VOL. 5, NO. 10

OMI Research Newsletter – July 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 6th International Conference on Computational Finance (ICCF), organised jointly by the Mathematical Institute and the OMI, will be held at St Catherine's College on 1-4 September 2026. For the schedule and registration details, follow this link.

Highlights

Key contributions in this month's edition include the following. Duan and Pelger (2026) develop a method-of-moments framework that uses a small, strategically selected human-labelled sample to correct bias in AI-generated covariates, while optimally combining outputs from multiple models and prompts to deliver valid and efficient downstream inference. Cartea, Cucuringu and Shi (2026) show that firm-specific lead-lag relationships remain pronounced at intraday frequencies after systematic risk is removed, and these relationships can be exploited through network-based trading strategies with substantial predictive power. Dai, Jia and Yu (2026) show that ultra-short asset-price prediction contracts allow traders to manipulate the underlying asset near settlement, thereby impairing price discovery and transferring wealth from retail traders, while longer contract horizons largely eliminate this problem. Liao, Neuhierl and Todorov (2026) construct a model-free portfolio of short-dated options that isolates shocks to expected market volatility and show that, although these shocks command a significant negative premium, the premium is largely attributable to dynamic market exposure rather than to a distinct cross-sectional volatility factor. Bondarenko et al. (2026) show that standard option-implied moments are generally not point-identified in incomplete option markets and develop robust alternatives with bounded and empirically small valuation uncertainty. Cabrera, Kolokolova and Zhang (2026) show that the announcement of an Analyst/Investor day predicts a sustained pre-event stock-price run-up, which is amplified by additional voluntary disclosures and is often followed by a post-event reversal. Di Nosse and Lillo (2026) develop an agent-based model of Uniswap v3 and show that dynamic fees can improve liquidity providers' hedged profitability, primarily by increasing fee income during periods of stale-price risk and toxic order flow.

Statistics, Econometrics & Machine Learning

Inference with AI-Generated Covariates

J. DUAN, M. PELGER
STANFORD UNIVERSITY

Key contributions in this month's edition include the following. Duan and Pelger (2026) develop a method-of-moments framework that uses a small, strategically selected human-labelled sample to correct bias in AI-generated covariates and optimally combine outputs across models and prompts. Cartea, Cucuringu and Shi (2026) show that firm-specific lead–lag relationships remain pronounced at intraday frequencies after removing systematic risk and can be exploited through network-based trading strategies. Dai, Jia and Yu (2026) show that ultra-short asset-price prediction contracts facilitate settlement manipulation, impair price discovery, and transfer wealth from retail traders, while longer horizons largely eliminate the problem. Liao, Neuhierl and Todorov (2026) construct a model-free short-dated option portfolio that isolates shocks to expected market volatility and show that their negative premium is largely explained by dynamic market exposure rather than a distinct volatility

factor. Bondarenko et al. (2026) show that standard option-implied moments are generally not point-identified in incomplete markets and develop robust alternatives with bounded uncertainty. Cabrera, Kolokolova and Zhang (2026) show that announcing an Analyst/Investor day predicts a sustained pre-event price run-up, amplified by additional voluntary disclosures and often followed by reversal. Finally, Di Nosse and Lillo (2026) develop an agent-based model of Uniswap v3 and show that dynamic fees improve liquidity providers' hedged profitability mainly by raising fee income during periods of stale-price risk and toxic order flow.

Causal Inference Using Factor Models

J. Bai, P. Wang
Columbia University, HKUST

We develop a factor-model framework for causal inference in panels with policy interventions. Treatment effects are represented as structural changes in treated units' exposure to latent common shocks and, in extensions, changes in the factor process itself. The approach does not impose the standard parallel-trends restriction, accommodates one or many treated units, and targets systematic effects when unit-time idiosyncratic effects are not point identified. We provide estimation and inference under both fixed and treatment-dependent factor processes. Simulations show coverage close to nominal levels. In applications to California tobacco control and German reunification, the method produces estimates broadly consistent with synthetic control while delivering formal confidence intervals.

Identification, Estimation and Inference Based on Structural Error Projection

C. Dong, J. Gao, O. Linton, B. Peng
Zhongnan University of Economics and Law, Monash University, University of Cambridge

This paper proposes to project and expand the conditional mean function of the structural error given the regressors in an endogenous regression under consideration. As the projection process is semiparametric, we define this procedure as a semiparametric projection (SP) method to address endogeneity in regression models by internally constructed instrumental variables. The SP method is applicable to many classes of regression models associated with endogeneity, such as linear, nonlinear, and non- and semi-parametric models, and provides a simple and computationally tractable alternative to conventional instrumental variable approaches available from the existing literature. This paper establishes identification conditions and derives the asymptotic properties of the resulting estimators. It then proposes a simple LASSO selection method to examine the finite-sample performance of both the proposed method and the established theory by simulated and real data examples.

Split-Session Cluster GARCH for Overnight and Intraday Returns: The Role of Tail Heterogeneity

X. Chen, P. R. Hansen, C. Tong
Xiamen University, University of North Carolina at Chapel Hill

We propose the Split-Session Cluster GARCH model for heavy-tailed multivariate dependence among asset returns decomposed into overnight and intraday components. The model uses convolution-t distributions to allow tail behavior to differ across clusters defined by trading sessions and, within each session, by economic sectors. It also accommodates block-structured conditional correlation matrices, preserving parsimony and scalability in high-dimensional settings. The resulting likelihood remains tractable and yields a score-driven specification for dynamic correlations. We apply the model to U.S. equity returns in six-asset and 100-asset applications. The results reveal pronounced tail heterogeneity between overnight and intraday returns. Model comparisons show that session-specific tail parameters substantially improve fit relative to a common multivariate-t specification, while sector-level tail partitioning delivers additional gains concentrated mainly in the overnight component. In the 100-asset application, asset-level tail heterogeneity delivers the strongest out-of-sample likelihood and global minimum-variance (GMV) portfolio performance.

Finance and Economics

Intraday Lead-Lag Relationships in Idiosyncratic Stock Returns

Á. Cartea, M. Cucuringu, D. Shi
Oxford-Man Institute, University of Oxford, University of California Los Angeles

We examine intraday lead-lag relationships in high-frequency stock returns that arise from firm-specific interactions. After removing systematic variation using common risk factors, we model lagged cross-dependencies in idiosyncratic returns with Kendall's correlation networks. Using S&P 500 constituent stocks, we document that lead-lag effects are both stronger and more persistent at higher sampling frequencies. In contrast to earlier evidence based on lower-frequency data, we find that these effects largely disappear in daily returns but remain pronounced at intraday horizons. A network-based trading strategy further shows that lead-lag signals extracted from high-frequency idiosyncratic returns have substantial predictive power and deliver significantly higher out-of-sample Sharpe ratios than signals based on low-frequency

returns or returns unadjusted for systematic risk.

Settlement Manipulation in Prediction Markets

D. Dai, R. Jia, S. Yu

Stanford University, Singapore Management University

Prediction markets increasingly list contracts settling on an asset price that holders can move by trading the underlying. We build a model showing that such contracts transfer wealth from prediction-market liquidity traders to manipulators and harm price discovery in the underlying, even as it becomes more liquid. After the launch of Polymarket's five-minute Bitcoin contract, settlement-time spot order flow spikes, causing large price reversals after settlement. Manipulators capture a large amount of profit, mostly from retail. Manipulation is largely absent in the fifteen-minute contracts: lengthening the contract horizon removes it, providing the market-design remedy our model and evidence support.

Is Expected Market Volatility a Priced Risk Factor?

Y. Liao, A. Neuhierl, V. Todorov

University of Iowa, Purdue University, Northwestern University

We study which sources of risk in expected market volatility are priced by investors. We construct a portfolio of index options spanning short-term forward variance, whose return identifies shocks to expected volatility while remaining immune to price jump risk. Its average return is negative and significant. We decompose this return into three orthogonal components: dynamic market and market-neutral cross-sectional stock positions hedging variance and a stock-unrelated residual. Most of the premium comes from the dynamic market component and the residual. Thus, once dynamic market exposure is accounted for, expected market volatility is not a separate cross-sectional risk factor.

What Can You Really Tell from Option Prices?

O. Bondarenko, Y. Dillschneider, P. Schneider, F. Trojani

University of Illinois at Chicago, University of Amsterdam, Tinbergen Institute, USI Lugano, Swiss Finance Institute, University of Geneva, University of Turin

Option-implied risk-neutral moments are widely used in empirical finance to measure moment risk premia, derive bounds on expected returns, and serve as inputs for models of conditional physical moments. Their identification, however, implicitly relies on the assumption of complete option markets. We develop novel model-free bounds that characterize the valuation uncertainty that an observed option crosssection leaves about a given risk-neutral moment under market incompleteness. For the VIX, as well as for essentially all standard measures of variance, skewness, and kurtosis, this uncertainty is unbounded. Consequently, predictive regressions and cyclicality tests based on these quantities can generate almost arbitrary parameter estimates, rendering their empirical conclusions largely uninformative. As a remedy, we introduce a family of robust risk-neutral moments whose valuation uncertainty remains bounded and empirically small, while integrating into existing empirical frameworks with only minimal modification.

Save The Date: Analyst/Investor Days as a Trading Signal

G. Cabrera, O. Kolokolova, S. S. Zhang

University of Manchester, Lancaster University

Firms increasingly use Analyst/Investor (A/I) days as discretionary disclosure channels. We show that announcing an A/I day leads to positive abnormal returns and a sustained pre-event price run-up. The run-up is more pronounced when firms issue additional pre-event voluntary disclosures and is followed by return reversals after the event. In contrast, events without any incremental disclosures exhibit positive post-event returns driven by event informativeness and changes in risk expectations. Exploiting these patterns, we design implementable long-short trading strategies that earn economically and statistically significant abnormal returns. A strategy based solely on the announcement yields positive alphas of around 0.095% before and 0.08% after transaction and borrowing costs. A deep-reinforcement-learning optimal stopping rule further improves performance, generating a positive alpha of 0.13% before costs, highlighting the importance of A/I days as a voluntary disclosure channel and the role of announcement and pre-event behaviour in shaping market prices.

Factors or Fake? A New Look at Anomalies

S. Chib, S. Xiao, L. Zhao

Washington University in St. Louis, City University of Hong Kong, Peking University

Based on the insight that declaring an anomaly as not spanned requires knowledge of the true stochastic discount factor (SDF), we reverse conventional testing and suppose that, under the alternative, the anomaly is spanned. Under this novel framework, we formulate a procedure that computes posterior probabilities for single anomalies, pairs, and larger sets, integrating over the distribution of possible SDFs, while controlling the expected false discovery rate at each stage. In the U.S.

data, 132 of 153 anomalies are classified as spanned. The evidence highlights redundancy and masking in the anomaly zoo.

Do Asset Pricing Models Change Over Time?

L. Cui, G. Feng, Y. Hong, J. Yang
City University of Hong Kong, Cornell University, Chinese Academy of Sciences, University of Chinese Academy of Sciences

Conditional asset pricing allows risk prices to vary with the state but treats the set of factors entering the SDF as fixed. We propose Sparse Fused GMM (SFGMM), a single-step estimator that allows this set, the support of the SDF loadings, to be time-varying through dynamic sparsity and heterogeneous structural breaks. Applied to 20 widely used U.S. factors, the estimates are difficult to reconcile with fully static specifications: the set of factors contributing to the SDF rotates over time, and estimated breaks cluster around recessions and inflation spikes. This is consistent with conditional asset pricing, in the sense that factors' marginal contribution to spanning the SDF varies across regimes. We also document correlation-driven substitution, whereby newer factors take over pricing roles once held by older ones. Emphasizing a factor's marginal contribution to the SDF, rather than its average premium, is associated with stronger cross-sectional fit and out-of-sample performance.

Volatility Disagreement in the Options Market

T. G. Bali, B. T. Kelly, M. Morke
Georgetown University, Yale University, ESCP Business School

We construct a novel stock-level measure of volatility disagreement as the cross-sectional dispersion of realized variance forecasts built from heterogeneous information sets and machine learning models. Sorting single-name delta-hedged straddles on this measure yields a long-short return of-5.14% per month, robust to known option return predictors and not subsumed by disagreement about first moments or the variance risk premium. A one-standard-deviation increase in volatility disagreement is associated with a 30% rise in option position opening. The cross-sectional patterns of our disagreement measure align with recent theoretical advances of beliefs and pricing of variance claims. Evidence on attention, ownership, and arbitrage costs is more consistent with mispricing than risk compensation.

Mathematical and Computational Finance

Mitigating Adverse Selection in Concentrated Liquidity AMMs with Dynamic Fees: An Agent-Based Model Approach

D. M. Di Nosse, F. Lillo
Scuola Normale Superiore

Automated Market Makers based on concentrated liquidity, such as Uniswap v3, significantly improve capital efficiency but expose Liquidity Providers (LPs) to adverse selection costs, formalized as Loss-Versus-Rebalancing (LVR). While theoretical literature quantifies these costs, the interplay between realistic blockchain microstructure and endogenous pricing mechanisms remains under-explored. This paper develops a granular Agent-Based Model of a Uniswap v3 pool interacting with a stochastic reference market governed by Heston volatility dynamics. The framework incorporates discrete block propagation, mempool latency, and a heterogeneous population of agents, including latency-sensitive arbitrageurs, smart routers, Maximal Extractable Value searchers, and active LPs benchmarked against a frictionless rebalancing strategy. We propose and evaluate dynamic fee schedules driven by volatility and order-flow toxicity proxies intended to compensate LPs for adverse-selection losses. Our simulations investigate the conditions under which LPs can achieve positive hedged Profit and Loss (fees minus LVR). The analysis suggests that dynamic fee adjustments can improve hedged LP profitability mainly by increasing fee income in states associated with stale-price risk. Depending on the configuration, these rules may also affect realized LVR, but the current aggregate results support compensation for LVR more directly than a reduction of LVR itself.

Volatility in Prediction Markets: A Structural Approach

W. Xi, C. C. Moallemi, M. Pai, S. Wang
Columbia University, Tempo Labs

Forward-looking volatility forecasts are central inputs to derivatives pricing, market making, risk management, and volatility-linked trading strategies, with ARCH and GARCH models serving as the canonical workhorses. Such models are natural in standard asset markets, where prices are positive-valued stochastic processes and volatility is typically inferred from return dynamics. Prediction markets have a different structure: prices are bounded probabilities, payoffs are binary, and contracts resolve at known deadlines. We develop and estimate a volatility model tailored to binary prediction markets. The model combines two economic mechanisms: a Wright-Fisher deadline-resolution component, capturing how remaining binary uncertainty is forced to resolve over time, and a Glosten-Milgrom order-flow component, capturing volatility from informed trading as reflected in spreads and volume. Using a large panel of Kalshi contracts, we show that these structural variables carry substantial forecasting power. Plain ARCH/GARCH benchmarks are dominated by structural specifications; combining the structural model with residual GARCH dynamics gives the best

overall forecasts. The model also provides an interpretable measurement framework: volatility is highest near fifty-fifty prices, rises near resolution, and varies across categories with the timing and discreteness of information arrival. Economics contracts are closer to smooth deadline-resolution dynamics, while sports contracts exhibit more event-concentrated, jump-like behavior. Across major categories, category-specific fitting does not systematically improve out-of-sample performance, suggesting that the structural specification transfers beyond the pooled headline result.

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

Bai, Jushan, and Peng Wang. 2026. "Causal inference using factor models." *arXiv preprint arXiv:2606.29691*.

Bali, Turan G, Bryan T Kelly, and Mathis Moerke. 2026. "Volatility Disagreement in the Options Market." *Available at SSRN 6961900*.

Bondarenko, Oleg, Yannick Dillschneider, Paul Schneider, and Fabio Trojani. 2026. "What Can You Really Tell from Option Prices?" *Available at SSRN 7017219*.

Cabrera, Gabriel, Olga Kolokolova, and S. Sarah Zhang. 2026. "Save The Date: Analyst/Investor Days as a Trading Signal." *Available at SSRN 6666218*.

Cartea, Álvaro, Mihai Cucuringu, and Danni Shi. 2026. "Intraday Lead-Lag Relationships in Idiosyncratic Stock Returns." *Available at SSRN 6811380*.

Chen, Xinxian, Peter Reinhard Hansen, and Chen Tong. 2026. "Split-Session Cluster GARCH for Overnight and Intraday Returns: The Role of Tail Heterogeneity." *arXiv preprint arXiv:2607.03669*.

Chib, Siddhartha, Shuhua Xiao, and Lingxiao Zhao. 2026. "Factors or Fake? A New Look at Anomalies." *Available at SSRN 6494339*.

Cui, Liyuan, Guanhao Feng, Yongmiao Hong, and Jiangshan Yang. 2026. "Do Asset Pricing Models Change Over Time?" *Available at SSRN 4431543*.

Dai, David, Ruizhe Jia, and Shihao Yu. 2026. "Settlement Manipulation in Prediction Markets." *arXiv preprint arXiv:2606.31675*.

Di Nosse, Daniele Maria, and Fabrizio Lillo. 2026. "Mitigating Adverse Selection in Concentrated Liquidity AMMs with Dynamic Fees: An Agent-Based Model Approach." *arXiv preprint arXiv:2606.23070*.

Dong, Chaohua, Jiti Gao, Oliver Linton, and Bin Peng. 2026. "Identification, Estimation and Inference Based on Structural Error Projection." *arXiv preprint arXiv:2607.05699*.

Duan, Junting, and Markus Pelger. 2026. "Inference with AI-Generated Covariates." *Available at SSRN 7053778*.

Liao, Yuan, Andreas Neuhierl, and Viktor Todorov. 2026. "Is Expected Market Volatility a Priced Risk Factor?" *Available at SSRN 6939978*.

Xi, Weiye, Ciamac C Moallemi, Mallesh Pai, and Shouqiao Wang. 2026. "Volatility in Prediction Markets: A Structural Approach." *arXiv preprint arXiv:2607.08199*.