VOL. 6, NO. 2

OMI Research Newsletter – November 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 include the following. Bandi, Korsaye and Su (2026) show that market timing alone can beat a zero forecast for individual stock returns. Against a historical panel mean, only flexible nonlinear models deliver consistent gains, concentrated among small, value and illiquid stocks. Chernov et al. (2026) recover the contribution of small covariance eigenvalues to the efficient frontier and estimate a maximum annualised Sharpe ratio for US equities of 7.4, falling to 5.0 when microcaps are excluded. Kötehne and Patton (2026) find that midday returns receive the greatest weight when forecasting correlations, whereas the weights for volatility forecasts rise through the trading day. Barbon and Drissi (2026) show that as traders delegate analysis to AI, prices can remain informative even as human expertise declines. An unexpected outage reveals the resulting loss of resilience. Caner et al. (2026) combine AI screening of company fundamentals and news with quantitative portfolio weights, improving on the best unscreened Sharpe ratio in their 2021–24 backtest. Riva, Targino and da Costa (2026) trace the cost of pursuing higher expected returns under a risk budget and find lower turnover and concentration than in conventional mean–variance portfolios. Bollerslev, Li and Luan (2026) use federal funds futures to show that FOMC announcements resolve uncertainty across the year ahead, while payroll releases chiefly clarify near-term policy decisions. Wiese, Buehler and Horvath (2026) model the joint evolution of option prices across strikes and maturities under no arbitrage. For short-dated S&P index options, over 95% of the initial model prices lie inside observed bid–ask spreads. Celary et al. (2026) characterise the kernels that keep fitted discount curves within the same finite-dimensional space as they evolve under the risk-neutral drift.

Econometrics, Statistics & Machine Learning

Noisy Financial Machine Learning

F. M. BANDI, S. A. KORSAYE, Y. SU
JOHNS HOPKINS UNIVERSITY

A model can beat a zero forecast for individual stock returns by timing movements in the overall market. How much does that tell us about its ability to rank stocks? The authors propose the historical average return across the panel as a benchmark and use a finite-sample analysis of ridge regression to explain why it is difficult to improve on when signals are weak. In their empirical comparison, only flexible nonlinear models generate consistent gains over the benchmark, concentrated among small, value and illiquid stocks.

The Hidden Sharpe Ratio

M. CHERNOV, B. T. KELLY, S. MALAMUD, J. SCHWAB
UCLA ANDERSON, YALE UNIVERSITY, AQR CAPITAL MANAGEMENT, EPFL, SWISS FINANCE INSTITUTE, CEPR, NBER

Regularising a large covariance matrix gives little weight to directions with small eigenvalues, even when those directions contribute to the efficient frontier. The authors estimate the maximum Sharpe ratio directly, progressively recovering this contribution and constructing bounds and confidence intervals without having to estimate the tangency portfolio. For US equities, they estimate a maximum annualised Sharpe ratio of 7.4, falling to 5.0 when microcaps are excluded. These estimates describe the population frontier and quantify how far portfolios based on regularised weights can fall short.

Bespoke Realized Correlations for Multivariate Volatility Prediction

B. Kötehne, A. J. Patton
Duke University, Singapore Management University

Which part of the trading day is most informative about future correlations? Using intraday returns for around 360 S&P 500 constituents from 2000 to 2023, the authors estimate weights tailored to the forecasting model. For correlations, the time-of-day weights rise from the open until around midday and then decline. For volatility, the weights rise through the close. Letting correlation weights depend on time of day, realised variance or trading volume significantly improves out-of-sample forecasts relative to weighting intraday returns equally.

An Itô-Type Formula: High Frequency Estimation under Rough Volatility

H. Guo, D. Chen, C. Li, L. Zhang, P. A. Mykland
Peking University, Singapore Management University, University of Illinois Chicago, University of Chicago

The authors develop an Itô-type formula for rough volatility and use it to correct the bias in high-frequency estimates of leverage, the dependence between price and volatility shocks. The way estimated leverage changes with the time lag identifies the Hurst exponent and yields a rate-optimal estimator with feasible standard errors. In S&P 500 intraday transaction data from 2018 to 2025, the estimated Hurst exponent lies below the semimartingale benchmark of one half for much of the sample. The test rejects that benchmark for most of the period from mid-2020 onwards.

Power Enhancement via Cross-Fit Variance Estimation: Applications to Specification, Overidentification, and Many-Restriction Testing

K. Sunada, Y. Matsushita, T. Otsu
Aarhus University, Hitotsubashi University, London School of Economics and Political Science

A quadratic-form test can lose power when the departure it is designed to detect also inflates its estimated variance. The authors construct a cross-fit estimator that filters this systematic component out of the residuals used to scale the statistic. For large departures concentrated in a few observations, the conventional test's rejection probability can approach zero while the cross-fit test's approaches one. The framework covers specification tests, overidentifying restrictions and joint tests of regression coefficients. An application uses data from the Oregon Health Insurance Experiment.

Finance and Economics

Traders on GPT: AI and Cognitive Capital in Financial Markets

A. Barbon, F. Drissi
University of St. Gallen, Oxford-Man Institute, University of Oxford

What can market prices reveal about traders' ability to interpret information? In a rational expectations model, traders divide information processing between their own analysis and AI. The authors assume that delegating analysis erodes expertise. When AI becomes sufficiently cheap, traders allow their skills to decline even though they anticipate the cost. Prices remain equally informative as expertise declines while AI is available. An unexpected outage exposes the loss of skill, leaving prices less informative than in an economy that never adopted AI.

Designing Agentic AI-Based Screening for Portfolio Investment

M. Caner, A. Capponi, N. Sun, J. Y. Tan
North Carolina State University, Columbia University

Two language-model agents screen S&P 500 stocks, one using company fundamentals and the other assessing news sentiment. The agents combine their recommendations, and an estimate of the inverse covariance matrix determines portfolio weights for the selected stocks. The authors establish conditions under which errors in screening still permit consistent estimation of the portfolio's squared Sharpe ratio. In their October 2021–April 2024 backtest with transaction costs of 10 basis points, the best two-agent specification achieves an annualised Sharpe ratio of 1.09, compared with 1.01 for the best unscreened quantitative strategy.

Risk-Budgeted Mean-Variance Portfolios

R. Riva, R. Targino, B. F. P. da Costa
Getulio Vargas Foundation, Universidade Federal do Rio de Janeiro

How much must a portfolio depart from its target risk contributions to attain a higher expected return? The authors add a return floor and a volatility ceiling to the risk-budgeting problem. The resulting convex programme has a unique solution whenever the constraints are feasible. Varying the return target traces the cost in deviations from the original risk budget. In out-of-sample tests on US stocks, industry portfolios and multi-asset ETFs, the portfolios earn higher average returns than risk parity, with lower turnover and concentration than conventional mean–variance portfolios.

Real-Time News and the Resolution of Monetary Policy Uncertainty

T. Bollerslev, S. Z. Li, Z. Luan
Duke University, Rutgers University, NBER

Using one-minute federal funds futures, the authors measure uncertainty about the policy rate separately for each of the next eight scheduled FOMC meetings. FOMC announcements reduce uncertainty along the full path, while payroll releases mainly clarify near-term decisions. Inflation releases are most informative when markets expect rates to fall. Federal Reserve Chair speeches generally reduce uncertainty, although testimony signalling a change in the expected path can increase it. Greater news-driven resolution of uncertainty also predicts higher stock market returns over the following month.

Time-Varying U.S. Equity Market Integration

F. Jedmo, A. J. Patton
Copenhagen Business School, Duke University, Singapore Management University

Do equity portfolios earn the same premium for exposure to a common factor? Using daily returns on 234 US portfolios from 1963 to 2024, the authors group portfolios by their estimated prices of factor risk. The distance between groups measures market segmentation. Recessions, heavy trading and illiquidity coincide with greater segmentation, which tends to persist when funding liquidity is low. Allowing for group-specific risk prices improves portfolio Sharpe ratios most when segmentation is high. The gains are small and statistically insignificant at lower levels of segmentation.

Limits to Arbitrage and Prediction Market Efficiency: The Case of PredictIt

E. Zitzewitz
Dartmouth College, NBER

Between 2016 and 2022, PredictIt attracted active trading on political events while limiting each participant's position in a contract to $850. Zitzewitz finds that bids across mutually exclusive outcomes frequently summed to more than $1.10 against a total payout of $1. The resulting arbitrage opportunities survived fees and could persist despite active trading. Longshots were also more overpriced than in other prediction markets studied. The evidence suggests that position limits can impair the accuracy of market-implied event probabilities by restricting arbitrage.

Delegate Pricing to Algorithms: When Slow and Steady Wins the Race

G. Ashkenazi-Golan, E. Plumb, C. Possnig, Y. Zhang
London School of Economics and Political Science, University of Waterloo, Imperial College London

Two firms choose how quickly their pricing algorithms learn, taking account of the profits they earn while prices adjust. When their actions are strategic substitutes, each firm prefers the fastest feasible adjustment. When prices are strategic complements, cutting prices quickly can bring forward a rival's response and erode interim profits. For a range of payoffs and initial conditions, both firms choose slower algorithms. Prices then remain above the competitive level for longer, even though the algorithms ultimately converge to the competitive equilibrium.

Signaling via Money

O. Bos, M. Pollrich
Université Paris-Saclay

A buyer's payment at auction can reveal private information to potential counterparties in a subsequent resale market. The authors ask what a mechanism designer loses when payments become public. They construct transfers that usually communicate the intended signal to outside observers, with occasional adjustments that preserve expected payments. The construction matches allocation probabilities and interim expected transfers exactly. Even when participants care about the beliefs of outside observers, the designer can approximate the original payoffs and incentives for truthful reporting arbitrarily closely.

How Learning Selects Equilibria and When It Leads to Collusion

H. Wang, M. Wu
University of California, Berkeley

In a repeated game, players test beliefs about their opponents and explore alternatives more readily when their payoffs are low. The authors show how this learning rule selects among equilibria by favouring those with the highest minimum transformed payoff across players. In repeated Cournot competition, sufficiently patient firms with enough memory and sufficiently accurate public monitoring earn more than the Cournot–Nash benchmark in every stochastically stable outcome. When firms are symmetric, the selected profits can approach the monopoly benchmark.

Mathematical and Computational Finance

GROMM: Generative Risk-Neutral Option Market Models

M. Wiese, H. Buehler, B. Horvath
Mathematical Institute and Oxford-Man Institute, University of Oxford

The authors model how an entire surface of European option prices evolves across strikes and maturities. Their model interpolates smoothly between contracts, excludes both static and dynamic arbitrage, and admits a unique strong solution. They estimate the model using S&P index options from January 2022 to August 2025. For the short-dated contracts considered, with expiries up to 25 business days, more than 95% of the initial model prices lie inside the observed bid–ask spreads.

Drift-Invariant Kernels for Discount Curve Models

A. Celary, Z. Eksi-Altay, D. Filipović, P. Krühner
Princeton University, Vienna University of Economics and Business, EPFL

Kernel ridge regression fits a discount curve in a finite-dimensional space determined by the kernel and the instruments' cash-flow dates. For the fitted curves to serve as states of a continuous-time interest-rate model, the risk-neutral drift must preserve this space. For kernels whose reproducing kernel Hilbert space has dimension greater than two, the authors characterise exactly when this condition holds for every grid of cash-flow dates. For the resulting family, fitted curves are affine combinations of discount curves with constant continuously compounded yields.

Deep Learning for Operator Maps in Optimal Liquidation

R. Donnelly, L. Galimberti, B. Healy, S. Zhang
King's College London, Decision Science Limited

A trader liquidating a position must specify how quickly its price impact decays. The authors show that the optimal trading schedule varies continuously with the impact kernel and the fundamental price process. They use this property to train a recurrent neural network to compute schedules across different impact models and trading signals. In their simulations, an exponential kernel that fits the true decay curve more closely can lead to a larger loss from following the resulting schedule.

Calibration of Passive Market Impact with a Queue-Sensitive Propagator Model

Y. Hafsi, F. Lillo, G. Szymanski
École Polytechnique, Deepfin Research, Scuola Normale Superiore, University of Luxembourg

How does an order that adds liquidity at the best quote affect subsequent prices? The authors express price changes in terms of signed order flow and a correction that depends on the depth available at the best bid and ask. In E-mini S&P 500 futures, trades move prices further when they consume a thin queue. Adding queue depth improves out-of-sample fit and the model's estimates of the subsequent price response to passive limit orders.

References

Ashkenazi-Golan, Galit, Edward Plumb, Clemens Possnig, and Yufei Zhang. 2026. "Delegate Pricing to Algorithms: When Slow and Steady Wins the Race." *arXiv preprint arXiv:2610.11675*.

Bandi, Federico M., Sofonias Alemu Korsaye, and Yinan Su. 2026. "Noisy Financial Machine Learning." *Available at SSRN 7573698*.

Barbon, Andrea, and Fayçal Drissi. 2026. "Traders on GPT: AI and Cognitive Capital in Financial Markets." *Available at SSRN 7661969*.

Bollerslev, Tim, Sophia Zhengzi Li, and Zeyao Luan. 2026. "Real-Time News and the Resolution of Monetary Policy Uncertainty." *Available at SSRN 7539278*.

Bos, Olivier, and Martin Pollrich. 2026. "Signaling via Money." *arXiv preprint arXiv:2610.03146*.

Caner, Mehmet, Agostino Capponi, Nathan Sun, and Jonathan Y. Tan. 2026. "Designing Agentic AI-Based Screening for Portfolio Investment." *arXiv preprint arXiv:2603.23300*.

Celary, Andreas, Zehra Eksi-Altay, Damir Filipović, and Paul Krühner. 2026. "Drift-invariant kernels for discount curve models." *arXiv preprint arXiv:2611.09751*.

Chernov, Mikhail, Bryan T. Kelly, Semyon Malamud, and Johannes Schwab. 2026. "The Hidden Sharpe Ratio." National Bureau of Economic Research Working Paper 35988.

Donnelly, Ryan, Luca Galimberti, Brian Healy, and Siqi Zhang. 2026. "Deep Learning for Operator Maps in Optimal Liquidation." *arXiv preprint arXiv:2611.14245*.

Guo, Haijian, Dachuan Chen, Chenxu Li, Lan Zhang, and Per A. Mykland. 2026. "An Itô-type Formula: High Frequency Estimation under Rough Volatility." *Available at SSRN 7615661*.

Hafsi, Yadh, Fabrizio Lillo, and Grégoire Szymanski. 2026. "Calibration of Passive Market Impact with a Queue-Sensitive Propagator Model." *arXiv preprint arXiv:2610.38442*.

Jedmo, Filip, and Andrew J. Patton. 2026. "Time-Varying U.S. Equity Market Integration." *Available at SSRN 7473221*.

Kötehne, Berna, and Andrew J. Patton. 2026. "Bespoke Realized Correlations for Multivariate Volatility Prediction." *Available at SSRN 7572359*.

Riva, Raul, Rodrigo Targino, and Bernardo Freitas Paulo da Costa. 2026. "Risk-Budgeted Mean-Variance Portfolios." *Available at SSRN 7672259*.

Sunada, Keita, Yukitoshi Matsushita, and Taisuke Otsu. 2026. "Power enhancement via cross-fit variance estimation: Applications to specification, overidentification, and many-restriction testing." *arXiv preprint arXiv:2610.06119*.

Wang, Hanzhang, and Manxi Wu. 2026. "How Learning Selects Equilibria and When It Leads to Collusion." *arXiv preprint arXiv:2610.36306*.

Wiese, Magnus, Hans Buehler, and Blanka Horvath. 2026. "GROMM: Generative Risk-neutral Option Market Models." *arXiv preprint arXiv:2610.16732*.

Zitzewitz, Eric. 2026. "Limits to Arbitrage and Prediction Market Efficiency: the Case of PredictIt." National Bureau of Economic Research Working Paper 35845.