VOL. 5, NO. 7

OMI Research Newsletter – April 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.*

Announcements

The Workshop on Microstructure & Financial Economics will be held at OMI on 7–8 May 2026. For the schedule and registration details, follow this link.

The OMI Machine Learning and Finance Conference will be held at the Mathematical Institute on 11–12 June 2026. For the schedule and registration details, follow this link.

The OMI–SBS Conference on FinTech and Blockchain Economics will be held at Saïd Business School and OMI on 25–26 June 2026. For the schedule and registration details, follow this link.

Highlights

Key contributions in this month's edition: Kolokolov and Yu (2026) develop a penalised high-frequency regression framework for estimating and selecting sparse time-varying factor exposures when both responses and covariates are semimartingales with jumps, and apply it to the "zoo" of high-frequency asset-pricing factors. Chen, Didisheim and Somoza (2026) develop an entropy-based framework that uses LLMs' inner probabilities to quantify predictive uncertainty through an "inner confidence" measure. Ding et al. (2026) propose a high-dimensional predictive factor model for stock jump intensities and show that exploiting cross-sectional jump dependence yields meaningful gains in prediction and risk assessment. Cartea et al. (2026) show that LLM-based AI agents participate in both irrational and rational speculative bubbles, and that greater reasoning capacity reduces, but does not eliminate, such bubbles or improve market efficiency. Oh and Patton (2026) propose a decomposition of mutual fund abnormal returns into "skill" and "efficiency", document significant heterogeneity in both across U.S. equity mutual funds, and show that efficiency is more persistent and more informative for predicting future abnormal returns. Muhle-Karbe et al. (2026) develop a unified market microstructure framework in which several core empirical regularities of order flow, volatility, and market impact emerge jointly from a common underlying mechanism. Bandi et al. (2026) propose a parsimonious model for ultra-short-term option surfaces that can simultaneously match the irregular at-the-money implied-volatility term structure across very short maturities and the implied-volatility smile at each maturity.

Statistics, Econometrics & Machine Learning

Realized Regularized Regressions

A. KOLOKOLOV, S. YU
NEW ECONOMIC SCHOOL, OXFORD-MAN INSTITUTE, UNIVERSITY OF OXFORD

Financial datasets are expanding rapidly, with an increasing number of variables observed at high frequencies. This raises challenges for estimation and variable selection in high-frequency regressions within a high-dimensional environment. We develop a continuous-time penalized regression framework for the estimation of time-varying coefficients when both the response and covariates are Itô semimartingales with jumps. The coefficient paths are approximated by spline basis expansions and estimated via least squares from truncated high-frequency increments. In a finite-dimensional setting, we establish consistency and derive a feasible asymptotic distribution for the integrated coefficient estimator under infill asymptotics. We then extend the framework to high-dimensional

settings in which the number of candidate covariates diverges, and show that a group-wise penalized estimator with a truncated ℓ_1-penalty attains the oracle property, which delivers both consistent model selection and coefficient estimation. An empirical application to a large panel of more than two hundred high-frequency factors investigates sparse factor structure across a large cross-section of stocks and industry portfolios.

Out of the Black Box: Uncertainty Quantification for LLMs via Conditional Probabilities

H. Chen, A. Didisheim, L. A. Somoza

Massachusetts Institute of Technology, ESSEC Business School, University of Lausanne

Autoregressive LLMs generate text by sampling from estimated probability distributions over the next token, conditional on prior context. We use these probabilities to construct an entropy-based measure of prediction uncertainty, termed inner confidence. In news classification, LLM predictions with higher inner confidence are systematically more accurate. To evaluate the measure's economic relevance, we form long-short portfolios based on LLM predictions. The portfolio based on high-confidence predictions achieves a Sharpe ratio about 20% higher than the unconditional benchmark, while the one based on low-confidence predictions yields no excess returns. In contrast, self-declared confidence exhibits significant decoding biases and provides no comparable performance gains.

Triple/Double-Debiased Lasso

D. Chetverikov, J. R.-V. Sørensen, A. Tsyvinski

University of California, Los Angeles, University of Copenhagen, Yale University, Aarhus Center for Econometrics (ACE)

In this paper, we propose a triple (or double-debiased) Lasso estimator for inference on a low-dimensional parameter in high-dimensional linear regression models. The estimator is based on a moment function that satisfies not only first- but also second-order Neyman orthogonality conditions, thereby eliminating both the leading bias and the second-order bias induced by regularization. We derive an asymptotic linear representation for the proposed estimator and show that its remainder terms are never larger and are often smaller in order than those in the corresponding asymptotic linear representation for the standard double Lasso estimator. Because of this improvement, the triple Lasso estimator often yields more accurate finite-sample inference and confidence intervals with better coverage. Monte Carlo simulations confirm these gains. In addition, we provide a general recursive formula for constructing higher-order Neyman orthogonal moment functions in Z-estimation problems, which underlies the proposed estimator as a special case.

Predictive Factor Model for Jump Intensities

Y. Ding, S. Hu, Y. Li, X. Zheng

University of Macau, Hong Kong University of Science and Technology

We propose a parsimonious factor model for jump intensities for a large number of stocks, which we call the Jump Factor Intensity Thinning (J-FIT) Model. The model utilizes a jump factor process to capture the complete grid of potential jumps, and a thinning procedure to capture individual stock jumps. The jump factor process features temporal dependence, which leads to predictability. We develop estimators and inference theories for the model parameters, enabling applications such as jump intensity prediction and constructing prediction intervals for the jump risk exposures of individual stocks and sectors. Empirically, we provide solid evidence that jumps are predictable with the J-FIT model, whose predictive performance substantially surpasses various benchmark models. We also demonstrate that J-FIT is useful for tail risk management.

Finance and Economics

AI Bubbles with Large Language Models

Á. Cartea, P. Chang, N. Chen, M. Zhong

Oxford-Man Institute, University of Oxford, The Chinese University of Hong Kong

AI agents participate in both rational and irrational bubbles. In a sequential bubble game, AI agents generate speculative trades even though there is a unique no-trade equilibrium. Increasing reasoning capacity reduces, but does not eliminate, such irrational bubbles. Chain-of-thought analysis shows that AI agents participate in irrational bubbles because they form incorrect beliefs through simplified expectations of the market, and they are sensitive to framing effects when forming their beliefs. When the environment admits both a no-trade equilibrium and a speculative bubble equilibrium, AI agents recognize this multiplicity and coordinate on the speculative equilibrium. Our results suggest that greater AI sophistication does not necessarily eliminate bubbles or improve market efficiency.

Skill and Efficiency in the U.S. Mutual Fund Industry

D. H. Oh, A. J. Patton

Federal Reserve Board, Duke University, Singapore Management University

We propose a new measure of mutual fund manager ability: "efficiency" is the ability to accrue the risk premium associated with a risk factor. The familiar

abnormal return, or alpha, is shown to be the sum of two distinct measures of ability: "aggregate efficiency" which is the beta-weighted sum of the fund's (in)efficiencies across risk factors, and "skill," the component that is unrelated to factor exposures. Using a panel of U.S. equity mutual fund returns from 1999-2023, we document significant heterogeneity in mutual fund manager skill and efficiency. We employ regression trees and their extensions to capture this heterogeneity. We find that efficiency is more persistent than skill, and we show that future abnormal returns can be better predicted by decomposing lagged abnormal returns into skill and efficiency.

The News-Contaminated Market: Local Factors and Asset Pricing

C. He

Northwestern University

Much of the information arriving in financial markets has a local structure: It directly affects subsets of firms while also generating a market-wide effect. Observationally, this can be modeled as a factor structure with a global factor and local factors that need not be orthogonal to each other. Because the value-weighted market factor aggregates both components, it conflates the global factor with the local components, giving rise to the news-contaminated market. We propose a decomposition of the market factor into a global component and an orthogonal residual, the local news factor (LNF), leveraging high-frequency data to capture the time-varying factor structure. Under this decomposition, market beta reflects exposure to both global and local components, giving rise to an omitted variable bias in beta-sorted portfolios. This bias flattens the security market line (SML), as the raw local factors are negatively correlated with the global factor, while the residual LNF is only weakly priced. In particular, the LNF is strongly priced overnight but exhibits excess volatility and negative returns during the day. Controlling for local exposures restores an upward-sloping SML, even during intraday periods and non-special days when the SML is particularly flat. The two-factor decomposition helps span low-risk anomalies, including betting-against-beta and idiosyncratic volatility, and explains the shrinking beta dispersion during the trading day. Together, these findings identify the contamination of the market factor by local news shocks as a common source of the flat SML, low-risk anomalies, and intraday beta dynamics.

Latency and the Look-Ahead Bias in Trade and Quote Data

R. H. Battalio, C. W. Holden†, M. Pierson, J. J. Shim, J. Wu

Indiana University, University of Notre Dame, University of Pennsylvania

The NYSE Trade and Quote (TAQ) dataset reports events sequentially out of order. Trades and quote changes are clustered but recorded with differential latency, partly due to geography, effectively mixing up the order of events. Consequently, quote changes after, and in response to, a trade are regularly reported before it, creating a look-ahead bias. This bias creates errors in signing and spread measurement when using the TAQ NBBO. Errors intensify with latency: around 20% of high-latency trades are signed incorrectly and effective spreads are understated by over 40%. We propose a more accurate latency-free signing methodology, grounded by exchange rules.

Breaking News

Lars Winkelmann, Wenying Yao

Freie Universität Berlin, University of Melbourne

This paper examines how regulatory interventions in high-frequency financial markets affect price discovery. We focus on Breaking news, where dynamic circuit breakers trigger trading halts immediately after the release of macroeconomic fundamentals. Within a high-frequency signal-in-noise model, we show that triggering rules complicate statistical inference for the price impact of news, rendering conventional non-parametric jump estimators inconsistent. Building on this insight, we develop a regression-based test for fundamental pricing that accounts for non-vanishing transition times. The test compares transition price changes to efficient jumps implied by observable factors. Our empirical analysis of CME E-mini S&P 500 futures shows that Breaking news are associated with systematic deviations from fundamental pricing, predominantly in the form of overshooting. Our findings highlight a regulatory trade-off: the appeal of simple and transparent circuit breaker rules must be weighed against their cost of preventing fundamentals from being priced contemporaneously, thereby creating adverse incentives and introducing distortions.

The AI Layoff Trap

B. H. Falk, G. Tsoukalas

University of Pennsylvania, Boston University

If AI displaces human workers faster than the economy can reabsorb them, it risks eroding the very consumer demand firms depend on. We show that knowing this is not enough for firms to stop it. In a competitive task-based model, demand externalities trap rational firms in an automation arms race, displacing workers well beyond what is collectively optimal. The resulting loss harms both workers and firm owners. More competition and "better" AI amplify the excess; wage adjustments and free entry cannot eliminate it. Neither can capital income taxes, worker equity participation, universal basic income, upskilling, or Coasian bargaining. Only a Pigouvian automation tax can. The results suggest that policy should

address not only the aftermath of AI labor displacement but also the competitive incentives that drive it.

Tokenomics and Blockchain Fragmentation

H. S. Shin

Bank for International Settlements

Money is a coordination device underpinned by strong network e§ects: the more others accept a form of money, the more I wish to adopt it too. The decentralisation agenda of public permissionless blockchains undercuts these network e§ects and leads to fragmentation of the monetary landscape. Validators who maintain the blockchain need to be rewarded to play their role with the necessary reward increasing in the degree of dependence on other validatorsíactions to sustain consensus. Since these rewards must ultimately be borne by users through congestion rents, capacity constraints are a feature, not a bug, especially for blockchains with more stringent standards for consensus. New blockchains with less stringent thresholds for consensus enter the market to serve users priced out of incumbent chains. The resulting fragmentation undercuts the very network e§ects that give money its social value. Stablecoins inherit this fragmentation from the blockchains on which they reside. The analysis has broader implications for the future of the monetary system.

Mathematical and Computational Finance

A Unified Theory of Order Flow, Market Impact, and Volatility

J. Muhle-Karbe, Y. O. Chahdi, M. Rosenbaum, G. Szymanski

Imperial College London, CentraleSupélec, Université Paris Dauphine-PSL, University of Luxembourg

We propose a microstructural model for the order flow in financial markets that distinguishes between *core orders* and *reaction flow*, both modeled as Hawkes processes. This model has a natural scaling limit that reconciles a number of salient empirical properties: persistent signed order flow, rough trading volume and volatility, and power-law market impact. In our framework, all these quantities are pinned down by a single statistic H_0, which measures the persistence of the core flow. Specifically, the signed flow converges to the sum of a fractional process with Hurst index H_0 and a martingale, while the limiting traded volume is a rough process with Hurst index $H_0 - 1/2$. No-arbitrage constraints imply that volatility is rough, with Hurst parameter $2H_0 - 3/2$, and that the price impact of trades follows a power law with exponent $2 - 2H_0$. The analysis of signed order flow data yields an estimate $H_0 \approx 3/4$. This is not only consistent with the square-root law of market impact, but also turns out to match estimates for the roughness of traded volumes and volatilities remarkably well.

Ultra-Short-Term Volatility Surfaces

F. M. Bandi, N. Fusari, G. Gazzani, R. Renò

Johns Hopkins University, University of Verona, ESSEC Business School

Options with maturities below one week, hereafter ultra-short-term options, have seen a sharp increase in trading activity in recent years. Yet, these instruments are difficult to price jointly using classical pricing models due to the pronounced oscillations observed in the at-the-money implied-volatility term structure across ultra-short-term tenors. We propose Edgeworth++, a parsimonious jump-diffusion model featuring a nonparametric stochastic volatility component, which provides flexibility in capturing the implied-volatility smiles for each tenor, combined with a deterministic shift extension, which allows the model to fit rich at-the-money impliedvolatility term structures across tenors. We derive a local (in tenor) expansion of the process characteristic function suited to value ultra-short-term options. The expansion leads to fast and accurate option pricing in closed form via standard Fourier inversion. We discuss the benefits of the proposed approach relative to natural benchmarks.

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 members at the institute: Patrick Chang, Fayçal Drissi, Christopher Hyland, Jiawen Liang, and Kevin Schneider.

References

Bandi, Federico M, Nicola Fusari, Guido Gazzani, and Roberto Renò. 2026. "Ultra-Short-Term Volatility Surfaces." *Available at SSRN 6500958*.

Battalio, Robert H, Craig W Holden, Matthew Pierson, John J Shim, and Jun Wu. 2026. "Latency and the Look-Ahead Bias in Trade and Quote Data." *Available at SSRN 5907665*.

Cartea, Alvaro, Patrick Chang, Nan Chen, and Mingyue Zhong. 2026. "AI Bubbles with Large Language Models." *Available at SSRN 6426761*.

Chen, Hui, Antoine Didisheim, and Luciano A Somoza. 2026. "Out of the black box: Uncertainty quantification for llms via conditional probabilities." NBER Working Paper No.w34965.

Chetverikov, Denis, Jesper R-V Sørensen, and Aleh Tsyvinski. 2026. "Triple/Double-Debiased Lasso." *arXiv preprint arXiv:2603.20134*.

Ding, Yi, Shiman Hu, Yingying Li, and Xinghua Zheng. 2026. "Predictive Factor Model for Jump Intensities." *Available at SSRN 6496119*.

Falk, Brett Hemenway, and Gerry Tsoukalas. 2026. "The AI Layoff Trap." *arXiv preprint arXiv:2603.20617*.

Hafner, Christian, Andrew C Harvey, and Linqi Wang. 2026. "Modeling Prices from Speculative Markets: Bursting Bubbles or Deflating Balloons?" *Available at SSRN 6427883*.

He, Chun. 2026. "The News-Contaminated Market: Local Factors and Asset Pricing." *Available at SSRN 6480458*.

Kolokolov, Aleksey, and Shifan Yu. 2026. "Realized Regularized Regressions." *Available at SSRN 6439163*.

Muhle-Karbe, Johannes, Youssef Ouazzani Chahdi, Mathieu Rosenbaum, and Grégoire Szymanski. 2026. "A Unified Theory of Order Flow, Market Impact, and Volatility." *arXiv preprint arXiv:2601.23172*.

Oh, Dong Hwan, and Andrew J Patton. 2026. "Skill and Efficiency in the U.S. Mutual Fund Industry." *Available at SSRN 6407798*.

Shin, Hyun Song. 2026. "Tokenomics and blockchain fragmentation." *Bank for International Settlements Working Paper*.

Winkelmann, Lars, and Wenying Yao. 2026. "Breaking News." *arXiv preprint arXiv:2603.22835*.