VOL. 2, NO. 11

OMI Research Newsletter – August 2023

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

News. The Conference on Learning in Games and Algorithmic Collusion Workshop will take place on 12 and 13 October 2023 at the University of Oxford. The conference aims at bringing together both junior and senior leading scholars from leading research institutions. Theorists, empiricists, and legal experts will discuss topics related to learning in games, algorithmic pricing, regulation, and detecting collusion.

The DeFox seminar will start again on 18 October 2023 on a weekly basis with talks on the mathematics and economics of decentralised finance. The list of speakers and details are here.

The Conference on Machine Learning and Financial Econometrics will take place on 12 and 13 December 2023 at the Oxford-Man Institute, more details will follow.

Research focus. Noteworthy contributions in this month's edition include works from the OMI, Technische Universität München, Imperial College London, Princeton University, London School of Economics, University of Illinois. In this work, the authors propose a nonparametric method to simulate a limit order book with GANs that learn the conditional distribution of the order book state. Here, the authors solve non-Markovian control problems with neural rough differential equations. The authors of this work presents a multi-agent RL method for the simultaneous execution of several orders. Finally, here the authors develop a hybrid Quantum Wasserstein GAN to load arbitrary distributions of assets and improve option pricing.

Machine Learning & Statistics

Limit Order Book Simulation with Generative Adversarial Networks

R. CONT, M. CUCURINGU, J. KOCHEMS, F. PRENZEL
OXFORD-MAN INSTITUTE, UNIVERSITY OF OXFORD

We propose a nonparametric method for simulating the dynamics of a limit order book using Generative Adversarial Networks (GAN) to learn the conditional distribution of the future state of the order book given its current state from time series of the limit order book. Our method yields a scenario generator for limit order books which captures a range of stylized facts and salient properties of limit order book transitions. We show that the trained generator is also able to correctly reproduce some key properties observed in empirical studies on market impact. In particular, the model exhibits a decaying marginal impact of trade size, higher impact of aggressive orders, as well as a decreasing relation between impact and order book depth.

A Neural RDE Approach for Continuous-time Non-Markovian Stochastic Control Problems

M. HÖGLUND, E. FERRUCCI, C. HERNÁNDEZ, A. M. GONZALEZ, C. SALVI, L. SÁNCHEZ-BETANCOURT, Y. ZHANG

OXFORD-MAN INSTITUTE, IMPERIAL COLLEGE LONDON, UNIVERSITY OF OXFORD, PRINCETON UNIVERSITY, LONDON SCHOOL OF ECONOMICS

We propose a novel framework for solving continuous-time non-Markovian stochastic optimal problems by means of neural rough differential equations (Neural RDEs) introduced in Morrill et al. (2021). Non-Markovianity naturally arises in control problems due to

the time delay effects in the system coefficients or the driving noises, which leads to optimal control strategies depending explicitly on the historical trajectories of the system state. By modelling the control process as the solution of a Neural RDE driven by the state process, we show that the control-state joint dynamics are governed by an uncontrolled, augmented Neural RDE, allowing for fast Monte-Carlo estimation of the value function via trajectories simulation and memory- efficient back-propagation. We provide theoretical underpinnings for the proposed algorithmic framework by demonstrating that Neural RDEs serve as universal approximators for functions of random rough paths. Exhaustive numerical ex- periments on non-Markovian stochastic control problems are presented, which reveal that the proposed framework is time-resolution-invariant and achieves higher accuracy and better stability in irregular sampling compared to existing RNN- based approaches.

Non-parametric Online Market Regime Detection and Regime Clustering for Multidimensional and Path-dependent Data Structures

B. Horvath, Z. Issa

Oxford-Man Institute, King's College London

In this work we present a non-parametric online market regime detection method for multidimensional data structures using a path-wise two-sample test derived from a maximum mean discrepancy-based similarity metric on path space that uses rough path signatures as a feature map. The latter similarity metric has been developed and applied as a discriminator in recent generative models for small data environments, and has been optimised here to the setting where the size of new incoming data is particularly small, for faster reactivity.

On the same principles, we also present a path-wise method for regime clustering which extends our previous work. The presented regime clustering techniques were designed as ex-ante market analysis tools that can identify periods of approximatively similar market activity, but the new results also apply to path-wise, high dimensional-, and to non-Markovian settings as well as to data structures that exhibit autocorrelation.

We demonstrate our clustering tools on easily verifiable synthetic datasets of increasing complexity, and also show how the outlined regime detection techniques can be used as fast on-line automatic regime change detectors or as outlier detection tools, including a fully automated pipeline. Finally, we apply the fine-tuned algorithms to real-world historical data including high-dimensional baskets of equities and the recent price evolution of crypto assets, and we show that our methodology swiftly and accurately indicated historical periods of market turmoil.

A Hybrid Quantum Wasserstein GAN with Applications to Option Pricing

F. Fuchs, B. Horvath

Oxford-Man Institute, Technische Universität München

This paper conducts an in-depth exploration of harnessing Machine Learning techniques in the context of Quantum methods for Option Pricing. We develop and implement a novel Hybrid Quantum Wasserstein GAN for loading arbitrary distributions from a classical state into a quantum state, which is of interest beyond its financial applications. In particular, our hybrid methodology eliminates several potential sources of instability and has superior performance compared to fully quantum generative models if the target distribution is classical.

Our qWGAN can be used to capture the probability distribution of an asset at maturity, and transmute it into a quantum state as demonstrated on synthetic as well as real data experiments. In an Option Pricing context we present a full pipeline using this methodology and leverage the Iterative Quantum Amplitude Estimation algorithm to derive the expected option payoff, ensuring a quadratic enhancement in error scaling compared to traditional methods.

Learning Multi-Agent Intention-Aware Communication for Optimal Multi-Order Execution in Finance

Y. Fang, Z. Tang, K. Ren, W. Liu, L. Zhao, J. Bian, D. Li, W. Zhang, Y. Yu, T-Y. Liu

Shanghai Jiao Tong University, University of Illinois, Microsoft Research Asia

Order execution is a fundamental task in quantitative finance, aiming at finishing acquisition or liquidation for a number of trading orders of the specific assets. Recent advance in model-free reinforcement learning (RL) provides a data-driven solution to the order execution problem. However, the existing works always optimize execution for an individual order, overlooking the practice that multiple orders are specified to execute simultaneously, resulting in suboptimality and bias. In this paper, we first present a multi-agent RL (MARL) method for multi-order execution considering practical constraints. Specifically, we treat every agent as an individual operator to trade one specific order, while keeping communicating with each other and collaborating for maximizing the overall profits. Nevertheless, the existing MARL algorithms often incorporate communication among agents by exchanging only the information of their partial observations, which is inefficient in complicated financial market. To improve collaboration, we then propose a learnable multi-round communication protocol, for the agents communicating the intended actions with each other and refining accordingly. It is optimized through a novel action

value attribution method which is provably consistent with the original learning objective yet more efficient. The experiments on the data from two real-world markets have illustrated superior performance with significantly better collaboration effectiveness achieved by our method.

SaGess: Sampling Graph Denoising Diffusion Model for Scalable Graph Generation

S. Limnios, P. Selvaraj, M. Cucuringu, C. Maple, G. Reinert, A. Elliott

Oxford-Man Institute, The Alan Turing Institute, University of Oxford, University of Warwick, University of Glasgow

Over recent years, denoising diffusion generative models have come to be considered as state-of-the-art methods for synthetic data generation, especially in the case of generating images. These approaches have also proved successful in other applications such as tabular and graph data generation. However, due to computational complexity, to this date, the application of these techniques to graph data has been restricted to small graphs, such as those used in molecular modeling. In this paper, we propose SaGess, a discrete denoising diffusion approach, which is able to generate large real-world networks by augmenting a diffusion model (DiGress) with a generalized divide-and-conquer framework. The algorithm is capable of generating larger graphs by sampling a covering of subgraphs of the initial graph in order to train DiGress. SaGess then constructs a synthetic graph using the subgraphs that have been generated by DiGress. We evaluate the quality of the synthetic data sets against several competitor methods by comparing graph statistics between the original and synthetic samples, as well as evaluating the utility of the synthetic data set produced by using it to train a task-driven model, namely link prediction. In our experiments, SaGess, outperforms most of the one-shot state-of-the-art graph generating methods by a significant factor, both on the graph metrics and on the link prediction task.

Conditional Generators for Limit Order Book Environments: Explainability, Challenges, and Robustness

A. Coletta, J. Jerome, R. Savani, S. Vyetrenko

J.P.Morgan AI Research, University of Liverpool, Alan Turing Institute

Limit order books are a fundamental and widespread market mechanism. This paper investigates the use of conditional generative models for order book simulation. For developing a trading agent, this approach has drawn recent attention as an alternative to traditional backtesting due to its ability to react to the presence of the trading agent. Using a state-of-the-art CGAN (from Coletta et al. (2022)), we explore its dependence upon input features, which highlights both strengths and weaknesses. To do this, we use "adversarial attacks" on the model's features and its mechanism. We then show how these insights can be used to improve the CGAN, both in terms of its realism and robustness. We finish by laying out a roadmap for future work.

Sig-Splines: Universal Approximation and Convex Calibration of Time Series Generative Models

M. Wiese, P. Murray, R. Korn

University of Kaiserlautern, Imperial College London

We propose a novel generative model for multivariate discrete-time time series data. Drawing inspiration from the construction of neural spline flows, our algorithm incorporates linear transformations and the signature transform as a seamless substitution for traditional neural networks. This approach enables us to achieve not only the universality property inherent in neural networks but also introduces convexity in the model's parameters.

Deep Reinforcement Learning for Optimal Portfolio Allocation: A Comparative Study with Mean-Variance Optimization

S. Sood, K. Papasotiriou, M. Vaiciulis, T. Balch

J.P. Morgan AI Research, Oxford-Man Institute

Portfolio Management is the process of overseeing a group of investments, referred to as a portfolio, with the objective of achieving predetermined investment goals and objectives. Portfolio Optimization is a key component that involves allocating the portfolio assets so as to maximize returns while minimizing risk taken. Portfolio optimization is typically carried out by financial professionals who use a combination of quantitative techniques and investment expertise to make decisions about the portfolio allocation. Recent applications of Deep Reinforcement Learning (DRL) have shown promising results when used to optimize portfolio allocation by training model-free agents on historical market data. Many of these methods compare their results against basic benchmarks or other state-of-the-art DRL agents but often fail to compare their performance against traditional methods used by financial professionals in practical settings. One of the most commonly used methods for this task is Mean-Variance Portfolio Optimization (MVO), which uses historical timeseries information to estimate expected asset returns and covariances, which are then used to optimize for an investment objective. Our work is a thorough comparison between model-free DRL and MVO for optimal portfolio allocation. We

detail the specifics of how to make DRL for portfolio optimization work in practice, also noting the comparable tweaks needed to get MVO operational. Backtest results display strong improved performance of the DRL agent in terms of many performance metrics, including Sharpe ratio, maximum drawdowns, and absolute returns.

Mathematical Finance

Contagious McKean–Vlasov Problems with Common Noise: from Smooth to Singular Feedback through Hitting Times

B. Hambly, A. Petronilia, C. Reisinger, S. Rigger, A. Søjmark

University of Oxford, University of Vienna, London School of Economics

We consider a family of McKean-Vlasov equations arising as the large particle limit of a system of interacting particles on the positive half-line with common noise and feedback. Such systems are motivated by structural models for systemic risk with contagion. This contagious interaction is such that when a particle hits zero, the impact is to move all the others toward the origin through a kernel which smooths the impact over time. We study a rescaling of the impact kernel under which it converges to the Dirac delta function so that the interaction happens instantaneously and the limiting singular McKean–Vlasov equation can exhibit jumps. Our approach provides a novel method to construct solutions to such singular problems that allows for more general drift and diffusion coefficients and we establish weak convergence to relaxed solutions in this setting. With more restrictions on the coefficients we can establish an almost sure version showing convergence to strong solutions. Under some regularity conditions on the contagion, we also show a rate of convergence up to the time the regularity of the contagion breaks down. Lastly, we perform some numerical experiments to investigate the sharpness of our bounds for the rate of convergence.

Approximately Optimal Trade Execution Strategies under Fast Mean-Reversion

D. Evangelista, Y. Thamsten

Escola de Matemática Aplicada, Universidade Federal Fluminens

In a fixed time horizon, appropriately executing a large amount of a particular asset – meaning a considerable portion of the volume traded within this frame – is challenging. Especially for illiquid or even highly liquid but also highly volatile ones, the role of "market quality" is quite relevant in properly designing execution strategies. Here, we model it by considering uncertain volatility and liquidity; hence, moments of high or low price impact and risk vary randomly throughout the trading period. We work under the central assumption: although there are these uncertain variations, we assume they occur in a fast mean-reverting fashion. We thus employ singular perturbation arguments to study approximations to the optimal strategies in this framework. By using high-frequency data, we provide estimation methods for our model in face of microstructure noise, as well as numerically assess all of our results.

Estimation of an Order Book Dependent Hawkes Process for Large Datasets

L. Mucciante, A. Sancetta

Royal Holloway University of London

A point process for event arrivals in high frequency trading is presented. The intensity is the product of a Hawkes process and high dimensional functions of covariates derived from the order book. Conditions for stationarity of the process are stated. An algorithm is presented to estimate the model even in the presence of billions of data points, possibly mapping covariates into a high dimensional space. The large sample size can be common for high frequency data applications using multiple liquid instruments. Convergence of the algorithm is shown, consistency results under weak conditions is established, and a test statistic to assess out of sample performance of different model specifications is suggested. The methodology is applied to the study of four stocks that trade on the New York Stock Exchange (NYSE). The out of sample testing procedure suggests that capturing the nonlinearity of the order book information adds value to the self exciting nature of high frequency trading events.

Macroscopic Market Making

I. Guo, S. Jin, K. Nam

Monash University

We propose the macroscopic market making model à la Avellaneda-Stoikov, using continuous processes for orders instead of discrete point processes. The model intends to bridge a gap between market making and optimal execution problems, while shedding light on the influence of order flows on the strategy. We demonstrate our model through three problems. The study provides a comprehensive analysis from Markovian to non-Markovian noises and from linear to non-linear intensity functions, encompassing both bounded and unbounded coefficients. Mathematically, the contribution lies in the existence and uniqueness of the optimal control, guaranteed by the well-posedness of the Hamilton-Jacobi-Bellman equation or the (non-)Lipschitz forward-backward stochastic differential equation.

Economics

Algorithms, Incentives, and Democracy

E. M. Penn, J. Patty
Emory University

Classification algorithms are increasingly used in areas such as housing, credit, and law enforcement in order to make decisions affecting peoples' lives. These algorithms can change individual behavior deliberately (a fraud prediction algorithm deterring fraud) or inadvertently (content sorting algorithms spreading misinformation), and they are increasingly facing public scrutiny and regulation. Some of these regulations, like the elimination of cash bail in some states, have focused on *lowering the stakes of certain classifications*. In this paper we characterize how optimal classification by an algorithm designer can affect the distribution of behavior in a population – sometimes in surprising ways. We then look at the effect of democratizing the rewards and punishments, or stakes, to algorithmic classification to consider how a society can potentially stem (or facilitate!) predatory classification. Our results speak to questions of algorithmic fairness in settings where behavior and algorithms are interdependent, and where typical measures of fairness focusing on statistical accuracy across groups may not be appropriate.

Markov Persuasion Processes with Endogenous Agent Beliefs

K. Iyer, H. Xu, Y. Zu
University of Minnesota, University of Chicago

We consider a dynamic Bayesian persuasion setting where a single long-lived sender persuades a stream of "short-lived" agents (receivers) by sharing information about a payoff-relevant state. The state transitions are Markovian and the sender seeks to maximize the long-run average reward by committing to a (possibly history-dependent) signaling mechanism. While most previous studies of Markov persuasion consider exogenous agent beliefs that are independent of the chain, we study a more natural variant with endogenous agent beliefs that depend on the chain's realized history. A key challenge to analyze such settings is to model the agents' partial knowledge about the history information. We analyze a Markov persuasion process (MPP) under various information models that differ in the amount of information the receivers have about the history of the process. Specifically, we formulate a general partial-information model where each receiver observes the history with an ℓ period lag. Our technical contribution start with analyzing two benchmark models, i.e., the full-history information model and the no-history information model. We establish an ordering of the sender's payoff as a function of the informativeness of agent's information model (with no-history as the least informative), and develop efficient algorithms to compute optimal solutions for these two benchmarks. For general ℓ, we present the technical challenges in finding an optimal signaling mechanism, where even determining the right dependency on the history becomes difficult. To bypass the difficulties, we use a robustness framework to design a "simple" *history-independent* signaling mechanism that approximately achieves optimal payoff when $\ell is reasonably large$.

Strategic Budget Selection in a Competitive Autobidding World

Y. Feng, B. Lucier, A. Slivkins
University of Chicago, Microsoft Research

We study a game played between advertisers in an online ad platform. The platform sells ad impressions by first-price auction and provides autobidding algorithms that optimize bids on each advertiser's behalf. Each advertiser strategically declares a budget constraint (and possibly a maximum bid) to their autobidder. The chosen constraints define an "inner" budget-pacing game for the autobidders, who compete to maximize the total value received subject to the constraints. Advertiser payoffs in the constraint-choosing "metagame" are determined by the equilibrium reached by the autobidders. Advertisers only specify budgets and linear values to their autobidders, but their true preferences can be more general: we assume only that they have weakly decreasing marginal value for clicks and weakly increasing marginal disutility for spending money. Our main result is that despite this gap between general preferences and simple autobidder constraints, the allocations at equilibrium are approximately efficient. Specifically, at any pure Nash equilibrium of the metagame, the resulting allocation obtains at least half of the liquid welfare of any allocation and this bound is tight. We also obtain a 4-approximation for any mixed Nash equilibrium, and this result extends also to Bayes-Nash equilibria. These results rely on the power to declare budgets: if advertisers can specify only a (linear) value per click but not a budget constraint, the approximation factor at equilibrium can be as bad as linear in the number of advertisers.

A Belief-Based Characterization of Reduced-Form Auctions

X. Lang
Southwest University of Finance and Economics

We study games of chance (e.g., pokers, dices, horse races) in the form of agents' first-order posterior beliefs about game outcomes. We ask for any profile of agents' posterior beliefs, is there a game that can generate these beliefs? We completely characterize all

feasible joint posterior beliefs from these games. The characterization enables us to find a new variant of Border's inequalities (Border, 1991), which we call a belief-based characterization of Border's inequalities. It also leads to a generalization of Aumann's Agreement Theorem. We show that the characterization results are powerful in bounding the correlation of agents' joint posterior beliefs.

Asset Embeddings

X. Gabaix, R. S. J. Koijen, M. Yogo
Harvard University, University of Chicago, Princeton University

Firm characteristics are ubiquitously used in economics. These characteristics are often based on readily-available information such as accounting data, but those reflect only a part of investors' information set. We show that useful information about firm characteristics is embedded in investors' holdings data and, via market clearing, in prices, returns, and trading data. Based on insights from the recent artificial intelligence (AI) and machine learning (ML) literature, in which unstructured data (e.g., words or speech) are represented as continuous vectors in a potentially high-dimensional space, we propose to learn asset embeddings from investors' holdings data. Indeed, just as documents arrange words that can be used to uncover word structures via embeddings, investors organize assets in portfolios that can be used to uncover firm characteristics that investors deem important via asset embeddings. This broad theme provides a natural bridge to connect recent advances in the fields of AI and ML to finance and economics. Specifically, we show how language models, including transformer models that feature prominently in large language models such as BERT and GPT, can handle numerical information, and in particular holdings data to estimate asset embeddings. We provide initial evidence on the value added of asset embeddings through a series of applications in the context of firm valuations, return comovement, and uncovering asset substitution patterns. As a by-product, the models generate investor embeddings, which can be used to measure investor similarity. We propose a programmatic list of potential applications of asset and investor embeddings to finance and economics more generally.

Acknowledgement

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

References

Coletta, Andrea, Joseph Jerome, Rahul Savani, and Svitlana Vyetrenko. 2023. "Conditional Generators for Limit Order Book Environments: Explainability, Challenges, and Robustness." *arXiv preprint arXiv:2306.12806*.

Cont, Rama, Mihai Cucuringu, Jonathan Kochems, and Felix Prenzel. 2023. "Limit Order Book Simulation with Generative Adversarial Networks." *Available at SSRN 4512356*.

Evangelista, David, and Yuri Thamsten. 2023. "Approximately optimal trade execution strategies under fast mean-reversion." *arXiv preprint arXiv:2307.07024*.

Fang, Yuchen, Zhenggang Tang, Kan Ren, Weiqing Liu, Li Zhao, Jiang Bian, Dongsheng Li, Weinan Zhang, Yong Yu, and Tie-Yan Liu. 2023. "Learning Multi-Agent Intention-Aware Communication for Optimal Multi-Order Execution in Finance." *arXiv preprint arXiv:2307.03119*.

Feng, Yiding, Brendan Lucier, and Aleksandrs Slivkins. 2023. "Strategic Budget Selection in a Competitive Autobidding World." *arXiv preprint arXiv:2307.07374*.

Fuchs, Felix, and Blanka Horvath. 2023. "A Hybrid Quantum Wasserstein GAN with Applications to Option Pricing." *Available at SSRN 4514510*.

Gabaix, Xavier, Ralph SJ Koijen, and Motohiro Yogo. 2023. "Asset Embeddings." *Available at SSRN*.

Guo, Ivan, Shijia Jin, and Kihun Nam. 2023. "Macroscopic Market Making." *arXiv preprint arXiv:2307.14129*.

Hambly, Ben, Aldaïr Petronilia, Christoph Reisinger, Stefan Rigger, and Andreas Søjmark. 2023. "Contagious McKean–Vlasov problems with common noise: from smooth to singular feedback through hitting times." *arXiv preprint arXiv:2307.10800*.

Hoglund, Melker, Emilio Ferrucci, Camilo Hernandez, Aitor Muguruza Gonzalez, Cristopher Salvi, Leandro Sanchez-Betancourt, and Yufei Zhang. 2023. "A Neural RDE approach for continuous-time non-Markovian stochastic control problems." *arXiv preprint arXiv:2306.14258*.

Horvath, Blanka, and Zach Issa. 2023. "Non-parametric Online Market Regime Detection and Regime Clustering for Multidimensional and Path-Dependent Data Structures." *Available at SSRN 4493344*.

Iyer, Krishnamurthy, Haifeng Xu, and You Zu. 2023. "Markov Persuasion Processes with Endogenous Agent Beliefs." *arXiv preprint arXiv:2307.03181*.

Lang, Xu. 2023. "A Belief-Based Characterization of Reduced-Form Auctions." *arXiv preprint arXiv:2307.04070*.

Limnios, Stratis, Praveen Selvaraj, Mihai Cucuringu, Carsten Maple, Gesine Reinert, and Andrew Elliott. 2023. "SaGess: Sampling Graph Denoising Diffusion Model for Scalable Graph Generation." *arXiv preprint arXiv:2306.16827*.

Mucciante, Luca, and Alessio Sancetta. 2023. "Estimation of an Order Book Dependent Hawkes Process for Large Datasets." *arXiv preprint arXiv:2307.09077*.

Penn, Elizabeth Maggie, and John W Patty. 2023. "Algorithms, Incentives, and Democracy." *arXiv preprint arXiv:2307.02319*.

Sood, Srijan, Kassiani Papasotiriou, Marius Vaiciulis, and Tucker Balch. 2023. "Deep Reinforcement Learning for Optimal Portfolio Allocation: A Comparative Study with Mean-Variance Optimization." *FinPlan 2023*, 21.

Wiese, Magnus, Phillip Murray, and Ralf Korn. 2023. "Sig-Splines: universal approximation and convex calibration of time series generative models." *arXiv preprint arXiv:2307.09767*.