OMI Research Newsletter – November 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.*

Highlights

Key contributions in this month's edition: Lai et al. (2025) summarize the core principles that have guided the development of diffusion models, tracing their origins and showing how diverse formulations arise from shared mathematical ideas. Chen and Feng (2025) propose a universal factor model that unifies mean and quantile factor analysis, and recovers factors even when they are weak at most quantiles. Simsek (2025) introduce a candlestick-based framework for estimating spot covariances in high frequency. Horvath et al. (2025) review the rise of "Market Generators" for synthetic financial data, detail finance-specific challenges and evaluation pitfalls, and highlight signature-based path-wise methods as effective discriminators powering a broader shift toward domain-tailored generative AI. Cui et al. (2025) propose a cross-sectional asset pricing model that incorporates structural breaks and regime-dependent momentum trends in factor premia. Kılıç (2025) find evidence for a limited "virtue of complexity" for exchange rate forecasting using nonlinear Ridge regressions with Random Fourier Features. Fusari and Lamichhane (2025) develop a structural model for equity option pricing that incorporates stochastic volatility and time-varying jumps. Di Nosse et al. (2025) empirically study how the emergence of DeFi modifies the well-studied stylized facts of prices, liquidity, and order flow observed in traditional markets.

Statistics, Econometrics & Machine Learning

The Principles of Diffusion Models: From Origins to Advances

C.-H. LAI, Y. SONG, D. KIM, Y. MITSUFUJI, S. ERMON

SONY AI, OPENAI, STANFORD UNIVERSITY

This monograph presents the core principles that have guided the development of diffusion models, tracing their origins and showing how diverse formulations arise from shared mathematical ideas. Diffusion modeling starts by defining a forward process that gradually corrupts data into noise, linking the data distribution to a simple prior through a continuum of intermediate distributions. The goal is to learn a reverse process that transforms noise back into data while recovering the same intermediates. We describe three complementary views. The variational view, inspired by variational autoencoders, sees diffusion as learning to remove noise step by step. The score-based view, rooted in energy-based modeling, learns the gradient of the evolving data distribution, indicating how to nudge samples toward more likely regions. The flow-based view, related to normalizing flows, treats generation as following a smooth path that moves samples from noise to data under a learned velocity field. These perspectives share a common backbone: a time-dependent velocity field whose flow transports a simple prior to the data. Sampling then amounts to solving a differential equation that evolves noise into data along a continuous trajectory. On this foundation, the monograph discusses guidance for controllable generation, efficient numerical solvers, and diffusion-motivated flow-map models that learn direct mappings between arbitrary times. It provides a conceptual and mathematically grounded understanding of diffusion models for readers with basic deep-learning knowledge.

Universal Factor Models

S. CHEN, J. FENG

ZHEJIANG UNIVERSITY, HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

We propose a new factor analysis framework and estimators of the factors and loadings that are robust to weak factors in a large N and large T setting. Our framework, by simultaneously considering all quantile levels of the outcome variable, induces standard mean and quantile factor models, but the factors can have an arbitrarily weak influence on the outcome's mean or quantile at most quantile levels. Our method estimates the factor space at the $\sqrt{N}$-rate without requiring the knowledge of weak factors' presence or strength, and achieves $\sqrt{N}$- and $\sqrt{T}$-asymptotic normality for the factors and loadings based on a novel sample splitting approach that handles incidental nuisance parameters. We also develop a weak-factor-robust estimator of the number of factors and consistent selectors of factors of any tolerated level of influence on the outcome's mean or quantiles. Monte Carlo simulations demonstrate the effectiveness of our method.

Beyond Returns: A Candlestick-Based Approach to Spot Covariance Estimation

Y. Simsek

Duke University

Spot covariance estimation is commonly based on high-frequency open-to-close return data over short time windows, but such approaches face a trade-off between statistical accuracy and localization. In this paper, I introduce a new estimation framework using high-frequency candlestick data, which include open, high, low, and close prices, effectively addressing this trade-off. By exploiting the information contained in candlesticks, the proposed method improves estimation accuracy relative to benchmarks while preserving local structure. I further develop a test for spot covariance inference based on candlesticks that demonstrates reasonable size control and a notable increase in power, particularly in small samples. Motivated by recent work in the finance literature, I empirically test the market neutrality of the iShares Bitcoin Trust ETF (IBIT) using 1-minute candlestick data for the full year of 2024. The results show systematic deviations from market neutrality, especially in periods of market stress. An event study around FOMC announcements further illustrates the new method's ability to detect subtle shifts in response to relatively mild information events.

Predictively Oriented Posteriors

Y. McLatchie, B.-E. Cherief-Abdellatif, D. T. Frazier, J. Knoblauch

University College London, Sorbonne Université, Monash University

We advocate for a new statistical principle that combines the most desirable aspects of both parameter inference and density estimation. This leads us to the predictively oriented (PrO) posterior, which expresses uncertainty as a consequence of predictive ability. Doing so leads to inferences which predictively dominate both classical and generalised Bayes posterior predictive distributions: up to logarithmic factors, PrO posteriors converge to the predictively optimal model average at rate $n^{-1/2}$. Whereas classical and generalised Bayes posteriors only achieve this rate if the model can recover the data-generating process, PrO posteriors adapt to the level of model misspecification. This means that they concentrate around the true model at rate $n^{1/2}$ in the same way as Bayes and Gibbs posteriors if the model can recover the data-generating distribution, but do not concentrate in the presence of non-trivial forms of model misspecification. Instead, they stabilise towards a predictively optimal posterior whose degree of irreducible uncertainty admits an interpretation as the degree of model misspecification – a sharp contrast to how Bayesian uncertainty and its existing extensions behave. Lastly, we show that PrO posteriors can be sampled from by evolving particles based on mean field Langevin dynamics, and verify the practical significance of our theoretical developments on a number of numerical examples.

TabMGP: Martingale Posterior with TabPFN

K. Ng, E. Fong, D. T. Frazier, J. Knoblauch, S. Wei

Monash University, University of Hong Kong, University College London

Bayesian inference provides principled uncertainty quantification but is often limited by challenges of prior elicitation, likelihood misspecification, and computational burden. The martingale posterior (MGP, Fong et al., 2023) offers an alternative, replacing prior-likelihood elicitation with a predictive rule - namely, a sequence of one-step-ahead predictive distributions - for forward data generation. The utility of MGPs depends on the choice of predictive rule, yet the literature has offered few compelling examples. Foundation transformers are well-suited here, as their autoregressive generation mirrors this forward simulation and their general-purpose design enables rich predictive modeling. We introduce TabMGP, an MGP built on TabPFN, a transformer foundation model that is currently state-of-the-art for tabular data. TabMGP produces credible sets with near-nominal coverage and often outperforms both existing MGP constructions and standard Bayes.

Finance and Economics

Generative Models in Finance: Market Generators, a Paradigm Shift in Financial Modeling

B. Horvath, J. Plenk, M. Vuletić, R. Saqur

Oxford-Man Institute for Quantitative Finance, University of Oxford, University of Toronto

Market Generators are a rapidly evolving class of neural-network-based models to simulate financial market behavior, offering a powerful alternative to classical stochastic models. These deep learning models are trained to encode the underlying distribution of financial data and generate new synthetic market scenarios from the learned distribution. Although the expressions "Market Generator" and its related form "Market Simulator" have only entered the vocabulary of financial modeling around 2019, today, the modeling techniques related to them have already grown into an area of its own right. This growing interest is matched by a dramatic rise in research, publications, and an accelerating rate of innovation in the ambient technological arena of generative modeling. Recently, this trend has culminated in the emergence of large generative networks, particularly GPT-type language models, which are proving to be driving one of the biggest disruptions in the history of technology.

Breaks and Trends in Factor Premia

L. Cui, G. Feng, J. Ma, Y. Su

City University of Hong Kong, University of Warwick, Johns Hopkins University

We propose a cross-sectional asset pricing model that incorporates structural breaks and regime-dependent momentum trends in factor premia. To estimate a time-varying slope factor model, we employ total-variation regularization (TV), which allows for flexible stock return prediction. This approach produces piecewise linear trend lines at varying frequencies, capturing either long-term or fine-grained factor momentum trends depending on the TV penalty, which is calibrated using a machine learning framework. The method demonstrates large-sample theoretical properties, including consistency and asymptotic normality. Out-of-sample trend-following factor-timing strategies outperform traditional buy-and-hold or fixed-window momentum strategies. Furthermore, dynamic factor selection and time-varying premia estimates lead to improved asset pricing factor model specification compared to time-invariant models. This framework provides a robust tool for capturing dynamic market conditions and offers significant advantages in both predictive accuracy and portfolio performance.

Virtue or Mirage? Complexity in Exchange Rate Prediction

R. Kiliç

Federal Reserve Board

This paper investigates whether the “virtue of complexity” (VoC), documented in equity return prediction, extends to exchange rate forecasting. Using nonlinear Ridge regressions with Random Fourier Features (Ridge–RFF), we compare the predictive performance of complex models against linear regression and the robust random walk benchmark. Forecasts are constructed across three sets of economic fundamentals—traditional monetary, expanded monetary and non-monetary, and Taylor-rule predictors—with nominal complexity varied through rolling training windows of 12, 60, and 120 months. Our results offer a cautionary perspective. Complexity delivers only modest, localized gains: in very small samples with rich predictor sets, Ridge–RFF can outperform linear regression. Yet these improvements never translate into systematic gains over the random walk. As training windows expand, Ridge–RFF quickly loses ground, while linear regression increasingly dominates, at times even surpassing the random walk under expanded fundamentals. Market-timing analyses reinforce these findings: complexity-based strategies yield occasional short-sample gains but are unstable and prone to sharp drawdowns, whereas simpler linear and random walk strategies provide more robust and consistent economic value. By incorporating formal forecast evaluation tests—including Clark–West and Diebold–Mariano—we show that apparent gains from complexity are fragile and rarely statistically significant. Overall, our evidence points to a limited virtue of complexity in FX forecasting: complexity may help under narrowly defined conditions, but parsimony and the random walk benchmark remain more reliable across samples, predictor sets, and economic evaluations.

The Impact of Active Managers on the Pricing of Underlying Assets in ETFs

C. Trzcinka, Z. Zhao

Indiana University, University of Lausanne and Swiss Finance Institute

We investigate the impact of active managers on the information efficiency of the underlying assets in passive ETF portfolios. Specifically, we explore how the increasing popularity of ETFs prompts active mutual fund managers to execute informed trades that generate alpha. Using trade-level data, we test whether trades by skilled active managers more accurately predict future abnormal stock returns as ETF ownership in these stocks rises. By leveraging the annual reclassification of stocks from the Russell 1000 to the Russell 2000 as an exogenous variation, we find that high-alpha mutual funds can mitigate the pricing inefficiency typically associated with ETFs.

Structural Equity Option Pricing: Implications for Credit Risk

N. Fusari, S. Lamichhane

Johns Hopkins University, International Monetary Fund

We develop a structural model for equity option pricing that incorporates stochastic volatility and time-varying

asset jumps. The asset dynamics belong to the broad class of affine models, but the corresponding equity dynamics are non-affine, making traditional Fourier-based valuation methods computationally intensive. We address this challenge by training a deep neural network to approximate the option pricing function. Using liquid, exchange-traded equity options, we estimate key firm-level characteristics, specifically, asset variance and market leverage, which are used to derive an option-implied probability of default (PD). This PD co-moves with over-the-counter CDS spreads but also exhibits persistent residuals that we link to intermediary constraints. Our approach delivers a real-time, forward-looking, equity market-based measure of credit risk that is, by construction, unaffected by the illiquidity and frictions inherent in credit markets.

Clustering and Similarity Learning in Financial Markets: A Tutorial for the Practitioners

D. Mehta, J. R. J. Thompson, H. Lee, Y. Lee

BlackRock Inc., University of British Columbia, Ulsan National Institute of Science and Technology

Clustering and similarity learning are increasingly indispensable for structuring heterogeneous financial data and supporting real-world decision making. Traditional heuristics such as industry codes, static style boxes, or return correlations, offer only coarse and rigid notions of peer groups. Recent advances in metric learning, graph methods, and large language models now make it possible to build adaptive neighborhoods of securities, funds, companies, and investors that align more closely with actual risk, liquidity, and thematic exposures. This tutorial synthesizes these methodological developments and demonstrates their use across major asset classes. Case studies show how supervised proximities improve bond substitution, how fund similarity systems reconcile category reproducibility with outlier detection, how multimodal pipelines refine company comparables for valuation and strategy, and how investor clustering enhances personalization and Know-Your-Client (KYC) analytics. We emphasize modeling choices that make clustering and similarity auditable and robust under regime shifts. We also outline their evaluation protocols such as neighborhood stability, substitution fidelity, and segment utility, etc., which align with investment, compliance, and fiduciary objectives. Overall, the central message for practitioners is pragmatic: similarity systems have moved beyond experimental prototypes and now stand as deployable techniques within real investment workflows.

Attention Factors for Statistical Arbitrage

E. L. Epstein, R. Wang, J. Choi, M. Pelger

Stanford University, Nanyang Technological University

Statistical arbitrage exploits temporal price differences between similar assets. We develop a framework to jointly identify similar assets through factors, identify mispricing and form a trading policy that maximizes risk-adjusted performance after trading costs. Our Attention Factors are conditional latent factors that are the most useful for arbitrage trading. They are learned from firm characteristic embeddings that allow for complex interactions. We identify time-series signals from the residual portfolios of our factors with a general sequence model. Estimating factors and the arbitrage trading strategy jointly is crucial to maximize profitability after trading costs. In a comprehensive empirical study we show that our Attention Factor model achieves an out-of-sample Sharpe ratio above 4 on the largest U.S. equities over a 24-year period. Our one-step solution yields an unprecedented Sharpe ratio of 2.3 net of transaction costs. We show that weak factors are important for arbitrage trading.

The Problem of Business Evil: Corporate Fraud, Rogue Trading and Pharmaceutical Misconduct

J. E. Thanassoulis

University of Warwick

We offer a new theory as to why ethical people engage in corporate misconduct such as accounting fraud, rogue trading, and suppression of drug side-effects. Modelling ethics as other-regarding preferences we show all agents, however ethical, have an in-too-deep threshold. If private knowledge that an activity they have begun is harmful arrives after a threshold time then the agent will choose to hide this fact and continue the activity, knowingly harming third parties – until publicly caught. We show that unethical agents can sell cash-flows linked to their activities at higher prices. That unethical agents inhibit the rate at which the market learns the value of their activities. That ethical agents are more likely to lose innovation races. And yet ethical agents are welfare enhancing. We show the in-too-deep threshold is always robust to patents and that only punitive punishments after public revelation can mitigate the problem.

Algorithmic Exclusion by Large Language Models

A. Nikandrova, A. Parekh

City St George's, University of London

This paper explores whether large language models (LLMs) can learn predatory strategies in dynamic environments in which an incumbent faces repeated entry threat. Using Ope-nAI's GPT-4.1 as decision-making agents, we find that LLMs learn to predate when both predation and accommodation are theoretically viable,

and adopt aggressive strategies when only accommodation is theoretically viable. Further, profit optimization is limited, highlighting both strategic learning and its limitations.

When and What to Learn in a Changing World

C. Barilla

University of Oxford

A decision-maker periodically acquires information about a changing state, controlling both the timing and content of updates. I characterize optimal policies using a decomposition of the dynamic problem into optimal stopping and static information acquisition. Eventually, information acquisition either stops or follows a simple cycle in which updates occur at regular intervals to restore prescribed levels of relative certainty. This enables precise analysis of long run dynamics across environments. As fixed costs of information vanish, belief changes become lumpy: it is optimal to either wait or acquire information so as to exactly confirm the current belief until rare news prompts a sudden change. The long run solution admits a closed-form characterization in terms of the "virtual flow payoff". I highlight an illustrative application to portfolio diversification.

Mathematical and Computational Finance

Deviations from Tradition: Stylized Facts in the Era of DeFi

D. M. Di Nosse, F. Gatta, F. Lillo, S. Jaimungal

Scuola Normale Superiore, Università di Bologna, University of Toronto, Oxford-Man Institute for Quantitative Finance

Decentralized Exchanges (DEXs) are now a significant component of the financial world where billions of dollars are traded daily. Differently from traditional markets, which are typically based on Limit Order Books, DEXs typically work as Automated Market Makers, and, since the implementation of Uniswap v3, feature concentrated liquidity. By investigating the twenty-four most active pools in Uniswap v3 during 2023 and 2024, we empirically study how this structural change in the organization of the markets modifies the well-studied stylized facts of prices, liquidity, and order flow observed in traditional markets. We find a series of new statistical regularities in the distributions and cross-autocorrelation functions of these variables that we are able to associate either with the market structure (e.g., the execution of orders in blocks) or with the intense activity of Maximal Extractable Value searchers, such as Just-in-Time liquidity providers and sandwich attackers.

Deep Reinforcement Learning for Optimal Trading with Partial Information

A. Macrì, S. Jaimungal, F. Lillo

Scuola Normale Superiore, University of Toronto, Oxford-Man Institute for Quantitative Finance, Università di Bologna

Reinforcement Learning (RL) applied to financial problems has been the subject of a lively area of research. The use of RL for optimal trading strategies that exploit latent information in the market is, to the best of our knowledge, not widely tackled. In this paper we study an optimal trading problem, where a trading signal follows an Ornstein-Uhlenbeck process with regime-switching dynamics. We employ a blend of RL and Recurrent Neural Networks (RNN) in order to make the most at extracting underlying information from the trading signal with latent parameters. The latent parameters driving mean reversion, speed, and volatility are filtered from observations of the signal, and trading strategies are derived via RL. To address this problem, we propose three Deep Deterministic Policy Gradient (DDPG)-based algorithms that integrate Gated Recurrent Unit (GRU) networks to capture temporal dependencies in the signal. The first, a one -step approach (hid-DDPG), directly encodes hidden states from the GRU into the RL trader. The second and third are two-step methods: one (prob-DDPG) makes use of posterior regime probability estimates, while the other (reg-DDPG) relies on forecasts of the next signal value. Through extensive simulations with increasingly complex Markovian regime dynamics for the trading signal's parameters, as well as an empirical application to equity pair trading, we find that prob-DDPG achieves superior cumulative rewards and exhibits more interpretable strategies. By contrast, reg-DDPG provides limited benefits, while hid-DDPG offers intermediate performance with less interpretable strategies. Our results show that the quality and structure of the information supplied to the agent are crucial: embedding probabilistic insights into latent regimes substantially improves both profitability and robustness of reinforcement learning-based trading strategies.

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: Patrick Chang, Fayçal Drissi, Gerardo Duran-Martin, Christopher Hyland, Jiawen Liang, and Kevin Schneider.

References

Barilla, César. 2025. "When and What to Learn in a Changing World." 916–916.

Chen, Songnian, and Junlong Feng. 2025. "Universal Factor Models." *arXiv preprint arXiv:2501.15761*.

Cui, Liyuan, Guanhao Feng, Jianxin Ma, and Yinan Su. 2025. "Breaks and Trends in Factor Premia." *Available at SSRN 5330449*.

Di Nosse, Daniele Maria, Federico Gatta, Fabrizio Lillo, and Sebastian Jaimungal. 2025. "Deviations from Tradition: Stylized Facts in the Era of DeFi." *arXiv preprint arXiv:2510.22834*.

Epstein, Elliot L, Rose Wang, Jaewon Choi, and Markus Pelger. 2025. "Attention Factors for Statistical Arbitrage." *arXiv preprint arXiv:2510.11616*.

Fusari, Nicola, and Sujan Lamichhane. 2025. "Structural Equity Option Pricing: Implications for Credit Risk." *Available at SSRN 5393252*.

Horvath, Blanka, Jonathan Plenk, Milena Vuletić, and Raeid Saqur. 2025. "Generative Models in Finance: Market Generators, a Paradigm Shift in Financial Modeling." *Available at SSRN 5284313*.

Kılıç, Rehim. 2025. "Virtue or Mirage? Complexity in Exchange Rate Prediction." *FEDS Working Paper No. 2025-89*.

Lai, Chieh-Hsin, Yang Song, Dongjun Kim, Yuki Mitsufuji, and Stefano Ermon. 2025. "The Principles of Diffusion Models: From Origins to Advances." *arXiv preprint arXiv:2510.21890*.

Macrì, Andrea, Sebastian Jaimungal, and Fabrizio Lillo. 2025. "Deep Reinforcement Learning for Optimal Trading with Partial Information." *arXiv preprint arXiv:2511.00190*.

McLatchie, Yann, Badr-Eddine Cherief-Abdellatif, David T Frazier, and Jeremias Knoblauch. 2025. "Predictively Oriented Posteriors." *arXiv preprint arXiv:2510.01915*.

Mehta, Dhagash, John RJ Thompson, Hoyoung Lee, and Yongjae Lee. 2025. "Clustering and Similarity Learning in Financial Markets: A Tutorial for the Practitioners." *Available at SSRN 5587353*.

Ng, Kenyon, Edwin Fong, David T Frazier, Jeremias Knoblauch, and Susan Wei. 2025. "TabMGP: Martingale Posterior with TabPFN." *arXiv preprint arXiv:2510.25154*.

Nikandrova, Arina, and Anushree Parekh. 2025. "Algorithmic Exclusion by Large Language Models." *Available at SSRN 5550761*.

Simsek, Yasin. 2025. "Beyond Returns: A Candlestick-Based Approach to Spot Covariance Estimation." *arXiv preprint arXiv:2510.12911*.

Thanassoulis, John E. 2025. "The Problem of Business Evil: Corporate Fraud, Rogue Trading and Pharmaceutical Misconduct." *Available at SSRN 5428157*.

Trzcinka, Charles, and Ziwei Zhao. 2025. "The Impact of Active Managers on the Pricing of Underlying Assets in ETFs." *Working Paper*.