VOL. 4, NO. 11

OMI Research Newsletter – August 2025



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 OMI is seeking to appoint one full-time Postdoctoral Research Assistant in Blockchain Economics and Decentralised Finance. Please find more details here.

Highlights

Key contributions in this month's edition: Calvo-Ordoñez et al. (2025) introduce a computationally efficient conditional flow matching method that uses simple pairwise Gibbs-kernel weights. Massacci et al. (2025) propose a new randomised test for zero alpha in linear asset pricing factor models that bypasses covariance matrix estimation, accommodates large panels, and remains robust to heteroskedasticity, non-Gaussian and strong cross-sectionally dependent error terms. Kolokolov (2025) propose a new nonparametric test for short-lived locally explosive trends in pure-jump processes, which is tailored to detect intraday flash crashes in cryptocurrency. Chernov, Dahlquist and Lochstoer (2025) show that characteristic-based asset-pricing factors embed sizable unpriced risk that pushes them off the mean–variance frontier, and decomposing them into priced and unpriced components and hedging the latter raises Sharpe ratios and isolates the true pricing factor. Bianchi, Jiao and Ma (2025) incorporate transaction costs directly into stochastic discount factor (SDF) estimation via adversarial neural networks, and show that trading frictions structurally shape the pricing kernel. Brownlees and Souza (2025) turn long–short portfolio construction into a statistical decision problem—buy stocks whose characteristics make them likely "winners," sell likely "losers," with tunable penalties for mistakes—and show that it outperforms standard expected-return sorts after transaction costs.

Statistics, Econometrics & Machine Learning

Weighted Conditional Flow Matching

S. CALVO-ORDOÑEZ, M. MEUNIER, Á. CARTEA, C. REISINGER, Y. GAL, J. M. HERNANDEZ-LOBATO
UNIVERSITY OF OXFORD, UNIVERSITY OF CAMBRIDGE

Conditional flow matching (CFM) has emerged as a powerful framework for training continuous normalizing flows due to its computational efficiency and effectiveness. However, standard CFM often produces paths that deviate significantly from straight-line interpolations between prior and target distributions, making generation slower and less accurate due to the need for fine discretization at inference. Recent methods enhance CFM performance by inducing shorter and straighter trajectories but typically rely on computationally expensive mini-batch optimal transport (OT). Drawing insights from entropic optimal transport (EOT), we propose Weighted Conditional Flow Matching (W-CFM), a novel approach that modifies the classical CFM loss by weighting each training pair (x, y) with a Gibbs kernel. We show that this weighting recovers the entropic OT coupling up to some bias in the marginals, and we provide the conditions under which the marginals remain nearly unchanged. Moreover, we establish an equivalence between W-CFM and the minibatch OT method in the large-batch limit, showing how our method overcomes computational and performance bottlenecks linked to batch size. Empirically, we test our method on unconditional generation on various synthetic and real datasets, confirming that W-CFM achieves comparable or superior sample quality, fidelity, and diversity to other alternative baselines while maintaining the computational efficiency of vanilla CFM.

A General Randomized Test for Alpha

D. Massacci, L. Sarno, L. Trapani, P. Vallarino

King's College London, University of Cambridge, Universita' di Pavia, University of Leicester, Erasmus University Rotterdam

We propose a methodology to construct tests for the null hypothesis that the pricing errors of a panel of asset returns are jointly equal to zero in a linear factor asset pricing model – that is, the null of "zero alpha". We consider, as a leading example, a model with observable, tradable factors, but we also develop extensions to accommodate for non-tradable and latent factors. The test is based on equation-by-equation estimation, using a randomized version of the estimated alphas, which only requires rates of convergence. The distinct features of the proposed methodology are that it does not require the estimation of any covariance matrix, and that it allows for both N and T to pass to infinity, with the former possibly faster than the latter. Further, unlike extant approaches, the procedure can accommodate conditional heteroskedasticity, non-Gaussianity, and even strong cross-sectional dependence in the error terms. We also propose a de-randomized decision rule to choose in favor or against the correct specification of a linear factor pricing model. Monte Carlo simulations show that the test has satisfactory properties and it compares favorably to several existing tests. The usefulness of the testing procedure is illustrated through an application of linear factor pricing models to price the constituents of the S&P 500.

Drift Bursts in Pure Jumps: Detection and Application to Bitcoin

A. Kolokolov

New Economic School

This paper proposes a new nonparametric test for detecting short-lived locally explosive trends (drift bursts) in pure-jump processes. The new test is designed specifically to detect intraday flash crashes and gradual jumps in cryptocurrency prices recorded at a high frequency. Empirical analysis shows that drift bursts in bitcoin price occur, on average, every second day. Their economic importance is highlighted by showing that hedge funds holding cryptocurrency in their portfolios are exposed to a risk factor associated with the intensity of bitcoin crashes. On average, hedge funds do not profit from intraday bitcoin crashes and do not hedge against the associated risk.

Yurinskii's Coupling for Martingales

M. D. Cattaneo, R. P. Masini, W. G. Underwood

Princeton University, University of California, Davis, University of Cambridge

Yurinskii's coupling is a popular theoretical tool for non-asymptotic distributional analysis in mathematical statistics and applied probability, offering a Gaussian strong approximation with an explicit error bound under easily verifiable conditions. Originally stated in ℓ_2-norm for sums of independent random vectors, it has recently been extended both to the ℓ_p-norm, for $1 \leq p \leq \infty$, and to vector-valued martingales in ℓ_2-norm, under some strong conditions. We present as our main result a Yurinskii coupling for approximate martingales in ℓ_p-norm, under substantially weaker conditions than those previously imposed. Our formulation further allows for the coupling variable to follow a more general Gaussian mixture distribution, and we provide a novel third-order coupling method which gives tighter approximations in certain settings. We specialize our main result to mixingales, martingales, and independent data, and derive uniform Gaussian mixture strong approximations for martingale empirical processes. Applications to nonparametric partitioning-based and local polynomial regression procedures are provided, alongside central limit theorems for high-dimensional martingale vectors.

Finance and Economics

Unpriced Risks: Rethinking Cross-Sectional Asset Pricing

M. Chernov, M. Dahlquist, L. A. Lochstoer

University of California, Los Angeles, Stockholm School of Economics

Characteristic-based factors embed large unpriced components that depress Sharpe ratios and deviate from the mean-variance efficient (MVE) frontier. We discuss how to decompose tradable factor returns into priced (MVE) and unpriced components, showing that hedging unpriced variation realigns factors with efficiency. We outline theoretical conditions for characteristic portfolios to span the MVE and describe practical hedge-portfolio construction. In some asset classes—currencies and sovereign bonds—real-time estimation of the MVE is feasible. In the case of equities, one can hedge unpriced risks from characteristic-based factors. Empirically, unpriced risks account for 30–99% of factor return variance, and hedging can more than double Sharpe ratios.

Transaction Costs and the Stochastic Discount Factor

D. Bianchi, T. Jiao, H. Ma

Queen Mary University of London

We show that transaction costs fundamentally reshape the stochastic discount factor (SDF) by determining which characteristics matter in equilibrium. Using deep neural networks, we incorporate transaction costs directly into robust SDF estimation rather than treating

them as post-optimization adjustments or arbitrary investment constraints. Transaction cost-aware SDFs yield substantially higher net Sharpe ratios and superior cross-sectional pricing through endogenous portfolio reallocation: increased diversification, reduced turnover, and lower exposure to costly high-turnover characteristics. These effects persist across sample configurations, market regimes, neural network specifications, and alternative definitions of transaction costs, demonstrating that trading frictions are structural determinants of equilibrium asset prices.

Can Equity Option Returns Be Explained by a Factor Model? IPCA Says Yes

A. Goyal, A. Saretto

Swiss Finance Institute, University of Lausanne, Federal Reserve Bank of Dallas

A number of delta-hedged equity option strategies exhibit very large average returns. We show that much of the profitability of these strategies can be explained by an IPCA factor model. The economic magnitude of the return-adjustment produced by IPCA is impressive: even before transaction costs, the average IPCA alpha of 46 long-short trading strategies constructed on previously discovered signals, is close to zero and contrasts with average realized returns of over 80 basis points per month. Our IPCA model can be used as a benchmark for assessing the performance of other option portfolios.

How To Bet On Winners (and Losers)

C. Brownlees, A. B. M. Souza

Universitat Pompeu Fabra, Barcelona School of Economics, ESADE Business School

We cast the construction of long-short portfolios as a statistical decision problem in which the investor seeks to buy the top-performing stocks (the "winners") and sell the worst-performing ones (the "losers") on the basis of stock characteristics. We derive the optimal portfolio selection rule implied by a loss function that accounts for different types of misclassification errors in portfolio construction. This approach leads to a return classification problem and the optimal rule buys or sells stocks based on their probabilities of being winners or losers, conditional on the stock characteristics. When returns are generated by an additive regression model and misclassification costs satisfy a symmetry condition, the optimal rule simplifies to the conventional sorting procedure based on expected returns. An empirical application using U.S. stock data shows that portfolios constructed using the optimal rule achieve higher Sharpe ratios compared to those built using conventional methods. Our results demonstrate that predictive signals in the cross-section of stock returns go beyond expected returns, and that properly optimized portfolio selection rules based on these signals can generate substantial economic value for investors.

Volatility Jumps Are Hard to Identify, Even With Options Data

Y. Aït-Sahalia, C. X. Li, C. Li

Princeton University, NBER, Renmin University of China, Peking University

We develop a closed-form expansion for the implied volatility surface and its shape characteristics when the underlying asset follows a general stochastic volatility model including jumps in both price and volatility. We focus in particular on the potential of the implied volatility surface to reveal information about volatility jumps, which are challenging to identify from price data alone. We show that volatility jumps arriving simultaneously with price jumps have a larger impact on the shape of the implied volatility surface than idiosyncratic volatility jumps. But even with rich data such as the entire implied volatility surface, volatility jumps remain challenging to identify both theoretically and empirically.

Uncertainty in Empirical Economics

F. Schorfheide, Z. You

University of Pennsylvania

Econometricians invest substantial effort in constructing standard errors that yield valid inference under a hypothetical data-generating process. This paper asks a fundamental question: Are the uncertainty statements reported by applied researchers consistent with empirical frequencies? The short answer is no. Drawing on the forecasting literature, we predict estimates from “new” studies using estimates from corresponding baseline studies. By doing this across a large number of study groups and linking parameters through a hierarchical model, we compare stated probabilities to observed empirical frequencies. Alignment occurs only under limited external validity, namely, that the studies estimate different parameters.

Mathematical and Computational Finance

Optimal Decisions for Liquid Staking: Allocation and Exit Timing

R. Ma, Z. Cai, W. Tang, D. Yao

In this paper, we study an investor’s optimal entry and exit decisions in a liquid staking protocol (LSP) and an automated market maker (AMM), primarily from the standpoint of the investor. Our analysis focuses on two key investor actions: the initial allocation decision at time $t = 0$, and the optimal timing of exit. First, we derive an optimal allocation strategy that enables the investor to

distribute risk across the LSP, AMM, and direct holding. Our results also offer insights for LSP and AMM designers, identifying the necessary and sufficient conditions under which the investor is incentivized to stake through an LSP, and further, to provide liquidity in addition to staking. These conditions include a lower bound on the transaction fee, for which we propose a fee mechanism that attains the bound. Second, given a fixed protocol design, we model the optimal exit timing of an individual investor using Laplace transforms and free-boundary techniques. We analyze scenarios with and without transaction fees. In the absence of fees, we decompose the investor's payoff into impermanent loss and opportunity cost, and provide theoretical results characterizing the investor's payoff and the optimal exit threshold. With transaction fees, we conduct numerical analyses to examine how fee accumulation influences exit strategies. Our results reveal that in both settings, a stop-loss strategy often maximizes the investor's expected payoff, driven by opportunity gains and the accumulation of fees where fees are present. Our analyses rely on various tools from stochastic processes and control theory, as well as convex optimization and analysis. We further support our theoretical insights with numerical experiments and explore additional properties of the investor's value function and optimal behavior.

Volatility Modeling with Rough Paths: A Signature-Based Alternative to Classical Expansions

E. Alòs, Ò. Burés, R. de Santiago, J. Vives

Universitat Pompeu Fabra, Barcelona School of Economics, Financera i Actuarial, IESE Business School, Universitat de Barcelona

We compare two methodologies for calibrating implied volatility surfaces: a second-order asymptotic expansion method derived via Malliavin calculus, and a data-driven approach based on path signatures from rough path theory. The former, developed in Alòs et al. (2015), yields efficient and accurate calibration formulas under the assumption that the asset price follows a Heston-type stochastic volatility model. The latter models volatility as a linear functional of the signature of a primary stochastic process, enabling a flexible approximation without requiring a specific parametric form. Our numerical experiments show that the signature-based method achieves calibration accuracy comparable to the asymptotic approach when the true dynamics are Heston. We then test the model in a more general setting where the asset follows a rough Bergomi volatility process-a regime beyond the scope of the asymptotic expansion-and show that the signature approach continues to deliver accurate results. These findings highlight the model-independence, robustness and adaptability of signature-based calibration methods in settings where volatility exhibits rough or non-Markovian features.

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, and Gerardo Duran-Martin.

References

Alòs, Elisa, Òscar Bur'es, Rafael de Santiago, and Josep Vives. 2025. "Volatility Modeling with Rough Paths: A Signature-Based Alternative to Classical Expansions." *arXiv preprint arXiv:2507.23392*.

Aït-Sahalia, Yacine, Chen Xu Li, and Chenxu Li. 2025. "Volatility Jumps Are Hard to Identify, Even With Options Data." *Available at SSRN 4919244*.

Bianchi, Daniele, Teng Jiao, and Hao Ma. 2025. "Transaction Costs and the Stochastic Discount Factor." *Available at SSRN 5365375*.

Brownlees, Christian, and André B M Souza. 2025. "How To Bet On Winners (and Losers)." *Available at SSRN 5352643*.

Calvo-Ordoñez, Sergio, Matthieu Meunier, Álvaro Cartea, Christoph Reisinger, Yarin Gal, and Jose Miguel Hernandez-Lobato. 2025. "Weighted Conditional Flow Matching." *arXiv preprint arXiv:2507.22270*.

Cattaneo, Matias D, Ricardo P Masini, and William G Underwood. 2022. "Yurinskii's Coupling for Martingales." *arXiv preprint arXiv:2210.00362*.

Chernov, Mikhail, Magnus Dahlquist, and Lars A Lochstoer. 2025. "Unpriced Risks: Rethinking Cross-Sectional Asset Pricing." NBER Working Paper.

Goyal, Amit, and Alessio Saretto. 2025. "Can equity option returns be explained by a factor model? IPCA says yes." *The Review of Financial Studies*, 38(6): 1783–1821.

Kolokolov, Aleksey. 2025. "Drift Bursts in Pure Jumps: Detection and Application to Bitcoin." *Journal of Business & Economic Statistics*, forthcoming.

Ma, Ruofei, Zhebiao Cai, Wenpin Tang, and David Yao. 2025. "Optimal Decisions for Liquid Staking: Allocation and Exit Timing." *arXiv preprint arXiv:2507.14810*.

Massacci, Daniele, Lucio Sarno, Lorenzo Trapani, and Pierluigi Vallarino. 2025. "A General Randomized Test for Alpha." *arXiv preprint arXiv:2507.17599*.

Schorfheide, Frank, and Zhiheng You. 2025. "Uncertainty in Empirical Economics." *NBER Working Paper No. w33962*.