VOL. 5, NO. 11

OMI Research Newsletter – August 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.*

Announcement

The OMI Financial Econometrics Workshop will be held at the Oxford-Man Institute on 22–23 October 2026. For the schedule and registration details, follow this link.

Highlights

Key contributions in this month's edition include the following. Chen and Wang (2026) connect transformers, foundation models, and diffusion forecasters to the econometric tradition of vector autoregressions, clarifying both their predictive gains and their current inferential limitations. Christensen and Norlyk (2026) derive rough microstructure noise from a tick-level model of permanent and fleeting price changes and develop a formal test for roughness. Borri, Liu and Tsyvinski (2026) use realized consumption across more than 400 large language models to construct an AI factor and document a sizeable, heterogeneous equity premium. Galiani and Pantano (2021) review structural modelling in empirical microeconomics, emphasizing its value for counterfactual analysis and how experimental variation can strengthen identification and validation. Chaudhry and Davis (2026) argue that weak substitution across individual stocks provides an economic explanation for the factor zoo. Bauer, Czarnota and Klein (2026) show that financial constraints materially weaken firms' price responses to monetary policy. Ellis, Piccione and Zhang (2026) show that widespread use of a machine-learning algorithm can reduce, rather than improve, equilibrium information aggregation. Horvath et al. (2026) establish minimax-optimal approximation rates for signature-based path regression and carry the rates through several statistical learning procedures.

Statistics, Econometrics & Machine Learning

From Vector Autoregressions to AI-Based Time Series Forecasting: A Review

L. CHEN, W. WANG

WASHINGTON UNIVERSITY IN ST. LOUIS, UNIVERSITY OF BRISTOL

This review organizes three major developments in AI-based time-series forecasting—transformers, large pretrained models for zero-shot prediction, and diffusion-based generative forecasters—around the conditional distribution of the future given the past, the same object that underlies the vector autoregression. It evaluates the approaches through the longstanding challenges of high dimensionality, nonstationarity, and nonlinearity. Modern methods expand the classical forecasting template through more flexible dynamics, larger information sets and training corpora, and richer predictive distributions. At the same time, they often lack the inferential and structural tools needed for hypothesis testing, explanation, and policy analysis. The paper concludes by identifying open problems where econometric methods remain especially valuable.

Asset-Specific Limit Order Microstructure Noise: Parameter Estimation and Empirical Evidence

M. BIBINGER, A. GRÜBER, M. JIRAK

UNIVERSITY OF WÜRZBURG, UNIVERSITY OF VIENNA

The paper extends the one-sided noise model for high-frequency limit-order-book quotes by allowing the lower tail of the noise distribution to vary across assets. The resulting tail parameter is central to inference and can be estimated from moments of high-frequency returns under

a parametric noise specification. For Gamma noise, the authors establish a central limit theorem for observations from the associated Gamma-difference distribution and show that it continues to hold when the observations also contain a general semimartingale log-price component. The analysis demonstrates how the tail parameter changes the attainable accuracy of efficient-price and volatility estimation. Evidence from recent NASDAQ limit-order-book data confirms that asset-specific tail behaviour is empirically important.

Factor Dipping

Y. Zeng
University of Cambridge

Principal-component factor estimation relies on a sufficiently large spectral gap, an assumption that can be fragile in macroeconomic and financial panels. The paper shows that weak factors, combined with common linear preprocessing choices, can create substantial factor-estimation uncertainty that conventional asymptotic intervals understate. Factor dipping re-estimates the factors after bounded, invertible transformations of the data and then reverses those transformations. Under the standard factor model, the dipped and original estimates remain nearly invariant; departures from this benchmark therefore support a specification test and a representation-sensitivity band. Applications to Stock–Watson factors, factor-augmented VAR impulse responses, three-pass risk-premium estimation, and quantile factor models show that the diagnostic can materially alter downstream conclusions.

Microstructural Foundations of Rough Noise

P. K. Christensen, A. M. Norlyk
Aarhus University

The authors propose a tick-by-tick price model that separates permanent price changes from fleeting changes generated by microstructure noise. At the macroscopic scale, the model converges to a standard semimartingale efficient-price process plus a continuous rough-noise component, thereby supplying a microstructural foundation for rough noise. They develop a GMM estimator for tick data and a formal roughness test, and establish the relevant asymptotic results. Simulations indicate good finite-sample behaviour. Applied to 2024 tick data for Dow Jones Industrial Average constituents, the method uncovers marked day-to-day variation: rough noise is present but not universal, its estimated roughness is usually mild, and detections are concentrated on days characterized by pronounced short-run price reversals.

Bootstrap Inference in Autoregressive Duration Models

G. Cavaliere, T. Mikosch, A. Rahbek, F. Vilandt
University of Bologna, University of Exeter, University of Copenhagen

This paper develops bootstrap inference for autoregressive conditional duration models observed over a fixed calendar span, where the number of recorded durations is itself random. It compares recursive schemes that hold fixed either calendar time or the realized event count. For the fixed-count bootstrap, consistency holds when the duration tail index is at least one. With infinite-mean durations and a tail index below one, the estimator has a mixed-normal limit and classical bootstrap consistency fails, but the bootstrap still reproduces the conditional Gaussian component. As a result, basic percentile intervals remain first-order valid and bootstrap t-statistics are asymptotically standard normal. Simulations cover finite- and infinite-mean regimes, and an application studies transaction durations for cryptocurrency exchange-traded funds.

Finance and Economics

AI Premium

N. Borri, Y. Liu, A. Tsyvinski
LUISS University, University of Rochester, Yale University

Using 380 trillion tokens of realized consumption across more than 400 large language models, the authors construct a high-frequency AI factor from growth in tokens, spending, and users, and estimate firm-level AI betas from stock-return comovement. Firms with higher AI betas subsequently earn higher returns: a value-weighted high-minus-low strategy earns 64.1 basis points per week. The premium is strongest for exposure to intensive, frontier-oriented use—closed-source models, paying and experienced users, and longer prompts—rather than casual or open-weight use. It extends beyond technology firms and is present in developed but not emerging markets. The paper also maps market-implied AI exposure to occupations and skills and documents the rapid rise of agentic AI consumption.

Structural Models: Inception and Frontier

S. Galiani, J. Pantano
University of Maryland, University of Arizona

The paper reviews the origins, evolution, and frontier of structural empirical microeconomics, with particular emphasis on labour, population, and human-capital applications. It first explains the central features of the structural approach for readers less familiar with the methodology, including how to determine when a structural model is needed and how economic primitives, behavioural restrictions, and equilibrium conditions

should be formulated. It then discusses identification, estimation, validation, and counterfactual policy analysis, drawing lessons from the field's development since the 1970s and 1980s. The review also considers how experiments, richer data, improved computation, and machine-learning methods can complement structural modelling while preserving its focus on mechanisms and policy-invariant primitives.

The Origins of the Factor Zoo: Investors Weakly Substitute Across Stocks

A. Chaudhry, C. Davis

The Ohio State University

The authors argue that investors view individual stocks as only weak substitutes and that this weak substitution explains why expected returns depend on a large "factor zoo." Classical asset-pricing models imply strong reallocation across securities that covary, concentrating expected-return differences in a small number of dominant covariance factors. Using cross-sectional return predictability, cross-price multipliers, and cross-price elasticities, the paper instead finds empirically weak substitution. In equilibrium, this allows even weak factors that explain little covariance to affect expected returns, while aggregate Sharpe ratios remain moderate. The same mechanism implies that misspecifying substitution patterns produces relatively small biases in estimated price elasticities and multipliers.

Holding Stocks, Trading Bonds

P. Guasoni, M. H. Weber

Dublin City University, University of Bologna, National University of Singapore

Why do rational investors trade, and which assets do they use? The paper studies a dynamic economy with stochastic growth and long-lived agents who differ in risk aversion, time preference, and income. Agents trade stocks, a long-term bond, and short-term loans under common information and beliefs. In closed-form equilibrium, they hold but do not trade stocks, even though stock returns remain time-varying and predictable. Instead, they continuously trade the long-term bond to hedge the intertemporal effects of growth shocks on interest rates, dividends, and personal income. If investors also receive access to bond funds that add no spanning power because they can be dynamically replicated, trade ceases. Long-term bonds therefore emerge as the natural instruments for hedging growth risk.

Active Mutual Fund Co-Holdings and the Buy-Side Peer Network

G. Feng, J. He, Y. Wang, S. Xiao

City University of Hong Kong

Rather than imposing peer groups using industries or characteristics, the paper infers them from active mutual funds' portfolio choices. For U.S. equities from 1990 to 2024, returns of nearby stocks in holdings space predict the focal stock's return in the following month. Predictability is strongest among the closest peers, disappears for dissimilar stocks, and survives controls for industry, common analysts, common ownership, and characteristics-based peers. Peer fundamentals and earnings news also forecast the focal firm's fundamentals. A value-weighted long–short strategy earns 0.67% per month with significant factor-adjusted returns. Continuation rather than reversal, stronger effects among high-friction stocks, and a supporting model point to limited investor coverage and gradual price adjustment.

Prices and Monetary Policy: The Role of Financial Constraints

M. Bauer, A. Czarnota, M. Klein

Federal Reserve Bank of San Francisco, Sveriges Riksbank

Using detailed microdata for Swedish public and private firms together with high-frequency surprises around Riksbank announcements, the authors show that firm-level financial constraints are quantitatively important for the transmission of monetary policy to prices. Smaller, more constrained firms adjust prices substantially less than larger firms after monetary-policy shocks, materially dampening the aggregate producer-price inflation response. A customer-markets mechanism with financial frictions explains the pattern: after a contraction raises external-finance premia, constrained firms preserve current cash flow by cutting prices less, at the cost of future market share. Heterogeneous responses of sales, debt, marginal costs, and markups support this channel, while several alternative explanations do not account for the main result.

The Quarter-Hour Effect: Periodic Algorithmic Trading and Return Predictability in Cryptocurrency Futures

C. Kim, P. R. Hansen

Korea Information Society Development Institute, University of North Carolina at Chapel Hill

Cryptocurrency futures exhibit recurring bursts of volume and volatility at one-minute, five-minute, and quarter-hour boundaries. Using trade data for six Binance perpetual contracts, the paper links these bursts to algorithmic participation through a sharp decline in trade-size roundness during the relevant windows. A new Autocorrelation Map resolves dependence by clock phase

and reveals serial patterns in order flow and returns at quarter-hour openings that conventional autocorrelation measures conceal. Opening returns are predictable out of sample, and order imbalance at the opening forecasts returns over the next four to twelve hours; the effects are much weaker at finer clock-time frequencies. The evidence documents distinct cross-frequency signatures of periodic algorithmic trading.

Auction Design with a Bit of Information

I. Ashlagi, S. Dobzinski, J. D. Leshno, S. Oren

Stanford University, Weizmann Institute of Science, University of Chicago, Ben-Gurion University

The paper asks which single binary signal is most valuable to a revenue-maximizing seller facing two bidders. It compares three information classes: demand signals about the overall level of values, ranking signals about which bidder is stronger, and anonymous competitiveness signals about how dispersed the two values are. For each class, the authors characterize the optimal signal and its associated auction mechanism. Competitiveness information is especially valuable: under suitable regularity conditions, the optimal competitiveness signal generates at least as much revenue as any demand or ranking signal. Moreover, among signals whose induced allocation is monotone, the optimal competitiveness signal weakly dominates every other binary signal.

Equilibrium Information Aggregation under Machine Learning

A. Ellis, M. Piccione, S. Zhang

London School of Economics, Carnegie Mellon University

The authors study the equilibrium consequences of machine learning when agents process information using the Chow–Liu tree algorithm, a tractable method that approximates a joint distribution with a tree-structured network. Embedding the algorithm in a Hellwig-style asset market with dispersed information generates a feedback loop between the price distribution used for learning and the beliefs and demands produced by the algorithm. Although algorithmic access benefits an agent holding prices fixed, market equilibrium is less informative than the rational-expectations benchmark. Prices fail to aggregate the information extracted by the algorithm even approximately. The equilibrium also typically features heterogeneous world models, demands, and utilities despite agents being identical before receiving information.

Mathematical and Computational Finance

How Fast Do Signatures Learn? Statistical Theory and Applications for Path Regression

B. Horvath, Wen Su, Wu Su, B. Wang, R. Zhang

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

Path signatures provide canonical finite-dimensional features for path-valued covariates, but their universal approximation theorem does not quantify how quickly truncation error vanishes. This paper establishes an L^2 approximation rate for smooth functionals of Itô diffusions and proves that the rate is minimax optimal. It then carries the truncation error through three learning procedures: Signature-OLS, Signature-LASSO, and Signature-Logistic, deriving consistency and related inferential or recovery guarantees. The theory makes explicit how the signature level and feature dimension may grow with the sample. Applications to foreign-exchange realized-volatility forecasting, battery end-of-life prediction, and epileptic-seizure detection show that signature features can outperform handcrafted summaries across continuous, sparse, and binary-response settings.

Mandate without Managers: Automated Market Makers as Verifiable Portfolio Products

Z. Feinstein, I. Florescu, S. O'Leary

Stevens Institute of Technology

The paper interprets an automated market maker not only as a decentralized exchange but also as a portfolio technology that enforces an observable economic mandate. For a geometric-mean market maker designed to maintain target portfolio weights, the authors introduce a multi-asset fee structure under which competitive arbitrage implements band rebalancing and bounds portfolio mis-weighting in advance. Compliance can therefore be verified directly from the pool's holdings. Simulated AMM portfolios are compared with VBIAX, EQL, and EDOW using annualized returns and tracking error relative to each portfolio mandate. Across these historical case studies, and even with arbitrage-only order flow, appropriately chosen fee ranges allow the AMM portfolios to outperform the incumbent funds on both dimensions.

Multi-Currency AMMs for Decentralized FOREX Markets: Feasibility & Optimal Design

R. K. X. Li, A. Park, A. Veneris, S. F. Singh

University of Toronto

Multi-currency automated market makers can pool liquidity across currency pairs and permit direct cross-currency trading that would otherwise route through a dominant vehicle currency. The paper studies cost-minimizing design under a constant-mean AMM, where deeper consolidated liquidity reduces price impact but joint exchange-rate risk raises liquidity providers'

impermanent loss. The authors derive closed-form trading costs, characterize optimal pool weights, and identify conditions under which an optimized multi-currency pool dominates vehicle-currency routing. They then formulate the system-level problem of partitioning currencies into pools and solve it with hierarchical agglomerative clustering. Using exchange-rate and trade data for 43 currencies from 2008 to 2023, the resulting pools reduce realized costs by about 13%, with gains that remain stable during periods of global financial stress.

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.

References

Ashlagi, Itai, Shahar Dobzinski, Jacob D. Leshno, and Sigal Oren. 2026. "Auction Design with a Bit of Information." *arXiv preprint arXiv:2608.06623*.

Bauer, Michael, Alexander Czarnota, and Mathias Klein. 2026. "Prices and Monetary Policy: The Role of Financial Constraints." Federal Reserve Bank of San Francisco Working Paper 2026-13.

Bibinger, Markus, Adrian Grüber, and Moritz Jirak. 2026. "Asset-Specific Limit Order Microstructure Noise: Parameter Estimation and Empirical Evidence." *arXiv preprint arXiv:2608.05838*.

Borri, Nicola, Yukun Liu, and Aleh Tsyvinski. 2026. "AI Premium." National Bureau of Economic Research Working Paper 35451.

Cavaliere, Giuseppe, Thomas Mikosch, Anders Rahbek, and Frederik Vilandt. 2026. "Bootstrap Inference in Autoregressive Duration Models." *arXiv preprint arXiv:2607.28294*.

Chaudhry, Aditya, and Carter Davis. 2026. "The Origins of the Factor Zoo: Investors Weakly Substitute Across Stocks." *Available at SSRN 6545318*.

Chen, Likai, and Weining Wang. 2026. "From Vector Autoregressions to AI-Based Time Series Forecasting: A Review." *arXiv preprint arXiv:2607.14279*.

Christensen, Peter Korsbakke, and Anders Norlyk. 2026. "Microstructural Foundations of Rough Noise." *arXiv preprint arXiv:2607.29442*.

Ellis, Andrew, Michele Piccione, and Shengxing Zhang. 2026. "Equilibrium Information Aggregation under Machine Learning." *arXiv preprint arXiv:2607.13670*.

Feinstein, Zachary, Ionut Florescu, and Sean O'Leary. 2026. "Mandate without Managers: Automated Market Makers as Verifiable Portfolio Products." *arXiv preprint arXiv:2608.02917*.

Feng, Guanhao, Jingyu He, Yirou Wang, and Shuhua Xiao. 2026. "Active Mutual Fund Co-Holdings and the Buy-Side Peer Network." *Available at SSRN 7000179*.

Galiani, Sebastian, and Juan Pantano. 2021. "Structural Models: Inception and Frontier." National Bureau of Economic Research Working Paper 28698.

Guasoni, Paolo, and Marko Hans Weber. 2026. "Holding Stocks, Trading Bonds." *Available at SSRN 7019618*.

Horvath, Blanka, Wen Su, Wu Su, Binnan Wang, and Ruixun Zhang. 2026. "How Fast Do Signatures Learn? Statistical Theory and Applications for Path Regression." *arXiv preprint arXiv:2607.17865*.

Kim, Chan, and Peter Reinhard Hansen. 2026. "The Quarter-Hour Effect: Periodic Algorithmic Trading and Return Predictability in Cryptocurrency Futures." *arXiv preprint arXiv:2607.09426*.

Li, Reina Ke Xin, Andreas Park, Andreas Veneris, and Srisht Fateh Singh. 2026. "Multi-Currency AMMs for Decentralized FOREX Markets: Feasibility & Optimal Design." *arXiv preprint arXiv:2607.26405*.

Zeng, Yinfeng. 2026. "Factor Dipping." *Available at SSRN 7188518*.