tT}}{\partial T}\,. \tag{36}$$

This can easily be checked. We have

$$r_t = \lim_{t\to T}\frac{\sigma_T^2}{\sqrt{2\pi}\Sigma_{tT}^{3/2}}\exp\left(-\frac{\xi_t^2}{2\Sigma_{tT}}\right). \tag{37}$$

Since $\lim_{t\to T}\Sigma_{tT} = 0$, the exponential term suppresses the right side to give $r_t = 0$ for all $t \geq 0$. Alternatively, by use of (32), a calculation shows that

$$\mathrm{d}P_{tT} = \lambda_t\Omega_{tT}P_{tT}\,\mathrm{d}t + \Omega_{tT}P_{tT}\,\mathrm{d}W_t\,, \tag{38}$$

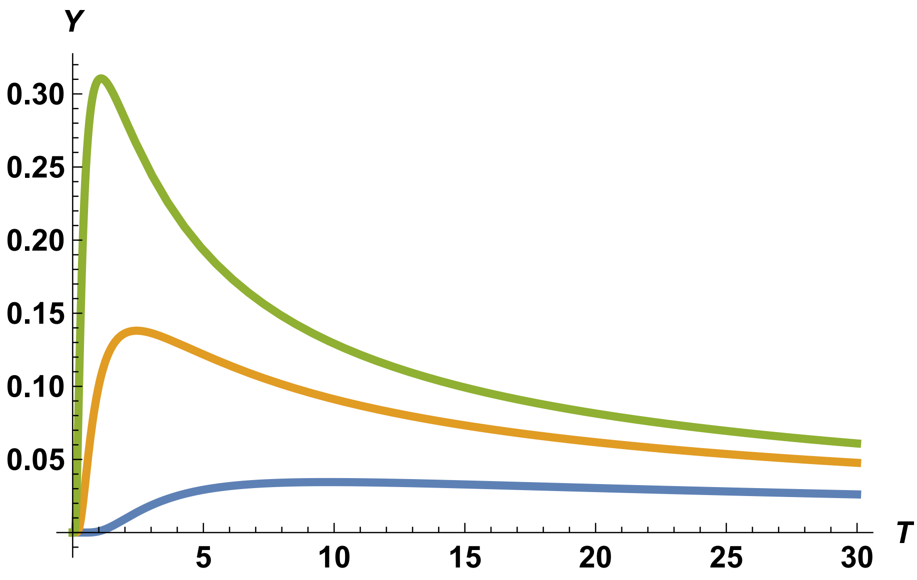

FIG. 2: *The initial yield curve for constant volatility.* Since $r_0 = 0$, the yield at $T = 0$ vanishes. For constant σ, the yield curve peaks and then decays to zero, where $Y_T \sim -\log(\sqrt{2/\pi\sigma T})/T$ as $T \to \infty$. We have plotted the yield for $\sigma = 0.3$ (blue), $\sigma = 0.6$ (orange) and $\sigma = 0.9$ (green).

where $\lambda_t = \sigma_t \pi_t$ is the market price of risk and

$$\Omega_{tT} = \frac{2\sigma_t}{P_{tT}\sqrt{2\pi\Sigma_{tT}}} \exp\left(-\frac{1}{2\Sigma_{tT}\pi_t^2}\right) \tag{39}$$

is the discount bond volatility. We observe that the contribution $r_t P_{tT}$ normally arising from the short rate in the drift is absent in (38), which confirms that the short rate indeed vanishes. We note also that $\lim_{t\to T}\Omega_{tT} = 0$, in accordance with the idea that the volatility of the bond should dwindle as the maturity date is approached.

But the fact that the short rate is zero does not imply that other rates are necessarily zero. For instance, as a consequence of the definition

$$f_{tT} = -\frac{\partial \log P_{tT}}{\partial T}, \tag{40}$$

a calculation shows that the instantaneous forward rates are given by

$$f_{tT} = \frac{\sigma_T^2 \, \exp\left(-\xi_t^2/2\Sigma_{tT}\right)}{\sqrt{2\pi}\Sigma_{tT}^{3/2} \, \operatorname{erf}\left(\xi_t/\sqrt{2\Sigma_{tT}}\right)}, \tag{41}$$

and we see that $\lim_{t\to T} f_{tT} = 0$. Similarly, for the initial yield curve $\{Y_T\}_{T\geq 0}$ we obtain

$$Y_T = -\frac{1}{T}\log\left[\operatorname{erf}\left(\frac{1}{\sqrt{2\Sigma_{0T}}}\right)\right]. \tag{42}$$

Here Y_T denotes the continuously compounded rate at time 0 for maturity T, defined as usual for all $T \geq 0$ by

$$P_{0T} = \exp\left[-T\,Y_T\right]. \tag{43}$$

Thus, initial yield curve data can be used to calibrate the freedom in the function $\{\sigma_t\}$. A typical set of yield curves arising from constant $\{\sigma_t\}$ is sketched in Figure 2. The calibration scheme that we have proposed is equivalent to that suggested in [17], section 8.3.4, using a time-change. Here, we can go a step further and derive a formula for $\{\sigma_t\}$ in terms of the initial discount function. Write erf$'$ for the derivative of erf, and erf^{-1} for the inverse of erf, so $\mathrm{erf}^{-1}(\mathrm{erf}(x)) = x$ for all $x \in \mathbb{R}$. Let P'_{0t} denote the derivative of the initial discount function at time t. Recall that if a function $f : \mathbb{R} \to \mathbb{R}$ is invertible, with inverse f^{-1}, and admits a continuous first derivative f', then for all $y \in \mathbb{R}$ it holds that

$$\frac{\mathrm{d}f^{-1}(y)}{\mathrm{d}y} = \frac{1}{f'(f^{-1}(y))} \,. \tag{44}$$

It then follows as a consequence of the relation

$$P_{0t} = \mathrm{erf}\left(\frac{\xi_t}{\sqrt{2\,\Sigma_{0t}}}\right) \tag{45}$$

that

$$(\sigma_t)^2 = \frac{-P'_{0t}}{\mathrm{erf}'(\mathrm{erf}^{-1}(P_{0t}))(\mathrm{erf}^{-1}(P_{0t}))^3} \,. \tag{46}$$

4. Bond options

We proceed to consider the pricing of options on discount bonds. To begin, we look at a European-style digital call with strike K and expiration t on a discount bond that matures at time T. Thus, the payout of an option that delivers one unit of cryptocurrency at time t in the event that $P_{tT} > K$ is the indicator function

$$H_t = \mathbb{1}\{P_{tT} > K\}, \tag{47}$$

and it follows that the price $D_0(K,t,T)$ of such an option at time zero is given by

$$D_0(K,t,T) = \mathbb{E}\left[\pi_t\,\mathbb{1}\{P_{tT} > K\}\right] = \mathbb{E}\left[\frac{1}{\xi_t}\mathbb{1}\left\{\mathrm{erf}\left(\frac{\xi_t}{\sqrt{2\Sigma_{tT}}}\right) > K\right\}\right], \tag{48}$$

where $\pi_t = \xi_t^{-1}$. Now, the error function is strictly increasing, so we find that there is a critical value ξ^* of the Bessel(3) process, given by

$$\xi^* = \sqrt{2\,\Sigma_{tT}}\ \mathrm{erf}^{-1}(K), \tag{49}$$

such that the option expires in-the-money if and only if $\xi_t > \xi^*$. Therefore, if we switch to a spherical representation for the volume element in $\mathbb{R}^3$, a calculation similar to that presented in (26) shows that

$$D_0(K,t,T) = \frac{1}{2}\left[\mathrm{erf}\left(\frac{\xi^*+1}{\sqrt{2\Sigma_{0t}}}\right) - \mathrm{erf}\left(\frac{\xi^*-1}{\sqrt{2\Sigma_{0t}}}\right)\right]. \tag{50}$$

More generally, we can consider the price process $\{D_s(K,t,T)\}_{0\leq s<t}$ of a digital call with strike K and expiration t on a discount bond with maturity T. This is given by

$$D_s(K,t,T) = \mathbb{1}_{\{0\leq s<t\}}\frac{1}{\pi_s}\mathbb{E}_s[\pi_t H_t]\,. \tag{51}$$

Noticing that conditional on $\mathcal{F}_s$ we have $X_t \sim N(X_s - a, \Sigma_{st})$, $Y_t \sim N(Y_s - b, \Sigma_{st})$ and $Z_t \sim N(Z_s - c, \Sigma_{st})$, one finds that a calculation analogous to that considered in (26) leads to the formula

$$D_s(K,t,T) = \frac{1}{2\xi_s}\left[\mathrm{erf}\left(\frac{\xi^* + \xi_s}{\sqrt{2\,\Sigma_{st}}}\right) - \mathrm{erf}\left(\frac{\xi^* - \xi_s}{\sqrt{2\,\Sigma_{st}}}\right)\right]. \tag{52}$$

We turn next to look at the pricing of an in-arrears caplet, for which the payout at time T can be taken to be of the form

$$H_T = X(L_{tT} - R)^+\,, \tag{53}$$

where X is the notional, R is the cap rate, and the term rate L_{tT} is defined by

$$L_{tT} = \frac{1}{T-t}\left(\frac{1}{P_{tT}} - 1\right). \tag{54}$$

As usual, we write $(x)^+$ for $\max(x,0)$. The value of such a caplet is given by

$$\kappa_0(R,t,T) = X\,\mathbb{E}\left[\pi_T\left(\frac{1}{T-t}\left(\frac{1}{P_{tT}} - 1\right) - R\right)^+\right]. \tag{55}$$

Since the caplet is paid "in arrears", meaning that the payoff is set at t and paid at T, and since L_{tT} is known at t, it follows by a familiar argument that the caplet is equivalent to a security that pays $H_t = P_{tT}H_T$ at the earlier time t, with

$$H_t = N\left(K - P_{tT}\right)^+\,, \tag{56}$$

where K and N are given by

$$K = \frac{1}{1+R(T-t)} \quad \text{and} \quad N = \frac{X[1+R(T-t)]}{T-t}\,. \tag{57}$$

Thus, as is well known, a position in an in-arrears caplet is equivalent to a position in N puts on a unit discount bond, where the strike K is the value of a discount bond with simple yield R, where R is the cap rate. Making use of (4) we deduce that

$$\begin{aligned}\kappa_0(R,t,T) &= N\,\mathbb{E}\left[\pi_t\left(K - P_{tT}\right)^+\right] \\ &= N\,\mathbb{E}\left[\frac{1}{\xi_t}\left(K - \mathrm{erf}\left(\frac{\xi_t}{\sqrt{2\Sigma_{tT}}}\right)\right)^+\right]. \end{aligned}\tag{58}$$

If we switch to the spherical representation for the volume element in $\mathbb{R}^3$, a calculation analogous to that presented in (26) shows that

$$\kappa_0(R,t,T) = NP_0(K,t,T)\,, \tag{59}$$

where the put price can be represented in terms of the following Gaussian integrals:

$$\begin{aligned} P_0(K,t,T) &= \frac{1}{\sqrt{2\pi\Sigma_{0t}}}\int_0^{\xi^*}\left[K-\mathrm{erf}\left(\frac{R}{\sqrt{2\Sigma_{tT}}}\right)\right]\left(\mathrm{e}^{-\frac{1}{2\Sigma_{0t}}(R-1)^2}-\mathrm{e}^{-\frac{1}{2\Sigma_{0t}}(R+1)^2}\right)\mathrm{d}R \\ &= \frac{1}{\sqrt{\pi}}\int_{\frac{1}{\sqrt{2\Sigma_{0t}}}}^{\frac{\xi^*+1}{\sqrt{2\Sigma_{tT}}}}\mathrm{e}^{-u^2}\mathrm{erf}\left(\tfrac{\sqrt{2\Sigma_{0t}}\,u-1}{\sqrt{2\Sigma_{tT}}}\right)\mathrm{d}u-\frac{1}{\sqrt{\pi}}\int_{\frac{-1}{\sqrt{2\Sigma_{0t}}}}^{\frac{\xi^*-1}{\sqrt{2\Sigma_{tT}}}}\mathrm{e}^{-u^2}\mathrm{erf}\left(\tfrac{\sqrt{2\Sigma_{0t}}\,u+1}{\sqrt{2\Sigma_{tT}}}\right)\mathrm{d}u \\ &\quad +K\left[\mathrm{erf}\left(\tfrac{1}{\sqrt{2\Sigma_{0t}}}\right)-\tfrac{1}{2}\left(\mathrm{erf}\left(\tfrac{\xi^*+1}{\sqrt{2\Sigma_{0t}}}\right)-\mathrm{erf}\left(\tfrac{\xi^*-1}{\sqrt{2\Sigma_{0t}}}\right)\right)\right]. \end{aligned} \tag{60}$$

While there appears to be no simpler representation for the Gaussian integrals appearing here, numerical evaluation is straightforward.

5. Models based on higher-order Bessel processes

Bessel processes of order four or more also give rise to cryptobond models, analogous to the one we have already investigated. In particular, if for $n \geq 3$ we consider a collection of Gaussian processes $\{X_t^k\}_{t\geq 0}$ for $k = 1, 2, \ldots, n$ of the type given by (10), then we can model the pricing kernel by setting

$$\pi_t = \left[(X_t^1)^2 + \cdots + (X_t^n)^2\right]^{(2-n)/2}. \tag{61}$$

A calculation shows that there exists a standard Brownian motion $\{W_t\}_{t\geq 0}$ such that

$$\mathrm{d}\pi_t = -(n-2)\,\sigma_t\,\pi_t^{(n-1)/(n-2)}\mathrm{d}W_t\,, \tag{62}$$

from which it follows on account of a test discussed in [18–20] that $\{\pi_t\}$ is a strict local martingale. In more detail, let the function $\beta : \mathbb{R}+ \to \mathbb{R}+$ be such that for all a, b satisfying $0 < a < b < \infty$ it holds that

$$0 < \inf_{a\leq x\leq b}\beta(x) < \sup_{a\leq x\leq b}\beta(x) < \infty\,. \tag{63}$$

Let $\{X_t\}$, with $X_0 = 1$, satisfy a stochastic differential equation of the form

$$\mathrm{d}X_t = \beta(X_t)\,\mathrm{d}W_t\,, \tag{64}$$

where the process is stopped when it first reaches the origin. Then the test states that $\{X_t\}$ is a strict local martingale if for some $c > 0$ it holds that

$$\int_c^\infty\left[\frac{x}{\beta^2(x)}\right]\mathrm{d}x < \infty\,, \tag{65}$$

and otherwise $\{X_t\}$ is a martingale. One can check by use of a time-change that if a deterministic time-dependent coefficient is included in (64) of the type one finds in (62) then a similar conclusion is applicable. A short calculation then shows that the test is passed for $n \geq 3$. Alternatively, for $n \geq 3$ one can construct a localizing sequence to confirm that if $\{\pi_t\}$ satisfies (62) then $\{\pi_t\}$ is a local martingale. A direct calculation allows one to verify that $\mathbb{E}[\pi_t] < \pi_0$ for all $t > 0$, and hence that $\{\pi_t\}$ is a strict local martingale.

As an illustration, we present a model for cryptocurrencies based on the reciprocal of the Bessel process in four dimensions. See [23] for properties of the Bessel(4) process. In our model we take

$$\pi_t = \frac{1}{(X_t^1 - a)^2 + (X_t^2 - b)^2 + (X_t^3 - c)^2 + (X_t^4 - d)^2}\,, \tag{66}$$

where the $\{X_t^k\}_{k=1,\dots,4}$ are independent Gaussian processes of the form (10), and the constants a, b, c, d are chosen such that $a^2 + b^2 + c^2 + d^2 = 1$. A calculation shows that the pricing kernel satisfies the stochastic differential equation

$$\mathrm{d}\pi_t = -2\,\sigma_t\,\pi_t^{3/2}\mathrm{d}W_t\,, \tag{67}$$

which corresponds to (62) for $n = 4$. We can work out the discount bond price in this model, which, following the logic of (24), with $\boldsymbol{\xi}_t = (X_t^1 - a, X_t^2 - b, X_t^3 - c, X_t^4 - d)$, is given by

$$P_{tT} = \frac{1}{\pi_t}\frac{1}{(2\pi\Sigma_{tT})^2}\int_{\mathbb{R}^4}\frac{1}{R^2}\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,|\boldsymbol{R}-\boldsymbol{\xi}_t|^2}\mathrm{d}^4\boldsymbol{R}\,. \tag{68}$$

We switch to a spherical representation. In four dimensions, we set $x = R\sin\theta\,\sin\varphi\,\cos\phi$, $y = R\sin\theta\,\sin\varphi\,\sin\phi$, $z = R\sin\theta\,\cos\varphi$, and $w = R\cos\theta$, with the volume element

$$\mathrm{d}^4\boldsymbol{R} = R^3\sin^2\theta\,\sin\varphi\,\mathrm{d}R\,\mathrm{d}\theta\,\mathrm{d}\varphi\,\mathrm{d}\phi\,. \tag{69}$$

Note that $\theta, \varphi \in [0, \pi]$ and $\phi \in [0, 2\pi]$. Since the vector $\boldsymbol{\xi}_t$ is fixed, and because of the spherical symmetry, we may without loss of generality choose $\boldsymbol{\xi}_t$ to be in the direction of the w-axis. A similar assumption was made in the three-dimensional case in (26), where $\boldsymbol{\xi}_t$ was taken to be in the z-direction. Then we have

$$\boldsymbol{R}\cdot\boldsymbol{\xi}_t = R\xi_t\cos\theta, \tag{70}$$

a choice that simplifies the calculation somewhat. Integration over ϕ gives 2π, whereas

$$\int_0^\pi \sin\varphi\,\mathrm{d}\varphi = 2, \tag{71}$$

so after performing the integration over these variables we obtain

$$P_{tT} = \frac{1}{\pi_t}\frac{1}{\pi\Sigma_{tT}^2}\int_0^\infty R\int_0^\pi \sin^2\theta\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,(R^2-2R\xi_t\cos\theta+\xi_t^2)}\mathrm{d}\theta\,\mathrm{d}R\,, \tag{72}$$

where $\{\xi_t\}$ represents the Bessel process in four dimensions, so $\pi_t = \xi_t^{-2}$. To proceed, we note the identity

$$\int_0^\pi \sin^2\theta\,\mathrm{e}^{\nu\cos\theta}\mathrm{d}\theta = \frac{\pi}{\nu}\,I_1(\nu). \tag{73}$$

This follows since $\sin^2\theta\,\mathrm{e}^{\nu\cos\theta} = (\sin\theta)(\sin\theta\,\mathrm{e}^{\nu\cos\theta})$ and $\sin\theta\,\mathrm{e}^{\nu\cos\theta} = -\nu^{-1}\partial_\theta\,\mathrm{e}^{\nu\cos\theta}$, which shows that we can integrate by parts to reduce the integrand to $\cos\theta\,\mathrm{e}^{\nu\cos\theta}$. But then we notice that $\cos\theta\,\mathrm{e}^{\nu\cos\theta} = \partial_\nu\,\mathrm{e}^{\nu\cos\theta}$, so moving ∂_ν outside the integration we see that the integrand reduces further to $\mathrm{e}^{\nu\cos\theta}$. But this gives rise to a Bessel function, and we have

$$\int_0^\pi \mathrm{e}^{\nu\cos\theta}\mathrm{d}\theta = \pi\,I_0(\nu). \tag{74}$$

Differentiating and using the differential identity $\partial_\nu I_0(\nu) = I_1(\nu)$, we arrive at the conclusion. Alternatively, if we recall the definition

$$I_n(\nu) = \frac{1}{\pi}\int_0^\pi \mathrm{e}^{\nu\cos\theta}\cos(n\theta)\mathrm{d}\theta \tag{75}$$

for the generalized Bessel function of the first kind, we arrive at the same conclusion more expediently. In any case, we deduce that

$$\int_0^\pi \sin^2\theta\, \mathrm{e}^{R\,\xi_t\,\Sigma_{tT}^{-1}\,\cos\theta}\mathrm{d}\theta = \frac{\pi\Sigma_{tT}}{R\xi_t}\, I_1(R\xi_t/\Sigma_{tT})\,. \tag{76}$$

Thus using $\pi_t = \xi_t^{-2}$ we obtain

$$P_{tT} = \frac{\xi_t}{\Sigma_{tT}}\int_0^\infty \mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,(R^2+\xi_t^2)} I_1(R\xi_t/\Sigma_{tT})\,\mathrm{d}R\,. \tag{77}$$

If we set $u = R/\sqrt{\Sigma_{tT}}$ and $\eta_t = \xi_t/\sqrt{\Sigma_{tT}}$, the expression simplifies to

$$P_{tT} = \eta_t\,\mathrm{e}^{-\frac{1}{2}\eta_t^2}\int_0^\infty \mathrm{e}^{-\frac{1}{2}u^2} I_1(\eta_t u)\,\mathrm{d}u\,. \tag{78}$$

Now we use the identity

$$\int_0^\infty \mathrm{e}^{-\frac{1}{2}u^2} I_1(\eta u)\,\mathrm{d}u = \frac{1}{\eta}\left(\mathrm{e}^{\frac{1}{2}\eta^2} - 1\right), \tag{79}$$

which can be established by use of the Taylor series expansion of the Bessel function

$$I_n(\nu) = \sum_{k=0}^\infty \frac{1}{2^{2k+n}\,k!\,\Gamma(n+k+1)}\,\nu^{2k+n} \tag{80}$$

along with the expression

$$\int_0^\infty \mathrm{e}^{-\frac{1}{2}u^2} u^{2k+n}\mathrm{d}u = 2^{(2k+n-1)/2}\,\Gamma\left(\frac{2k+n+1}{2}\right) \tag{81}$$

for the Gaussian moments. Specifically, substituting (80) for $n = 1$ and $\nu = \eta u$ in the left side of (79) and using (81) for $n = 1$, we obtain

$$\begin{aligned}\int_0^\infty \mathrm{e}^{-\frac{1}{2}u^2} I_1(\eta u)\,\mathrm{d}u &= \sum_{k=0}^\infty \frac{\eta^{2k+1}}{2^{2k+1}\,k!\,\Gamma(k+2)}\,2^k\Gamma(k+1) \\ &= \sum_{k=0}^\infty \frac{(\eta/2)^{k+1}}{(k+1)!} \\ &= \frac{1}{\eta}\left(\sum_{k=0}^\infty \frac{(\eta/2)^k}{k!} - 1\right),\end{aligned} \tag{82}$$

and this establishes (79). Putting these together we arrive at the bond price

$$P_{tT} = 1 - \exp\left(-\frac{1}{2}\frac{\xi_t^2}{\Sigma_{tT}}\right), \tag{83}$$

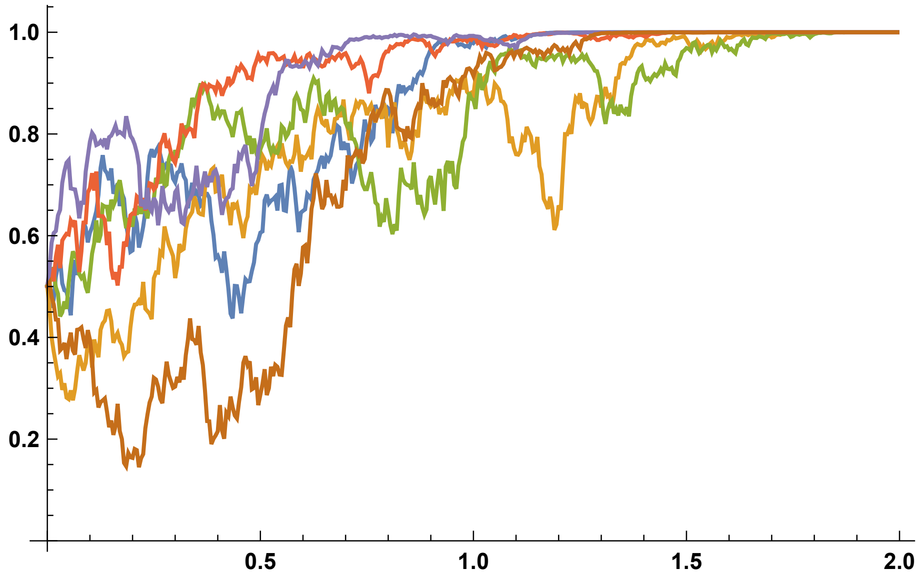

FIG. 3: *The discount bond price process with* $\sigma_t = 0.6$ *and* $T = 2$. Six sample paths are displayed, illustrating the qualitative behaviour of the bond price in a model based on the Bessel(4) process.

which turns out to be surprisingly simple.

For illustration we have shown in Figure 3 some sample paths for the bond price process. Note that $\lim_{t\to T} P_{tT} = 1$, whereas, assuming that $\sigma_t > 0$ for all $t \geq 0$, we have $\lim_{T\to\infty} \Sigma_{tT} = \infty$, from which it follows that for fixed t we have $\lim_{T\to\infty} P_{tT} \to 0$. The initial bond price is

$$P_{0T} = 1 - \exp\left(-\frac{1}{2}\frac{1}{\Sigma_{0T}}\right), \tag{84}$$

from which we deduce that the initial yield curve takes the form

$$Y_T = -\frac{1}{T}\log\left[1 - \exp\left(-\frac{1}{2}\frac{1}{\Sigma_{0T}}\right)\right]. \tag{85}$$

This relation can be used to calibrate the volatility function to market data. Specifically, we find that

$$\sigma_t^2 = -\frac{1}{2}\frac{(Y_t + t\,Y_t')\,\mathrm{e}^{-tY_t}}{(1-\mathrm{e}^{-tY_t})\,(\log(1-\mathrm{e}^{-tY_t}))^2}. \tag{86}$$

Next we examine the dynamics of the bond price. If we start with

$$\mathrm{d}\xi_t = \frac{3\sigma_t^2}{2\xi_t}\,\mathrm{d}t + \sigma_t\,\mathrm{d}W_t\,, \tag{87}$$

an application of Ito's formula gives

$$\mathrm{d}P_{tT} = \lambda_t\Omega_{tT}P_{tT}\,\mathrm{d}t + \Omega_{tT}P_{tT}\,\mathrm{d}W_t\,, \tag{88}$$

where

$$\lambda_t = 2\sigma_t\xi_t^{-1} \qquad \text{and} \qquad \Omega_{tT} = \frac{\sigma_t\xi_t}{P_{tT}\Sigma_{tT}}\,\mathrm{e}^{-\frac{1}{2}\Sigma_{tT}^{-1}\,\xi_t^2}. \tag{89}$$

This result confirms the fact that $r_t = 0$ for all $t \geq 0$ in the present model.

Let us now consider the valuation of a call option on a discount bond. The payoff takes the form

$$H_t = (P_{tT} - K)^+, \tag{90}$$

where K is the strike, t is the expiration date of the option, and $T > t$ is the maturity date of the bond. We assume that $0 < K < 1$. Since the bond price is an increasing function of ξ_t, we find that there is a critical value

$$\xi^* = \sqrt{-2\,\Sigma_{tT}\log(1-K)} \tag{91}$$

such that $H_t = 0$ if $\xi_t \leq \xi^*$. After we perform the integration over the (ϕ, φ) variables, we find that the price of the option is given by the integral

$$C_0(K,t,T) = \frac{1}{\pi\Sigma_{0t}^2}\int_{\xi^*}^{\infty} R\left((1-K) - \mathrm{e}^{-\frac{1}{2\Sigma_{tT}}R^2}\right)\mathrm{e}^{-\frac{1}{2\Sigma_{0t}}(R^2+1)}\int_0^{\pi}\sin^2\theta\,\mathrm{e}^{\frac{R}{\Sigma_{0t}}\cos\theta}\mathrm{d}\theta\,\mathrm{d}R\,. \tag{92}$$

Performing the θ integration, we thus have

$$C_0(K,t,T) = \frac{1}{\Sigma_{0t}}\int_{\xi^*}^{\infty}\left((1-K) - \mathrm{e}^{-\frac{1}{2\Sigma_{tT}}R^2}\right) I_1\left(\frac{R}{\Sigma_{0t}}\right)\mathrm{e}^{-\frac{1}{2\Sigma_{0t}}(R^2+1)}\,\mathrm{d}R\,. \tag{93}$$

Similarly, for a put option with payout

$$H_t = (K - P_{tT})^+, \tag{94}$$

we obtain

$$P_0(K,t,T) = \frac{1}{\Sigma_{0t}}\int_0^{\xi^*}\left((K-1) + \mathrm{e}^{-\frac{1}{2\Sigma_{tT}}R^2}\right) I_1\left(\frac{R}{\Sigma_{0t}}\right)\mathrm{e}^{-\frac{1}{2\Sigma_{0t}}(R^2+1)}\,\mathrm{d}R\,, \tag{95}$$

from which we deduce that

$$C_0(K,t,T) - P_0(K,t,T) = \frac{1}{\Sigma_{0t}}\int_0^{\infty}\left((K-1) + \mathrm{e}^{-\frac{1}{2\Sigma_{tT}}R^2}\right) I_1\left(\frac{R}{\Sigma_{0t}}\right)\mathrm{e}^{-\frac{1}{2\Sigma_{0t}}(R^2+1)}\,\mathrm{d}R\,. \tag{96}$$

Using (79) we can integrate the right side of (96) explicitly to obtain the put-call parity relation:

$$\begin{aligned} C_0(K,t,T) - P_0(K,t,T) &= (1-K)\left(1 - \mathrm{e}^{-\frac{1}{2\Sigma_{0t}}}\right) + \mathrm{e}^{-\frac{1}{2\Sigma_{0t}}} - \mathrm{e}^{-\frac{1}{2\Sigma_{0t}}\left[1-\frac{\Sigma_{tT}}{\Sigma_{0T}}\right]} \\ &= P_{0T} - KP_{0t}\,, \end{aligned} \tag{97}$$

where we have made use of the fact that $\Sigma_{0t} + \Sigma_{tT} = \Sigma_{0T}$.

The indefinite Gaussian integrals of the Bessel function for the option prices have to be evaluated numerically. It is interesting to note that despite the simplicity of the model for the bond price, the option price cannot be expressed in closed form in terms of known functions. Nevertheless, numerical valuation is straightforward. To see this, we use the Taylor series expansion (80) for the Bessel function to obtain

$$C_0(K,t,T) = \frac{1}{\Sigma_{0t}} \mathrm{e}^{-\frac{1}{2\Sigma_{0t}}} \sum_{k=0}^{\infty} \frac{(1/2\Sigma_{0t})^{2k+1}}{k!(k+1)!} \int_{\xi^*}^{\infty} R^{2k+1} \left[(1-K) - \mathrm{e}^{-\frac{1}{2\Sigma_{tT}}R^2} \right] \mathrm{e}^{-\frac{1}{2\Sigma_{0t}}R^2} \mathrm{d}R\,. \tag{98}$$

Then, changing the integration variable by setting $u = R^2$, we find that the integration reduces to that of an incomplete gamma function

$$\Gamma(a,z) = \int_z^{\infty} u^{a-1} \mathrm{e}^{-u} \mathrm{d}u, \tag{99}$$

and we thus obtain the option price in the form of a series:

$$\begin{aligned} C_0(K,t,T) \;=\; & \mathrm{e}^{-\frac{1}{2\Sigma_{0t}}} \sum_{k=0}^{\infty} \frac{(1/2\Sigma_{0t})^{k+1}}{k!(k+1)!} \Bigg[(1-K)\,\Gamma\left(k+1, -\frac{\Sigma_{tT}}{\Sigma_{0t}} \log(1-K)\right) \\ & - \left(\frac{\Sigma_{tT}}{\Sigma_{0T}}\right)^{k+1} \Gamma\left(k+1, -\frac{\Sigma_{0T}}{\Sigma_{0t}} \log(1-K)\right) \Bigg]. \end{aligned} \tag{100}$$

On account of the appearance of the double factorial in the denominator in the summand in the expression above, the series converges rapidly, making it a useful expression for numerical valuation of the option price.

In particular, truncating the sum at, say, $k = 20$, we can obtain option prices very rapidly by use of standard numerical tools: the difference of the result thus obtained and the result of a standard numerical valuation of the integral (98) is of the order 10^{-16}.

More generally, in higher dimensions it should be evident that by use of the spherical representation for calculating the expectation $\mathbb{E}_t[\pi_T]$ it is always possible to set the direction of the vector $\boldsymbol{\xi}_t$ such that $\boldsymbol{R}\cdot\boldsymbol{\xi}_t = R\,\xi_t\,\cos\theta$. Thus the only nontrivial integration concerns the variables θ and R. Performing the θ integration we arrive at a linear combination of Bessel functions if the dimension n of the Bessel process is even, which then has to be integrated with respect to a Gaussian measure to obtain an expression for the bond price; whereas if n is odd, the θ integration gives rise to a linear combination of exponential functions, which again has to be integrated with respect to a Gaussian measure to obtain an expression for the bond price. Thus, depending on whether n is even or odd, for $n \geq 3$ we obtain two different types of cryptocurrency interest rate models.

Finally, one might wonder what the situation is in the case of a Bessel process with $n = 2$. Following the notation that we have been using in the previous examples, one sets

$$\xi_t = \sqrt{(X_t - a)^2 + (Y_t - b)^2}\,. \tag{101}$$

Then we find that the process $\{\pi_t\}_{t\geq 0}$ defined by $\pi_t = \log \xi_t$ is indeed a strict local martingale: first one introduces a localizing sequence to prove that it is a local martingale, and then one shows that $\mathbb{E}[\pi_t] > \pi_0$. However, in this case $\{\pi_t\}$ can take negative as well as positive values, so it is unsuitable as a model for a pricing kernel in our scheme. In fact, for $n = 2$ we find that $\{\pi_t\}$ is a submartingale.

6. Complex extensions of the model

The model associated with the reciprocal of the Bessel process in three dimensions can be extended in an alternative manner to allow for parametric degrees of freedom to be incorporated. This can be achieved if we allow the parameters a, b and c appearing in (11) to be complex numbers. The real part of the resulting complexified process $\{\pi_t^{\mathbb{C}}\}$ then defines an admissible model for the pricing kernel, with vanishing short rate. The reason for this is that when the parameters a, b, and c are complex, then both real and imaginary parts of the function $u(x, y, z)$ satisfy Laplace's equation, and the real part is strictly positive. The additional freedom thus arising can be used, for instance, to calibrate the model not only against the yield curve but also against option prices.

To proceed, let us therefore write $a = a_0 + \mathrm{i}a_1$, $b = b_0 + \mathrm{i}b_1$ and $c = c_0 + \mathrm{i}c_1$. Additionally, let us write $\tilde{x} = x - a_0$, $\tilde{y} = y - b_0$ and $\tilde{z} = z - c_0$. Then we have

$$u(x,y,z) = \sqrt{\frac{\tilde{x}^2 + \tilde{y}^2 + \tilde{z}^2 - a_1^2 - b_1^2 - c_1^2 + 2\mathrm{i}(\tilde{x}a_1 + \tilde{y}b_1 + \tilde{z}c_1)}{\left(\tilde{x}^2 + \tilde{y}^2 + \tilde{z}^2 - a_1^2 - b_1^2 - c_1^2\right)^2 + 4\left(\tilde{x}a_1 + \tilde{y}b_1 + \tilde{z}c_1\right)^2}}\,. \tag{102}$$

This function is a solution of Laplace's equation away from the "ring" singularity defined by the intersection of the two-sphere Σ of radius $\sqrt{a_1^2 + b_1^2 + c_1^2}$ centred at the point (a_0, b_0, c_0) and the two-plane Π defined by $a_1 x + b_1 y + c_1 z = a_0 a_1 + b_0 b_1 + c_0 c_1$ which passes through the point (a_0, b_0, c_0) and hence cuts Σ in an equatorial circle. We recall the formula

$$\sqrt{A + \mathrm{i}B} = \sqrt{\frac{A + \sqrt{A^2 + B^2}}{2}} + \mathrm{i}\frac{B}{|B|}\sqrt{\frac{-A + \sqrt{A^2 + B^2}}{2}} \tag{103}$$

for the real and the imaginary parts of the principal square-root of a complex number for which $B \neq 0$. It follows that $\mathrm{Re}\,(u) > 0$ on $\mathbb{R}^3 \backslash \{\Sigma \cap \Pi\}$, whereas $\mathrm{Im}\,(u) = 0$ on $\Pi \backslash \{\Sigma \cap \Pi\}$. With these results at hand, we introduce a new crypto-rate model by setting

$$\pi_t = \mathrm{Re}\,\left(\pi_t^{\mathbb{C}}\right). \tag{104}$$

Then writing $\tilde{X}_t = X_t - a_0$, $\tilde{Y}_t = Y_t - b_0$, and $\tilde{Z}_t = Z_t - c_0$, we obtain the following expression:

$$\pi_t = \sqrt{\frac{\tilde{X}_t^2 + \tilde{Y}_t^2 + \tilde{Z}_t^2 - a_1^2 - b_1^2 - c_1^2 + \sqrt{\left(\tilde{X}_t^2 + \tilde{Y}_t^2 + \tilde{Z}_t^2 - a_1^2 - b_1^2 - c_1^2\right)^2 + 4\left(\tilde{X}_t a_1 + \tilde{Y}_t b_1 + \tilde{Z}_t c_1\right)^2}}{2\left[\left(\tilde{X}_t^2 + \tilde{Y}_t^2 + \tilde{Z}_t^2 - a_1^2 - b_1^2 - c_1^2\right)^2 + 4\left(\tilde{X}_t a_1 + \tilde{Y}_t b_1 + \tilde{Z}_t c_1\right)^2\right]}}\,. \tag{105}$$

The normalization $\pi_0 = 1$ imposes one constraint, whereas rotational symmetry can be used to eliminate two further parameters. Thus we are left with a model with three exogenously specifiable parameters that can be used to fit option prices.

To obtain an expression for the bond price, we need to work out the conditional expectation $\mathbb{E}_t[\pi_T]$. Rather than using expression (105) for the pricing kernel, which makes the computation cumbersome, we can take advantage of the fact that

$$\mathbb{E}_t\left[\mathrm{Re}\left(\pi_t^{\mathbb{C}}\right)\right] = \mathrm{Re}\left(\mathbb{E}_t\left[\pi_T^{\mathbb{C}}\right]\right). \tag{106}$$

Then we can use the simpler formula (11) for the pricing kernel, with complex parameters a, b, and c, and calculate its expectation, taking the real part of the result. Using the spherical representation, we have

$$\mathbb{E}_t[\pi_T^{\mathbb{C}}] = \frac{1}{(\sqrt{2\pi\Sigma_{tT}})^3}\int_{\mathbb{R}^3}\frac{1}{R}\,\mathrm{e}^{-\frac{1}{2\Sigma_{tT}}|\boldsymbol{R}-(\boldsymbol{\xi}_t-\mathrm{i}\boldsymbol{\delta})|^2}R^2\sin\theta\,\mathrm{d}R\,\mathrm{d}\theta\,\mathrm{d}\phi, \tag{107}$$

where $\boldsymbol{\xi}_t = (X_t - a_0, Y_t - b_0, Z_t - c_0)$ and $\boldsymbol{\delta} = (a_1, b_1, c_1)$.

It turns out that the calculation leading to (32) is applicable for complex parameters a, b, and c. To see this, we perform the integration explicitly. Recall that in the real case where $\boldsymbol{\delta} = \mathbf{0}$ we have one fixed vector $\boldsymbol{\xi}_t$ in the exponent of the integrand, so by using the spherical symmetry we choose this vector to point in the z direction, resulting in the simple expression $\boldsymbol{R}\cdot\boldsymbol{\xi}_t = R\,\xi_t\cos\theta$, which was used in the calculation of (32).

In the present case, we have two fixed vectors $\boldsymbol{\xi}_t$ and $\boldsymbol{\delta}$, so we can use the rotational symmetry to let the two vectors lie on the x-y plane, symmetrically placed about the x-axis. We let $2\alpha_t$ denote the angle between the two vectors $\boldsymbol{\xi}_t$ and $\boldsymbol{\delta}$. In other words, we have $\boldsymbol{\xi}_t\cdot\boldsymbol{\delta} = \xi_t\,\delta\cos(2\alpha_t)$, where $\xi_t^2 = \boldsymbol{\xi}_t\cdot\boldsymbol{\xi}_t$ and $\delta^2 = \boldsymbol{\delta}\cdot\boldsymbol{\delta}$. Thus, the angle between $\boldsymbol{\xi}_t$ and the x-axis is α_t, and similarly the angle between $\boldsymbol{\delta}$ and the x-axis is $-\alpha_t$. With this choice of coordinates we have

$$\boldsymbol{R}\cdot(\boldsymbol{\xi}_t - \mathrm{i}\boldsymbol{\delta}) = \big(R(\xi_t-\mathrm{i}\delta)\sin\theta\cos\alpha_t\big)\cos\phi + \big(R(\xi_t+\mathrm{i}\delta)\sin\theta\sin\alpha_t\big)\sin\phi. \tag{108}$$

We are now in a position to perform the integration over the variable ϕ. To this end we recall the identity

$$\int_0^{2\pi}\mathrm{e}^{p\cos\phi+q\sin\phi}\mathrm{d}\phi = 2\pi I_0\left(\sqrt{p^2+q^2}\right). \tag{109}$$

This can be seen by viewing the exponent of the integrand as an inner product between the vector (p,q) and the unit vector placed at an angle ϕ from the vector (p,q). Then the exponent is equivalent to $\sqrt{p^2+q^2}\cos\phi$, and the result follows. In the present case we have $p = R(\xi_t - \mathrm{i}\delta)\sin\theta\cos\alpha_t$ and $q = R(\xi_t+\mathrm{i}\delta)\sin\theta\sin\alpha_t$, so $p^2+q^2 = R^2\omega_t^2\sin^2\theta$, where

$$\omega_t^2 = |(\boldsymbol{\xi}_t - \mathrm{i}\boldsymbol{\delta})|^2 = \xi_t^2 - \delta^2 - 2\mathrm{i}\xi_t\delta\cos(2\alpha_t). \tag{110}$$

We thus deduce that

$$\mathbb{E}_t\left[\pi_T^{\mathbb{C}}\right] = \frac{2\pi}{(\sqrt{2\pi\Sigma_{tT}})^3}\int_0^\infty\int_0^\pi R\,\mathrm{e}^{-\frac{1}{2\Sigma_{tT}}(R^2+\omega_t^2)}\sin\theta I_0\left(\frac{R\omega_t}{\Sigma_{tT}}\sin\theta\right)\mathrm{d}\theta\mathrm{d}R. \tag{111}$$

To perform the integration over the variable θ we note from (80) that

$$I_0(\nu\sin\theta) = \sum_{k=0}^\infty\frac{\nu^{2k}}{2^{2k}(k!)^2}(\sin\theta)^{2k}. \tag{112}$$

Thus, since

$$\int_0^\pi(\sin\theta)^{2k+1}\mathrm{d}\theta = \frac{2^{2k+1}(k!)^2}{(2k+1)!}\,, \tag{113}$$

and taking into account the Taylor series expansion

$$\frac{2\sinh(\nu)}{\nu} = 2\sum_{k=0}^{\infty} \frac{\nu^{2k}}{(2k+1)!}\,, \tag{114}$$

we deduce the identity

$$\int_0^{\pi} \sin\theta\, I_0(\nu\sin\theta)\mathrm{d}\theta = \frac{1}{\nu}\left(\mathrm{e}^{\nu} - \mathrm{e}^{-\nu}\right), \tag{115}$$

from which it follows that

$$\mathbb{E}_t\left[\pi_T^{\mathbb{C}}\right] = \frac{\omega_t^{-1}}{\sqrt{2\pi\Sigma_{tT}}}\int_0^{\infty} \mathrm{e}^{-\frac{1}{2\Sigma_{tT}}(R^2+\omega_t^2)}\left(\mathrm{e}^{\frac{R\omega_t}{\Sigma_{tT}}} - \mathrm{e}^{-\frac{R\omega_t}{\Sigma_{tT}}}\right)\mathrm{d}R = \omega_t^{-1}\,\mathrm{erf}\left(\frac{\omega_t}{\sqrt{2\Sigma_{tT}}}\right). \tag{116}$$

Noting that $\pi_t^{\mathbb{C}} = \omega_t^{-1}$ we thus validate the claim that the calculation leading to (32) is applicable for complex parameters a, b, and c. In particular, for the bond price we have

$$P_{tT} = \frac{1}{\mathrm{Re}(\omega_t^{-1})}\,\mathrm{Re}\left(\omega_t^{-1}\,\mathrm{erf}\left(\frac{\omega_t}{\sqrt{2\Sigma_{tT}}}\right)\right), \tag{117}$$

where ω_t is defined by (110). It should be apparent that in the real case for which $\boldsymbol{\delta} = \mathbf{0}$, we recover from (117) the previous expression (32) for the bond price.

Hence, by complexification of a model based on the Bessel(3) process we obtain a genuine parametric extension of the resulting term structure model. The complexification method that we have applied here is reminiscent of an analogous technique that has been used in physical applications [21, 22].

7. Discussion

The notion that strict local martingales should play a role in finance has been considered in various contexts by a number of authors. One can mention, in particular, the so-called benchmark approach of Platen and his collaborators, the Föllmer-Jarrow-Protter theory of price bubbles, and the Cox-Hobson analysis of option pricing in markets with bubbles as examples that have attracted considerable attention [16, 23–28]. In the present paper, we have put forward an altogether different proposal for the application of local martingales in the theory of finance — namely, the idea that strict local martingales can be used as a basis for modelling the pricing kernel in a cryptocurrency economy where there is no money market account. The familiar rules of risk-neutral pricing no longer apply, since the money market account is not available to act as a numeraire. Nevertheless, as we have shown, the existence of a market price of risk is sufficient to ensure nontrivial discounting, despite the vanishing of the short rate. Our approach to cryptocurrency interest rates has been developed in some detail in models based on Bessel processes of order three and order four. These models have the advantage that explicit formulae, or semi-explicit expressions involving Gaussian integrals, can be obtained for the prices of a variety of derivative contracts.

More generally, one can envisage a market admitting numerous decentralized currencies. Now, in a friction-free market with n cryptocurrencies, if we write S_t^{ij} $(i, j = 1, 2, 3, \ldots, n)$ for the price at time t of one unit of currency i quoted in units of currency j, then we have

$$S_t^{ij} = \pi_t^i / \pi_t^j\,, \tag{118}$$

where $\{\pi_t^i\}$ denotes the pricing kernel for assets denominated in units of currency i [29, 30]. In the present modelling framework, we consider a situation in which the pricing kernels each take the form (11), with initial values

$$\pi_0^i = \frac{1}{\sqrt{(a_i^2 + b_i^2 + c_i^2)}}\,, \tag{119}$$

and respective $\{\sigma_t^i\}$ functions. Some of the Brownian motions are shared throughout the crypto economy, representing systematic risk, while others may apply, perhaps, only to one or two cryptocurrencies, representing idiosyncratic risk. A call option on the cryptocurrency exchange rate with strike K and maturity T has the payout

$$H_T = (\pi_T^i/\pi_T^j - K)^+\,, \tag{120}$$

and hence the value in units of currency j of a call option on currency i is given by

$$C_0^{ij}(K,T) = \frac{1}{\pi_0^j}\,\mathbb{E}\left[(\pi_T^i - K\pi_T^j)^+\right], \tag{121}$$

which can easily be computed numerically on account of the Gaussian nature of the setup.

In a similar vein, we can examine the problem of pricing a call option on the exchange rate between cryptocurrency i and a sovereign currency, say, USD. Then the dollar price of an option to purchase one unit of cryptocurrency i at the strike price K at time T is

$$C_0^{i\$}(K,T) = \frac{1}{\pi_0^\$}\,\mathbb{E}\left[(\pi_T^i - K\pi_T^\$)^+\right]. \tag{122}$$

As an illustration, suppose that we have a geometric Brownian motion model

$$\pi_t^\$ = \pi_0^\$\,\mathrm{e}^{-rt-\lambda B_t - \frac{1}{2}\lambda^2 t} \tag{123}$$

for the dollar pricing kernel, where the short rate r is constant. For the cryptocurrency (say, bitcoin), we consider a pricing kernel of the form (11) with initial value

$$\pi_0^{฿} = \frac{1}{\sqrt{(a^2 + b^2 + c^2)}}\,, \tag{124}$$

and such that $\{B_t\}$ and $\{W_t\}$ are independent. The option price is then given by

$$C_0^{฿\$}(K,T) = \frac{1}{\pi_0^\$}\,\mathbb{E}\left[(\pi_T^{฿} - K\pi_T^\$)^+\right], \tag{125}$$

which is nonzero only if

$$B_T > -\frac{1}{\lambda}\left[\log\left(\pi_T^{฿}/K\pi_0^\$\right) + rT + \frac{1}{2}\lambda^2 T\right]. \tag{126}$$

A straightforward calculation then shows that the option price is given by

$$C_0^{฿\$}(K,T) = \frac{1}{\pi_0^\$}\,\mathbb{E}\left[\pi_T^{฿} N(g_+) - K\mathrm{e}^{-rT}\pi_0^\$ N(g_-)\right], \tag{127}$$

where

$$g_{\pm} = \frac{\log\left(\pi_T^{฿}/K\pi_0^{\$}\right) + rT \pm \frac{1}{2}\lambda^2 T}{\lambda\sqrt{T}}\,. \tag{128}$$

Hence, the cryptocurrency exchange-rate option prices can easily be computed numerically. Although for simplicity we have taken the dollar term structure to be exponential with a constant short rate, it is straightforward to extend the model to allow for calibration to the initial dollar term structure, while preserving the overall tractability of the results. Likewise one can work out the prices of options on exchange-rate futures contracts. In this respect, our analysis of options on cryptocurrency exchange rates can be contrasted with the pioneering work of Madan, Reyners & Schoutens [31], where the dollar and bitcoin interest rates are taken to be constant (in fact, zero). With the development of interest rate models applicable to bitcoin and other cryptocurrencies, financial institutions will be in a position to trade in cryptocurrency interest rate products. It could be that distributed ledger technologies will eventually find a way of rewarding the holders of positions in virtual currencies with interest on a continuous basis. In the meantime, there should be a role for models of no interest.

Acknowledgements The authors wish to express their gratitude to J. R. Boland, M. Kecman, A. Rafailidis and seminar participants at Tokyo Metropolitan University for helpful comments. We are also grateful for the feedback of two anonymous reviewers.

[1] Nakamoto, S. (2008) Bitcoin: A peer-to-peer electronic cash system. Bitcoin white paper, https://bitcoin.org.

[2] Bank of International Settlements (2018) BIS Statistical Bulletin, December 2018, Monetary and Economic Department. www.bis.org/statistics/bulletin1812.htm.

[3] Hughston, L. P., ed. (1996) *Vasicek and Beyond* (London: Risk Publications).

[4] Hughston, L. P., ed. (2000) *The New Interest Rate Models* (London: Risk Publications).

[5] Heath, D., Jarrow, R. & Morton, A. (1992) Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica* **60**, 77-105.

[6] Constantinides, G. M. (1992) A theory of the nominal term structure of interest rates. *Review of Financial Studies* **5**, 531-552.

[7] Flesaker, B. & Hughston, L. P. (1996) Positive interest. *Risk* **9** (1), 46-49.

[8] Brody, D. C. & Hughston, L. P. (2004) Chaos and coherence: a new framework for interest rate modelling. *Proceedings of the Royal Society London* A**460**, 85-110.

[9] Hughston, L. P. & Rafailidis, A. (2005) A chaotic approach to interest rate modelling. *Finance and Stochastics* **9**, 43-65.

[10] Brody, D. C. & Hughston, L. P. (2018) Social discounting and the long rate of interest. *Mathematical Finance* **28**, 306-334.

[11] Jobert, A. & Rogers, L. C. G. (2002) Valuations and dynamic convex risk measures. *Mathematical Finance* **18**, 1-22.

[12] Johnson, G. & Helms, L. L. (1963) Class D supermartingales. *Bulletin of the American Mathematical Society* **6**9, 59-62.

[13] Revuz, D. & Yor, M. (1999) *Continuous Martingales and Brownian Motion*, third edition (Berlin: Springer).

[14] Madan, D. B. & Yor, M. (2006) Ito's integrated formula for strict local martingales. In: M. Émery & M. Yor (eds.) *In Memoriam Paul-André Meyer*, Séminaire de probabilités XXXIX (Berlin: Springer).

[15] Pal, S. & Protter, P. (2010) Analysis of continuous strict local martingales via h-transforms. *Stochastic Processes and Their Applications* **120**, 1424-1443.

[16] Föllmer, H. & Protter, P. (2011) Local martingales and filtration shrinkage. *ESAIM: PS* **15**, S25-S38.

[17] Hunt, P. J. & Kennedy, J. E. (2004) *Financial Derivatives in Theory and Practice*, revised edition (Chichester: John Wiley & Sons).

[18] Delbaen, F. & Shirakawa, H. (2002) No arbitrage condition for positive diffusion processes. *Asia-Pacific Financial Markets* **9**, 159-168.

[19] Blei, S. & Engelbert H. J. (2009) On exponential local martingales associated with strong Markov continuous local martingales. *Stochastic Processes and Their Applications* **119**, 2859-2880.

[20] Mijatović, A. & Urusov, M. (2012) The martingale property of certain local martingales. *Probability Theory and Related Fields* **152**, 1-30.

[21] Newman, E. T., Couch, E., Chinnapared, K., Exton, A., Prakash, A. & Torrence, R. J. (1965) Metric of a rotating charged mass. *J. Math. Phys.* **6** (6), 918-919.

[22] Synge, J. L. (1972) A special class of solutions of the Schrödinger equation for a free particle. *Foundations of Physics*, **2** (1), 35-40.

[23] Platen, E. & Heath, D. (2009) *A Benchmark Approach to Quantitative Finance* (Berlin: Springer).

[24] Jarrow, R. & Protter, P. (2009) Forward and futures prices with bubbles. *International Journal of Theoretical and Applied Finance* **12**, 901-924.

[25] Jarrow, R. (2015) Asset price bubbles. *Annual Review of Financial Economics* **7**, 201-218.

[26] Jarrow, R. (2016) Bubbles and multiple-factor asset pricing models. *International Journal of Theoretical and Applied Finance* **19**, 1650007.

[27] Jarrow, R. & Protter, P. (2019) A rational asset pricing model for premiums and discounts on closed-end funds: the bubble theory. *Mathematical Finance* **29**, 1151-1170.

[28] Cox, A. M. G. & Hobson, D. (2005) Local martingales, bubbles and option prices. *Finance and Stochastics* **9**, 477-492.

[29] Flesaker, B & Hughston, L. P. (1997) International models for interest rates and foreign exchange. *Net Exposure*, **3**, 55-79. Reprinted in: L. P. Hughston (ed.) *The New Interest Rate Models*, London: Risk Publications (2000), 217-235.

[30] Rogers, L. C. G. (1997) The potential approach to the term structure of interest rates and foreign exchange rates. *Mathematical Finance* **7**, 157-176.

[31] Madan, D. B., Reyners, S. & Schoutens, W. (2019) Advanced model calibration on bitcoin options. *Digital Finance*, https://doi.org/10.1007/s42521-019-00002-1.

Special Sessions in memory of Prof. Tomas Bjork

Time Inconsistency and Stochastic Optimal Control

Agatha Murgoci

Abstract: Whenever we need to make an optimal decision in a dynamic context, the standard toolkit consists of the Bellman optimality principle and dynamic programming. However, the toolkit fails in a number of important economic situations. If we have an explicit dependence on the initial point of our problem, either for the state variable or for the exact time, the Bellman optimality principle fails. Tomas Björk was one of the main actors proposing a game theoretic approach to time inconsistency. In this talk, we review the contributions of Tomas and also bring forward some recent developments in the field.

Special Sessions in memory of Prof. Tomas Bjork

Time-inconsistent stopping

Speaker: Kristoffer Lindensj¨o

Abstract: Standard Markovian optimal stopping problems are consistent in the sense that the optimal stopping rule "stop the first time that the state process enters the stopping region", is independent of the starting value of the state process. Problems without this consistency property are known as time-inconsistent. We have developed a game-theoretic framework for time-inconsistent stopping problems, in the spirit of Tomas Bjork's and his co-authors' research on time-inconsistent stochastic control.

Special Sessions *in memory of* Tomas Björk

11th World Congress of the Bachelier Finance Society 2022

Organisers

Raquel M. Gaspar

Mariana Khapko

Tomas Björk: the man behind the name

– On his career and beyond –

Session 1: PhD Students and Co-authors

Agatha Murgoci

Time Inconsistency and Stochastic Optimal Control

Abstract: Whenever we need to make an optimal decision in a dynamic context, the standard toolkit consists of the Bellman optimality principle and dynamic programming. However, the toolkit fails in a number of important economic situations. If we have an explicit dependence on the initial point of our problem, either for the state variable or for the exact time, the Bellman optimality principle fails. Tomas Björk was one of the main actors proposing a game theoretic approach to time inconsistency. In this talk, we review the contributions of Tomas and also bring forward some recent developments in the field.

Kristoffer Lindensjö

Time-inconsistent stopping

Abstract: Standard Markovian optimal stopping problems are consistent in the sense that the optimal stopping rule "stop the first time that the state process enters the stopping region", is independent of the starting value of the state process. Problems without this consistency property are known as time-inconsistent. We have developed a game-theoretic framework for time-inconsistent stopping problems, in the spirit of Tomas Björk's and his co-authors' research on time-inconsistent stochastic control.

Francesca Biagini

Reduced form framework under model uncertainty

Abstract:In this talk we present a reduced form framework under model uncertainty. In particular we introduce a sublinear conditional operator with respect to a family of possibly nondominated probability measures in presence of a single or multiple ordered default times, and use it for the valuation of credit portfolio derivatives under model uncertainty. In addition, we extend this framework to include mortality intensities following an affine process under parameter uncertainty.
This talk is based on joint work with several co-authors: A. Mazzon, K.Oberpriller, and Y. Zhang.

Wolfgang Runggaldier

Recalling some of Tomas Björk's contributions to the mathematical theory of interest rates

Abstract: Tomas Björk may well be considered as one of the main actors in the development of the mathematical theory of the term structure of interest rates. I have had the privilege of working with him on this topic from very early onwards. In this talk I try to recall some of the main issues in this area that Tomas had been working on, partly jointly with me, and what were some of his contributions there.

Damir Filipovic

Stripping the discount curve - a robust machine learning approach

Abstract: We introduce a robust, flexible and easy-to-implement method for estimating the yield curve from Treasury securities. This method is non-parametric and optimally learns basis functions in reproducing Hilbert spaces with an economically motivated smoothness reward. We provide a closed-form solution of our machine learning estimator as a simple kernel ridge regression, which is straightforward and fast to implement. We show in an extensive empirical study on U.S. Treasury securities, that our method strongly dominates all parametric and non-parametric benchmarks. Our method achieves substantially smaller out-of-sample yield and pricing errors, while being robust to outliers and data selection choices. We attribute the superior performance to the optimal trade-off between flexibility and smoothness, which positions our method as the new standard for yield curve estimation.
Joint with M.Pelger and Y. Ye.

Ali Lazrak

Voting and decentralized trading of payment promises

Abstract: We examine a model where a binary decision is made by a committee based on voting. Committee members disagree on the decision but they can freely make credible and enforceable payment promises contingent on the committee decision. The payment promises are zero-sum but can involve coalitions of any size ranging from a pair to the entire committee. We show that trading payment promises enable the committee to choose the socially optimal decision and remove any inefficiencies resulting from voting externalities. Equilibrium payment promises are indeterminate even when restricted to generate a minimal total payment. When the voting outcome is inefficient, payment promises restore efficiency because intense minorities compensate the members of the majority to change their vote. When the voting outcome is efficient, payment promises can be required to preempt coalitions to form, get a better payoff and yet lead the group to adopt the inefficient decision.
Joint work with Jianfeng Zhang.

Eckhard Platen

Principles for modeling long-term market dynamics

Abstract: The paper derives eight principles that allow explaining the long-term dynamics of large stock markets, the typical distribution of the market capitalization of stocks, the risk premia for stock portfolios, the key role of optimal portfolios at the growth efficient frontier, the least expensive pricing and hedging of long-term payoffs, and other market features. By applying the concepts of entropy maximization and energy conservation in a richer modeling world than typically considered, most of these market properties follow rather directly. Furthermore, popular fundamental tools for portfolio and risk management, including the intertemporal capital asset pricing model and the preferred pricing rule, become revised.

Walter Schachermayer

Martingale transport, DeMarch-Touzi paving, and stretched Brownian motion in R^d

Abstract: In the classical optimal transport, the contributions of Benamou-Breiner and McCann regarding the time-dependent version of the problem are cornerstones of the filed and form the basis for a variety of applications in other mathematical areas. For μ, ν probability measures on R^d, increasing in convex order, stretched Brownian motion, Backho-Veraguas et al. provide an analogue for the martingale version of this problem. In dimension $d = 1$ it was shown in Beiglböck and Juillet that any martingale transport decomposes into at most countably many invariant intervals and that this decomposition is universal. Extensions for $d \geq 2$ have been studied DeMarch and Touzi and others. We show that the dual optimization problems attached to a stretched Brownian motion induces the universal DeMarch-Touzi paving of R^d.
Joint work with M.Beiglböck, J. Backhoff, and B. Tschiderer.

Josef Teichmann

Geometry of term structures

Abstract: Tomas Björk was a leading researcher in introducing fine geometric reasonings to the analysis of term structure equations. We review some of his main contributions including some personal thoughts about his admired didactical approaches and we show some recent developments in the field.

Schedule Restrictions:

1) Wolfgang Runggaldier cannot on Monday
2) Damir Filipovic cannot Monday, Tuesday nor Thursday

Suggestion: Is it possible to do it Wednesday or Friday?

Toward A General Framework for Modelling Roll-Over Risk

Mesias Alfeus[a]
[a] Department of Statistics and Actuarial Science
Stellenbosch University
mesias@sun.ac.za

Abstract

Quantitative Finance underwent significant development over the past decade. For example, modelling of the term structure of interest rates after the GFC poses a unique challenge. The persistent phenomena of market basis spreads are an indication that markets are pricing various risks which are not captured in classical models. We pioneer a roll-over risk modelling framework to provide empirical evidence to the observed basis spreads, i.e., a spread between LIBOR of different tenors and LIBOR-OIS spread. This roll-over risk consists of two components, a credit risk component due to the possibility of being downgraded and thus facing a higher credit spread when attempting to roll over short–term borrowing, and a component reflecting the (systemic) possibility of being unable to roll over short–term borrowing at the reference rate (e.g., LIBOR) due to an absence of liquidity in the market. The modelling framework is of "reduced form" in the sense that the source of credit risk is not modelled (nor is the source of liquidity risk). We show how such model can be calibrated to market data, and used for relative pricing of interest rate derivatives, including bespoke tenor frequencies not liquidly traded in the market.

Keywords: Roll–Over Risk, Basis Spread, LIBOR–OIS, Funding Liquidity risk, Multiple Tenors

Category: **Mathematics of Finance**

References

[1] **Alfeus, M., Grasselli, M. and E. Schlögl** (2020). A Consistent Stochastic Model of the Term Structure of Interest Rates for Multiple Tenors. *Journal of Economic Dynamics and Control*, Elsevier, vol. 114(C), `https://doi.org/10.1016/j.jedc.2020.103861`

[2] **Backwell, A., Macrina, A., Schlögl, E. Skovmand, D.** (2021), Term Rates, Multicurve Term Structures and Overnight Rate Benchmarks: a Roll-Over Risk Approach, `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3399680`

[3] **Grbac, Z. and W. J. Runggaldier** (2015). Interest Rate Modeling: Post–Crisis Challenges and Approaches. *Springer Briefs in Quantitative Finance*

Trading constraints in continuous-time Kyle models

Jin Hyuk Choi [1] , Heeyoung Kwon [1] and Kasper Larsen [2]

1) *Mathematical Sciences, Ulsan National Institute of Science and Technology (UNIST), Ulsan, South Korea*

2) *Department of Mathematics, Rutgers University, New Brunswick, USA*

Corresponding Author : Kasper Larsen, KL756@math.rutgers.edu

ABSTRACT

In a continuous-time Kyle setting, we prove global existence of an equilibrium when the insider faces a terminal trading constraint. We prove that our equilibrium model produces output consistent with several empirical stylized facts such as autocorrelated aggregate holdings, decreasing price impacts over the trading day, and U shaped optimal trading patterns.

REFERENCES

1. Almgren, R. (2003): *Optimal execution with nonlinear impact functions and trading-enhanced risk*, Applied Mathematical Finance **10**, 1–18.
2. Almgren, R., and N. Chriss (1999): *Value under liquidation*, Risk **12**, 61–63.
3. Almgren, R., and N. Chriss (2000): *Optimal execution of portfolio transactions*, Journal of Risk **3**, 5–39.
4. Back, K. (1992): *Insider trading in continuous time*, Review of Financial Studies, **5**, 387–409.
5. Back, K., and S. Baruch (2004): *Information in securities markets: Kyle meets Glosten and Milgrom*, Econometrica **72**, 433–465.
6. Back, K., P. Collin-Dufresne, V. Fos, T. Li, and A. Ljungqvist (2018): *Activism, Strategic trading, and liquidity*, Econometrica **86**, 1431–1463.
7. Barardehi, Y. H., and D. Bernhardt (2021): *Uncovering the impacts of endogenous liquidity consumption in intraday trading patterns*, working paper.
8. Barger, W., and R. Donnelly (2020): *Insider trading with temporary price impact*, to appear in International Journal of Theoretical and Applied Finance.
9. Baruch, S. (2002): *Insider trading and risk aversion*, Journal of Financial Markets **5**, 451–464
10. Basak, S., and D. Cuoco (1998): *An equilibrium model with restricted stock market participation*, Review of Financial Studies **11**, 309-341.
11. Bouchard, B., M. Fukasawa, M. Herdegen, and J. Muhle-Karbe (2018): *Equilibrium returns with transaction costs*, Finance and Stochastics **22**, 569–601.
12. Brunnermeier, M. K., and L. H. Pedersen (2005): *Predatory trading*, Journal of Finance **60**, 1825–1863.

13. Caldentey, R., and E. Stacchetti (2010): *Insider trading with a random deadline*, Econometrica **7**8, 245–283.
14. Carlin, B., M. Lobo, and S. Viswanathan (2007): Episodic liquidity crises: Cooperative and predatory trading. Journal of Finance, 62, 2235–2274.
15. Çetin, U. (2018): *Financial equilibrium with asymmetric information and random horizon*, Finance and Stochastics **22**, 97–126.
16. Çetin, U., and A. Danilova (2016): *Markovian Nash equilibrium in financial markets with asymmetric information and related forward–backward systems*, Annals of Applied Probability **26**, 1996–2029.
17. Çetin, U., and A. Danilova (2020): *On pricing rules and optimal strategies in general Kyle-Back models*, to appear in SIAM J. Control and Optimization.
18. Cho, K. H. (2003): *Continuous auctions and insider trading: uniqueness and risk aversion*, Finance and Stochastics **7**, 47–71.
19. Choi, J. H., K. Larsen, and D. Seppi (2019): *Information and trading targets in a dynamic market equilibrium*, Journal of Financial Economics **132**, 22–49.
20. Choi, J. H., K. Larsen, and D. Seppi (2021): *Equilibrium effects of intraday order-splitting benchmarks*, Mathematics and Financial Economics **15**, 315–352.
21. Collin-Dufresne, P. and F. Vyacheslav (2016): *Insider trading, stochastic liquidity, and equilibrium prices*, Econometrica **84**, 1441–1475.
22. Cuoco, D., and H. He (1994): *Dynamic equilibrium in infinite-dimensional economies with incomplete financial markets*, working paper.
23. Degryse, H., F. de Jong, and V. van Kervel (2014): *Does order splitting signal uninformed order flow?*, working paper.
24. Gârleanu, N., and L. H. Pedersen (2016): *Dynamic portfolio choice with frictions*, Journal of Economic Theory **165**, 487–516.
25. Hugonnier, J. (2012): *Rational asset pricing bubbles and portfolio constraints*. Journal of Economic Theory **147**, 2260-2302.
26. van Kervel, V., A. Kwan, and P. Westerholm (2020): *Order splitting and interacting with a counterparty*, working paper.
27. Kyle, A. (1985): *Continuous auctions and insider trading*, Econometrica **53**, 1315–1336.
28. Lipster, R. S., and A. N. Shiryaev (2001): *Statistics of random processes I*, Springer.
29. Prieto, R. (2010): *Dynamic equilibrium with heterogeneous agents and risk constraints*, working paper.
30. Protter, P. (2005): *Stochastic integration and differential equations*, Springer (2.1 Ed.).
31. Schied, A., and T. Schöneborn (2009): *Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets*, Finance and Stochastics **13**, 181–204.

Trading Signals In VIX Futures*

M. Avellaneda†, T. N. Lǐ‡, A. Papanicolaou§, G. Wáng¶

December 9, 2021

Abstract

We propose a new approach for trading VIX futures. We assume that the term structure of VIX futures follows a Markov model. Our trading strategy selects a position in VIX futures by maximizing the expected utility for a day-ahead horizon given the current shape and level of the term structure. Computationally, we model the functional dependence between the VIX futures curve, the VIX futures positions, and the expected utility as a deep neural network with five hidden layers. Out-of-sample backtests of the VIX futures trading strategy suggest that this approach gives rise to reasonable portfolio performance, and to positions in which the investor will be either long or short VIX futures contracts depending on the market environment.

Keywords: Contango, Cross Validation, Deep Learning, Feedforward Neural Networks, Trading Signals, VIX Futures.

AMS Subject Codes: 62P05, 68T05, 91B28.

*The authors would like to thank Brian Healy and Xunyang Wu for their feedback and support.

†Department of Mathematics, New York University. 251 Mercer Street, New York, NY, 10012. *avellane@cims.nyu.edu*

‡Department of Mathematics, New York University. 251 Mercer Street, New York, NY, 10012. *thomli@cims.nyu.edu*

§Department of Mathematics, North Carolina State University. Campus Box 8205, Raleigh, NC 27695. *apapani@ncsu.edu*. The author is partially supported by NSF grant DMS-1907518.

¶Department of Mathematics, Columbia University. 2990 Broadway New York, NY 10027. *gw2376@columbia.edu*

[1] Here, "quickly" means relative to other curves such as crude oil or treasuries.

Special Sessions in memory of Prof. Mark H.A. Davis

Trading, Pricing, and Hedging with Transaction Costs

Johannes Muhle-Karbe

Abstract: In this talk, we review some of the fundamental contributions Mark Davis has made to optimal trading, option pricing, and hedging in models that account for transaction costs. We then discuss some of the developments in the recent literature that build on these pioneering achievements.

Universal Approximation of Path Space Functionals

Christa Cuchiero* Patrick Mijatovic†
Philipp Schmocker‡ Josef Teichmann§

Abstract

The universal approximation property of neural networks on Euclidean spaces was proven for different families of functions. However, in many cases these classical results cannot be directly extended to infinite dimensional spaces, which are needed e.g. for applications in functional data analysis. For this purpose, we introduce so-called *functional neural networks* (FNNs) defined on an infinite dimensional weighted space and show their universal approximation property in the space of smooth non-anticipative functionals introduced in [5], [3] and [4]. This novel machine learning technique can then be used within the framework of Stochastic Portfolio Theory (SPT) summarized in [6] to learn an optimal portfolio, which outperforms the market portfolio of the S&P500 universe in a pathwise setting.

Topic Areas:

- Machine Learning including reinforcement learning and deep learning
- Artificial Intelligence in Finance

Summary

We introduce the generalization of neural networks to infinite dimensional spaces, which we shall call *functional neural networks (FNN)*, and can therefore be counted to the machine learning methods of supervised learning (see [9]). Neural networks on Euclidean spaces were originally discovered by Warren McCulloch and Walter Pitts in their seminal work [8] and mimick the functionality of a human brain. In order to prove the universal approximation property of functional neural networks defined on an infinite dimensional space, we rely on classical approximation theory including Stone-Weierstrass theorems and results of Nachbin type, see [13] and [10]. Inspired by Leopoldo Nachbin's weighted

*Vienna University, Department of Statistics and Operations Research, Data Science @ Uni Vienna, Kolingasse 14-16, 1090 Vienna, Austria, christa.cuchiero@univie.ac.at

†ETH Zürich, Departement of Mathematics, Rämistrasse 101, 8092 Zurich, Switzerland, mipatric@student.ethz.ch

‡Nanyang Technological University, Division of Mathematical Sciences, 21 Nanyang Link, Singapore 637371, Singapore, philippt001@e.ntu.edu.sg

§ETH Zürich, Departement of Mathematics, Rämistrasse 101, 8092 Zurich, Switzerland, jteichma@math.ethz.ch

The first author gratefully acknowledges financial support through grant Y 1235 of the START-program.

approximation problem in [11], we use the structure of a weighted space to obtain global universal approximation results beyond compact input spaces. Similar generalizations of neural networks to infinite dimensional spaces have been found in [2] and [12], and recently in [7] by using feature maps and in [1] for functionals on a Fréchet space. Moreover, we show that smooth functional neural networks are not only able to approximate the values of a given functional, but also its derivatives, which can be used to approximate a smooth non-anticipative functional introduced in [5], [3] and [4].

Within the framework of Stochastic Portfolio Theory (SPT), we apply the classical universal approximation result and the infinite dimensional extension to learn an optimal portfolio, which outperforms the market portfolio of the S&P500 universe in a pathwise setting. More precisely, we consider for a neural network $\varphi^{\mathrm{NNGP}} : \overline{\Delta}^d \to \mathbb{R}$ the *neural network-generated portfolio (NNGP)* in the sense of [6, Chapter 3] with portfolio weights

$$\pi_i^{\mathrm{NNGP}}(t) = \left(D_i \varphi^{\mathrm{NNGP}}(\mu(t)) + 1 - \sum_{j=1}^{d} \mu_j(t) D_j \varphi^{\mathrm{NNGP}}(\mu(t)) \right) \mu_i(t),$$

for $t \in [0,T]$ and $i = 1, ..., d$, where D_i denotes the ith partial derivative, and where $\mu(t) \in \overline{\Delta}^d := \{x \in [0,1]^d;\, x_1 + ... + x_d = 1\}$ represents the vector of market weights. On the other hand, we define for a vector-valued neural network $\varphi : \overline{\Delta}^d \to \mathbb{R}^d$ the *neural network portfolios (NNPs)* with portfolio weights given by

$$\pi_i^{\mathrm{NNP}}(t) = \left(\varphi_i^{\mathrm{NNP}}(\mu(t)) + 1 - \sum_{j=1}^{d} \mu_j(t) \varphi_j^{\mathrm{NNP}}(\mu(t)) \right) \mu_i(t)$$

for $t \in [0,T]$ and $i = 1, ..., d$. Then, by learning an optimal NNGP and NNP in the stock market universe of the S&P500 with respect to the logarithmic return $\log\left(\frac{Z_\pi(T)}{Z_\mu(T)}\right)$, we observe that the NNGP does not perform significantly worse than the NNP, even though the NNP class is much larger than the set of NNGPs.

Moreover, the setting can be extended to the path-dependent (non-Markovian) case by using non-anticipative functionals, where the neural networks depends on the past trajectory of the market weights. Hereby, we use the universal approximation result on path spaces to learn an optimal portfolio with the help of a functional neural network.

References

[1] F. E. Benth, N. Detering, and L. Galimberti. Neural networks in Fréchet spaces, 2021.

[2] T. Chen and H. Chen. Approximation capability to functions of several variables, nonlinear functionals, and operators by radial basis function neural networks. *IEEE Transactions on Neural Networks*, 6(4):904–910, 1995.

[3] R. Cont and D.-A. Fournié. Change of variable formulas for non-anticipative functionals on path space. *Journal of Functional Analysis*, 259(4):1043–1072, 2010.

[4] R. Cont and D.-A. Fournié. Functional Itô calculus and stochastic integral representation of martingales. *Annals of Probability*, 41(1):109–133, 01 2013.

[5] B. Dupire. Functional Itô calculus. Technical report, Bloomberg, 2009. Bloomberg Portfolio Research Paper No. 2009-04-FRONTIERS.

[6] R. E. Fernholz. *Stochastic Portfolio Theory*. Springer Science + Business Media, Berlin Heidelberg, 1st edition edition, 2002.

[7] A. Kratsios and I. Bilokopytov. Non-Euclidean universal approximation, 2020.

[8] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5:115–133, 1943.

[9] T. M. Mitchell. *Machine Learning*. McGraw-Hill series in computer science. WCB McGraw-Hill, Boston MA, 1997.

[10] L. Nachbin. Sur les algèbras denses de fonctions différentiables sur une variété. *Comptes rendus de l'Académie des Sciences de Paris*, 228:1549–1551, 1949.

[11] L. Nachbin. Weighted approximation for algebras and modules of continuous functions: Real and self-adjoint complex cases. *Annals of Mathematics*, 81(2):289–302, 1965.

[12] M. B. Stinchcombe. Neural network approximation of continuous functionals and continuous functions on compactifications. *Neural Networks*, 12(3):467–477, 1999.

[13] M. H. Stone. The generalized Weierstrass approximation theorem. *Mathematics Magazine*, 21(4):167–184, 1948.

Universal approximation theorems for continuous functionals of càdlàg paths and Lévy-type signature models

Francesca Primavera

based on joint work with Christa Cuchiero and Sara Svaluto-Ferro

Signature-based models have recently entered the field of Mathematical Finance, opening the door to more data-driven and thus more robust model selection mechanisms, see e.g. Arribas et al. (2020). The fundamental motivation for choosing the signature as the main building block for stochastic models can be explained by a *universal approximation theorem* (UAT) according to which continuous functionals of *continuous* paths can be approximated by linear functions of the time extended signature. This powerful result is however only proved for continuous paths and therefore leaves open the question of approximating continuous functionals of the more general set of *càdlàg* paths, which have particular relevance when it comes to financial modeling.

Based on recent results on the signature of càdlàg paths by Friz and Shekhar (2017), we prove two versions of a universal approximation theorem, one with respect to the Skorokhod J_1-topology and the other one with respect to (a rough path version of) the Skorokhod M_1-topology, introduced in Chevyrev and Friz (2019).

Our main motivation to treat this question comes from signature-based models for finance that allow for the inclusion of jumps. Indeed, all the signature models for asset prices that have been proposed so far have only dealt with continuous trajectories.

Relying on the càdlàg versions of the UAT, we then present a new class of signature models with jumps and show their universality among all classical jump diffusion models in finance.

The approach that we follow consists in parameterizing the model itself or its characteristics as linear functions of the signature of a time-augmented Lévy process X, that we call *market's primary process*, defined via

$$X_t : t \mapsto \left(t, W_t, \int_0^t \int_{\mathbb{R}} x \, (\mu - \nu)(ds, dx), \int_0^t \int_{\mathbb{R}} x^2 \, \mu(ds, dx), \dots, \int_0^t \int_{\mathbb{R}} x^N \, \mu(ds, dx) \right).$$

Here, μ denotes a Poisson random measure with compensator ν and W a one-dimensional Brownian motion. As introduced in Friz and Shekhar (2017) for general semimartingales, the signature $\mathbb{X}$ of the Lévy process X is defined as solution of the Marcus canonical equation (see Kurtz et al. (1995) for more details) in the extended tensor algebra over $\mathbb{R}^d$,

$$d\mathbb{X} = \sum_{i=1}^{d} \mathbb{X} \otimes \diamond \, dX^i$$
$$\mathbb{X}_0 = (1, 0, 0, \dots) \in T((\mathbb{R}^d)).$$

We thus define *Lévy-type signature models*, under some risk neutral measure $\mathbb{Q}$, as follows

$$S(\boldsymbol{\ell})_t = S_0 + \int_0^t \left(\sum_{|J| \leqslant n} \ell_W^J \langle \epsilon_J, \mathbb{X}_{s-} \rangle \right) dW_s + \int_0^t \int_{\mathbb{R}} \left(\sum_{|J| \leqslant n} \ell_\nu^J \langle \epsilon_J, \mathbb{X}_{s-} \rangle \right) y \, (\mu - \nu)(ds, dy),$$

for some parameters $\ell_W^J, \ell_\nu^J \in \mathbb{R}$.

To analyze tractability properties of these models, we first show that the signature process of a generic multivariate Lévy process is a polynomial process on the extended tensor algebra and derive its expected value via the so-called moment formula, thus solving a linear ODE; see Cuchiero et al. (2012); Filipović and Larsson (2020). We then exploit this result to get explicit formulas of the expected signature of S in terms of the expected signature of X, which makes pricing and calibration to market data very efficient.

Not only pricing can be traced back to the signature of X. Indeed, we also solve the hedging problem by computing the explicit form of the local risk minimizing strategy in terms of $\mathbb{X}$.

Finally, we discuss equivalent measure changes that allow to keep the tractable structure of Lévy-type signature models also under the physical measure $\mathbb{P}$.

References

Imanol Perez Arribas, Cristopher Salvi, and Lukasz Szpruch. Sig-sdes model for quantitative finance, 2020.

Ilya Chevyrev and Peter K. Friz. Canonical RDEs and general semimartingales as rough paths. *The Annals of Probability*, 47(1):420 – 463, 2019. doi: 10.1214/18-AOP1264. URL https://doi.org/10.1214/18-AOP1264.

Christa Cuchiero, Martin Keller-Ressel, and Josef Teichmann. Polynomial processes and their applications to mathematical finance. *Finance and Stochastics*, 16(4):711–740, 2012. URL https://EconPapers.repec.org/RePEc:spr:finsto:v:16:y:2012:i:4:p:711-740.

Damir Filipović and Martin Larsson. Polynomial jump-diffusion models. *Stochastic Systems*, 10, 02 2020. doi: 10.1287/stsy.2019.0052.

Peter K. Friz and Atul Shekhar. General rough integration, Lévy rough paths and a Lévy–Kintchine-type formula. *The Annals of Probability*, 45(4):2707 – 2765, 2017. doi: 10.1214/16-AOP1123. URL https://doi.org/10.1214/16-AOP1123.

Thomas G. Kurtz, Étienne Pardoux, and Philip Protter. Stratonovich stochastic differential equations driven by general semimartingales. *Annales de l'I.H.P. Probabilités et statistiques*, 31(2):351–377, 1995. URL http://www.numdam.org/item/AIHPB_1995__31_2_351_0/.

Universal Regular Conditional Distributions

Anastasis Kratsios anastasis.kratsios@math.ethz.ch
Department of Mathematics
Eidgenössische Technische Hochschule Zürich (ETH)
Rämistrasse 101, 8092 Zürich, ZH, Switzerland

June 20, 2021

Abstract

We introduce a general framework for approximating regular conditional distributions (RCDs). Our approximations of these RCDs are implemented by a new class of geometric deep learning models with inputs in $\mathbb{R}^d$ and outputs in the Wasserstein-1 space $\mathcal{P}_1(\mathbb{R}^D)$. We find that the models built using our framework can approximate any continuous functions from $\mathbb{R}^d$ to $\mathcal{P}_1(\mathbb{R}^D)$ uniformly on compacts, quantitatively. We identify two methods in which the curse of dimensionality can be avoided. The first solution describes functions in $C(\mathbb{R}^d, \mathcal{P}_1(\mathbb{R}^D))$ which can be efficiently approximated on any compact subset of $\mathbb{R}^d$. Conversely, the second approach describes sets in $\mathbb{R}^d$, on which any function in $C(\mathbb{R}^d, \mathcal{P}_1(\mathbb{R}^D))$ can be efficiently approximated. Our framework is used to obtain an affirmative answer to the open conjecture of Bishop 1994; namely: *mixture density networks are universal regular conditional distributions*. Experiments are performed in the context of ELMs, model uncertainty, and heteroscedastic regression. All the results are obtained for more general input and output spaces; thus, our analysis applies to geometric deep learning contexts.

This project is part of the following geometric deep learning program:

- Non-Euclidean Universal Approximation - **NeurIPS - 2020** - *A. Kratsios et. al*,
- NEU: A Meta-Algorithm for Universal UAP-Invariant Feature Representation - **Journal of Machine Learning Research - 2021** - *A. Kratsios et. al*,
- Universal Approximation Theorems for Geometric Deep Learning - *submitted* - *A. Kratsios et. al.*

Measure-Valued Neural Networks, Universal Regular Conditional Distributions, Conditional Expectation, Optimal Transport, Geometric Deep Learning.

2

Unspanned stochastic variance for VIX

Chenxu Li*

Guanghua School of Management

Peking University

Xin Zang†

School of Mathematical Sciences

Peking University

Abstract

Most existing models for volatility derivatives imply that variance swaps or the CBOE VIX span the whole space of volatility risk. However, we find that variance swaps and the VIX have limited explanatory power for VVIX, the CBOE VIX of VIX. We term this feature as unspanned stochastic variance (USV) for VIX. Thus, we present a new class of canonical-form affine models incorporating such USV factors. For practical applications, we study the novel two-factor and three-factor USV models and demonstrate their superior empirical performance in fitting the market by comparing with several existing popular models based on the estimation results using unscented Kalman filter.

Keywords: Unspanned stochastic volatility, market incompleteness, affine model, variance swaps, VIX.

*Corresponding author. Address: Guanghua School of Management, Peking University, Beijing, 100871, P.R.China. E-mail address: `cxli@gsm.pku.edu.cn`. Li's research is supported by the Guanghua School of Management, the Center for Statistical Sciences, and the Key Laboratory of Mathematical Economics and Quantitative Finance (Ministry of Education) at Peking University, as well as the National Natural Science Foundation of China (Project 11201009).

†Address: School of Mathematical Sciences, Peking University, Beijing, 100871, P.R.China. E-mail address: `xzang@pku.edu.cn`. Zang's research is supported by the Key Laboratory of Mathematical Economics and Quantitative Finance (Ministry of Education) at Peking University.

Using the COS Method to Calculate Risk Measures and Risk Contributions in Multifactor Copula Models

Xiaoyu Shen*† Fang Fang‡† Chujun Qiu§

January 26, 2020

Abstract

Measurement of a portfolio's credit risk in factor-copula models underlies various practical applications, from pricing credit derivatives such as collateralised debt obligations (CDOs) to quantifying regulatory capital (RC) and economical capital (EC) for both banking book and trading book positions. An re-emerging interest in such a topic arises from the Default Risk Charge (DRC) requirement, one of the capital charge components set out by the Basel Committee on Banking Supervision (BCBS) in the Fundamental Review of the Trading Book (FRTB). DRC under the internal model approach is defined as Value-at-Risk (VaR) at the 99.9% quantile of the default loss distribution of the trading book positions that are subject to issuer risk. Under multifactor copula models, the widely adopted method in industry is Monte Carlo (MC) simulation when it comes to quantification of such credit risks. However, MC simulation is not only time consuming but also inaccurate, especially in risk attribution calculations, whereby a risk measure such as VaR or Expected Shortfall (ES) is allocated among individual obligors or sub-portfolios so as to identify risk concentrations. In this paper, we present a semi-analytical method to efficiently calculate both risk measures and risk attribution of portfolio credit losses in the set-up of multifactor copula models. In essence, it directly recovers the cumulative distribution function of the total portfolio loss via the COS method, a method based on the Fourier-cosine series expansion. Fourier coefficients are extracted from the characteristic function (ch.f.) of the portfolio loss and the ch.f. can be pre-calculated by means

*ING Bank, the Netherlands; xiaoyu.shen@ing.com

†The views expressed herein are our own and do not necessarily reflect those of ING Bank.

‡ING Bank, the Netherlands; FF Quant Advisory B.V.; FangFang@fsquaredquant.nl

§Tsinghua University, China; qcj573@sina.com

of numerical integration. We demonstrate that both the computational speed and the accuracy of this method are much superior to MC simulation, using the examples of a multifactor Gaussian and a Gaussian-t hybrid model. Note that in case that the default threshold has no analytical expression, such as the case of the Gaussian-t-hybrid model, the default threshold can again be recovered semi-analytically by employing the COS method.

Variable Clustering via Distributionally Robust Nodewise Regression

Kaizheng Wang, Xiao Xu, Xun Yu Zhou

Abstract

We study a multi-factor block model for variable clustering and connect it to the subspace clustering problem. We then derive a distributionally robust version of nodewise regression to solve the subspace clustering problem. To solve the distributionally robust nodewise regression, we obtain a convenient convex relaxation, provide guidance on informing the choice of tuning parameter from the data, and propose an ADMM algorithm for efficient implementation. We validate our method in extensive simulation studies. We also apply our method to stock return data, utilizing clustering analysis to facilitate portfolio selection, and obtain promising portfolio backtesting results.

Definition 0.1 (Multi-factor block model)**.** Consider a d-dimensional random vector $X = (X_1, \ldots, X_d)$, and an underlying partition $G := \{G_1, \ldots, G_K\}$ of the indices $\{1, \ldots, d\}$. Denote by m_k the size of cluster G_k. For each $k = 1, \ldots, K$, let F_G^k be a d_k-dimensional random vector that represents the factors controlling the k-th cluster, and without loss of generality, assume that $\mathrm{Cov}(F_G^k) = \boldsymbol{I}$. We also assume that $m_k > d_k$, i.e., there are more variables than factors in each cluster. For each $i = 1, \ldots, d$, denote by $z(i) \in 1, \ldots, K$ the index of the cluster that X_i belongs to. Let $\beta_i \in \mathbb{R}^{d_k}$ be the loadings of the i-th variable on the factors $F_G^{z(i)}$. Under the multi-factor block model, for each i, the random variable X_i satisfies $X_i = (F_G^{z(i)})^\top \beta_i + U_i$, where U_i is a one-dimensional random variable that represents the idiosyncratic part that satisfies $\mathrm{Cov}(U_i, U_j) = 0$ for $i \neq j$.

Let $\boldsymbol{X} \in \mathbb{R}^{n \times d}$ be the data matrix of n observations of X. Suppose that the (unobserved) realizations of the latent factors are $\boldsymbol{F}_G \in \mathbb{R}^{n \times D}$, then $\boldsymbol{X}$ can be decomposed as $\boldsymbol{X} = \boldsymbol{Y} + \boldsymbol{U}$, where $\boldsymbol{Y} = \boldsymbol{F}_G \boldsymbol{A}$ is the group-specific component controlled by the factors, and $\boldsymbol{U}$ is the idiosyncratic components. The columns of $\boldsymbol{Y}$ lie in the union of K subspaces spanned by the factors. The task of variable clustering is now to identify the subspaces based on $\boldsymbol{X}$. A natural approach is nodewise regression, where each column of $\boldsymbol{X}$ is regressed against all other columns. In matrix form, nodewise regression is formulated as

$$\min_{\boldsymbol{B}\in\mathbb{R}^{d\times d}} \|\boldsymbol{X}-\boldsymbol{X}\boldsymbol{B}\|_F^2 \quad \text{s.t.} \quad \operatorname{diag}(\boldsymbol{B})=0. \tag{0.1}$$

The hope is that the resulting regression coefficients $\boldsymbol{B}$ will mainly connect vectors that are in the same cluster, i.e., $|b_{ij}|\approx 0$ if $z(i)\neq z(j)$. Then, the clusters can be easily recovered by performing, for example, spectral clustering on the symmetrized matrix $\boldsymbol{C}:=\boldsymbol{B}_{abs}^\top+\boldsymbol{B}_{abs}$, where $(\boldsymbol{B}_{abs})_{ij}=|b_{ij}|$. Our study focuses on the nodewise regression part.

Specifically, we propose to solve a distributionally-robust version of nodewise regression:

$$\underset{\boldsymbol{B}\in\mathbb{R}^{d\times d},\operatorname{diag}(\boldsymbol{B})=0}{\text{minimize}} \quad \sup_{\mathbb{P}:\ W_2(\mathbb{P},\mathbb{P}_n)\leq\delta} \mathbb{E}_{\mathbb{P}}\left[\left\|X-\boldsymbol{B}^\top X\right\|_2^2\right]. \tag{0.2}$$

Here W_2 is the Wasserstein-2 distance and $\mathbb{P}_n$ is the empirical distribution of the data. To facilitate computation, we obtain a convenient upper bound on the above objective function:

$$\sup_{\mathbb{P}:\ W_2(\mathbb{P},\mathbb{P}_n)\leq\delta} \mathbb{E}_{\mathbb{P}}\left[\left\|X-\boldsymbol{B}^\top X\right\|_2^2\right] \leq \left(\frac{1}{\sqrt{n}}\|\boldsymbol{X}-\boldsymbol{X}\boldsymbol{B}\|_F+\sqrt{\delta}\,\|\boldsymbol{I}-\boldsymbol{B}\|_2\right)^2, \qquad \forall \boldsymbol{B}\in\mathbb{R}^{d\times d}.$$

Then, we show that the program (0.2) can be relaxed to

$$\underset{\boldsymbol{B}\in\mathbb{R}^{d\times d},\operatorname{diag}(\boldsymbol{B})=0}{\text{minimize}} \left\{\frac{1}{\sqrt{n}}\|\boldsymbol{X}-\boldsymbol{X}\boldsymbol{B}\|_F+\sqrt{\delta}\,\|\boldsymbol{I}-\boldsymbol{B}\|_2\right\}, \tag{0.3}$$

which is a convex program.

We propose an efficient algorithm based on the alternating direction method of multipliers (ADMM). We also provide a simple, actionable recipe for choosing the tuning parameter δ. Finally, we conduct numerical experiments to test the efficacy of the proposed method.

Volatility Ambiguity in Dual

Kyunghyun Park*　　Hoi Ying Wong†

Abstract

Consider a robust consumption-investment problem for a risk- and ambiguity-averse investor who is concerned about return ambiguity in risky asset prices. When the investor aims to maximize the worst-case scenario of his/her consumption-investment objective, we propose a dual approach to the robust optimization problem in a dual economy with volatility ambiguity. Using the G-expectation framework, we establish the duality theorem to bridge between the primal problem with return ambiguity and the dual problem with volatility ambiguity, and hence characterize the robust strategy for a general class of utility functions subject to the non-negative consumption rate and wealth constraints. The volatility ambiguity in the dual problem induces correlation ambiguity when the primal economy comprises multiple risky assets with return ambiguity. By analyzing the dual economy, we show that the robust investment strategy favors a sparse portfolio, in addition to its usual feature—having the least exposure to ambiguity. We also extend the analysis to robust consumption-investment with retirement option. Using the duality approach, the robust retirement timing is characterized through a G-stopping.

Keywords: return ambiguity in primal, volatility ambiguity in dual, retirement, G-expectation, optimal G-stopping time, consumption-investment, robust strategy

*E-mail: khpark@cuhk.edu.hk Department of Statistics, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong
†E-mail: hywong@cuhk.edu.hk Department of Statistics, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong

Volatility forecasting with machine learning and intraday commonality

Chao Zhang*[1], Yihuang Zhang*[2,4], Mihai Cucuringu[1,2,3], and Zhongmin Qian[1,4]

[1]Department of Statistics, University of Oxford, Oxford, UK
[2]Mathematical Institute, University of Oxford, Oxford, UK
[3]The Alan Turing Institute, London, UK
[4]Oxford Suzhou Centre for Advanced Research

Abstract

We apply machine learning models to forecast intraday realized volatility (RV), by exploiting commonality in intraday volatility via pooling stock data together, and by incorporating a proxy for the market volatility. Neural networks dominate linear regressions and tree models in terms of performance, due to their ability to uncover and model complex latent interactions among variables. Our findings remain robust when we apply trained models to new stocks that have not been included in the training set, thus providing new empirical evidence for a *universal volatility* mechanism among stocks. Finally, we propose a new approach to forecasting one-day-ahead RVs using past intraday RVs as predictors, and highlight interesting diurnal effects that aid the forecasting mechanism. The results demonstrate that the proposed methodology yields superior out-of-sample forecasts over a strong set of traditional baselines that only rely on past daily RVs.

Introduction

Forecasting and modeling stock return volatility has been of interest to both academics and practitioners over the past years. Recent advances in high-frequency trading (HFT) highlight the need for robust and accurate intraday volatility forecasts. Intraday volatility estimates are important for pricing derivatives, managing risk, and devising quantitative strategies.

To the best of our knowledge, unlike daily volatility forecasting models, forecasting intraday volatility has received scant attention in the research literature [5]. It is pointed out that conventional parametric models, such as GARCH and stochastic volatility models, may be inadequate for modeling intraday returns [1]. In recent works [2, 4], high-frequency data are used to estimate daily **realized volatility (RV)** by summing squared intraday returns. However, these methods are potentially restrictive and are often difficult to apply when forecasting intraday volatility.

Main results

In the present paper, we study and analyze various non-parametric machine learning models for forecasting *multi-asset intraday and daily volatilities* by using high-frequency data from the U.S. equity market. We demonstrate that, by taking advantage of *commonality* in intraday volatility, the model's forecasting performance can be significantly be improved.

*Equal contribution. Correspondence to: Chao Zhang <chao.zhang@stats.ox.ac.uk>

More specifically, we first propose a measure to evaluate the commonality in intraday volatility and give a statistical analysis of the commonality based on historical data. The measure is defined as the adjusted R-squared value from linear regressions of a given stock's RVs against the market RVs. It is demonstrated that commonality over the daily horizon is turbulent over time, although commonality in intraday RVs is strong and stable. On the other hand, the analysis of the high-frequency data from the real market reveals that during a trading session, commonality achieves a peak near the close, in contrast to the diurnal volatility pattern.

Second, in order to assess the benefits of incorporating commonality into models aimed at predicting intraday volatility, we train multiple machine learning algorithms (including HAR [4], OLS, LASSO, XGBoost, MLP, and LSTM) under three different schemes: (a) **Single**: training specific models for each asset; (b) **Universal**: training one model with pooled data for all assets; (c) **Augmented**: training one model with pooled data and adding an additional predictor which takes account the impact of market realized volatility. We find that for most models, the incorporation of commonality (**Augmented**) leads to better out-of-sample performance through pooling data together and adding market volatility as additional features.

In addition, the empirical results demonstrate that neural networks (NNs) are, in general, superior to other techniques. New empirical evidence is provided, in order to demonstrate the capability of NNs for handling complex interactions among predictors. Furthermore, to alleviate the concerns of overfitting, we conduct a stringent out-of-sample test, using the trained models to forecast the volatility of completely new stocks that are not included in the training sample. Our results reveal that NNs still outperform other approaches (including the OLS models trained for each new stock), thus presenting empirical evidence for a universal volatility mechanism among stocks.

Finally, we propose a new approach for predicting daily volatility, in which the past intraday volatilities rather than the past daily volatilities are used as predictors. This approach fully utilizes the available high-frequency data, and contributes to the improvement over traditional methods of modeling daily volatilities with past daily volatilities (e.g. HAR [4], SHAR [6], HARQ [3]). To the best of our knowledge, this is the first line of work that studies the effectiveness of past intraday time-of-day dependent volatilities in forecasting future daily volatility.

References

[1] Torben G Andersen and Tim Bollerslev. Intraday periodicity and volatility persistence in financial markets. *Journal of empirical finance*, 4(2-3):115–158, 1997.

[2] Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Paul Labys. Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625, 2003.

[3] Tim Bollerslev, Andrew J Patton, and Rogier Quaedvlieg. Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18, 2016.

[4] Fulvio Corsi. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196, 2009.

[5] Robert F Engle and Magdalena E Sokalska. Forecasting intraday volatility in the us equity market. multiplicative component garch. *Journal of Financial Econometrics*, 10(1):54–83, 2012.

[6] Andrew J Patton and Kevin Sheppard. Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics*, 97(3):683–697, 2015.

Special Sessions in memory of Prof. Peter Carr

Volatility is (Mostly) Path-Dependent

Speaker: Julien Guyon

Abstract: Under We learn from data that volatility is mostly path-dependent: more than 90% of the variance of the implied volatility of equity indexes is explained endogenously by past index returns, and more than 70% for future daily realized volatility. The path-dependency that we uncover is remarkably simple: a linear combination of a weighted sum of past daily returns and the square root of a weighted sum of past daily squared returns with different time-shifted power-law weights. This suggests a simple continuous-time path-dependent volatility model that may be fed historical or risk-neutral parameters. The weights can be approximated by a superposition of exponential kernels to produce a Markovian model.

Special Sessions in memory of Prof. Tomas Bjork

Voting and decentralized trading of payment promises

Speaker: Ali Lazrak

Abstract: We examine a model where a binary decision is made by a committee based on voting. Committee members disagree on the decision but they can freely make credible and enforceable payment promises contingent on the committee decision. The payment promises are zero-sum but can involve coalitions of any size ranging from a pair to the entire committee. We show that trading payment promises enable the committee to choose the socially optimal decision and remove any inefficiencies resulting from voting externalities. Equilibrium payment promises are indeterminate even when restricted to generate a minimal total payment. When the voting outcome is inefficient, payment promises restore efficiency because intense minorities compensate the members of the majority to change their vote. When the voting outcome is efficient, payment promises can be required to preempt coalitions to form, get a better payoff and yet lead the group to adopt the inefficient decision. Joint work with Jianfeng Zhang.

What happen when we stop a Reflected backward stochastic differential equations at arbitrary random time?

Safa' Alsheyab[1]

[1]*Department of Mathematical and Statistical Sciences, University of Alberta, Central Academic Building, Edmonton, AB T6G 2G1, Canada*

February 23, 2022

Abstract: This paper addresses reflected backward stochastic differential equations (RBSDE hereafter) that take the form of

$$\begin{cases} dY_t = f(t, Y_t, Z_t)d(t \wedge \tau) + Z_t dW_t^{\tau} + dM_t - dK_t, \quad Y_\tau = \xi, \\ Y \geq S \quad \text{on} \quad [\![0, \tau[\![, \quad \displaystyle\int_0^{\tau} (Y_{s-} - S_{s-})dK_s = 0 \quad P\text{-a.s..} \end{cases}$$

Here τ is an arbitrary random time that might not be a stopping time for the filtration $\mathbb{F}$ generated by the Brownian motion W. On the one hand, this RBSDE is a natural extension of the cases treated in the literature and which assume that either τ is a stopping time or "somehow independent" of the filtration W. On the other hand, this class of RBSDEs is directly connected to the exponential hedging problem of defaultable claims or under random horizon in general. To study the RBSDE introduced above, we consider the filtration $\mathbb{G}$ resulting from the progressive enlargement of $\mathbb{F}$ with τ where this becomes a stopping time. Precisely, the following problems are among the questions that we deal with:

1. What are the sufficient minimal conditions on the data (f, ξ, S, τ) that guarantee the existence of the solution of the $\mathbb{G}$-RBSDE in L^p $(p > 1)$?
2. How can we estimate the solution in norm using the triplet-data (f, ξ, S)? What are the adequate spaces and the adequate norms?
3. Is there an RBSDE under $\mathbb{F}$ that is intimately related to the current one and how their solutions are related to each other?

This paper answers all these questions deeply and beyond. Importantly, we prove that for any random time, having a positive Azéma supermartingale, there exists a positive discount factor $\widetilde{\mathcal{E}}$ –a positive and non-increasing $\mathbb{F}$-adapted and RCLL process– that is vital in answering our questions without assuming any further assumption on τ, and determining the space for the triplet-data (f, ξ, S) and the space for the solution of the RBSDE as well. Furthermore, we found that the conditions for the $\mathbb{G}$-RBSDE are weaker than the conditions for its $\mathbb{F}$-RBSDE counterpart when the horizon is unbounded. Our approach sounds very robust, as it relies on sharp martingale inequalities that hold no matter what is the filtration, and on explicit description of how Snell envelope evolves with τ. Thus, we treat both the linear and general case of RBSDEs for bounded and unbounded horizon. If time permits, we will talk about some extensions and how the solutions of the resulting RBSDEs evolve.

This talk is based on joint paper with Tahir Choulli (my PhD supervisor)

What if we knew what the future brings, even on short term notice?

Peter Bank

We study optimal investment problems for an insider who can peek some time units into the future, but cannot arbitrarily take advantage of this knowledge because of quadratic transaction costs. In the Bachelier setting with exponential utility, we give an explicit solution to this control problem with intrinsically infinite-dimensional state variable. This is made possible by solving the dual problem where we make use of the theory of Gaussian Volterra integral equations. For the limiting case of very short term notice on price changes, we consider a model with jumps on which our investor receives a possibly noisy signal. Here dynamic programming allows us to compute an explicit optimal investment ratio which allows us to quantify the signal's monetary value.

The first part of the talk is based on joint work with Yan Dolinsky and Miklos Rasonyi, the second is joint work with Laura Koerber.

When Do Models Make Money? A Path from Return Predictability to Profitability*

Yufan Chen† and Ruixun Zhang‡

March 14, 2022

Abstract

Although machine learning models have been increasingly adopted to forecast asset prices and discover market inefficiencies, it remains difficult to bring these models from research to the investment management practice. In the context of a single-asset market, we reconcile this disconnect by proposing an analytical framework to link the profitability of a model-driven investment strategy to the predictability of the underlying model. Our framework accounts for trading costs, applies to general return-forecast models with arbitrary precision–recall curves, and yields an explicit solution to strategy profit as a function of model performance metrics. This provides a definition of *arbitrage-free model accuracy*—the level of model predictability for investment strategies to break even in profits—which is shown to be much higher than commonly believed. In addition, our model provides an explanation for the asymmetry between opening and closing positions, a popular heuristic among investment professionals. We demonstrate the applicability of our model in the context of algorithmic trading in the stock market, with tick-level data that accounts for the bid-ask spread and actual liquidity in the limit order book. Empirical results show that the model-implied analytical profitability aligns with estimates of true profitability. These results provide a low-cost and efficient model profitability estimate for both researchers and practitioners. More generally, they highlight that machine learning models should meet a high standard as evidence for market inefficiency and profitability, especially for high-frequency investing that is sensitive to bid-offer spread and trading costs.

Keywords: Quantitative trading; High-frequency trading; Profitability; Machine learning; Statistical models

JEL Classification: C02, G14, G24

*We thank Andrew Lo, Lan Wu, Xun Yu Zhou, and seminar participants at the Peking University Laboratory for Mathematical Economics and Quantitative Finance for very helpful comments and discussion.

†Department of Financial Mathematics, School of Mathematical Sciences, Peking University

‡Assistant Professor and Boya Young Fellow, Peking University School of Mathematical Sciences, Center for Statistical Science, and Laboratory for Mathematical Economics and Quantitative Finance. `zhangruixun@pku.edu.cn`.

When do you Stop Supporting your Bankrupt Subsidiary?

Maxim Bichuch* Nils Detering†

January 29, 2020

1 Extended abstract

Most global banks consist of dozens if not hundreds of subsidiaries. For corporations becoming a holding with subsidiaries has several advantages: 1. It allows them to limit the spillover risk if one business line is in trouble. In this case, the holding acts as a pure shareholder and enjoys limited liability (type B). 2. In other instances it allows corporations to defer taxable business income and use it for other business opportunities or, by channeling income to low tax countries reduce taxes altogether (type A). For example a bank holding may naturally divide into subsidiaries based on location (i.e. Europe, US) and/or business lines (i.e. equity trading, fixed income trading). Type B is the standard holding structure in the US while type A is common in Germany and is accompanied by a profit transfer agreement ensuring fiscal unity. Motivation for type A is when subsidiaries have credit lines with other subsidiaries through their holding, with subsidiaries defaulting after exhausting the limit. Motivation for type B is holdings supporting subsidiaries to some level, and giving up on them only when faced with substantial losses. While for the corporation themselves, the advantages are imminent, whether holdings are beneficial from a society perspective is less clear. In particular it is not well understood how shocks propagate in a network of bank holdings and what influence the holding structure has on the contagion process.

Therefore we consider a network of n banks. Each bank is structured as a holding with two subsidiaries – subsidiary 1 and 2. We assume that subsidiaries only trade with subsidiaries of the same kind, i.e. subsidiary 1 of bank i might trade with subsidiary 1 of bank j but not with subsidiary 2. The capital of the holding is the sum of the capitals of the subsidiaries, because the holding itself has no business activity. For each holding structures (A and B) we then consider a contagion process triggered by a financial shock and propagates throughout the network according to the following rules:

1. Subsidiary defaults if it has non-positive capital unless its holding supports it (i.e. is being supported by the other subsidiary).

2. The holding will support a subsidiary until its capital reaches $x < 0$ if its other subsidiary has the capital to do so, otherwise the subsidiary defaults.

*Department of Applied Mathematics & Statistics, Johns Hopkins University, Baltimore, MD 21218, USA Email: mbichuch@jhu.edu.

†Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106, USA. Email: detering@pstat.ucsb.edu

3. For holding structure of type A (type B respectively) a defaulted subsidiary reduces the capital of the holding's other subsidiary (does not change the capital of its other subsidiary respectively).

All holdings will support their subsidiaries up to the same level x when possible. In the limit as $x \to -\infty$ type A network of holdings is similar in nature to having no subsidiaries, but a multilayer network of banks where each layer represents a different business relationship.

For fixed n we build a random network by connecting subsidiary k of bank i to subsidiary k of bank j with probability p_k/n, $k = 1, 2,\ 1 \leq i, j \leq n$. We then consider propagation of shocks via default contagion of subsidiaries and holdings asymptotically as $n \to \infty$.

Contributions: Given this network parameters and x our first main result allows to calculate the spillover effects from any initial shock to the system. It allows us to calculate the number of defaulted subsidiaries and holdings at the end of the contagion process for each holding type A or B. We show that comprising a firms business activity in a holding can have both, positive or negative effects on systemic stability. We analyse to what extent voluntary support benefits society and/or the holding itself. We observe that an increased commitment of the holding to its subsidiaries (i.e. decreasing x) can have both, positive and negative effects on systemic risk depending on the type of the holding. We then find the optimal level of support x from the perspective of reducing the risk of the entire system and from the perspective of one individual bank with a given capital structure. We compare the two types of holding structures and compare it to having no holdings at all, but indivisible atom banks with no holding structure.

From a technical perspective our model setup provides a first instance of a multilayered random network in which contagion spreads through different channels and where interactions between layers is through nodes. This complicates the analysis compared to a one layer network as the state of the nodes is not determined by a one dimensional quantity, but rather lies in a multidimensional domain.

Special Sessions in memory of Prof. Mark H.A. Davis

Working with Msc students on the cutting edge - Dynamically Controlled Kernel Estimation

Gordon Lee

Abstract: We introduce a data driven and model free approach for computing conditional expectations. The new method is based on classical techniques combined with machine learning methods. In particular, we consider kernel density estimation based on simulated risk factors combined with a control variate. This is used in a Gaussian process regression for finally approximating the conditional expectation. In this way we increase not only the stability of the estimator but we also need a significantly lower amount of simulations due to the variance reduction. Since we apply Gaussian process regression, we do not only get a point estimate, but also the full distribution. It turns out that the optimal coefficient for the control variate is the minimal variance delta. Thus, in this way we obtain model free and purely data driven hedges. Finally, we apply our method to several examples from option pricing including exotic option payoffs and payoffs with multiple underlyings for different models including the rough Bergomi model. A discussion on the challenges to extend the method to a large dimensional settings is provided and is partially solved by using Quasi random number sequences.

Authors Index