OMI Research Newsletter – November 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.*

Announcements

The Workshop on New Challenges and Risks in Finance will take place on 14 November 2024 at the Man Group (London). Click here for a schedule and to register.

The Oxford–Harvard Conference on Decentralised Finance will take place on 20 and 21 January 2025 at the Oxford-Man Institute. Click here for a schedule and to register.

Highlights

Key contributions in this month's edition include a comprehensive review of the impact of generative AI, such as ChatGPT, on firm valuation and financial research methodologies (Eisfeldt and Schubert, 2024), an AI-assisted approach to identifying instrumental variables in social science research (Han, 2024), a differential calculus framework for analysing generalisation errors in transfer learning models (Aminian, Szpruch and Cohen, 2024), a scalable approach to distribution regression (Alden, Ventre and Horvath, 2024), option-based lower bounds for stock crash probabilities (Martin and Shi, 2024), momentum spillovers in option returns across peer firms (Li and Yang, 2024), and inefficiencies in decentralised finance markets (Bodisz, Hautsch and Voigt, 2024), etc.

Machine Learning & Statistics

Mining Causality: AI-Assisted Search for Instrumental Variables

S. HAN
UNIVERSITY OF BRISTOL

The instrumental variables (IVs) method is a leading empirical strategy for causal inference. Finding IVs is a heuristic and creative process, and justifying its validity (especially exclusion restrictions) is largely rhetorical. We propose using large language models (LLMs) to search for new IVs through narratives and counterfactual reasoning, similar to how a human researcher would. The stark difference, however, is that LLMs can accelerate this process exponentially and explore an extremely large search space. We demonstrate how to construct prompts to search for potentially valid IVs. We argue that multi-step prompting is useful and role-playing prompts are suitable for mimicking the endogenous decisions of economic agents. We apply our method to three well-known examples in economics: returns to schooling, production functions, and peer effects. We then extend our strategy to finding (i) control variables in regression and difference-in-differences and (ii) running variables in regression discontinuity designs.

Understanding Transfer Learning via Mean-field Analysis

G. AMINIAN, Ł. SZPRUCH, S. N. COHEN
ALAN TURING INSTITUTE, UNIVERSITY OF OXFORD, UNIVERSITY OF EDINBURGH

We propose a novel framework for exploring generalization errors of transfer learning through the lens of differential calculus on the space of probability measures. In particular, we consider two main transfer learning scenarios, α-ERM and fine-tuning with the KL-regularized empirical risk minimization and establish generic conditions under which the generalization error and the population risk convergence rates for these scenarios are studied. Based on our theoretical results, we show the benefits of transfer learning with a one-hidden-

layer neural network in the mean-field regime under some suitable integrability and regularity assumptions on the loss and activation functions.

Scalable Signature-Based Distribution Regression via Reference Sets

A. Alden, C. Ventre, B. Horvath
King's College London, University of Oxford

Distribution Regression (DR) on stochastic processes describes the learning task of regression on collections of time series. Path signatures, a technique prevalent in stochastic analysis, have been used to solve the DR problem. Recent works have demonstrated the ability of such solutions to leverage the information encoded in paths via signature-based features. However, current state of the art DR solutions are memory intensive and incur a high computation cost. This leads to a trade-off between path length and the number of paths considered. This computational bottleneck limits the application to small sample sizes which consequently introduces estimation uncertainty. In this paper, we present a methodology for addressing the above issues; resolving estimation uncertainties whilst also proposing a pipeline that enables us to use DR for a wide variety of learning tasks. Integral to our approach is our novel distance approximator. This allows us to seamlessly apply our methodology across different application domains, sampling rates, and stochastic process dimensions. We show that our model performs well in applications related to estimation theory, quantitative finance, and physical sciences. We demonstrate that our model generalises well, not only to unseen data within a given distribution, but also under unseen regimes (unseen classes of stochastic models).

Unsupervised-to-Online Reinforcement Learning

J. Kim, S. Park, S. Levine
KAIST, University of California, Berkeley

Offline-to-online reinforcement learning (RL), a framework that trains a policy with offline RL and then further fine-tunes it with online RL, has been considered a promising recipe for data-driven decision-making. While sensible, this framework has drawbacks: it requires domain-specific offline RL pre-training for each task, and is often brittle in practice. In this work, we propose unsupervised-to-online RL (U2O RL), which replaces domain-specific supervised offline RL with unsupervised offline RL, as a better alternative to offline-to-online RL. U2O RL not only enables reusing a single pre-trained model for multiple downstream tasks, but also learns better representations, which often result in even better performance and stability than supervised offline-to-online RL. To instantiate U2O RL in practice, we propose a general recipe for U2O RL to bridge task-agnostic unsupervised offline skill-based policy pre-training and supervised online fine-tuning. Throughout our experiments in nine state-based and pixel-based environments, we empirically demonstrate that U2O RL achieves strong performance that matches or even outperforms previous offline-to-online RL approaches, while being able to reuse a single pre-trained model for a number of different downstream tasks.

When to Sense and Control? A Time-adaptive Approach for Continuous-Time RL

L. Treven, B. Sukhija, Y. As, F. Dörfler, A. Krause
ETH Zurich

Reinforcement learning (RL) excels in optimizing policies for discrete-time Markov decision processes (MDP). However, various systems are inherently continuous in time, making discrete-time MDPs an inexact modeling choice. In many applications, such as greenhouse control or medical treatments, each interaction (measurement or switching of action) involves manual intervention and thus is inherently costly. Therefore, we generally prefer a time-adaptive approach with fewer interactions with the system. In this work, we formalize an RL framework, Time-adaptive Control & Sensing (TaCoS), that tackles this challenge by optimizing over policies that besides control predict the duration of its application. Our formulation results in an extended MDP that any standard RL algorithm can solve. We demonstrate that state-of-the-art RL algorithms trained on TaCoS drastically reduce the interaction amount over their discrete-time counterpart while retaining the same or improved performance, and exhibiting robustness over discretization frequency. Finally, we propose OTaCoS, an efficient model-based algorithm for our setting. We show that OTaCoS enjoys sublinear regret for systems with sufficiently smooth dynamics and empirically results in further sample-efficiency gains.

NeoRL: Efficient Exploration for Nonepisodic RL

B. Sukhija, L. Treven, F. Dörfler, S. Coros, A. Krause
ETH Zurich

We study the problem of nonepisodic reinforcement learning (RL) for nonlinear dynamical systems, where the system dynamics are unknown and the RL agent has to learn from a single trajectory, i.e., without resets. We propose Nonepisodic Optimistic RL (NeoRL), an approach based on the principle of optimism in the face of uncertainty. NeoRL uses well-calibrated probabilistic models and plans optimistically w.r.t. the epistemic uncertainty about the unknown dynamics. Under continuity and bounded energy assumptions on the system, we provide a first-of-its-kind regret

bound of $(\beta_T\sqrt{TT_T})$ for general nonlinear systems with Gaussian process dynamics. We compare NeoRL to other baselines on several deep RL environments and empirically demonstrate that NeoRL achieves the optimal average cost while incurring the least regret.

Strategic Collusion of LLM Agents: Market Division in Multi-Commodity Competitions

R. Y. Lin, S. Ojha, K. Cai, M. F. Chen
California Institute of Technology

Machine-learning technologies are seeing increased deployment in real-world market scenarios. In this work, we explore the strategic behaviors of large language models (LLMs) when deployed as autonomous agents in multi-commodity markets, specifically within Cournot competition frameworks. We examine whether LLMs can independently engage in anti-competitive practices such as collusion or, more specifically, market division. Our findings demonstrate that LLMs can effectively monopolize specific commodities by dynamically adjusting their pricing and resource allocation strategies, thereby maximizing profitability without direct human input or explicit collusion commands. These results pose unique challenges and opportunities for businesses looking to integrate AI into strategic roles and for regulatory bodies tasked with maintaining fair and competitive markets. The study provides a foundation for further exploration into the ramifications of deferring high-stakes decisions to LLM-based agents.

Finance and Economics

AI and Finance

A. L. Eisfeldt, G. Schubert
UCLA Anderson School of Management

We provide evidence that the development and adoption of Generative AI is driving a significant technological shift for firms and for financial research. We review the literature on the impact of ChatGPT on firm value and provide directions for future research investigating the impact of this major technology shock. Finally, we review and describe innovations in research methods linked to improvements in AI tools, along with their applications. We offer a practical introduction to available tools and advice for researchers interested in using these tools.

Forecasting Crashes with a Smile

I. W. R. Martin, R. Shi
London School of Economics, University of Colorado Boulder

We use option prices to derive bounds on the probability of a crash in an individual stock, and argue that the lower bound should be close to the truth. Empirically, the lower bound is highly statistically and economically significant; on its own, it outperforms 15 stock characteristics proposed by the prior literature combined. In a multivariate regression, a one standard deviation increase in the bound raises the predicted crash probability by 3 percentage points, whereas a one standard deviation increase in the next most important predictor (a measure of short interest) raises the predicted probability by only 0.3 percentage points.

Cross-Sectional Dependence in Idiosyncratic Volatility

I. Kalnina, K. Tewou
North Carolina State University, National Bank of Canada, Canadian Derivatives Institute

This paper introduces an econometric framework for analyzing cross-sectional dependence in the idiosyncratic volatilities of assets using high frequency data. We first consider the estimation of standard measures of dependence in the idiosyncratic volatilities such as covariances and correlations. Naive estimators of these measures are biased due to the use of the error-laden estimates of idiosyncratic volatilities. We provide bias-corrected estimators and the relevant asymptotic theory. Next, we introduce an idiosyncratic volatility factor model, in which we decompose the variation in idiosyncratic volatilities into two parts: the variation related to the systematic factors such as the market volatility, and the residual variation. Again, naive estimators of the decomposition are biased, and we provide bias-corrected estimators. We also provide the asymptotic theory that allows us to test whether the residual (non-systematic) components of the idiosyncratic volatilities exhibit cross-sectional dependence. We apply our methodology to the S&P 100 index constituents, and document strong cross-sectional dependence in their idiosyncratic volatilities. We consider two different sets of idiosyncratic volatility factors, and find that neither can fully account for the cross-sectional dependence in idiosyncratic volatilities. For each model, we map out the network of dependencies in residual (non-systematic) idiosyncratic volatilities across all stocks.

Peer Option Momentum

S. Li, L. Yang
City University of Hong Kong

We document option momentum spillovers across peer firms with shared analyst coverage. Firms whose peers have higher past-12-month delta-hedged straddle returns tend to have higher future option returns. Peer option momentum has a magnitude similar to the option momentum documented by Heston et al. (2023), but it is distinct from option momentum. Past option returns of a firm's peers can predict the firm's realized variance even

after we control the firm's option-implied variance and its own past option returns. In addition, peer momentum is stronger for indirect linkages. The findings are consistent with limited investor attention on volatility information from peer firms. Compared with peer stock momentum, peer option momentum persists longer and exhibits spillovers via more economic linkages that cannot be subsumed by analyst-based linkages.

Dynamic Factor Allocation Leveraging Regime-Switching Signals

Y. Shu, J. M. Mulvey
Princeton University

This article explores dynamic factor allocation by analyzing the cyclical performance of factors through regime analysis. The authors focus on a U.S. equity investment universe comprising seven long-only indices representing the market and six style factors: value, size, momentum, quality, low volatility, and growth. Their approach integrates factor-specific regime inferences of each factor index's active performance relative to the market into the Black-Litterman model to construct a fully-invested, long-only multi-factor portfolio. First, the authors apply the sparse jump model (SJM) to identify bull and bear market regimes for individual factors, using a feature set based on risk and return measures from historical factor active returns, as well as variables reflecting the broader market environment. The regimes identified by the SJM exhibit enhanced stability and interpretability compared to traditional methods. A hypothetical single-factor long-short strategy is then used to assess these regime inferences and fine-tune hyperparameters, resulting in a positive Sharpe ratio of this strategy across all factors with low correlation among them. These regime inferences are then incorporated into the Black-Litterman framework to dynamically adjust allocations among the seven indices, with an equally weighted (EW) portfolio serving as the benchmark. Empirical results show that the constructed multi-factor portfolio significantly improves the information ratio (IR) relative to the market, raising it from just 0.05 for the EW benchmark to approximately 0.4. When measured relative to the EW benchmark itself, the dynamic allocation achieves an IR of around 0.4 to 0.5. The strategy also enhances absolute portfolio performance across key metrics such as the Sharpe ratio and maximum drawdown. These findings highlight the effectiveness of leveraging regime-switching signals to enhance factor allocation by capitalizing on factor cyclicality.

The DeFi Dilemma

A. Bodisz, N. Hautsch, S. Voigt
University of Vienna, University of Copenhagen

Blockchain technology enables financial interactions without central counterparties by distributing pending transactions in an open network and offering compensation fees to validators in order to ensure a reliable record of transaction histories. We show that these essential building blocks of decentralized settlement make decentralized finance (DeFi) applications inherently inefficient: Front-running risk causes a wasteful arms race for transaction fees and renders arbitrage prohibitively costly. Exploiting a novel, granular dataset of prices and transactions on decentralized exchanges, we show that cross-exchange price differences are substantial and persistent, while transaction fees almost entirely consume arbitrage profits. Overcoming these inefficiencies is possible only by reinstalling central counterparties, thus undermining the fundamental principle of DeFi.

Mathematical Finance

Statistical Arbitrage in Rank Space

Y.-F. Li, G. Papanicolaou
Stanford University

Equity market dynamics are conventionally investigated in name space where stocks are indexed by company names. In contrast, by indexing stocks based on their ranks in capitalization, we gain a different perspective of market dynamics in rank space. Here, we demonstrate the superior performance of statistical arbitrage in rank space over name space, driven by a robust market representation and enhanced mean-reverting properties of residual returns in rank space. Our statistical arbitrage algorithm features an intraday rebalancing mechanism for effective conversion between portfolios in name and rank space. We explore statistical arbitrage with and without neural networks in both name and rank space and show that the portfolios obtained in rank space with neural networks significantly outperform those in name space.

Equilibrium Control Theory for Kihlstrom-Mirman Preferences in Continuous Time

L. De Gennaro Aquino, S. Desmettre, Y. Havrylenko, M. Steffensen
Reykjavik University, Johannes Kepler University, University of Copenhagen, University of Ulm

In intertemporal settings, the multiattribute utility theory of Kihlstrom and Mirman suggests the application of a concave transform of the lifetime utility index. This construction, while allowing time and risk attitudes to be separated, leads to dynamically inconsistent preferences. We address this issue in a game-theoretic sense by formalizing an equilibrium control theory for continuous-time Markov processes. In these terms, we describe the equilibrium strategy and value function as the solution of an extended Hamilton-Jacobi-Bellman system of partial

differential equations. We verify that (the solution of) this system is a sufficient condition for an equilibrium and examine some of its novel features. A consumption-investment problem for an agent with CRRA-CES utility showcases our approach.

Optimal Execution with Deterministically Time Varying Liquidity: Well Posedness and Price Manipulation

G. Palmari, F. Lillo, Z. Eisler
Scuola Normale Superiore, Imperial College London, University of Bologna

We investigate the well-posedness in the Hadamard sense and the absence of price manipulation in the optimal execution problem within the Almgren-Chriss framework, where the temporary and permanent impact parameters vary deterministically over time. We present sufficient conditions for the existence of a unique solution and provide second-order conditions for the problem, with a particular focus on scenarios where impact parameters change monotonically over time. Additionally, we establish conditions to prevent transaction-triggered price manipulation in the optimal solution, i.e. the occurence of buying and selling in the same trading program. Our findings are supported by numerical analyses that explore various regimes in simple parametric settings for the dynamics of impact parameters.

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

Alden, Andrew, Carmine Ventre, and Blanka Horvath. 2024. "Scalable Signature-Based Distribution Regression via Reference Sets." *arXiv preprint arXiv:2410.09196*.

Aminian, Gholamali, Łukasz Szpruch, and Samuel N Cohen. 2024. "Understanding Transfer Learning via Mean-field Analysis." *arXiv preprint arXiv:2410.17128*.

Aquino, Luca De Gennaro, Sascha Desmettre, Yevhen Havrylenko, and Mogens Steffensen. 2024. "Equilibrium Control Theory for Kihlstrom-Mirman Preferences in Continuous Time." *arXiv preprint arXiv:2407.16525*.

Bodisz, Aron, Nikolaus Hautsch, and Stefan Voigt. 2024. "The DeFi Dilemma." *Working Paper*.

Eisfeldt, Andrea L., and Gregor Schubert. 2024. "AI and Finance." *Available at SSRN 4988553*.

Han, Sukjin. 2024. "Mining Causality: AI-Assisted Search for Instrumental Variables." *arXiv preprint arXiv:2409.14202*.

Kalnina, Ilze, and Kokouvi Tewou. 2024. "Cross-Sectional Dependence in Idiosyncratic Volatility." *arXiv preprint arXiv:2408.13437*.

Kim, Junsu, Seohong Park, and Sergey Levine. 2024. "Unsupervised-to-Online Reinforcement Learning." *arXiv preprint arXiv:2408.14785*.

Lin, Ryan Y, Siddhartha Ojha, Kevin Cai, and Maxwell F Chen. 2024. "Strategic Collusion of LLM Agents: Market Division in Multi-Commodity Competitions." *arXiv preprint arXiv:2410.00031*.

Li, Shuaiqi, and Lihai Yang. 2024. "Peer Option Momentum." *Available at SSRN 4964637*.

Li, Y-F, and G Papanicolaou. 2024. "Statistical Arbitrage in Rank Space." *arXiv preprint arXiv:2410.06568*.

Martin, Ian W. R., and Ran Shi. 2024. "Forecasting Crashes with a Smile." *Working Paper*.

Palmari, Gianluca, Fabrizio Lillo, and Zoltan Eisler. 2024. "Optimal Execution with Deterministically Time Varying Liquidity: Well Posedness and Price Manipulation." *arXiv preprint arXiv:2410.04867*.

Shu, Yizhan, and John M. Mulvey. 2024. "Dynamic Factor Allocation Leveraging Regime-Switching Signals." *Available at SSRN 4960484*.

Sukhija, Bhavya, Lenart Treven, Florian Dörfler, Stelian Coros, and Andreas Krause. 2024. "NeoRL: Efficient Exploration for Nonepisodic RL." *arXiv preprint arXiv:2406.01175*.

Treven, Lenart, Bhavya Sukhija, Yarden As, Florian Dörfler, and Andreas Krause. 2024. "When to Sense and Control? A Time-adaptive Approach for Continuous-Time RL." *arXiv preprint arXiv:2406.01163*.