VOL. 5, NO. 6

OMI Research Newsletter – March 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. Please find the schedule and register here.

The OMI-SBS Conference on FinTech and Blockchain Economics will be held at Saïd Business School and OMI on 25–26 June 2026. Please find the schedule here, and contact faycal.drissi@omi.ox.ac.uk for further information about registration.

Highlights

Key contributions in this month's edition: Chong and Todorov (2026) develop a nonparametric high-frequency measure of variance disagreement to test whether equity and options markets are fully integrated, and document recurring but moderate episodes of segmentation in major stocks and index ETFs. Chang et al. (2026) show that a tiny-ridge "Ridgelet" minimum-variance portfolio can generalise well out of sample in very high dimensions, relative to the standard ridgeless approach. Cohen, Lu and Nguyen (2026) use AI and machine learning to distinguish the predictable from the novel components of mutual fund managers' trading, showing that about 71% of trade directions are predictable and that less predictable managers and positions perform better. Hyland (2026) develops a New Keynesian macro-finance framework in which tail-risk shocks act like negative demand shocks and push down the neutral real rate, implying that monetary policy should respond to the full distribution of risks, not just volatility. Cartea, Cucuringu and Fang (2026) show that jump variation improves realised volatility forecasts when modelled as an idiosyncratic network spillover across stocks using graph-based methods. Hizmeri and Bevilacqua (2026) find that morning VVIX, measured at 10:00 EST during the U.S.–European market overlap, strongly predicts next-day variance-asset returns, consistent with delayed incorporation of volatility uncertainty. Bergault, Hafsi and Sánchez-Betancourt (2026) develop a mixed-control model for trading across centralised (CEX) and decentralised (DEX) exchanges with stochastic execution delays and priority fees, and show that optimal priority-fee choice can significantly improve trading performance. Aqsha, Bank and Sánchez-Betancourt (2026) propose a signature-based numerical method for multidimensional linear-quadratic stochastic control that reduces the problem to a finite-dimensional convex quadratic optimisation. Baggiani, Herdegen and Sánchez-Betancourt (2026) characterise an approximate Nash equilibrium in a game between DEXs competing for order flow through dynamic trading fees.

Statistics, Econometrics & Machine Learning

Realized Variance Disagreement

C. H. CHONG, V. TODOROV
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY, NORTHWESTERN UNIVERSITY

We propose realized variance disagreement as a nonparametric measure of segmentation between equity and options markets over a fixed time interval. Defined as the aggregate difference between diffusive variance components of high-frequency asset returns and their risk-neutral conditional short-term expectations recovered from options expiring at the end of the trading day, this statistic captures realized violations of the fundamental no-arbitrage condition that diffusive variance coincides under the physical and risk-neutral probability measures.

We derive a central limit theorem that enables feasible inference on integrated variance disagreement between equity and options markets. We further develop market integration tests that can detect spot and integrated variance disagreements. Empirically, we find evidence for episodes of disagreement between equity and options markets. These periods exhibit mild persistence and differ in magnitude across assets.

Zero Variance Portfolio

J. Chang, Y. Ding, Z. Shi, B. Zhang

Southwestern University of Finance and Economics, Chinese Academy of Sciences, University of Macau, The Chinese University of Hong Kong, University of Science and Technology of China

When the number of assets is larger than the sample size, the minimum variance portfolio interpolates the training data, delivering pathological zero in-sample variance. We show that if the weights of the zero variance portfolio are learned by a novel "Ridgelet" estimator, in a new test data this portfolio enjoys out-of-sample generalizability. It exhibits the double descent phenomenon and can achieve optimal risk in the overparametrized regime when the number of assets dominates the sample size. In contrast, a "Ridgeless" estimator which invokes the pseudoinverse fails in-sample interpolation and diverges away from out-of-sample optimality. Extensive simulations and empirical studies demonstrate that the Ridgelet method performs competitively in high-dimensional portfolio optimization.

Nonparametric Inference for Spot Volatility in Pure-Jump Semimartingales

C. Yan, D. Chen, J. Li

Singapore Management University

We provide a comprehensive analysis of spot volatility inference in pure-jump semimartingales under two asymptotic settings: fixed-k, where each local window uses a fixed number of observations, and large-k, where this number grows with sampling frequency. For both active- and possibly inactive-jump settings, we derive generally nonstandard, typically non-Gaussian limit distributions and establish valid inference, including when the jump-activity index is consistently estimated. Simulations show that fixed-k asymptotics offer markedly better finite-sample accuracy, underscoring their practical advantage for nonparametric spot volatility inference.

Online Learning in Semiparametric Econometric Models

X. Chen, E. Tamer, Q. Yao

Yale University, Harvard University, Louisiana State University

Data in modern economic and financial applications often arrive as a stream, requiring models and inference to be updated in real time – yet most semiparametric methods remain batch-based and computationally impractical in large-scale streaming settings. We develop an online learning framework for semiparametric monotone index models with an unknown monotone link function. Our approach uses a two-phase learning paradigm. In a warm-start phase, we introduce a new online algorithm for the finite-dimensional parameter that is globally stable, yielding consistent estimation from arbitrary initialization. In a subsequent rate-optimal phase, we update the finite-dimensional parameter using an orthogonalized score while learning the unknown link via an online sieve method; this phase achieves optimal convergence rates for both components. The procedure processes only the most recent data batch, making it suitable when data cannot be stored (e.g., memory, privacy, or security constraints), and its resulting parameter trajectories enable online inference such as confidence regions–on parameters including policy-effect analysis with negligible additional computation. Monte Carlo experiments on both simulated and real data show adequate performance especially relative to full sample methods.

Finance and Economics

Mimicking Finance

L. Cohen, Y. Lu, Q. H. Nguyen

Harvard University, University of Pennsylvania, DePaul University

We use frontier advancements in Artificial Intelligence and machine learning to extract and classify the part of key economic agents' behaviors that are predictable from past behaviors. Even the agents themselves might view these as novel (innovative) decisions; however, we show in strong contrast that a large percentage of these actions and behaviors can be predicted—and thus mimicked—in the absence of these individuals. In particular, we show that 71% of mutual fund managers' trade directions can be predicted in the absence of the agent making a single trade. For some managers, this increases to nearly all of their trades in a given quarter. Further, we find that manager behavior is more predictable and replicable for managers who have a longer history of trading and are in less competitive categories. The larger the ownership stake of the manager in the fund, the less predictable their behavior. Lastly, we show strong performance implications: less predictable managers strongly outperform their peers, while the most predictable managers significantly underperform. Even within each manager's portfolio, those stock positions that are more difficult to predict strongly outperform those that are easier to predict. Aggregating across the universe of fund managers each quarter, stocks

whose position changes are least predictable additionally significantly outperform stocks whose position changes are most predictable across the universe. Our framework allows researchers to delineate and classify the portion of financial agents' action sets which are predictable from those which are novel responses to stimuli – open to being evaluated for value creation or destruction.

Taming Black Swans: Monetary Policy in the Presence of Tail Risks

C. C. Hyland
University of Oxford

Central banks have long engaged in risk management, adjusting policy preemptively when tail risks rise. This paper sets out a formal framework for doing so. Using options data, I show empirically that tail-risk shocks act as negative aggregate demand shocks, with output and inflation falling together. I then develop a New Keynesian macro-finance framework in which all higher-order moments of the shock distribution enter asset prices, aggregate demand, and the neutral real rate. Tail risks are quantitatively important for the neutral rate: a one percentage-point increase in disaster probability lowers it by approximately 16 basis points. For policy implications, I find large welfare gains from central banks actively adjusting policy in response to the full shape of the shock distribution, not just its mean and variance. I then construct a counterfactual for the US during the Global Financial Crisis, finding that ignoring tail risk would have more than doubled the output gap decline in 2007Q3–Q4. Finally, I provide a new measure of the neutral rate of interest adjusted for risk. This risk-adjusted neutral rate documents that policy has been overly restrictive during periods of elevated tail risk, with sharp crisis-driven contractions absent from standard estimates.

Network-Augmented Volatility Forecasting with Idiosyncratic Jump Spillovers

Á. Cartea, M. Cucuringu, L. Fang
Oxford-Man Institute of Quantitative Finance, University of Oxford, University of California, Los Angeles

Empirical studies often conclude that jump variation contributes little to realized volatility (RV) forecasting in univariate settings. We revisit this question from a crosssectional perspective, asking whether jump shocks propagate across stocks through network channels and improve RV forecasts. Using high-frequency data for SP 100 stocks over 2007-2021, we construct idiosyncratic returns via factor regressions and estimate sparse dependency networks from these residuals. Combining these networks with HAR-type predictors in linear (GHAR-JV) and graph neural network (GNNHAR-JV) models, we find that idiosyncratic-network-based specifications reduce one-dayahead MSE by around 30% relative to a standard HAR benchmark, while jump models on raw or market-based networks deliver little benefit. We conclude that jump variation becomes economically important once viewed as an idiosyncratic network spillover mechanism, and that structured network learning enhances volatility forecasting performance, while preserving an economically interpretable cross-sectional structure.

Early Birds Get the Vol: Morning Volatility Uncertainty and Variance Risk Premium

R. Hizmeri, M. Bevilacqua
University of Liverpool

We document that morning volatility-of-volatility (VVIX), measured at 10:00 EST during the U.S.-European market overlap, strongly predicts next-day variance asset returns. Predictive power peaks with t-statistics reaching 5.6 and adjusted R^2 of 2.6%, but diminishes after 11:00 EST. Trading strategies exploiting this pattern generate Sharpe ratios exceeding 2.0, remaining economically meaningful even after transaction costs. Crucially, intraday returns display a flat relationship with morning VVIX on high volatility uncertainty days, while close-to-close returns exhibit strong predictability. These findings are difficult to reconcile with standard risk premium theories but are consistent with limited attention and slow-moving beliefs about volatility.

A Test of Lookahead Bias in LLM Forecasts

Z. Gao, W. Jiang, Y. Yan
The Chinese University of Hong Kong

We develop a statistical test to detect lookahead bias in economic forecasts generated by large language models (LLMs). Using state-of-the-art pre-training data detection techniques, we estimate the likelihood that a given prompt appeared in an LLM's training corpus, a statistic we term Lookahead Propensity (LAP). We formally show that a positive correlation between LAP and forecast accuracy indicates the presence and magnitude of lookahead bias, and apply the test to two forecasting tasks: news headlines predicting stock returns and earnings call transcripts predicting capital expenditures. Our test provides a cost-efficient, diagnostic tool for assessing the validity and reliability of LLM-generated forecasts.

Mathematical and Computational Finance

Trading in CEXs and DEXs with Priority Fees and Stochastic Delays

P. Bergault, Y. Hafsi, L. Sánchez-Betancourt

Université Paris Dauphine-PSL, École Polytechnique, Oxford-Man Institute of Quantitative Finance, University of Oxford

We develop a mixed control framework that combines absolutely continuous controls with impulse interventions subject to stochastic execution delays. The model extends current impulse control formulations by allowing (i) the controller to choose the mean of the stochastic delay of their impulses, and allowing (ii) for multiple pending orders, so that several impulses can be submitted and executed asynchronously at random times. The framework is motivated by an optimal trading problem between centralized (CEX) and decentralized (DEX) exchanges. In DEXs, traders control the distribution of the execution delay through the priority fee paid, introducing a fundamental trade-off between delays, uncertainty, and costs. We study the optimal trading problem of a trader exploiting trading signals in CEXs and DEXs. From a mathematical perspective, we derive the associated dynamic programming principle of this new class of impulse control problems, and establish the viscosity properties of the corresponding quasi-variational inequalities. From a financial perspective, our model provides insights on how to carry out execution across CEXs and DEXs, highlighting how traders manage latency risk optimally through priority fee selection. We show that employing the optimal priority fee has a significant outperformance over non-strategic fee selection.

Solving Linear-Quadratic Stochastic Control Problems with Signatures

A. Aqsha, P. Bank, L. Sánchez-Betancourt

Oxford-Man Institute of Quantitative Finance, University of Oxford, Technische Universität Berlin

We study a signature-driven numerical scheme to solve multi-dimensional linear-quadratic (LQ) stochastic control problems. Using that linear signature functionals are dense in the natural class of admissible controls, we show that our approach turns the original LQ problem into a deterministic convex quadratic polynomial optimisation. To underpin a numerical approach based on truncated signatures, we prove that the problem's value function can be approximated by finite-dimensional polynomial approximations when the truncation levels are chosen sufficiently high. Remarkably, our numerical experiments show very decent accuracy already for small truncation levels. Key tools for our analysis are (i) the algebraic representation of controlled stochastic differential equations and the associated cost function as linear functionals of the path signatures of the driving noise, (ii) the convergence of the truncated linear functionals, and (iii) the density of signature controls.

Competition between DEXs through Dynamic Fees

L. Baggiani, M. Herdegen, L. Sánchez-Betancourt

University of Warwick, University of Stuttgart, Oxford-Man Institute of Quantitative Finance, University of Oxford

We find an approximate Nash equilibrium in a game between decentralized exchanges (DEXs) that compete for order flow by setting dynamic trading fees. We characterize the equilibrium via a coupled system of partial differential equations and derive tractable approximate closed-form expressions for the equilibrium fees. Our analysis shows that the two-regime structure found in monopoly models persists under competition: pools alternate between raising fees to deter arbitrage and lowering fees to attract noise trading and increase volatility. Under competition, however, the switching boundary shifts from the oracle price to a weighted average of the oracle and competitors' exchange rates. Our numerical experiments show that, holding total liquidity fixed, an increase in the number of competing DEXs reduces execution slippage for strategic liquidity takers and lowers fee revenue per DEX. Finally, the effect on noise traders' slippage depends on market activity: they are worse off in low-activity markets but better off in high-activity ones.

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

Aqsha, Alif, Peter Bank, and Leandro Sánchez-Betancourt. 2026. "Solving Linear-Quadratic Stochastic Control Problems with Signatures." *arXiv preprint arXiv:2602.23473*.

Baggiani, Leonardo, Martin Herdegen, and Leandro Sánchez-Betancourt. 2026. "Competition between DEXs through Dynamic Fees." *arXiv preprint arXiv:2603.09669*.

Bergault, Philippe, Yadh Hafsi, and Leandro Sánchez-Betancourt. 2026. "Trading in CEXs and DEXs with

Priority Fees and Stochastic Delays." *arXiv preprint arXiv:2602.10798*.

Cartea, Álvaro, Mihai Cucuringu, and Lei Fang. 2026. "Network-Augmented Volatility Forecasting with Idiosyncratic Jump Spillovers." *Available at SSRN 6333798*.

Chang, Jinyuan, Yi Ding, Zhentao Shi, and Bo Zhang. 2026. "Zero Variance Portfolio." *arXiv preprint arXiv:2602.19462*.

Chen, Xiaohong, Elie Tamer, and Qingsong Yao. 2026. "Online Learning in Semiparametric Econometric Models." *arXiv preprint arXiv:2603.08614*.

Chong, Carsten H, and Viktor Todorov. 2026. "Realized Variance Disagreement." *Available at SSRN 6249218*.

Cohen, Lauren, Yiwen Lu, and Quoc H Nguyen. 2026. "Mimicking Finance." *NBER Working Paper No. 34849*.

Gao, Zhenyu, Wenxi Jiang, and Yutong Yan. 2025. "A Test of Lookahead Bias in LLM Forecasts." *arXiv preprint arXiv:2512.23847*.

Hizmeri, Rodrigo, and Mattia Bevilacqua. 2026. "Early Birds Get the Vol: Morning Volatility Uncertainty and Variance Risk Premium." *Available at SSRN 6212458*.

Hyland, Christopher. 2026. "Taming Black Swans: Monetary Policy in the Presence of Tail Risks." *Available at SSRN 6240078*.

Yan, Chengxin, Dachuan Chen, and Jia Li. 2026. "Non-parametric Inference for Spot Volatility in Pure-Jump Semimartingales." *arXiv preprint arXiv:2601.18371*.