VOL. 3, NO. 8

OMI Research Newsletter – June 2024

The Oxford-Man Institute of Quantitative Finance (OMI) is a world-leading centre for interdisciplinary research in financial markets. The OMI addresses fundamental problems in quantitative finance with a strong focus on machine learning and data driven models.

*The **OMI Research Newsletter** is a curated summary that covers the most recent research, insights, and tools related to financial decision making.*

Highlights

Research focus. Key contributions in this month's edition include works from the OMI, Imperial College London, Stevens Institute of Technology, and UC Berkeley. Duran-Martin et al. (2024) derive robust Bayesian update rules for online filtering in the presence of outliers and misspecified measurement models. Cao and Feinstein (2024) explores LLMs as tools to interpret complex financial regulation. Mastrolia and Xu (2024) design a periodic auction when both a continuous market and a periodic auction market are available to traders. Finally, Calvo-Ordonez et al. (2024) present a new family of infinitely deep neural networks with partial stochasticity to improve computational efficiency.

Machine Learning & Statistics

Outlier-robust Kalman Filtering through Generalised Bayes

G. DURAN-MARTIN, M. ALTAMIRANO, A. Y. SHESTOPALOFF, L. SÁNCHEZ-BETANCOURT, J. KNOBLAUCH, M. JONES, F-X. BRIOL, K. MURPHY

OXFORD-MAN INSTITUTE, QUEEN MARY UNIVERSITY, UNIVERSITY COLLEGE LONDON, UNIVERSITY OF COLORADO BOULDER

We derive a novel, provably robust, and closed-form Bayesian update rule for online filtering in state-space models in the presence of outliers and misspecified measurement models. Our method combines generalised Bayesian inference with filtering methods such as the extended and ensemble Kalman filter. We use the former to show robustness and the latter to ensure computational efficiency in the case of nonlinear models. Our method matches or outperforms other robust filtering methods (such as those based on variational Bayes) at a much lower computational cost. We show this empirically on a range of filtering problems with outlier measurements, such as object tracking, state estimation in high-dimensional chaotic systems, and online learning of neural networks.

ChatGPT is bullshit

M. T. HICKS, J. HUMPHRIES, J. SLATER

UNIVERSITY OF GLASGOW

Recently, there has been considerable interest in large language models: machine learning systems which produce human-like text and dialogue. Applications of these systems have been plagued by persistent inaccuracies in their output; these are often called "AI hallucinations". We argue that these falsehoods, and the overall activity of large language models, is better understood as bullshit in the sense explored by Frankfurt (On Bullshit, Princeton, 2005): the models are in an important way indifferent to the truth of their outputs. We distinguish two ways in which the models can be said to be bullshitters, and argue that they clearly meet at least one of these definitions. We further argue that describing AI misrepresentations as bullshit is both a more useful and more accurate way of predicting and discussing the behaviour of these systems.

Partially Stochastic Infinitely Deep Bayesian Neural Networks

S. Calvo-Ordonez, M. Meunier, F. Piatti, Y. Shi

Oxford-Man Institute, Imperial College London

In this paper, we present Partially Stochastic Infinitely Deep Bayesian Neural Networks, a novel family of architectures that integrates partial stochasticity into the framework of infinitely deep neural networks. Our new class of architectures is designed to improve the computational efficiency of existing architectures at training and inference time. To do this, we leverage the advantages of partial stochasticity in the infinite-depth limit which include the benefits of full stochasticity e.g. robustness, uncertainty quantification, and memory efficiency, whilst improving their limitations around computational complexity. We present a variety of architectural configurations, offering flexibility in network design including different methods for weight partition. We also provide mathematical guarantees on the expressivity of our models by establishing that our network family qualifies as Universal Conditional Distribution Approximators. Lastly, empirical evaluations across multiple tasks show that our proposed architectures achieve better downstream task performance and uncertainty quantification than their counterparts while being significantly more efficient.

Online Learning of Order Flow and Market Impact with Bayesian Change-Point Detection Methods

I-Y. Tsaknaki, F. Lillo, P. Mazzarisi

Scuola Normale Superiore, Università di Siena, Università di Bologna

Financial order flow exhibits a remarkable level of persistence, wherein buy (sell) trades are often followed by subsequent buy (sell) trades over extended periods. This persistence can be attributed to the division and gradual execution of large orders. Consequently, distinct order flow regimes might emerge, which can be identified through suitable time series models applied to market data. In this paper, we propose the use of Bayesian online change-point detection (BOCPD) methods to identify regime shifts in real-time and enable online predictions of order flow and market impact. To enhance the effectiveness of our approach, we have developed a novel BOCPD method using a score-driven approach. This method accommodates temporal correlations and time-varying parameters within each regime. Through empirical application to NASDAQ data, we have found that: (i) Our newly proposed model demonstrates superior out-of-sample predictive performance compared to existing models that assume i.i.d. behavior within each regime; (ii) When examining the residuals, our model demonstrates good specification in terms of both distributional assumptions and temporal correlations; (iii) Within a given regime, the price dynamics exhibit a concave relationship with respect to time and volume, mirroring the characteristics of actual large orders; (iv) By incorporating regime information, our model produces more accurate online predictions of order flow and market impact compared to models that do not consider regimes.

Large Language Model in Financial Regulatory Interpretation

Z. Cao, Z. Feinstein

Stevens Institute of Technology

This study explores the innovative use of Large Language Models (LLMs) as analytical tools for interpreting complex financial regulations. The primary objective is to design effective prompts that guide LLMs in distilling verbose and intricate regulatory texts, such as the Basel III capital requirement regulations, into a concise mathematical framework that can be subsequently translated into actionable code. This novel approach aims to streamline the implementation of regulatory mandates within the financial reporting and risk management systems of global banking institutions. A case study was conducted to assess the performance of various LLMs, demonstrating that GPT-4 outperforms other models in processing and collecting necessary information, as well as executing mathematical calculations. The case study utilized numerical simulations with asset holdings – including fixed income, equities, currency pairs, and commodities – to demonstrate how LLMs can effectively implement the Basel III capital adequacy requirements.

HLOB – Information Persistence and Structure in Limit Order Books

A. Briola, S. Bartolucci, T. Aste

University College London, London School of Economics

We introduce a novel large-scale deep learning model for Limit Order Book mid-price changes forecasting, and we name it 'HLOB'. This architecture (i) exploits the information encoded by an Information Filtering Network, namely the Triangulated Maximally Filtered Graph, to unveil deeper and non-trivial dependency structures among volume levels; and (ii) guarantees deterministic design choices to handle the complexity of the underlying system by drawing inspiration from the groundbreaking class of Homological Convolutional Neural Networks. We test our model against 9 state-of-the-art deep learning alternatives on 3 real-world Limit Order Book datasets, each including 15 stocks traded on the NASDAQ exchange, and we systematically characterize the scenarios where HLOB outperforms state-

of-the-art architectures. Our approach sheds new light on the spatial distribution of information in Limit Order Books and on its degradation over increasing prediction horizons, narrowing the gap between microstructural modeling and deep learning-based forecasting in high-frequency financial markets.

Generative Adversarial Bayesian Optimization for Surrogate Objectives

M. S. Yao, Y. Zeng, H. Bastani, J. Gardner, J. C. Gee, O. Bastani

University of Pennsylvania

Offline model-based policy optimization seeks to optimize a learned surrogate objective function without querying the true oracle objective during optimization. However, inaccurate surrogate model predictions are frequently encountered along the optimization trajectory. To address this limitation, we propose generative adversarial Bayesian optimization (GABO) using adaptive source critic regularization, a task-agnostic framework for Bayesian optimization that employs a Lipschitz-bounded source critic model to constrain the optimization trajectory to regions where the surrogate function is reliable. We show that under certain assumptions for the continuous input space prior, our algorithm dynamically adjusts the strength of the source critic regularization. GABO outperforms existing baselines on a number of different offline optimization tasks across a variety of scientific domains.

Finance and Economics

Clearing time randomization and transaction fees for auction market design

T. Mastrolia, T. Xu

UC Berkeley

Flaws of a continuous limit order book mechanism raise the question of whether a continuous trading session and a periodic auction session would bring better efficiency. This paper wants to go further in designing a periodic auction when both a continuous market and a periodic auction market are available to traders. In a periodic auction, we discover that a strategic trader could take advantage of the accumulated information available along the auction duration by arriving at the latest moment before the auction closes, increasing the price impact on the market. Such price impact moves the clearing price away from the efficient price and may disturb the efficiency of a periodic auction market. We thus propose and quantify the effect of two remedies to mitigate these flaws: randomizing the auction's closing time and optimally designing a transaction fees policy. Our results show that these policies encourage a strategic trader to send their orders earlier to enhance the efficiency of the auction market, illustrated by data extracted from Alphabet and Apple stocks.

Ponzi Funds

P. van der Beck, J-P. Bouchaud, D. Villamaina

Harvard Business School, Capital Fund Management, Ecole polytechnique

Many active funds hold concentrated portfolios. Flow-driven trading in these securities causes price pressure, which pushes up the funds' existing positions resulting in realized returns. We decompose fund returns into a price pressure (self-inflated) and a fundamental component and show that when allocating capital across funds, investors are unable to identify whether realized returns are self-inflated or fundamental. Because investors chase self-inflated fund returns at a high frequency, even short-lived impact meaningfully affects fund flows at longer time scales. The combination of price impact and return chasing causes an endogenous feedback loop and a reallocation of wealth to early fund investors, which unravels once the price pressure reverts. We find that flows chasing self-inflated returns predict bubbles in ETFs and their subsequent crashes, and lead to a daily wealth reallocation of 500 Million from ETFs alone. We provide a simple regulatory reporting measure – fund illiquidity – which captures a fund's potential for self-inflated returns.

Influencer Cartels

M. Hinnosaar, T. Hinnosaar

University of Nottingham

Social media influencers account for a growing share of marketing worldwide. We demonstrate the existence of a novel form of market failure in this advertising market: influencer cartels, where groups of influencers collude to increase their advertising revenue by inflating their engagement. Our theoretical model shows that influencer cartels can improve consumer welfare if they expand social media engagement to the target audience, or reduce welfare if they divert engagement to less relevant audiences. We validate the model empirically using novel data on influencer cartels combined with machine learning tools, and derive policy implications for how to maximize consumer welfare.

LEGO: Optimal Online Learning under Sequential Price Competition

Shukai Li, C. Shi, S. Mehrotra

Northwestern University

We consider price competition among multiple sellers over a selling horizon of T periods. In each period, sellers simultaneously offer their prices and subsequently

observe their respective demand that is unobservable to competitors. The realized demand of each seller depends on the prices of all sellers following a private unknown linear model. We propose a least-squares estimation then gradient optimization (LEGO) policy, which does not require sellers to communicate demand information or coordinate price experiments throughout the selling horizon. We show that our policy, when employed by all sellers, leads at a fast rate $O(\frac{1}{\sqrt{T}})$ to the Nash equilibrium prices that sellers would reach if they were fully informed. Meanwhile, each seller achieves an optimal order-of-$\sqrt{T}$ regret relative to a dynamic benchmark policy. Our analysis further shows that the unknown individual price sensitivity contributes to the major difficulty of dynamic pricing in sequential competition and forces regret to the order of $\sqrt{T}$ in the worst case. If each seller knows their individual price sensitivity coefficient, then a gradient optimization policy can achieve an optimal order-of-1/T convergence rate to Nash equilibrium as well as an optimal order-of-log T regret.

Power and Financial Interdependence

B. SETSER

COUNCIL ON FOREIGN RELATIONS

The link between financial self-reliance and geopolitical power has long been debated. The unbalanced Sino-American trade relationship has created asymmetric financial ties which generate potential sources of leverage for both parties and will not quickly disappear. Absent a clarifying major crisis, it will be difficult to definitively determine which party has greater leverage.

Many in the United States (US) are concerned about indebtedness to its primary strategic rival, and the risks posed by a sudden Chinese withdrawal from US financial markets. US policymakers actively sought to encourage China's top leadership not to withdraw financing from the market for US Agency securities in the run-up to the global financial crisis.

Yet China also sees risks in this unbalanced financial relationship. Chinese policymakers have expressed concern about the domestic political consequences of losses on either their Treasury or Agency holdings and actively have sought to diversify China's reserves – including by substituting the risk of lending to developing economies for the visibility associated with large holdings of Treasuries in US custodians. China increasingly worries that its dollar holdings and the dollar's global role increase its vulnerability to potential financial sanctions.

Both parties thus worry about the possibility that financial interdependence can be weaponized yet find it hard to extricate themselves from the inevitability of financial interdependence absent a clean break from an entrenched pattern of trade imbalances.

Designing Social Learning

A. SMIRNOV, E. STARKOV

UNIVERSITY OF COPENHAGEN, HSE UNIVERSITY

This paper studies strategic communication in the context of social learning. Product reviews are used by consumers to learn product quality, but in order to write a review, a consumer must be convinced to purchase the item first. When reviewers care about welfare of future consumers, this leads to a conflict: a reviewer today wants the future consumers to purchase the item even when this comes at a loss to them, so that more information is revealed for the consumers that come after. We show that due to this conflict, communication via reviews is inevitably noisy, regardless of whether reviewers can commit to a communication strategy or have to resort to cheap talk. The optimal communication mechanism involves truthful communication of extreme experiences and pools the moderate experiences together.

Mathematical Finance

Price-Aware Automated Market Makers: Models Beyond Brownian Prices and Static Liquidity

P. BERGAULT, L. BERTUCCI, D. BOUBA, O. GUÉANT, J. GUILBERT

PARIS DAUPHINE-PSL, UNIVERSITÉ PARIS 1 PANTHÉON-SORBONNE, INSTITUT LOUIS BACHELIER

In this paper, we introduce a suite of models for price-aware automated market making platforms willing to optimize their quotes. These models incorporate advanced price dynamics, including stochastic volatility, jumps, and microstructural price models based on Hawkes processes. Additionally, we address the variability in demand from liquidity takers through models that employ either Hawkes or Markov-modulated Poisson processes. Each model is analyzed with particular emphasis placed on the complexity of the numerical methods required to compute optimal quotes.

Fourier-Laplace transforms in polynomial Ornstein-Uhlenbeck volatility models

E. ABI JABER, S. LI, X. LIN

ECOLE POLYTECHNIQUE, UNIVERSITÉ PARIS 1 PANTHÉON-SORBONNE

We consider the Fourier-Laplace transforms of a broad class of polynomial Ornstein-Uhlenbeck (OU) volatility models, including the well-known Stein-Stein, Schöbel-Zhu, one-factor Bergomi, and the recently introduced Quintic OU models motivated by the SPX-VIX joint calibration problem. We show the connection between the joint Fourier-Laplace functional of the log-price and the integrated variance, and the solution of an infinite

dimensional Riccati equation. Next, under some non-vanishing conditions of the Fourier-Laplace transforms, we establish an existence result for such Riccati equation and we provide a discretized approximation of the joint characteristic functional that is exponentially entire. On the practical side, we develop a numerical scheme to solve the stiff infinite dimensional Riccati equations and demonstrate the efficiency and accuracy of the scheme for pricing SPX options and volatility swaps using Fourier and Laplace inversions, with specific examples of the Quintic OU and the one-factor Bergomi models and their calibration to real market data.

Acknowledgement

The OMI Research Newsletter is curated by Álvaro Cartea, Director of the Oxford-Man Institute, and Fayçal Drissi, postdoctoral researcher at the institute. Paper suggestions are from fellow members at the institute: Álvaro Arroyo, Patrick Chang, Gerardo Durán-Martín, Marcello Monga, and Harrison Waldon.

References

Bergault, Philippe, Louis Bertucci, David Bouba, Olivier Guéant, and Julien Guilbert. 2024. "Price-Aware Automated Market Makers: Models Beyond Brownian Prices and Static Liquidity." *arXiv preprint arXiv:2405.03496*.

Briola, Antonio, Silvia Bartolucci, and Tomaso Aste. 2024. "HLOB–Information Persistence and Structure in Limit Order Books." *arXiv preprint arXiv:2405.18938*.

Calvo-Ordonez, Sergio, Matthieu Meunier, Francesco Piatti, and Yuantao Shi. 2024. "Partially Stochastic Infinitely Deep Bayesian Neural Networks." *arXiv preprint arXiv:2402.03495*.

Cao, Zhiyu, and Zachary Feinstein. 2024. "Large Language Model in Financial Regulatory Interpretation." *arXiv preprint arXiv:2405.06808*.

Duran-Martin, Gerardo, Matias Altamirano, Alexander Y Shestopaloff, Leandro Sánchez-Betancourt, Jeremias Knoblauch, Matt Jones, François-Xavier Briol, and Kevin Murphy. 2024. "Outlier-robust Kalman Filtering through Generalised Bayes." *arXiv preprint arXiv:2405.05646*.

Hicks, Michael Townsen, James Humphries, and Joe Slater. 2024. "ChatGPT is bullshit." *Ethics and Information Technology*, 26(2): 38.

Hinnosaar, Marit, and Toomas Hinnosaar. 2024. "Influencer cartels." *arXiv preprint arXiv:2405.10231*.

Jaber, Eduardo Abi, Xuyang Lin, et al. 2024. "Fourier-Laplace transforms in polynomial Ornstein-Uhlenbeck volatility models." *arXiv preprint arXiv:2405.02170*.

Li, Shukai, Cong Shi, and Sanjay Mehrotra. 2024. "LEGO: Optimal Online Learning under Sequential Price Competition." *Available at SSRN 4803002*.

Mastrolia, Thibaut, and Tianrui Xu. 2024. "Clearing time randomization and transaction fees for auction market design." *arXiv preprint arXiv:2405.09764*.

Smirnov, Aleksei, and Egor Starkov. 2024. "Designing Social Learning." *arXiv preprint arXiv:2405.05744*.

Tsaknaki, Ioanna-Yvonni, Fabrizio Lillo, and Piero Mazzarisi. 2024. "Online learning of order flow and market impact with Bayesian change-point detection methods." *Quantitative Finance*, 1–16.

van der Beck, Philippe, Jean-Philippe Bouchaud, and Dario Villamaina. 2024. "Ponzi Funds." *arXiv preprint arXiv:2405.12768*.

Yao, Michael S, Yimeng Zeng, Hamsa Bastani, Jacob Gardner, James C Gee, and Osbert Bastani. 2024. "Generative Adversarial Bayesian Optimization for Surrogate Objectives." *arXiv preprint arXiv:2402.06532*.