VOL. 4, NO. 7

OMI Research Newsletter – April 2025

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 Decentralised Finance and Market Microstructure, co-organised by the University of Toronto, Columbia University and OMI, will be held on 8–9 May 2025 at the Fields Institute for Research in Mathematical Sciences, University of Toronto.

The Machine Learning in Quantitative Finance Conference 2025, organised by OMI, will be held on 5–6 June 2025 at the Mathematical Institute, University of Oxford. Registration will be available on 10 April 2025.

Highlights

Key contributions in this month's edition: Baker et al. (2025) offer a practitioner-oriented guide that unifies various Difference-in-Differences (DiD) designs for causal inference. Murphy (2025) provides a comprehensive overview of reinforcement learning and sequential decision making. Liang et al. (2025) introduce a large-scale city road network dataset and highlight long-range dependencies in graph neural networks. Liao et al. (2025) develop new methods to construct forecast confidence intervals for expected returns obtained from neural networks. Li et al. (2025) propose an innovative framework to construct confidence sets that account for proxy estimation error and enable real-time assessment of forecast "acceptability." Cartea, Jin and Shi (2025) challenge the "virtue of complexity" in return prediction by showing that, although adding more factors can improve return forecasts under ideal conditions, measurement noise in factors ultimately undermines those benefits, leading to an optimal—rather than unbounded—level of model complexity. Cartea and Jin (2025) find that only a small fraction of analysts consistently generate alpha beyond a naive buy-and-hold benchmark, and dynamically pooling these outperformers in a "fund-of-analysts" strategy delivers significant excess returns. Oliveira et al. (2025) propose a clustering-based framework for macroeconomic regime detection, which substantially improves tactical asset allocation performance over conventional benchmarks. Kroner (2025) shows that investor attention, heightened in high-inflation periods, strongly shapes how financial markets incorporate macroeconomic news. Cartea and Sánchez-Betancourt (2025) derive a closed-form optimal broker strategy for dealing with both informed ("toxic") and uninformed order flows. Finally, Aqsha, Bergault and Sánchez-Betancourt (2025) develop an optimal contract for automated market makers that offsets impermanent loss, incentivises strategic liquidity providers, and leverages noise trading to maximise order flow.

Statistics, Econometrics & Machine Learning

Difference-in-Differences Designs: A Practitioner's Guide

A. BAKER, B. CALLAWAY, S.CUNNINGHAM, A. GOODMAN-BACON, P. H. C. SANT'ANNA

UNIVERSITY OF CALIFORNIA, BERKELEY, UNIVERSITY OF GEORGIA, BAYLOR UNIVERSITY, FEDERAL RESERVE BANK OF MINNEAPOLIS, EMORY UNIVERSITY

Difference-in-Differences (DiD) is arguably the most popular quasi-experimental research design. Its canonical form, with two groups and two periods, is well-understood. However, empirical practices can be ad hoc when researchers go beyond that simple case. This article provides an organizing framework for discussing different

types of DiD designs and their associated DiD estimators. It discusses covariates, weights, handling multiple periods, and staggered treatments. The organizational framework, however, applies to other extensions of DiD methods as well.

Reinforcement Learning: A Comprehensive Overview

K. P. Murphy
Google

This manuscript gives a big-picture, up-to-date overview of the field of (deep) reinforcement learning and sequential decision making, covering value-based method, policy-gradient methods, model-based methods, and various other topics (e.g., multi-agent RL, RL+LLMs, and RL+inference).

Towards Quantifying Long-Range Interactions in Graph Machine Learning: A Large Graph Dataset and a Measurement

H. Liang, H. Sáez de Ocáriz Borde, B. Sripathmanathan, M. Bronstein, X. Dong
University of Oxford

Long-range dependencies are critical for effective graph representation learning, yet most existing datasets focus on small graphs tailored to inductive tasks, offering limited insight into long-range interactions. Current evaluations primarily compare models employing global attention (e.g., graph transformers) with those using local neighborhood aggregation (e.g., message-passing neural networks) without a direct measurement of long-range dependency. In this work, we introduce City-Networks, a novel large-scale transductive learning dataset derived from real-world city roads. This dataset features graphs with over 105 nodes and significantly larger diameters than those in existing benchmarks, naturally embodying long-range information. We annotate the graphs using an eccentricity-based approach, ensuring that the classification task inherently requires information from distant nodes. Furthermore, we propose a model-agnostic measurement based on the Jacobians of neighbors from distant hops, offering a principled quantification of long-range dependencies. Finally, we provide theoretical justifications for both our dataset design and the proposed measurement - particularly by focusing on over-smoothing and influence score dilution - which establishes a robust foundation for further exploration of long-range interactions in graph neural networks.

Deep Learning is Not So Mysterious or Different

A. G. Wilson
New York University

Deep neural networks are often seen as different from other model classes by defying conventional notions of generalization. Popular examples of anomalous generalization behaviour include benign overfitting, double descent, and the success of overparametrization. We argue that these phenomena are not distinct to neural networks, or particularly mysterious. Moreover, this generalization behaviour can be intuitively understood, and rigorously characterized using long-standing generalization frameworks such as PAC-Bayes and countable hypothesis bounds. We present soft inductive biases as a key unifying principle in explaining these phenomena: rather than restricting the hypothesis space to avoid overfitting, embrace a flexible hypothesis space, with a soft preference for simpler solutions that are consistent with the data. This principle can be encoded in many model classes, and thus deep learning is not as mysterious or different from other model classes as it might seem. However, we also highlight how deep learning is relatively distinct in other ways, such as its ability for representation learning, phenomena such as mode connectivity, and its relative universality.

The Uncertainty of Machine Learning Predictions in Asset Pricing

Y. Liao, X. Ma, A. Neuhierl, L. Schilling
Rutgers University, National University of Singapore, Washington University in St. Louis

Machine learning in asset pricing typically predicts expected returns as point estimates, ignoring uncertainty. We develop new methods to construct forecast confidence intervals for expected returns obtained from neural networks. We show that neural network forecasts of expected returns share the same asymptotic distribution as classic nonparametric methods, enabling a closed-form expression for their standard errors. We also propose a computationally feasible bootstrap to obtain the asymptotic distribution. We incorporate these forecast confidence intervals into an uncertainty-averse investment framework. This provides an economic rationale for shrinkage implementations of portfolio selection. Empirically, our methods improve out-of-sample performance.

Online Forecast Evaluation with High-Frequency Proxies

J. Li, Y. Li, A. J. Patton, T. Xie
Singapore Management University, University of Manchester, University of New South Wales, Duke University, Shanghai University of Finance and Economics

We present methods to evaluate risk forecasts "online," enabling the analyst to determine in real time whether a forecasting model is producing "statistically acceptable" forecasts. We use the asymptotic distribution of the volatility proxy used for forecast evaluation to construct a confidence set for the true, latent, risk measure and we use the inclusion, or not, of the forecast in that set as a formal measure of acceptability. We consider the construction of the confidence sets in both large- and small-sample settings, allowing us to consider forecasts of quantities like integrated variance and spot volatility. We illustrate the method by evaluating spot volatility forecasts for a newly listed stock, including on the day of its IPO.

Finance and Economics

The Limited Virtue of Complexity in a Noisy World

Á. Cartea, Q. Jin, Y. Shi
University of Oxford

Measurement error in factors affects the out-of-sample performance of portfolios and of return predictions. In a high-dimensional factor space, increasing model complexity under proper regularization can enhance the predictability of asset returns; however, the Sharpe ratio of a portfolio of assets and the R-squared of the prediction of the asset returns decrease monotonically and are convex as the noise level in factors increases. Furthermore, we show that when only a subset of factors is observed, there is an optimal level of complexity beyond which incorporating additional factors can degrade portfolio performance due to the effect of noise in the factors. Thus, the marginal benefits of increasing model complexity will, at some point, start to diminish as the number of observed factors increases. Our findings underscore a limited virtue of complexity in financial forecasting, where the performance of portfolios depends on the noise level in factors, and where more complex models do not necessarily lead to better performance when factors are not perfectly observed.

Alpha in Analysts

Á. Cartea, Q. Jin
University of Oxford

This paper examines the investment value in sell-side analyst price targets. We treat each analyst as a portfolio manager and use their price targets to construct 12-month implied return forecasts. We invest in self-financing long-short portfolios for individual analysts, where we go long on stocks with positive forecasts and go short on those with negative forecasts, and the weights in the portfolio are proportional to the magnitude of the implied returns. Our empirical analysis shows that while the average analyst does not generate statistically significant alpha relative to the returns of a long-only portfolio benchmark, a subset of analysts exhibits persistent alpha. Motivated by this heterogeneity, we introduce a "fund-of-analysts" framework that first predicts analyst performance and then dynamically allocates weights across analysts based on predicted analyst performances. Our results show that this meta-portfolio strategy can yield significant alpha over long-only benchmarks, providing new insights into the role of analyst heterogeneity in equity market pricing.

Tactical Asset Allocation with Macroeconomic Regime Detection

D. C. Oliveira, D. Sandfelder, A. Fujita, X. Dong, M. Cucuringu
University of São Paulo, University of Oxford, University of California Los Angeles

This paper extends the tactical asset allocation literature by incorporating regime modeling using techniques from machine learning. We propose a novel model that classifies current regimes, forecasts the distribution of future regimes, and integrates these forecasts with the historical performance of individual assets to optimize portfolio allocations. Utilizing a macroeconomic data set from the FRED-MD database, our approach employs a modified k-means algorithm to ensure consistent regime classification over time. We then leverage these regime predictions to estimate expected returns and volatilities, which are subsequently mapped into portfolio allocations using various sizing schemes. Our method outperforms traditional benchmarks such as equal-weight, buy-and-hold, and random regime models. Additionally, we are the first to apply a regime detection model from a large macroeconomic dataset to tactical asset allocation, demonstrating significant improvements in portfolio performance. Our work presents several key contributions, including a novel data-driven regime detection algorithm tailored for uncertainty in forecasted regimes and applying the FRED-MD data set for tactical asset allocation.

A Test of the Efficiency of a Given Portfolio in High Dimensions

M. Chernov, B. T. Kelly, S. Malamud, J. Schwab
University of California Los Angeles, Yale University, Swiss Finance Institute, Ecole Polytechnique Federale de Lausanne

We generalize the seminal Gibbons-Ross-Shanken test to the empirically relevant case where the number of test assets far exceeds the number of observations. In such a setting, one needs to use a regularized estimator of the covariance matrix of test assets, which leads to biases in the original test statistic. Random Matrix Theory allows us to account for these biases and to evaluate the test's power. Power increases with the number of test assets and reaches

the maximum for a broad range of local alternatives. These conclusions are supported by an extensive simulation study. We implement the test empirically for state-of-the-art candidate efficient portfolios and test assets.

Short-Selling Profitability, Stock Lending Fees, and Asset Pricing Anomalies

ZHI DA, C. FU, N. LIN, L. LU
UNIVERSITY OF NOTRE DAME, UNIVERSITY OF NORTHERN BRITISH COLUMBIA, ARKANSAS STATE UNIVERSITY, UNIVERSITY OF MANITOBA

We measure a stock's short-selling profitability (SSP) as its price sensitivity to short-selling activities over recent periods. Our findings show that short-selling strongly and negatively predicts future returns, particularly among high-SSP stocks. Furthermore, we identify SSP as a novel determinant of stock lending fees in the cross-section. While the profitability of anomalies decreases when accounting for short-selling fees, they remain exploitable among high-SSP stocks. These results support the presence of a stock lending market in which lenders allow short sellers to retain a portion of arbitrage profits. This suggests that short-selling constraints alone do not fully explain the persistence of anomalies, especially among high-SSP stocks.

How Markets Process Macro News: The Importance of Investor Attention

T. N. KRONER
FEDERAL RESERVE BOARD

I provide evidence that investors' attention allocation plays a critical role in how financial markets incorporate macroeconomic news. Using intraday data, I document a sharp increase in the market reaction to Consumer Price Index (CPI) releases during the 2021-2023 inflation surge. Bond yields, market-implied inflation expectations, and other asset prices exhibit significantly stronger responses to CPI surprises, while reactions to other macroeconomic announcements remain largely unchanged. The joint reactions of these asset prices point to an attention-based explanation - an interpretation I corroborate throughout the rest of the paper. Specifically, I construct a measure of CPI investor attention and find that: (1) attention was exceptionally elevated around CPI announcements during the inflation surge, and (2) higher pre-announcement attention robustly leads to stronger market reactions. Studying investor attention in the context of Employment Report releases and Federal Reserve announcements, I document a similar importance of attention allocation for market reactions. Lastly, I find that markets tend to overreact to announcements that attract high levels of attention.

Non-Bayesian Learning in Misspecified Models

S. BERVOETS, M. FAURE, L. RENOU
AIX-MARSEILLE UNIVERSITY, QUEEN MARY UNIVERSITY OF LONDON

Deviations from Bayesian updating are traditionally categorized as biases, errors, or fallacies, thus implying their inherent "sub-optimality." We offer a more nuanced view. We demonstrate that, in learning problems with misspecified models, non-Bayesian updating can outperform Bayesian updating.

Hate in the Time of Algorithms: Evidence on Online Behavior from a Large-Scale Experiment

A. KALRA
BROWN UNIVERSITY

Social media algorithms are thought to amplify variation in user beliefs, thus contributing to radicalization. However, quantitative evidence on how algorithms and user preferences jointly shape harmful online engagement is limited. I conduct an individually randomized experiment with 8 million users of an Indian TikTok-like platform, replacing algorithmic ranking with random content delivery. Exposure to "toxic" posts decreases by 27%, mainly due to reduced platform usage by users with higher interest in such content. Strikingly, these users increase engagement with toxic posts they find. Survey evidence indicates shifts to other platforms. Model-based counterfactuals highlight the limitations of blanket algorithmic regulation.

Value of Information in Social Learning

H. SATO, K. SHIMIZU
NAGOYA UNIVERSITY, UNIVERSITY OF TOKYO

This study extends Blackwell's (1953) comparison of information to a sequential social learning model, where agents make decisions sequentially based on both private signals and the observed actions of others. In this context, we introduce a new binary relation over information structures: an information structure is more socially valuable than another if it yields higher expected payoffs for all agents, regardless of their preferences. First, we establish that this binary relation is strictly stronger than the Blackwell order. Then, we provide a necessary and sufficient condition for our binary relation and propose a simpler sufficient condition that is easier to verify.

Mathematical and Computational Finance

A Simple Strategy to Deal with Toxic Flow

Á. CARTEA, L. SÁNCHEZ-BETANCOURT
UNIVERSITY OF OXFORD

We model the trading activity between a broker and her clients (informed and uninformed traders) as an infinite-horizon stochastic control problem. We derive the broker's optimal dealing strategy in closed form and use this to introduce an algorithm that bypasses the need to calibrate individual parameters, so the dealing strategy can be executed in real-world trading environments. Finally, we characterise the discount in the price of liquidity a broker offers clients. The discount strikes the optimal balance between maximising the order flow from the broker's clients and minimising adverse selection losses to the informed traders.

Equilibrium Reward for Liquidity Providers in Automated Market Makers

A. Aqsha, P. Bergault, L. Sánchez-Betancourt
University of Oxford, Université Paris I Panthéon-Sorbonne

We find the equilibrium contract that an automated market maker (AMM) offers to their strategic liquidity providers (LPs) in order to maximize the order flow that gets processed by the venue. Our model is formulated as a leader-follower stochastic game, where the venue is the leader and a representative LP is the follower. We derive approximate closed-form equilibrium solutions to the stochastic game and analyze the reward structure. Our findings suggest that under the equilibrium contract, LPs have incentives to add liquidity to the pool only when higher liquidity on average attracts more noise trading. The equilibrium contract depends on the external price, the pool reference price, and the pool reserves. Our framework offers insights into AMM design for maximizing order flow while ensuring LP profitability.

Short-Time Behavior of the At-The-Money Implied Volatility for the Jump-Diffusion Stochastic Volatility Bachelier Model

E. Alòs, Ò. Burés, J. Vives
1Universitat Pompeu Fabra, Barcelona School of Economics, Universitat de Barcelona

In this paper we use Malliavin Calculus techniques in order to obtain expressions for the short-time behavior of the at-the-money implied volatility (ATM-IV) level and skew for a jump-diffusion stock price. The diffusion part is assumed to be the stochastic volatility Bachelier model and the jumps are modeled by a pure-jump Lévy process with drift so that the stock price is a martingale. Regarding the level, we show that the short-time behavior of the ATM-IV level is the same for all pure-jump Lévy processes and, regarding the skew, we give conditions on the law of the jumps for the skew to exist. We also give several numerical examples of stochastic volatilities and Lévy processes that confirm the theoretical results found in the paper.

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, and Qi Jin.

References

Alòs, Elisa, Òscar Burés, and Josep Vives. 2025. "Short-Time Behavior of the At-The-Money Implied Volatility for the Jump-Diffusion Stochastic Volatility Bachelier Model." *arXiv preprint arXiv:2503.22282*.

Aqsha, Alif, Philippe Bergault, and Leandro Sánchez-Betancourt. 2025. "Equilibrium Reward for Liquidity Providers in Automated Market Makers." *SSRN Working Paper*.

Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2025. "Difference-in-Differences Designs: A Practitioner's Guide." *arXiv preprint arXiv:2503.13323*.

Bervoets, Sebastian, Mathieu Faure, and Ludovic Renou. 2025. "Non-Bayesian Learning in Misspecified Models." *arXiv preprint arXiv:2503.18024*.

Cartea, Álvaro, and Leandro Sánchez-Betancourt. 2025. "A Simple Strategy to Deal with Toxic Flow." *SSRN Working Paper*.

Cartea, Álvaro, and Qi Jin. 2025. "Alpha in Analysts." *SSRN Working Paper*.

Cartea, Álvaro, Qi Jin, and Yuantao Shi. 2025. "The Limited Virtue of Complexity in a Noisy World." *SSRN Working Paper*.

Chernov, Mikhail, Bryan T. Kelly, Semyon Malamud, and Johannes Schwab. 2025. "A Test of the Efficiency of a Given Portfolio in High Dimensions." *SSRN Working Paper*.

Da, Zhi, Chengbo Fu, Nanying Lin, and Lei Lu. 2025. "Short-selling Profitability, Stock Lending Fees, and Asset Pricing Anomalies." *SSRN Working Paper*.

Kalra, Aarushi. 2025. "Hate in the Time of Algorithms: Evidence on Online Behavior from a Large-Scale Experiment." *arXiv preprint arXiv:2503.06244*.

Kroner, T. Niklas. 2025. "How Markets Process Macro News: The Importance of Investor Attention." *SSRN Working Paper*.

Liang, Huidong, Haitz Sáez de Ocáriz Borde, Baskaran Sripathmanathan, Michael Bronstein, and Xiaowen Dong. 2025. "Towards Quantifying Long-Range Interactions in Graph Machine Learning: A Large Graph Dataset and a Measurement." *arXiv preprint arXiv:2503.09008*.

Liao, Yuan, Xinjie Ma, Andreas Neuhierl, and Linda Schilling. 2025. "The Uncertainty of Machine Learning Predictions in Asset Pricing." *SSRN Working Paper*.

Li, Jia, Yifan Li, Andrew J Patton, and Tian Xie. 2025. "Online Forecast Evaluation with High-Frequency Proxies." *Working Paper*.

Murphy, Kevin P. 2025. "Reinforcement Learning: A Comprehensive Overview." *arXiv preprint arXiv:2412.05265*.

Oliveira, Daniel Cunha, Dylan Sandfelder, André Fujita, Xiaowen Dong, and Mihai Cucuringu. 2025. "Tactical Asset Allocation with Macroeconomic Regime Detection." *arXiv preprint arXiv:2503.11499*.

Sato, Hiroto, and Konan Shimizu. 2025. "Value of Information in Social Learning." *arXiv preprint arXiv:2503.05015*.

Wilson, Andrew Gordon. 2025. "Deep Learning is Not So Mysterious or Different." *arXiv preprint arXiv:2503.02113*.