OMI Research Newsletter – October 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 advancements in interpretable language models for analysing financial narratives (Cao et al., 2024), robust reinforcement learning frameworks for risk management (Coache and Jaimungal, 2024), the implications of AI for scientific discovery (Ghafarollahi and Buehler, 2024) and its impact on consumer behaviour (Werner et al., 2024), new empirical findings on price jumps, news, and earnings announcements (Aït-Sahalia, Li and Li, 2024; Christensen, Timmermann and Veliyev, 2024), an innovative approach to characteristic-based factor portfolio construction (Kagkadis et al., 2024), and pricing of liquidity positions in automated market makers in decentralised finance (Nadkarni, Kulkarni and Viswanath, 2024), etc.

Machine Learning & Statistics

WhiteBox inside Blackbox: Using Interpretable Language Models to Analyze Financial Narratives

S. CAO, Y. KIM, A. WANG, H. XIAO
UNIVERSITY OF MARYLAND, SANTA CLARA UNIVERSITY, THE HONG KONG POLYTECHNIC UNIVERSITY, GEORGIA STATE UNIVERSITY

This paper applies technological innovation (interpretable language models as a component function of large language models (LLMs)) to financial innovation (extracting context-specific information from unstructured financial narrative data with a high level of interpretability). Rapid technological advances spur financial innovations based on the analysis of unstructured texts. While LLMs such as Google BERT and GPT can capture the sequence and relationship between words, they are black-box models. Users do not know the exact sequences or relationships they capture. This paper introduces an interpretable language model that incorporates complex syntactical features by integrating two machine learning techniques (i.e., neural-network-based natural language parsing and word embedding) in the training process. While our model is part of the LLM ecosystem, the interpretable component of it reveals the exact textual features, offering interpretability which is lacking in general LLMs. We demonstrate the usefulness of this interpretable language model by constructing a new sentiment measure of financial narratives in earnings conference calls that is specific to performance discussions and adjusted for complex contextual negations. We find that this performance-specific sentiment explains cross-sectional returns and future operating performance better than the commonly used umbrella sentiment proxies and simple rule-based measures. As a second application, an analysis of earnings-related forward-looking statements in conference calls also confirms the value of this new approach in better utilizing LLMs.

Robust Reinforcement Learning with Dynamic Distortion Risk Measures

A. COACHE, S. JAIMUNGAL
IMPERIAL COLLEGE LONDON, UNIVERSITY OF TORONTO, OXFORD-MAN INSTITUTE

In a reinforcement learning (RL) setting, the agent's optimal strategy heavily depends on her risk preferences and the underlying model dynamics of the training environment. These two aspects influence the agent's ability to make well-informed and time-consistent

decisions when facing testing environments. In this work, we devise a framework to solve robust risk-aware RL problems where we simultaneously account for environmental uncertainty and risk with a class of dynamic robust distortion risk measures. Robustness is introduced by considering all models within a Wasserstein ball around a reference model. We estimate such dynamic robust risk measures using neural networks by making use of strictly consistent scoring functions, derive policy gradient formulae using the quantile representation of distortion risk measures, and construct an actor-critic algorithm to solve this class of robust risk-aware RL problems. We demonstrate the performance of our algorithm on a portfolio allocation example.

SciAgents: Automating Scientific Discovery through Multi-Agent Intelligent Graph Reasoning

A. Ghafarollahi, M. J. Buehler
University of Zurich

A key challenge in artificial intelligence is the creation of systems capable of autonomously advancing scientific understanding by exploring novel domains, identifying complex patterns, and uncovering previously unseen connections in vast scientific data. In this work, we present SciAgents, an approach that leverages three core concepts: (1) the use of large-scale ontological knowledge graphs to organize and interconnect diverse scientific concepts, (2) a suite of large language models (LLMs) and data retrieval tools, and (3) multi-agent systems with in-situ learning capabilities. Applied to biologically inspired materials, SciAgents reveals hidden interdisciplinary relationships that were previously considered unrelated, achieving a scale, precision, and exploratory power that surpasses traditional human-driven research methods. The framework autonomously generates and refines research hypotheses, elucidating underlying mechanisms, design principles, and unexpected material properties. By integrating these capabilities in a modular fashion, the intelligent system yields material discoveries, critique and improve existing hypotheses, retrieve up-to-date data about existing research, and highlights their strengths and limitations. Our case studies demonstrate scalable capabilities to combine generative AI, ontological representations, and multi-agent modeling, harnessing a 'swarm of intelligence' similar to biological systems. This provides new avenues for materials discovery and accelerates the development of advanced materials by unlocking Nature's design principles.

LLM Processes: Numerical Predictive Distributions Conditioned on Natural Language

J. Requeima, J. Bronskill, D. Choi, R. E. Turner, D. Duvenaud
University of Toronto, University of Cambridge, Vector Institute

Machine learning practitioners often face significant challenges in formally integrating their prior knowledge and beliefs into predictive models, limiting the potential for nuanced and context-aware analyses. Moreover, the expertise needed to integrate this prior knowledge into probabilistic modeling typically limits the application of these models to specialists. Our goal is to build a regression model that can process numerical data and make probabilistic predictions at arbitrary locations, guided by natural language text which describes a user's prior knowledge. Large Language Models (LLMs) provide a useful starting point for designing such a tool since they 1) provide an interface where users can incorporate expert insights in natural language and 2) provide an opportunity for leveraging latent problem-relevant knowledge encoded in LLMs that users may not have themselves. We start by exploring strategies for eliciting explicit, coherent numerical predictive distributions from LLMs. We examine these joint predictive distributions, which we call LLM Processes, over arbitrarily-many quantities in settings such as forecasting, multi-dimensional regression, black-box optimization, and image modeling. We investigate the practical details of prompting to elicit coherent predictive distributions, and demonstrate their effectiveness at regression. Finally, we demonstrate the ability to usefully incorporate text into numerical predictions, improving predictive performance and giving quantitative structure that reflects qualitative descriptions. This lets us begin to explore the rich, grounded hypothesis space that LLMs implicitly encode.

Finance and Economics

So Many Jumps, So Few News

Y. Aït-Sahalia, C. X. Li, C. Li
Princeton University, Peking University, Renmin University of China

This paper relates jumps in high frequency stock prices to firm-level, industry and macroeconomic news, in the form of machine-readable releases from Thomson Reuters News Analytics. We find that most relevant news, both idiosyncratic and systematic, lead quickly to price jumps, as market efficiency suggests they should. However, in the reverse direction, the vast majority of price jumps do not have identifiable public news that can explain them, in a departure from the ideal of a fair, orderly and efficient market. Microstructure-driven variables have only limited predictive power to help distinguish between jumps with and without news.

Warp Speed Price Moves: Jumps after Earnings Announcements

K. Christensen, A. Timmermann, B. Veliyev
Aarhus University, University of California, San Diego

Corporate earnings announcements unpack large bundles of information that should, in efficient markets, trigger jumps in stock prices. Testing this implication is difficult in practice, as it requires noisy high-frequency data from after-hours markets, where most earnings announcements are released. With a unique dataset and a new microstructure noise-robust jump test, we show that earnings announcements almost always induce jumps in the stock price of announcing firms. They also significantly raise the probability of price co-jumps in non-announcing firms and the market. We find that returns from a post-announcement trading strategy are consistent with an efficient price formation after 2016.

Power Sorting

A. Kagkadis, H. Lohre, I. Nolte, s. Nolte, N. Vasilas
University of Liverpool, Robeco Quantitative Investments, Lancaster University, Ultramarin

We propose a novel approach for constructing characteristic-based equity factors, termed "power sorting". Power sorting exploits the non-linearities and asymmetries inherent in the characteristic-return relations, while it remains computationally simple and avoids excessive weights. We demonstrate that power sorting achieves consistently superior out-of-sample performance compared to traditional quantile sorting and other factor portfolio construction methods. Our results are pervasive across factors, robust through time, cannot be attributed to increased turnover or tail risk, and extend to multi-factor strategies. Finally, we show that power sorted versions of well-known asset pricing factor models outperform the original ones.

Predicting Anomalies

B. Bowles, A. V. Reed, M. C. Ringgenberg, J. Thornock
Texas A&M University, University of North Carolina, University of Utah, Brigham Young University

We show that stock returns follow predictable patterns before the publication of anomaly trading signals. Moreover, anomaly trading signals derived from financial data are themselves predictable, making it possible to trade before financial statements are released. A trading strategy based on predicted signals earns an annualized return of 3.65% in the quarter before the signal is released. In recent periods this predictability is concentrated in signals that are harder to forecast and returns are increasingly earned several quarters before signals are released. Our findings suggest anomalies are more anomalous than previously recognized.

Deep Learning to Play Games

D. Condorelli, M. Furlan
University of Warwick

We train two neural networks adversarially to play normal-form games. At each iteration, a row and column network take a new randomly generated game and output individual mixed strategies. The parameters of each network are independently updated via stochastic gradient descent to minimize expected regret given the opponent's strategy. Our simulations demonstrate that the joint behavior of the networks converges to strategies close to Nash equilibria in almost all games. For all 22 and in 80% of 33 games with multiple equilibria, the networks select the risk-dominant equilibrium. Our results show how Nash equilibrium emerges from learning across heterogeneous games.

Experimental Evidence That Conversational Artificial Intelligence Can Steer Consumer Behavior Without Detection

T. Werner, I. Soraperra, E. Calvano, D. C. Parkes, I. Rahwan
Max Planck Institute for Human Development, Luiss University, Harvard University

Conversational AI models are becoming increasingly popular and are about to replace traditional search engines for information retrieval and product discovery. This raises concerns about monetization strategies and the potential for subtle consumer manipulation. Companies may have financial incentives to steer users toward search results or products in a conversation in ways that are unnoticeable to consumers. Using a behavioral experiment, we show that conversational AI models can indeed significantly shift consumer preferences. We discuss implications and ask whether regulators are sufficiently prepared to combat potential consumer deception.

Mathematical Finance

Optimal Investment with Costly Expert Opinions

C. Knochenhauer, A. Merkel, Y. Zhang
Technical University of Munich, Technische Universität Berlin, Imperial College London

We consider the Merton problem of optimizing expected power utility of terminal wealth in the case of an unobservable Markov-modulated drift. What makes the model special is that the agent is allowed to purchase costly expert opinions of varying quality on the current

state of the drift, leading to a mixed stochastic control problem with regular and impulse controls involving random consequences. Using ideas from filtering theory, we first embed the original problem with unobservable drift into a full information problem on a larger state space. The value function of the full information problem is characterized as the unique viscosity solution of the dynamic programming PDE. This characterization is achieved by a new variant of the stochastic Perron's method, which additionally allows us to show that, in between purchases of expert opinions, the problem reduces to an exit time control problem which is known to admit an optimal feedback control. Under the assumption of sufficient regularity of this feedback map, we are able to construct optimal trading and expert opinion strategies.

Adaptive Curves for Optimally Efficient Market Making

V. Nadkarni, S. Kulkarni, P. Viswanath
Princeton University

Automated Market Makers (AMMs) are essential in Decentralized Finance (DeFi) as they match liquidity supply with demand. They function through liquidity providers (LPs) who deposit assets into liquidity pools. However, the asset trading prices in these pools often trail behind those in more dynamic, centralized exchanges, leading to potential arbitrage losses for LPs. This issue is tackled by adapting market maker bonding curves to trader behavior, based on the classical market microstructure model of Glosten and Milgrom. Our approach ensures a zero-profit condition for the market maker's prices. We derive the differential equation that an optimal adaptive curve should follow to minimize arbitrage losses while remaining competitive. Solutions to this optimality equation are obtained for standard Gaussian and Lognormal price models using Kalman filtering. A key feature of our method is its ability to estimate the external market price without relying on price or loss oracles. We also provide an equivalent differential equation for the implied dynamics of canonical static bonding curves and establish conditions for their optimality. Our algorithms demonstrate robustness to changing market conditions and adversarial perturbations, and we offer an on-chain implementation using Uniswap v4 alongside off-chain AI co-processors.

Market Simulation under Adverse Selection

L. Lalor, A. Swishchuk
University Of Calgary

In this paper, we study the effects of fill probabilities and adverse fills on the trading strategy simulation process. We specifically focus on a stochastic optimal control market-making problem and test the strategy on ES (E-mini SP 500), NQ (E-mini Nasdaq 100), CL (Crude Oil) and ZN (10-Year Treasury Note), which are some of the most liquid futures contract listed on the CME (Chicago Mercantile Exchange). We provide empirical evidence which shows how fill probabilities and adverse fills can significantly effect performance, and propose a more prudent simulation framework for dealing with this. Many previous works aim to measure different types of adverse selection in the limit order book, however, they often simulate price processes and market orders independently. This has the ability to largely inflate the performance of a short-term style trading strategy. Our studies show that using more realistic fill probabilities, and tracking adverse fills, in the strategy simulation process, more accurately portrays how these types of trading strategies would perform in reality.

How Does Liquidity Shape the Yield Curve?

V. Le Coz, I. Mastromatteo, M. Benzaquen
École polytechnique, Université Paris-Saclay, Quant AI Lab, Capital Fund Management

The phenomenology of the forward rate curve (FRC) can be accurately understood by the fluctuations of a stiff elastic string (Le Coz and Bouchaud, 2024). By relating the exogenous shocks driving such fluctuations to the surprises in the order flows, we elevate the model from purely describing price variations to a microstructural model that incorporates the joint dynamics of prices and order flows, accounting for both impact and cross-impact effects. Remarkably, this framework allows for at least the same explanatory power as existing cross-impact models, while using significantly fewer parameters. In addition, our model generates liquidity-dependent correlations between the forward rate of one tenor and the order flow of another, consistent with recent empirical findings. We show that the model also account for the non-martingale behavior of prices at short timescales.

Exploratory Optimal Stopping: A Singular Control Formulation

J. Dianetti, G. Ferrari, R. Xu
Bielefeld University, New York University

This paper explores continuous-time and state-space optimal stopping problems from a reinforcement learning perspective. We begin by formulating the stopping problem using randomized stopping times, where the decision maker's control is represented by the probability of stopping within a given time–specifically, a bounded, non-decreasing, càdlàg control process. To encourage exploration and facilitate learning, we introduce a regularized version of the problem by penalizing it with the cumulative residual entropy of the randomized stopping time. The regularized problem takes the form of an $(n+1)$-dimensional degenerate singular stochastic control with finite-fuel. We address this through the

dynamic programming principle, which enables us to identify the unique optimal exploratory strategy. For the specific case of a real option problem, we derive a semi-explicit solution to the regularized problem, allowing us to assess the impact of entropy regularization and analyze the vanishing entropy limit. Finally, we propose a reinforcement learning algorithm based on policy iteration. We show both policy improvement and policy convergence results for our proposed algorithm.

DeFi Arbitrage in Hedged Liquidity Tokens

M. Bichuch, Z. Feinstein
SUNY University at Buffalo, Stevens Institute of Technology

Empirically, the prevailing market prices for liquidity tokens of the constant product market maker (CPMM) – as offered in practice by companies such as Uniswap – readily permit arbitrage opportunities by delta hedging the risk of the position. Herein, we investigate this arbitrage opportunity by treating the liquidity token as a derivative position in the prices of the underlying assets for the CPMM. In doing so, not dissimilar to the Black-Scholes result, we deduce risk-neutral pricing and hedging formulas for these liquidity tokens. Furthermore, with our novel pricing formula, we construct a method to calibrate a volatility to data which provides an updated (non-market) price which would not permit arbitrage if quoted by the CPMM. We conclude with a discussion of novel AMM designs which would bring the pricing of liquidity tokens into the modern financial era.

Acknowledgement

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

References

Aït-Sahalia, Yacine, Chen Xu Li, and Chenxu Li. 2024. "So Many Jumps, So Few News." *NBER Working Paper*.

Bichuch, Maxim, and Zachary Feinstein. 2024. "DeFi Arbitrage in Hedged Liquidity Tokens." *arXiv preprint arXiv:2409.11339*.

Bowles, Boone, Adam V. Reed, Matthew C. Ringgenberg, and Jake Thornock. 2024. "Predicting Anomalies." *Available at SSRN 4939779*.

Cao, Sean, Yongtae Kim, Angie Wang, and Houping Xiao. 2024. "WhiteBox inside Blackbox: Using Interpretable Language Models to Analyze Financial Narratives." *Available at SSRN 3568504*.

Christensen, Kim, Allan Timmermann, and Bezirgen Veliyev. 2024. "Warp Speed Price Moves: Jumps after Earnings Announcements." *Available at SSRN 4422376*.

Coache, Anthony, and Sebastian Jaimungal. 2024. "Robust Reinforcement Learning with Dynamic Distortion Risk Measures." *arXiv preprint arXiv:2409.10096*.

Condorelli, Daniele, and Massimiliano Furlan. 2024. "Deep Learning to Play Games." *arXiv preprint arXiv:2409.15197*.

Dianetti, Jodi, Giorgio Ferrari, and Renyuan Xu. 2024. "Exploratory Optimal Stopping: A Singular Control Formulation." *arXiv preprint arXiv:2408.09335*.

Ghafarollahi, Alireza, and Markus J Buehler. 2024. "SciAgents: Automating Scientific Discovery through Multi-Agent Intelligent Graph Reasoning." *arXiv preprint arXiv:2409.05556*.

Kagkadis, Anastasios, Harald Lohre, Ingmar Nolte, Sandra Nolte, and Nikolaos Vasilas. 2024. "Power Sorting." *Available at SSRN 4552208*.

Knochenhauer, Christoph, Alexander Merkel, and Yufei Zhang. 2024. "Optimal Investment with Costly Expert Opinions." *arXiv preprint arXiv:2409.11569*.

Lalor, Luca, and Anatoliy Swishchuk. 2024. "Market Simulation under Adverse Selection." *arXiv preprint arXiv:2409.12721*.

Le Coz, Victor, Iacopo Mastromatteo, and Michael Benzaquen. 2024. "How Does Liquidity Shape the Yield Curve?" *arXiv preprint arXiv:2409.12282*.

Nadkarni, Viraj, Sanjeev Kulkarni, and Pramod Viswanath. 2024. "Adaptive Curves for Optimally Efficient Market Making." *arXiv preprint arXiv:2406.13794*.

Requeima, James, John Bronskill, Dami Choi, Richard E Turner, and David Duvenaud. 2024. "LLM Processes: Numerical Predictive Distributions Conditioned on Natural Language." *arXiv preprint arXiv:2405.12856*.

Werner, Tobias, Ivan Soraperra, Emilio Calvano, David C Parkes, and Iyad Rahwan. 2024. "Experimental Evidence That Conversational Artificial Intelligence Can Steer Consumer Behavior Without Detection." *arXiv preprint arXiv:2409.12143*.