VOL. 5, NO. 4

OMI Research Newsletter – January 2026

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: Abadie et al. (2025) survey modern differences-in-differences and event-study methods with time-varying treatment effects, with special attention to normalisation, heterogeneity, pre-trend tests, and logs-versus-levels choices. Xu and Fan (2025) propose a single-index, characteristics-augmented quantile factor model for time-varying factor loadings and document improved fit for U.S. corporate bond returns. Cartea and Shi (2025) prove that the "virtue of complexity" found in random Fourier-feature return forecasting can persist even when each rolling regression is fit on a single observation, implying the rising out-of-sample R^2 and Sharpe ratio with more features can be a mechanical artifact of random features rather than genuine learning. Brogaard et al. (2025) show that flickering quotes—limit orders cancelled within one second—have a nonzero, persistent price impact and account for about 20% of limit-order price discovery. Harvey, Sancetta and Zhao (2025) clarify how to do valid multiple-testing inference for factor models under dependence and selection bias, and recommends using local FDR and a minimum t-statistic threshold of about 3.0 for new factor discoveries. Moskowitz et al. (2025) argue that time-series momentum is best captured by a nonlinear, "rollback" (S-shaped) signal, which improves out-of-sample Sharpe ratios (especially in market downturns) across futures and factor portfolios. Drissi, Wu and Jaimungal (2025) develop an equilibrium model of liquidity provision on decentralised exchanges with risk-averse liquidity providers who hedge on a frictionless centralised venue, which implies that liquidity supply declines with risk aversion and fundamental volatility but increases with more elastic and less-informed demand.

Statistics, Econometrics & Machine Learning

Harvesting Differences-in-Differences and Event-Study Evidence

A. ABADIE, J. ANGRIST, B. FRANDSEN, J.-S. PISCHKE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY, NBER, BRIGHAM YOUNG UNIVERSITY, LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

This paper surveys econometric innovations related to differences-in-differences estimators and event-study models with time-varying treatment effects. Our discussion highlights tricky normalization issues, heterogeneous policy effects, the interpretation of exposure designs, pretrends pretesting, and the ever-bothersome question of logs versus levels. Key ideas are illustrated with applications.

Single-Index Quantile Factor Model with Observed Characteristics

R. XU, Q. FAN
THE CHINESE UNIVERSITY OF HONG KONG

We propose a characteristics-augmented quantile factor (QCF) model, where unknown factor loading functions are linked to a large set of observed individual-level (e.g., bond- or stock-specific) covariates via a single-index projection. The single-index specification offers a parsimonious, interpretable, and statistically efficient way to nonparametrically characterize the time-varying loadings, while avoiding the curse of dimensionality in flexible nonparametric models. Using a three-step sieve estimation procedure, the QCF model demonstrates high in-sample and out-of-sample accuracy in simulations. We establish asymptotic properties for estimators of the latent factor, loading functions, and index parameters. In an

empirical study, we analyze the dynamic distributional structure of U.S. corporate bond returns from 2003 to 2020. Our method outperforms the benchmark quantile Fama-French five-factor model and quantile latent factor model, particularly in the tails ($\tau = 0.05, 0.95$). The model reveals state-dependent risk exposures driven by characteristics such as bond and equity volatility, coupon, and spread. Finally, we provide economic interpretations of the latent factors.

Exponentially Weighted Estimands and the Exponential Family: Filtering, Prediction and Smoothing

S. D. van Heel, N. Shephard
Harvard University

We propose using a discounted version of a convex combination of the log-likelihood with the corresponding expected log-likelihood such that when they are maximized they yield a filter, predictor and smoother for time series. This paper then focuses on working out the implications of this in the case of the canonical exponential family. The results are simple exact filters, predictors and smoothers with linear recursions. A theory for these models is developed and the models are illustrated on simulated and real data.

BayesSum: Bayesian Quadrature in Discrete Spaces

S. S. Kang, F.-X. Briol, T. Karvonen, Z. Chen
University College London, Lappeenranta–Lahti University of Technology LUT

This paper addresses the challenging computational problem of estimating intractable expectations over discrete domains. Existing approaches, including Monte Carlo and Russian Roulette estimators, are consistent but often require a large number of samples to achieve accurate results. We propose a novel estimator, *BayesSum*, which is an extension of Bayesian quadrature to discrete domains. It is more sample efficient than alternatives due to its ability to make use of prior information about the integrand through a Gaussian process. We show this through theory, deriving a convergence rate significantly faster than Monte Carlo in a broad range of settings. We also demonstrate empirically that our proposed method does indeed require fewer samples on several synthetic settings as well as for parameter estimation for Conway-Maxwell-Poisson and Potts models.

Finance and Economics

The Mechanical Virtue of Random Fourier Features and Model Complexity

Á. Cartea, Y. Shi
Oxford-Man Institute of Quantitative Finance, University of Oxford

In this paper, we examine the empirical "virtue of complexity" in Kelly et al. (2024) in the extreme case where a random Fourier-feature model is estimated on a rolling window containing only one observation. We show that the out-of-sample R-squared of the return forecast, and the average return and Sharpe ratio of an investment market-timing strategy increase as model complexity increases (i.e., the number of features increases). To understand this behavior, we derive closed-form expressions for the expected out-of-sample R-squared of the return forecast and for the expected return of the market-timing strategy, where the equity return and macroeconomic data are fixed, so the expectations are taken only over the randomness in the random features. These formulae show that the relationship between model size and model performance arises mechanically from the structure of the random features rather than from learning. We further show that reversing the sign of returns will reverse the direction of the monotone relationship between model complexity and both the average return and the Sharpe ratio of the timing strategy, that is, as model complexity increases, the performance of the model decreases. Numerical simulations confirm these theoretical predictions in synthetic settings with pure momentum and pure reversal.

The Lasting Impact of Flickering Quotes

J. Brogaard, Á. Cartea, P. Chang, R. Graumans
University of Utah, Oxford-Man Institute of Quantitative Finance, University of Oxford

Flickering quotes are limit orders that are cancelled shortly after entry. Theoretically, the net price impact of flickers should be zero; yet, we find that the impact of flickers persists beyond the cancellation leg of the flicker. Flickers account for around 20% of the total price discovery of limit orders, even though they account for less than 0.3% of the total lifetime of all limit orders that are cancelled. Flickers are followed by reactions that result in a thicker book on the side of the flicker and a thinner book on the opposite side of the flicker, where the new shape on both sides of the book lasts beyond the cancellation leg of the flicker. Thus, flickers can lead to long-lived price discovery by inducing other orders to impound information into prices.

What Threshold Should be Applied to Tests of Factor Models?

C. R. Harvey, A. Sancetta, Y. Zhao
Duke University, Royal Holloway, University of London, University of Sussex

Researchers broadly acknowledge that statistical tests

must be adjusted when hundreds of factors and trading strategies have been examined. But how should these adjustments be made? Existing methods are often misunderstood or misapplied. We show that proper inference requires accounting for dependence across tests, correctly specifying the null distribution, and mitigating sample-selection bias. We develop a simple framework that avoids assumptions about the total number of tests run and yields a lower bound on valid significance thresholds - implying that researchers should employ a t-statistic cutoff of at least 3.0. In addition, we advocate using the local False Discovery Rate, which provides the probability that the null hypothesis is true for a given test-statistic realization - information that a conventional p-value cannot supply.

Nonlinear Time Series Momentum

T. J. Moskowitz, R. Sabbatucci, A. Tamoni, B. Uhl

Yale University, NBER, Stockholm School of Economics, University of Notre Dame, Universität Hamburg

We document a persistent nonlinear relationship between price trends and risk-adjusted returns across markets and asset classes that is consistent with asset pricing theory. Nonlinearities in time series momentum are consistent with past returns reflecting information about conditional expected returns, in line with investors using conditioning information to form efficient portfolios. Machine learning techniques are useful in uncovering these relationships and yield economically and statistically significant out-of-sample improvements in time series momentum strategies.

A Practitioner's Guide to AI+ML in Portfolio Investing

M. Caner, Q. Fan

North Carolina State University, The Chinese University of Hong Kong

In this review, we provide practical guidance on some of the main machine learning tools used in portfolio weight formation. This is not an exhaustive list, but a fraction of the ones used and have some statistical analysis behind it. All this research is essentially tied to precision matrix of excess asset returns. Our main point is that the techniques should be used in conjunction with outlined objective functions. In other words, there should be joint analysis of Machine Learning (ML) technique with the possible portfolio choice-objective functions in terms of test period Sharpe Ratio or returns. The ML method with the best objective function should provide the weight for portfolio formation. Empirically we analyze five time periods of interest, that are out-sample and show performance of some ML-Artificial Intelligence (AI) methods. We see that nodewise regression with Global Minimum Variance portfolio based weights deliver very good Sharpe Ratio and returns across five time periods in this century we analyze. We cover three downturns, and 2 long term investment spans.

Deep Parametric Portfolio Policies

F. Simon, S. Weibels, T. Zimmermann

LIQID, University of Cologne

We consider parametric portfolio policies of any complexity using deep neural networks to optimize investor utility. Risk aversion acts as an economic regularization mechanism, with higher risk aversion constraining model complexity. Empirically, Deep Parametric Portfolio Policies generate 43-102 basis points higher monthly certainty equivalent returns compared to linear policies. Looking beyond expected returns, non-linear portfolio policies better capture the complex relationship between investor preferences and firm characteristics but the benefits of using complex models vary with investor preferences. Results hold across different utility functions and remain robust to transaction costs and short-selling restrictions. Overall, economic regularization constrains model complexity much like statistical regularization but emerges endogenously from investor preferences.

Shrinkage Alignment in High-Dimensional Portfolios

S. Ke, M. Pourmohammadi

Yale University

We study how shrinkage affects portfolio efficiency when the number of assets approaches or exceeds the sample size. Standard methods, such as ridge, impose uniform shrinkage, treating all assets as ex-ante identical and creating inefficiency when profitability is heterogeneous. Empirically, this "one-size-fits-all" design produces a hump-shaped relationship between model complexity and out-of-sample Sharpe ratios: adding assets can paradoxically reduce performance. We introduce shrinkage alignment, showing that efficiency requires matching shrinkage strength to each asset's true profitability. Building on this insight, we propose Sharpe Ratio Shrinkage (SRS)-a data-driven approach that aligns shrinkage intensity with empirical Sharpe ratios. SRS outperforms conventional methods under profitability heterogeneity and restores the virtue of complexity in high-dimensional portfolio construction.

Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database

M. Acosta, A. Ajello, M. Bauer, F. Loria, S. Miranda-Agrippino

University of Wisconsin–Madison, Federal Reserve Board, University of Oxford, CPER

This paper introduces the U.S. Monetary Policy Event-Study Database (USMPD), a novel, public, and regularly updated dataset of financial market data around Federal Open Market Committee (FOMC) policy announcements, press conferences, and minutes releases. Using the rich high-frequency data in the USMPD, we document several new empirical findings. Large monetary policy surprises have made a comeback in recent years, and post-meeting press conferences have become the most important source of policy news. Monetary policy surprises have pronounced negative effects on breakeven inflation based on Treasury yields. Risk assets, including dividend derivatives, also respond strongly and negatively to monetary policy surprises, consistent with conventional channels of monetary transmission. Press conferences have stronger effects than FOMC statements on most asset prices. Finally, the term structure evidence shows peak effects on market-based inflation and dividend expectations at horizons of several years.

Dual Peer Effects and Cross-Stock Predictability

D. Avramov, S. Ge, S. Li, O. Linton

Reichman University, Nankai University, Peking University, University of Cambridge

This paper introduces a Peer Index (PI) constructed from economically motivated peer networks that summarizes (i) the strength of a firm's peers and (ii) the firm's position within its peer group. PI predicts stock returns and earnings surprises over short and long horizons. Machine-learning models based solely on firm-level characteristics do not subsume PI's predictive power, supporting the interpretation that it captures genuine cross-stock information. Lag-augmented local projections show that positive PI innovations are followed by higher next-month returns that gradually decay without reversal, consistent with slow diffusion of peer information into prices.

Monetary Policy Shocks: A New Hope — Large Language Models and Central Bank Communication

Ruben Fernandez-Fuertes

Bocconi University

I develop a multi-agent LLM framework that processes Federal Reserve communications to construct narrative monetary policy surprises. By analyzing Beige Books and Minutes released before each FOMC meeting, the system generates conditional expectations that yield less noisy surprises than market-based measures. These surprises produce theoretically consistent impulse responses where contractionary shocks generate persistent disinflationary effects and enable profitable yield curve trading strategies that outperform alternatives. By directly extracting expectations rather than cleaning surprises ex post, this approach demonstrates how multi-agent LLMs can implement narrative identification at scale without contamination in high-frequency measures.

Mathematical and Computational Finance

Equilibrium Liquidity and Risk Offsetting in Decentralised Markets

F. Drissi, X. Wu, S. Jaimungal

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

We develop an economic model of decentralised exchanges (DEXs) in which risk-averse liquidity providers (LPs) manage their exposure in a centralised exchange (CEX) based on risk preferences, private information, and trading costs. Rational, risk-averse LPs anticipate the frictions associated with replication and manage risk primarily by reducing the reserves supplied to the DEX. Greater risk aversion reduces the equilibrium viability of liquidity provision, resulting in thinner markets and lower trading volumes. Greater uninformed demand supports deeper liquidity, whereas higher fundamental price volatility erodes it. Finally, while moderate anticipated price changes can improve LP performance, larger changes require more intensive trading in the CEX, generate higher replication costs, and induce LPs to reduce liquidity supply.

Multi-Task Gaussian Processes for Tactical Execution Under Regime Shifts

Á. Cartea, F. Drissi, G. Palmari

Oxford-Man Institute of Quantitative Finance, University of Oxford

In electronic markets, broker execution algorithms comprise a strategic layer that sets a trading schedule and a tactical layer that optimise child-order execution. This paper introduces the Multi-Task Gaussian Process bandit with Likelihood Ratio test (MTGP-LR) for the tactical layer. MTGP-LR maps market features to trading actions, adapts to regime shifts through change point detection, and employs a kernel that transfers learning across related actions. In stationary markets, MTGP-LR achieves sublinear regret. In nonstationary markets, we show how to control inference error and provide guarantees for detection delay. Our main result derives high-probability regret bounds in regime-switching markets. We illustrate the framework in two applications: child market orders timing and child limit order placement. Without assumptions on dynamics or regime transitions, MTGP-LR is accurate, computationally efficient, and outperforms

control-based models when brokers misspecify model parameters.

On the Impossibility of Transparent and Decentralized DeFi Trading

H. Halaburda

New York University

Blockchain technology promises a transparent, decentralized, and fair architecture for peer-to-peer interaction, including finance. This paper explores the inherent tension between this decentralized architecture and the economic equilibrium that emerges on top of it. We argue that this equilibrium is centralized and extractive. We develop a generalized, protocol-agnostic model of strategic interaction between payoff-maximizing agents (e.g., block builders, miners) who compete to create blocks. We show that free entry, transparency, and latency-bounded settlement create a natural and powerful tendency towards centralization and rent extraction. This tendency is driven by three main forces: (1) transparency and latency-bounded settlement create predictable opportunities for MEV extraction, so even small differences in agents' extraction capabilities generate persistent heterogeneity in effective rewards; (2) block production across PoW, PoS, and PBS is a winner-take-all contest that strictly sorts agents by reward and magnifies small advantages, producing increasing concentration as contest decisiveness rises or reward dispersion widens; and (3) rational traders avoid extraction in the public mempool by routing privately to the most reliable execution agent, introducing strong network complementarities that make monopoly private order flow the only stable outcome under any positive multihoming cost, while profitable "seeding" by traders can eliminate routing equilibria altogether. Together, these forces generate an impossibility result: transparent and decentralized trading cannot coexist in equilibrium. All robust equilibria are centralized, while decentralized allocations require knife-edge conditions or fail to exist. Transparency—intended to eliminate intermediaries—systematically recreates and entrenches them.

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

Abadie, Alberto, Joshua Angrist, Brigham Frandsen, and Jörn-Steffen Pischke. 2025. "Harvesting Differences-in-Differences and Event-Study Evidence."

Acosta, Miguel, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino. 2025. "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database." *Available at SSRN 5944756*.

Avramov, Doron, Shuyi Ge, Shaoran Li, and Oliver B Linton. 2025. "Dual Peer Effects and Cross-Stock Predictability." *Available at SSRN 5910346*.

Brogaard, Jonathan, Álvaro Cartea, Patrick Chang, and Rob Graumans. 2025. "The Lasting Impact of Flickering Quotes." *Available at SSRN 5963035*.

Caner, Mehmet, and Qingliang Fan. 2025. "A Practitioner's Guide to AI+ML in Portfolio Investing." *Available at SSRN 5546200*.

Cartea, Álvaro, and Yuantao Shi. 2025. "The Mechanical Virtue of Random Fourier Features and Model Complexity." *Available at SSRN 5930257*.

Cartea, Álvaro, Fayçal Drissi, and Gianluca Palmari. 2025. "Multi-Task Gaussian Processes for Tactical Execution Under Regime Shifts." *Available at SSRN 5946898*.

Drissi, Fayçal, Xuchen Wu, and Sebastian Jaimungal. 2025. "Equilibrium Liquidity and Risk Offsetting in Decentralised Markets." *Available at SSRN*.

Fernandez-Fuertes, Ruben. 2025. "Monetary Policy Shocks: A New Hope — Large Language Models and Central Bank Communication." *Available at SSRN 5669212*.

Halaburda, Hanna. 2025. "On the Impossibility of Transparent and Decentralized DeFi Trading Preliminary-Comments Welcome." *Available at SSRN 5772982*.

Harvey, Campbell R, Alessio Sancetta, and Yuqian Zhao. 2025. "What Threshold Should be Applied to Tests of Factor Models?" *Available at SSRN 5925386*.

Kang, Sophia Seulkee, François-Xavier Briol, Toni Karvonen, and Zonghao Chen. 2025. "BayesSum: Bayesian Quadrature in Discrete Spaces." *arXiv preprint arXiv:2512.16105*.

Ke, Shikun, and Mohammad Pourmohammadi. 2025. "Shrinkage Alignment in High-Dimensional Portfolios." *Available at SSRN 5723922*.

Moskowitz, Tobias J, Riccardo Sabbatucci, Andrea Tamoni, and Björn Uhl. 2025. "Nonlinear Time Series Momentum." *Available at SSRN 5933974*.

Simon, Frederik, Sebastian Weibels, and Tom Zimmermann. 2025. "Deep Parametric Portfolio Policies." *Available at SSRN 4150292*.

van Heel, Simon Donker, and Neil Shephard. 2025. "Exponentially weighted estimands and the exponential family: filtering, prediction and smoothing." *arXiv preprint arXiv:2512.16745*.

Xu, Ruofan, and Qingliang Fan. 2025. "Single-Index Quantile Factor Model with Observed Characteristics." *arXiv preprint arXiv:2506.19586*.