VOL. 3, NO. 6

OMI Research Newsletter – April 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, University of Cambridge, Imperial College London, University College London, London School of Economics, Columbia University, and Santa Clara University. Calvo-Ordonez et al. (2023) design a denoising network for diffusion models of continuous dynamical systems. Briola, Bartolucci and Aste (2024) develop an open-source library to study the predictability of limit order book prices with state-of-the-art deep learning models. Finally, Fox, Glosten and Guan (2024) study the legality of arbitrage strategies designed to manipulate prices at which market makers execute trades for retail clients.

Machine Learning & Statistics

The Missing U for Efficient Diffusion Models

S. CALVO-ORDONEZ, C-W. CHENG, J. HUANG, L. ZHANG, G. YANG, C-B. SCHONLIEB, A. I AVILES-RIVERO
OXFORD-MAN INSTITUTE, UNIVERSITY OF CAMBRIDGE, IMPERIAL COLLEGE LONDON

Diffusion Probabilistic Models stand as a critical tool in generative modelling, enabling the generation of complex data distributions. This family of generative models yields record-breaking performance in tasks such as image synthesis, video generation, and molecule design. Despite their capabilities, their efficiency, especially in the reverse process, remains a challenge due to slow convergence rates and high computational costs. In this paper, we introduce an approach that leverages continuous dynamical systems to design a novel denoising network for diffusion models that is more parameter-efficient, exhibits faster convergence, and demonstrates increased noise robustness. Experimenting with Denoising Diffusion Probabilistic Models (DDPMs), our framework operates with approximately a quarter of the parameters, and ∼30% of the Floating Point Operations (FLOPs) compared to standard U-Nets in DDPMs. Furthermore, our model is notably faster in inference than the baseline when measured in fair and equal conditions. We also provide a mathematical intuition as to why our proposed reverse process is faster as well as a mathematical discussion of the empirical tradeoffs in the denoising downstream task. Finally, we argue that our method is compatible with existing performance enhancement techniques, enabling further improvements in efficiency, quality, and speed.

Deep Limit Order Book Forecasting

A. BRIOLA, S. BARTOLUCCI, T. ASTE
UNIVERSITY COLLEGE LONDON, LONDON SCHOOL OF ECONOMICS

We exploit cutting-edge deep learning methodologies to explore the predictability of high-frequency Limit Order Book mid-price changes for a heterogeneous set of stocks traded on the NASDAQ exchange. In so doing, we release 'LOBFrame', an open-source code base to efficiently process large-scale Limit Order Book data and quantitatively assess state-of-the-art deep learning models' forecasting capabilities. Our results are twofold. We demonstrate that the stocks' microstructural characteristics influence the efficacy of deep learning methods and that their high forecasting power does not necessarily correspond to actionable trading signals. We argue that traditional machine learning metrics fail to adequately assess the quality of forecasts in the Limit Order Book context. As an alternative, we propose

an innovative operational framework that evaluates predictions' practicality by focusing on the probability of accurately forecasting complete transactions. This work offers academics and practitioners an avenue to make informed and robust decisions on the application of deep learning techniques, their scope and limitations, effectively exploiting emergent statistical properties of the Limit Order Book.

Robust Utility Optimization via a GAN Approach

F. Krach, J. Teichmann, H. Wutte
ETH Zürich

Robust utility optimization enables an investor to deal with market uncertainty in a structured way, with the goal of maximizing the worst-case outcome. In this work, we propose a generative adversarial network (GAN) approach to (approximately) solve robust utility optimization problems in general and realistic settings. In particular, we model both the investor and the market by neural networks (NN) and train them in a mini-max zero-sum game. This approach is applicable for any continuous utility function and in realistic market settings with trading costs, where only observable information of the market can be used. A large empirical study shows the versatile usability of our method. Whenever an optimal reference strategy is available, our method performs on par with it and in the (many) settings without known optimal strategy, our method outperforms all other reference strategies. Moreover, we can conclude from our study that the trained path-dependent strategies do not outperform Markovian ones. Lastly, we uncover that our generative approach for learning optimal, (non-) robust investments under trading costs generates universally applicable alternatives to well known asymptotic strategies of idealized settings.

Stable Neural Stochastic Differential Equations in Analyzing Irregular Time Series Data

Y. Oh, D. Lim, S. Kim
UNIST

Irregular sampling intervals and missing values in real-world time series data present challenges for conventional methods that assume consistent intervals and complete data. Neural Ordinary Differential Equations (Neural ODEs) offer an alternative approach, utilizing neural networks combined with ODE solvers to learn continuous latent representations through parameterized vector fields. Neural Stochastic Differential Equations (Neural SDEs) extend Neural ODEs by incorporating a diffusion term, although this addition is not trivial, particularly when addressing irregular intervals and missing values. Consequently, careful design of drift and diffusion functions is crucial for maintaining stability and enhancing performance, while incautious choices can result in adverse properties such as the absence of strong solutions, stochastic destabilization, or unstable Euler discretizations, significantly affecting Neural SDEs' performance. In this study, we propose three stable classes of Neural SDEs: Langevin-type SDE, Linear Noise SDE, and Geometric SDE. Then, we rigorously demonstrate their robustness in maintaining excellent performance under distribution shift, while effectively preventing overfitting. To assess the effectiveness of our approach, we conduct extensive experiments on four benchmark datasets for interpolation, forecasting, and classification tasks, and analyze the robustness of our methods with 30 public datasets under different missing rates. Our results demonstrate the efficacy of the proposed method in handling real-world irregular time series data.

In deep reinforcement learning, a pruned network is a good network

J. Obando-Ceron, A. Courville, P. Samuel Castro
Québec AI Institute, Université de Montréal

Recent work has shown that deep reinforcement learning agents have difficulty in effectively using their network parameters. We leverage prior insights into the advantages of sparse training techniques and demonstrate that gradual magnitude pruning enables agents to maximize parameter effectiveness. This results in networks that yield dramatic performance improvements over traditional networks and exhibit a type of "scaling law", using only a small fraction of the full network parameters.

Log Neural Controlled Differential Equations: The Lie Brackets Make a Difference

B. Walker, A. D. McLeod, T. Qin, Y. Cheng, H. Li, T. Lyons
University of Oxford, City University of Hong Kong

The vector field of a controlled differential equation (CDE) describes the relationship between a control path and the evolution of a solution path. Neural CDEs (NCDEs) treat time series data as observations from a control path, parameterise a CDE's vector field using a neural network, and use the solution path as a continuously evolving hidden state. As their formulation makes them robust to irregular sampling rates, NCDEs are a powerful approach for modelling real-world data. Building on neural rough differential equations (NRDEs), we introduce Log-NCDEs, a novel and effective method for training NCDEs. The core component of Log-NCDEs is the Log-ODE method, a tool from the study of rough paths for approximating a CDE's solution. On a range of multivariate time series classification benchmarks, Log-

NCDEs are shown to achieve a higher average test set accuracy than NCDEs, NRDEs, and two state-of-the-art models, S5 and the linear recurrent unit.

Finance and Economics

Manipulating Citadel: Profiting at the Expense of Retail Stock Traders' Market Makers

M. B. Fox, L. R. Glosten, S. S. Guan
Columbia University, Santa Clara University

This Article considers whether securities market strategies designed to profit at the expense of so-called "internalizers" should properly be considered illegal manipulation. An internalizer acquires from a brokerage firm the right to be the market maker for the broker's full order flow from its retail customers, promising in return to execute each order at a price slightly better than the best price available on any exchange ("price improvement") as well as to pay the broker a fee for each executed order ("payment for order flow"). Almost all retail trading — about 29% of the country's total share volume — is executed in this fashion, amounting in 2021 to about $41 trillion in transactions, a figure almost twice the nation's GDP that year. The internalizer can run a viable business while promising both price improvement and payment for order flow because retail traders rarely possess information not already reflected in price. This makes the buy and sell orders internalizers receive less dangerous to fill than the more varied order flow going to exchanges. The internalizer's business model, though, has a vulnerability: a trader can influence what is the best price available on the exchanges and then profit by sending an order to an internalizer that, as a result, executes at a price more favorable to her.

Using a framework that derives its key results from microstructure and financial economics, this Article seeks answers to four questions: (1) Exactly what actions in the market can traders take that would allow them to profit in this fashion? (2) What are the consequences to the various players in the market from traders undertaking such actions? (3) Would it be socially desirable to use legal prohibitions to try to prevent traders from profiting in this fashion? (4) How are such practices treated under existing law, and what reforms, if any, are desirable? The usual rhetoric concerning the evils of manipulation stresses its unfairness and its distortion of prices. This Article, however, concludes that strategies aimed at profiting off internalizers raise no serious fairness issues. Equally surprisingly, it concludes that if these strategies were freely occurring, they would probably indirectly marginally improve price accuracy. It is unlikely, however, that this effect would be more socially valuable than the practices' socially negative impact on liquidity. This negative social welfare assessment becomes that much bigger when one adds in the resources consumed by traders engaging in these strategies and by internalizers to protect against them, resources that otherwise would have been available to produce valuable goods and services for society. The status of these strategies under current case law is uncertain. If they are ultimately adjudicated to be legal, their use would expand greatly. The language of Sections 9(a)(2) and 10(b) of the Securities Exchange Act of 1934 and SEC Rule 10b-5 leave room, however, for the development of a coherent doctrine that definitively extends the Act's prohibitions against manipulation to cover these strategies. The analysis in this Article gives the courts good reasons to do so.

Learning Optimal Behavior Through Reasoning and Experiences

C. Ilut, R. Valchev
Duke University, Boston College

We develop a novel framework of bounded rationality under cognitive frictions that studies learning over optimal behavior through both deliberative reasoning and accumulated experiences. Using both types of information, agents engage in Bayesian non-parametric estimation of the unknown action value function. Reasoning signals are produced internally through mental deliberation, subject to a cognitive cost. Experience signals are the observed utility outcomes at previous actions. Agents' subjective estimation uncertainty, which evolves through information accumulation, modulates the two modes of learning in a state- and history-dependent way. We discuss how the model draws on and bridges conceptual, methodological and empirical insights from both economics and the cognitive sciences literature on reinforcement learning.

Mathematical Finance

Optimal positioning in derivative securities in incomplete markets

T. Leung, M. Lorig, Y. Shirai
University of Washington

This paper analyzes a problem of optimal static hedging using derivatives in incomplete markets. The investor is assumed to have a risk exposure to two underlying assets. The hedging instruments are vanilla options written on a single underlying asset. The hedging problem is formulated as a utility maximization problem whereby the form of the optimal static hedge is determined. Among our results, a semi-analytical solution for the optimizer is found through variational methods for exponential, power/logarithmic, and quadratic utility. When vanilla options are available for each underlying asset, the optimal solution is related to the fixed points of

a Lipschitz map. In the case of exponential utility, there is only one such fixed point, and subsequent iterations of the map converge to it.

am-AMM: An Auction-Managed Automated Market Maker

A. Adams, C. Moallemi, S. Reynolds, D. Robinson

Columbia University

Automated market makers (AMMs) have emerged as the dominant market mechanism for trading on decentralized exchanges implemented on blockchains. This paper presents a single mechanism that targets two important unsolved problems for AMMs: reducing losses to informed orderflow, and maximizing revenue from uninformed orderflow. The "auction-managed AMM" works by running a censorship-resistant onchain auction for the right to temporarily act as "pool manager" for a constant-product AMM. The pool manager sets the swap fee rate on the pool, and also receives the accrued fees from swaps. The pool manager can exclusively capture some arbitrage by trading against the pool in response to small price movements, and also can set swap fees incorporating price sensitivity of retail orderflow and adapting to changing market conditions, with the benefits from both ultimately accruing to liquidity providers. Liquidity providers can enter and exit the pool freely in response to changing rent, though they must pay a small fee on withdrawal. We prove that under certain assumptions, this AMM should have higher liquidity in equilibrium than any standard, fixed-fee AMM.

Fill Probabilities in a Limit Order Book with State-Dependent Stochastic Order Flows

F. Lokin, F. Yu

TU Delft

This paper focuses on computing the fill probabilities for limit orders positioned at various price levels within the limit order book, which play a crucial role in optimizing executions. We adopt a generic stochastic model to capture the dynamics of the order book as a series of queueing systems. This generic model is state-dependent and also incorporates stylized factors. We subsequently derive semi-analytical expressions to compute the relevant probabilities within the context of state-dependent stochastic order flows. These probabilities cover various scenarios, including the probability of a change in the mid-price, the fill probabilities of orders posted at the best quotes, and those posted at a price level deeper than the best quotes in the book, before the opposite best quote moves. These expressions can be further generalized to accommodate orders posted even deeper in the order book, although the associated probabilities are typically very small in such cases. Lastly, we conduct extensive numerical experiments using real order book data from the foreign exchange spot market. Our findings suggest that the model is tractable and possesses the capability to effectively capture the dynamics of the limit order book. Moreover, the derived formulas and numerical methods demonstrate reasonably good accuracy in estimating the fill probabilities.

Agents' Behavior and Interest Rate Model Optimization in DeFi Lending

C. Bertucci, L. Bertucci, O. Guéant

École Polytechnique, Université Paris 1 Panthéon-Sorbonne, Institut Louis Bachelier,

In the past five years, Decentralized Finance (DeFi) lending has seen rapid growth, leading to a substantial market with dozens of billions of USD in various cryptocurrencies currently borrowed as of the start of 2024. Contrasting sharply with traditional money, bond and bond futures markets, where interest rates emerge organically from participant interactions, DeFi lending platforms employ rule-based interest rates that are algorithmically set. Thus, the selection of an effective interest rate model is paramount for the success of a lending protocol. This paper delves into the modeling of agents' behaviors on lending platforms, proposing a theoretical framework for formulating optimal interest rate models. Furthermore, we examine streamlined interest rate models, assessing their characteristics and practical implications.

A Mean-Field Game of Market Entry: Portfolio Liquidation with Trading Constraints

G. Fu, P. P. Hager, U. Horst

The Hong Kong Polytechnic University, Technische Universitat Berlin, Humboldt University Berlin,

We consider both N-player and mean-field games of optimal portfolio liquidation in which the players are not allowed to change the direction of trading. Players with an initially short position of stocks are only allowed to buy while players with an initially long position are only allowed to sell the stock. Under suitable conditions on the model parameters we show that the games are equivalent to games of timing where the players need to determine the optimal times of market entry and exit. We identify the equilibrium entry and exit times and prove that equilibrium mean-trading rates can be characterized in terms of the solutions to a highly non-linear higher-order integral equation with endogenous terminal condition. We prove the existence of a unique solution to the integral equation from which we obtain the existence of a unique equilibrium both in the mean-field and the N-player game.

Optimal Portfolio Choice with Cross-Impact Propagators

E. Abi Jaber, E. Neuman, S. Tuschmann
Ecole Polytechnique, Imperial College London

We consider a class of optimal portfolio choice problems in continuous time where the agent's transactions create both transient cross-impact driven by a matrix-valued Volterra propagator, as well as temporary price impact. We formulate this problem as the maximization of a revenue-risk functional, where the agent also exploits available information on a progressively measurable price predicting signal. We solve the maximization problem explicitly in terms of operator resolvents, by reducing the corresponding first order condition to a coupled system of stochastic Fredholm equations of the second kind and deriving its solution. We then give sufficient conditions on the matrix-valued propagator so that the model does not permit price manipulation. We also provide an implementation of the solutions to the optimal portfolio choice problem and to the associated optimal execution problem. Our solutions yield financial insights on the influence of cross-impact on the optimal strategies and its interplay with alpha decays.

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

Abi Jaber, Eduardo, Eyal Neuman, and Sturmius Tuschmann. 2024. "Optimal Portfolio Choice with Cross-Impact Propagators." *Available at SSRN*.

Adams, Austin, Ciamac Moallemi, Sara Reynolds, and Dan Robinson. 2024. "am-AMM: An Auction-Managed Automated Market Maker." *arXiv preprint arXiv:2403.03367*.

Bertucci, Charles, Louis Bertucci, and Olivier Guéant. n.d.. "Agents' Behavior and Interest Rate Model Optimization in DeFi Lending."

Briola, Antonio, Silvia Bartolucci, and Tomaso Aste. 2024. "Deep Limit Order Book Forecasting." *arXiv preprint arXiv:2403.09267*.

Calvo-Ordonez, Sergio, Jiahao Huang, Lipei Zhang, Guang Yang, Carola-Bibiane Schonlieb, and Angelica I Aviles-Rivero. 2023. "The Missing U for Efficient Diffusion Models." *arXiv preprint arXiv:2310.20092*.

Fox, Merritt B, Lawrence R Glosten, and Sue S Guan. 2024. "Manipulating Citadel: Profiting at the Expense of Retail Stock Traders' Market Makers." *Harv. Bus. L. Rev.*, 14: 1.

Fu, Guanxing, Paul P Hager, and Ulrich Horst. 2024. "A Mean-Field Game of Market Entry: Portfolio Liquidation with Trading Constraints." *arXiv preprint arXiv:2403.10441*.

Ilut, Cosmin, and Rosen Valchev. 2024. "Learning Optimal Behavior Through Reasoning and Experiences." *arXiv preprint arXiv:2403.18185*.

Krach, Florian, Josef Teichmann, and Hanna Wutte. 2024. "Robust Utility Optimization via a GAN Approach." *arXiv preprint arXiv:2403.15243*.

Leung, Tim, Matthew Lorig, and Yoshihiro Shirai. 2024. "Optimal positioning in derivative securities in incomplete markets." *arXiv preprint arXiv:2403.00139*.

Lokin, Felix, and Fenghui Yu. 2024. "Fill Probabilities in a Limit Order Book with State-Dependent Stochastic Order Flows." *arXiv preprint arXiv:2403.02572*.

Obando-Ceron, Johan, Aaron Courville, and Pablo Samuel Castro. 2024. "In deep reinforcement learning, a pruned network is a good network." *arXiv preprint arXiv:2402.12479*.

Oh, YongKyung, Dongyoung Lim, and Sungil Kim. 2024. "Stable Neural Stochastic Differential Equations in Analyzing Irregular Time Series Data." *arXiv preprint arXiv:2402.14989*.

Walker, Benjamin, Andrew D McLeod, Tiexin Qin, Yichuan Cheng, Haoliang Li, and Terry Lyons. 2024. "Log Neural Controlled Differential Equations: The Lie Brackets Make a Difference." *arXiv preprint arXiv:2402.18512*.