VOL. 5, NO. 12

OMI Research Newsletter – September 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. Chong et al. (2026) develop inference for the autocorrelation of intraday volatility changes that remains valid with spot-volatility estimation error, jumps, and market microstructure noise. Todorov et al. (2026) establish functional inference for intraday correlation periodicity and document time-varying patterns in stock–bond futures correlations. Gürkaynak, Swanson and Wright (2026) review high-frequency identification in financial event studies, vector autoregressions, and local projections, including its measurement and interpretation challenges. Keane, Neal and Vu (2026) show that choosing instruments after comparing first- and second-stage results can generate material bias towards ordinary least squares. Hassan, Mertens and Wang (2026) identify a broad tension between leading quantitative asset-pricing models and the large, persistent interest-rate differentials observed across risky and safe currencies. Chen and Welch (2026) find that published equity anomalies deliver little implementable alpha for non-micro-cap managers in post-2005 data once modest luck or trading costs are allowed for. Bandi and Su (2026) explain why standard cash-and-carry pricing fails for non-storable compute and provide early evidence of a positive compute risk premium. Wiegand and Ochoa (2026) decompose policy-path uncertainty into macroeconomic and Federal Reserve reaction-function components that imply opposite Treasury-yield responses. Leclère and Rosenbaum (2026) show how heterogeneous forecast horizons and aggregate market impact can cancel predictable price signals and restore martingale-like prices.

Econometrics, Statistics & Machine Learning

Intraday Volatility Dynamics

C. H. CHONG, M. HOFFMANN, M. ROSENBAUM, G. SZYMANSKI

HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY, UNIVERSITÉ PARIS DAUPHINE–PSL, ÉCOLE POLYTECHNIQUE, UNIVERSITY OF LUXEMBOURG

Return-based spot-volatility estimates are noisy and can be biased by jumps and market microstructure noise. The paper develops inference for the autocorrelation of intraday volatility changes that is robust to all three problems. It first shows that, over shrinking horizons, the limiting autocorrelation function of increments of a broad continuous-time process is parametrically determined and coincides with that of fractional Gaussian noise. It then derives infill asymptotics for realised autocovariances constructed from spot-variance estimators and combines the results in a feasible GMM procedure. Applied to SPY transaction data, the method finds widespread short-horizon predictability in intraday spot-volatility changes and highlights implications for standard high-frequency volatility estimators.

The Intraday Correlation Pattern

T. G. ANDERSEN, Y. TAN, Z. TAN, V. TODOROV, Z. ZHANG

NORTHWESTERN UNIVERSITY, HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY, SHANGHAI UNIVERSITY OF FINANCE AND ECONOMICS

The authors establish a functional central limit theorem for a nonparametric estimator of the periodic intraday component of instantaneous correlation between two assets, using high-frequency observations over a long span. The theory is developed in Skorokhod space because the correlation functional is not Hadamard differentiable under the conventional L^2 metric. It accommodates patterns of either sign and the benchmark case of zero integrated correlation. An estimator of the limiting covariance function supports an equivalence test across non-overlapping periods. Evidence from S&P 500 index and U.S. Treasury futures reveals pronounced intraday correlation periodicity whose shape evolves over time and helps illuminate longer-run variation in stock–bond correlation.

Optimal Transfer Learning for High-Frequency Prediction by Local Cross-Validation

D. Chen, J. Li, Y. Sun, Y. Tang

Singapore Management University, Academy of Mathematics and Systems Science, Chinese Academy of Sciences

High-frequency forecasters often face several plausible predictions based on different information sets, methods, or asset populations. The paper proposes Transfer-HFEL, a high-frequency ensemble-learning method that optimally combines forecasts from heterogeneous populations in a state-space framework using a local L-fold cross-validation criterion. The authors establish asymptotic optimality and consistency and assess the approach in simulations and an ETF application. Relative to competing methods, Transfer-HFEL improves predictive accuracy and produces an annualised return and Sharpe ratio that exceed the Fama–French three-factor benchmark by 79% and 65%, respectively. Transferring idiosyncratic variation is especially effective, suggesting that time-varying premia, misspecification, and microstructure frictions contain useful predictive information.

Fixed-Order PCA: Theory for Overestimated Factor Models

Y. Liao, X. Tong, W. Wang, D. Xiu

University of Iowa, National University of Singapore, University of Chicago, NBER

This paper develops asymptotic theory for principal component analysis in a high-dimensional factor model when the working dimension R is fixed and need only be an upper bound on the true factor count r. Using anisotropic local laws from random-matrix theory, the authors show that extra empirical eigencomponents beyond r are governed by noise, are incoherent, and are nearly orthogonal to the factor loadings. Expanded and compressed rotations then deliver consistent factor estimates. In a factor-augmented treatment-effect regression, $\sqrt{T}$-asymptotic normality holds for every fixed $R \geq r$. The results justify using a conservative upper bound when exact factor-number selection is unstable, while making clear that underestimating r can still have first-order consequences.

High-Frequency Identification

R. S. Gürkaynak, E. T. Swanson, J. H. Wright

Bilkent University, University of California Irvine, Johns Hopkins University, CEPR, NBER

Macroeconomic and financial variables respond jointly at monthly or quarterly frequencies, making causal effects hard to isolate. This review explains how sufficiently narrow event windows can instead render data releases and policy announcements plausibly predetermined. It covers high-frequency identification in financial-market event studies and its use as an external instrument in vector autoregressions and local projections for lower-frequency outcomes. The discussion addresses data construction, timing, measurement error, weak instruments, nonlinearity, predictability of surprises, and central-bank information effects. It also surveys applications beyond monetary policy, including commodity, fiscal, and political shocks. The paper provides a practical account of both the reach of high-frequency identification and the assumptions required to interpret its estimates causally.

Instrument Hacking

M. P. Keane, T. Neal, P. Vu

Johns Hopkins University, University of New South Wales, Brown University, NBER

Instrumental-variable researchers may examine several candidate instruments and report the one producing the most favourable first-stage strength or second-stage estimate. The authors show that this selection process, termed instrument hacking, biases the median IV estimate towards ordinary least squares. When instruments have equal population strength, the bias grows monotonically with the number considered. Heterogeneous strength can break that monotonicity in theory when a genuinely strong instrument is present, but calibrated simulations indicate that bias is commonly material even with only a few candidates. The paper therefore recommends avoiding post hoc instrument selection and instead using all credible instruments with weak-instrument-robust procedures suited to the assumed error structure.

Finance and Economics

A Currency Premium Puzzle

T. A. Hassan, T. M. Mertens, J. Wang

Boston University, Federal Reserve Bank of San Francisco, Renmin University of China, NBER

Quantitative asset-pricing models that reproduce the equity and risk-free-rate puzzles often fail in an open economy: they cannot generate the large and persistent interest-rate differentials between risky and safe currencies. In many such models, the mean of the log stochastic discount factor almost offsets its variance one-for-one, leaving little scope for cross-country interest-rate variation. Expected currency returns must then be generated chiefly by predictable exchange-rate movements, contrary to the data. The argument applies across a broad family of models, including incomplete-markets settings and models with exchange-rate disconnect. This currency-premium puzzle exposes a central tension linking exchange rates, interest rates, and international capital flows.

What Useful Alphas?

A. Y. Chen, I. Welch

Federal Reserve Board, University of California Los Angeles

The paper evaluates roughly 200 published long–short equity anomalies with two filters that matter for implementation: the calendar period and the investable stock universe. Across all stocks through 2005, the median zero-investment return is 48 basis points per month. It falls to 19 basis points after 2005, to 26 basis points when micro-caps are excluded, and to only 7 basis points when both restrictions apply. Modest allowances for data-mining luck or transaction costs eliminate even this residual. The results suggest that published anomalies have offered little useful alpha to non-micro-cap portfolio managers in the twenty-first century and underscore how historical samples and very small stocks drive much of the apparent anomaly premium.

(Early) AI Compute Asset Pricing

F. M. Bandi, Y. Su

Johns Hopkins University

Compute is a scarce, capital-intensive input whose expenditure and service flow already exceed 1% of U.S. GDP, while announced futures contracts are set to make its price risk tradable. The paper provides an early framework for this emerging asset class. Standard cash-and-carry relations fail because compute cannot be stored; synthetic futures constructed from existing term rentals are likely upper bounds on eventual futures prices; and, after financialisation, futures should equal expected expiry spot prices net of a risk premium. Using term-rental data, the authors build a first panel of synthetic compute-futures returns by GPU generation and maturity. Preliminary evidence points to a positive compute premium consistent with hedging pressure from service providers.

Asset Pricing When T Equals One

P. Zaffaroni

Imperial College London

The paper develops inference for the pricing ability of asset-pricing models, and for associated mean–variance portfolios, from a single large cross-section of returns: $T = 1$ and large N. The method imposes no structure on the time variation of risk premia or exposures and rests on new principal-component asymptotics for one time observation. It therefore tracks pricing and portfolio performance period by period rather than averaging across time, and identifies analytically what drives out-of-sample performance. The framework handles observed factors, latent factors, or both. Applications revisit the factor zoo, compare common cross-sectional models out of sample, and study the relation between security-market-line flatness and macroeconomic announcements.

Illiquidity at Risk

D. Lacava, P. Santucci de Magistris

University of Messina, LUISS University

The authors introduce Illiquidity-at-Risk (IlliQaR), a tail-risk measure for extreme market-liquidity dry-ups. It is built from realised Amihud illiquidity, measured with high-frequency data as realised volatility relative to trading volume. Linear and nonlinear forecasting models reveal that discontinuous jumps are essential for accurate tail coverage and stress-period forecasts; models with purely continuous dynamics systematically understate severe liquidity evaporation. In data for the S&P 500 and 25 large U.S. equities, incorporating jumps materially improves short-horizon predictions, while the benefits diminish as illiquidity mean-reverts. Violations at individual stocks cluster during index-level liquidity stress, indicating that extreme illiquidity is systemic and that the broad index can lead dry-ups in its constituents.

Decomposing Price Chart Signals in Return Prediction

J. T. Bøjstrup, B. Veliyev, J. N. Wulff

Aarhus University

Image models trained on stock charts predict returns out of sample even though a chart contains no information beyond the open, high, low, close, and volume data used to draw it. The paper decomposes their advantage over tabular models into three sources: function-class capacity, chart-specific encoding, and the way each model smooths noisy return targets. Model capacity explains little of the gap. Much of the predictive gain instead reflects the inductive structure created by the visual representation and its local smoothing of recent price and volume

patterns. Translating the learned visual structures back into numerical predictors closes much of the performance difference, offering a more interpretable account of why representation changes what a model can learn from the same raw data.

Monetary Policy Uncertainty: Sources and Consequences

C. Wiegand, M. Ochoa

University of Pennsylvania, Federal Reserve Board

Why do stocks and bonds sometimes hedge one another and sometimes move together? Using daily revisions in expected policy rates, the authors decompose a new market-based measure of policy-path uncertainty into macroeconomic uncertainty and uncertainty about the Federal Reserve's reaction function. The components have sharply different consequences. Macroeconomic uncertainty lowers output and inflation, induces policy easing, and reduces Treasury yields. Reaction-function uncertainty creates stagflationary pressure, induces tightening, and raises yields through higher real term premia. Both lower equity valuations. The changing mix therefore provides a structural explanation for sign switches in the stock–bond correlation: it is negative when macro uncertainty dominates and positive when policy-response uncertainty is more important.

Ex Machina: Financial Stability in the Age of Artificial Intelligence

K. Anand, S. Kazinnik, A. Leonello, E. Panetti

Deutsche Bundesbank, Stanford University, European Central Bank, University of Naples Federico II

The paper uses simulations to study artificial-intelligence investors in a mutual-fund redemption game with strategic complementarities and multiple equilibria. Different architectures generate systematically different outcomes. Q-learning agents coordinate effectively, but default risk leads them to redeem excessively and too early, amplifying fund fragility. Large-language-model investors internalise the equilibrium structure more closely, yet heterogeneity in their beliefs weakens coordination and makes aggregate redemptions harder to predict. These contrasts arise even when the economic environment is held fixed. The findings imply that the design of automated decision systems—not simply the extent of their adoption—is a first-order determinant of financial stability and should matter for stress testing and regulation.

Completions, Not Convictions: Can LLM Probabilities Proxy for Human Beliefs?

S. Nagel, C. Tseng, D. Xiu

University of Chicago

Large language models could in principle elicit beliefs at scale, but the paper finds that their probabilities differ systematically from human consensus in prediction markets and analyst earnings forecasts. As prompt continuations, model outputs react to tone, anchors, label order, exposure in the training corpus, and the direction of the outcome. More information, demographic aggregation, reasoning instructions, and larger models shift these gaps without reliably removing them. Replacing analyst beliefs with LLM probabilities also changes the relation with announcement returns and reverses the usual forecast-error-on-revision conclusion: analysts appear to underreact, while the LLM appears to overreact. Machine-elicited probabilities therefore require validation for the specific model, prompt, setting, and downstream use.

Exploration and Stopping

Y. Sannikov, W. Zhong

Stanford Graduate School of Business

The authors study a decision-maker who dynamically chooses what to learn before deciding when to act. They show that this control problem is equivalent to selecting the joint distribution of the stopped state and stopping time subject to a sequence of information-budget constraints, and they characterise exactly which distributions are attainable. The resulting convex programme has a strong dual that prices information over time; the optimal policy concavifies the stopping payoff net of these shadow prices. Time-preference curvature governs behaviour: convex preferences generate Poisson exploration, concave preferences confine stopping to a window and can force an initial phase of pure exploration, while the linear case lies at the boundary. Applications include real options, speed–accuracy trade-offs, and exploration contests.

Mathematical and Computational Finance

Rough Volatility Across Assets

S. Mouti

University of New Haven

Using a common pipeline, the paper measures volatility roughness for 3,926 U.S. equities, 34 CME futures roots spanning rates, currencies, and commodities, and options on 44 underlyings from 2010 to 2025. Realised volatility is rough in every asset class, with median Hurst estimates ranging from 0.05 for livestock to 0.20 for equity indices. Option-implied roughness is informative only when leverage produces a clean skew term structure: it works reasonably for equity indices but fails for rates and foreign exchange despite their rough realised volatility.

The author also quantifies mean-reversion contamination and additive measurement error, neither of which restores Brownian-scale estimates, and proposes a taxonomy of settings where rough-volatility methods succeed or fail.

Mean-Field Equilibrium of Heterogeneous Agents under Market Impact

J. Leclère, M. Rosenbaum
Université Paris Dauphine–PSL

Market participants share information but form forecasts over heterogeneous horizons, while their aggregate positions feed back into prices through market impact. The paper introduces a linear mean-field model in which the observed price combines a martingale, a predictable Volterra signal, and impact from aggregate positions. In a Gaussian–Volterra setting, equilibrium is characterised by a linear fixed-point equation with explicit existence and uniqueness conditions. A balance condition cancels direct transmission of the common signal. As agents increasingly account for aggregate impact, suitably scaled equilibria make the predictable signal and impact offset one another, so the observed price converges to its martingale component. For fractional signals and Gamma-distributed horizons, the authors derive local regularity bounds compatible with Brownian price paths.

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

Anand, Kartik, Sophia Kazinnik, Agnese Leonello, and Ettore Panetti. 2026. "Ex Machina: Financial Stability in the Age of Artificial Intelligence." *Available at SSRN 6728539*.

Bandi, Federico M., and Yinan Su. 2026. "(Early) AI Compute Asset Pricing." *arXiv preprint arXiv:2607.12156*.

Bøjstrup, Jonas Theill, Bezirgen Veliyev, and Jesper N. Wulff. 2026. "Decomposing Price Chart Signals in Return Prediction." *Available at SSRN 7119420*.

Chen, Andrew Y., and Ivo Welch. 2026. "What Useful Alphas?" *arXiv preprint arXiv:2607.06502*.

Chen, Dachuan, Jia Li, Yuying Sun, and Yue Tang. 2026. "Optimal Transfer Learning for High-Frequency Prediction by Local Cross-Validation." *Available at SSRN 7189238*.

Chong, Carsten H., Marc Hoffmann, Mathieu Rosenbaum, and Grégoire Szymanski. 2026. "Intraday Volatility Dynamics." *Available at SSRN 7344018*.

Gürkaynak, Refet S., Eric T. Swanson, and Jonathan H. Wright. 2026. "High-Frequency Identification." *Working Paper*.

Hassan, Tarek Alexander, Thomas M. Mertens, and Jingye Wang. 2026. "A Currency Premium Puzzle." National Bureau of Economic Research Working Paper 35572.

Keane, Michael P., Timothy Neal, and Patrick Vu. 2026. "Instrument Hacking." National Bureau of Economic Research Working Paper 35633.

Lacava, Demetrio, and Paolo Santucci de Magistris. 2026. "Illiquidity at Risk." *arXiv preprint arXiv:2609.00943*.

Leclère, Joseph, and Mathieu Rosenbaum. 2026. "Mean-Field Equilibrium of Heterogeneous Agents under Market Impact." *arXiv preprint arXiv:2609.03115*.

Liao, Yuan, Xin Tong, Wanjie Wang, and Dacheng Xiu. 2026. "Fixed-Order PCA: Theory for Overestimated Factor Models." *arXiv preprint arXiv:2605.18448*.

Mouti, Saad. 2026. "Rough Volatility Across Assets." *arXiv preprint arXiv:2608.16749*.

Nagel, Stefan, Claire Tseng, and Dacheng Xiu. 2026. "Completions, Not Convictions: Can LLM Probabilities Proxy for Human Beliefs?" *Available at SSRN 7265538*.

Sannikov, Yuliy, and Weijie Zhong. 2026. "Exploration and Stopping." *arXiv preprint arXiv:2608.10274*.

Todorov, Viktor, Torben G. Andersen, Zhiyuan Zhang, Yingwen Tan, and Zhensi Tan. 2026. "The Intraday Correlation Pattern." *Available at SSRN 7295718*.

Wiegand, Courtney, and Marcelo Ochoa. 2026. "Monetary Policy Uncertainty: Sources and Consequences." *Available at SSRN 7066998*.

Zaffaroni, Paolo. 2026. "Asset Pricing When T Equals One." *Available at SSRN 6522918*.