OMI Research Newsletter – October 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 Workshop on "New Challenges and Risks in Finance" will be held at the Man Group (London) on 13 November 2025. Please find the schedule here.

Highlights

Key contributions in this month's edition: Chen et al. (2025) develop a high-dimensional framework that learns a mean-variance-efficient stochastic discount factor from cross-sectional regression-based characteristic returns, which outperforms benchmark factor models and enables inference on which characteristic themes drive premia. Bouamara, Laurent and Shi (2025) introduce a stepwise Cauchy combination test that pinpoints which hypotheses drive rejections while controlling familywise error under dependence, with empirical applications to drift bursts and nonzero alphas. Fortin, Gagliardini and Scaillet (2025) construct optimal GMM-based tests of the key PCA assumption in short panels and show it fails for U.S. stock returns, calling routine PCA into question. Dew-Becker and Giglio (2025) document that the variance risk premium has largely vanished: traded index options' once strongly negative alphas have moved to about zero since around 2010, while synthetic replications never exhibited negative alphas. Avramov et al. (2025) propose "Schrödinger's Sparsity," a Bayesian conditional factor model that infers—rather than fixes—the degree of sparsity in alphas and betas, finding that alphas are typically sparser and that the approach outperforms dense and fixed-sparsity benchmarks. Bergault, Bieber and Sánchez-Betancourt (2025) model a liquidity provider's optimal exit from an automated market maker as a stochastic control problem and show numerically that exits tighten when oracle and pool prices diverge and depend on volatility, fee levels, and the activity of arbitrageurs and noise traders. Chilenje et al. (2025) model limit-order-book competition as a leader–follower game between a representative market maker and the aggregate rivals, derive equilibrium quoting rules, and show how spreads and inventory skew provide testable signals of competition versus cooperation.

Statistics, Econometrics & Machine Learning

Cross-Sectional Learning and Inference for the Stochastic Discount Factor

Z. CHEN, Y. DING, Y. LI, X. ZHENG
HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY, UNIVERSITY OF MACAU

We develop a statistical learning framework for constructing the stochastic discount factor (SDF) portfolio. To address the dimensionality challenge, we extend the MAXSER method (Ao et al., 2019) to allow for $N \gg T$; prove that it surely screens for useful characteristics; and establish asymptotic normality for the SDF loading estimates. Using 153 characteristics returns from cross-sectional regressions (Fama and French, 2020), our framework not only constructs an SDF with a high out-of sample Sharpe ratio that successfully prices the cross-section of expected returns, but also allows us to identify key characteristic themes and test the significance of their contributions.

A Stepwise Cauchy Combination Test for Multiple Testing Problems with Financial Applications

N. BOUAMARA, S. LAURENT, S. SHI
NATIONAL BANK OF BELGIUM, AIX-MARSEILLE UNIVERSITY, MACQUARIE UNIVERSITY

We propose the stepwise Cauchy combination test (StepC), a new procedure for multiple testing with dependent statistics and sparse signals. Unlike the global version, StepC pinpoints which p-values drive rejections, while maintaining strong familywise error control. It is less conservative under dependence and more powerful than conventional multiple testing corrections. In simulations and in applications to drift burst detection and testing for nonzero alphas, StepC consistently boosts power and yields more meaningful rejections, making it a practical alternative for large-scale financial datasets.

Optimal Maximin GMM Tests for Sphericity in Latent Factor Analysis of Short Panels

A.-P. Fortin, P. Gagliardini, O. Scaillet
University of Chicago

We derive optimal maximin tests for parametric hypotheses in short panels with latent common factors. We rely on a Generalized Method of Moments setting with optimal weighting under a large cross sectional dimension n and a fixed time series dimension T. We outline the asymptotic distributions of the estimators as well as the asymptotic maximin optimality of the Wald, Lagrange Multiplier, and Likelihood Ratio type tests. The characterisation of optimality relies on finding the limit Gaussian experiment in strongly identified GMM models under a block dependence structure and unobserved heterogeneity. We reject sphericity of idiosyncratic errors in an empirical application to a large cross-section of U.S. stocks, which casts doubt on the validity of routinely applying Principal Component Analysis to short panels of monthly financial returns.

Testing for Rough Volatility When Prices are Purely Discontinuous

C. H. Chong, V. Todorov
Hong Kong University of Science and Technology, Northwestern University

We consider the problem of nonparametric testing for rough volatility, using highfrequency data with a fixed time span, in a setting when the price is purely discontinuous. More specifically we analyze the asymptotic properties of a test we developed in Chong and Todorov (2025) in a pure-jump setting. We show that the asymptotic behavior of the test remains unchanged, regardless of whether the underlying price contains a diffusion or is of purely discontinuous type. This adaptive nature of the test is very convenient from an empirical point of view as it implies that no pre-test for the underlying features of the price process, i.e., whether it contains a diffusion, is needed when testing for rough volatility. A Monte Carlo study shows good finite sample behavior of the test in purely discontinuous price settings.

Smooth Tests for Normality in ANOVA

P. Jia, X. Song, H. Wei
Peking University, University of California, San Diego

The normality assumption for random errors is fundamental in the analysis of variance (ANOVA) models, yet it is seldom subjected to formal testing in practice. In this paper, we develop Neyman's smooth tests for assessing normality in a broad class of ANOVA models. The proposed test statistics are constructed via the Gaussian probability integral transformation of ANOVA residuals and are shown to follow an asymptotic Chi-square distribution under the null hypothesis, with degrees of freedom determined by the dimension of the smooth model. We further propose a data-driven selection of the model dimension based on a modified Schwarz's criterion. Monte Carlo simulations demonstrate that the tests maintain the nominal size and achieve high power against a wide range of alternatives. Our framework thus provides a systematic and effective tool for formally validating the normality assumption in ANOVA models.

Finance and Economics

Bond-Stock Comovements

J. Y. Campbell, C. Pflueger, L. M. Viceira
Harvard University, University of Chicago, NBER, CEPR

This paper documents that during the late 20th Century, nominal government bonds and stocks tended to comove positively, whereas during the first quarter of the 21st Century they have tended to comove negatively. A similar sign switch is observable for real government bonds and breakeven inflation rates. Recent macroeconomic events have caused short-lived changes in these comovements, and periods with high risk premia tend to be periods in which bond-stock comovements are large in absolute value. The paper surveys theoretical models of these phenomena.

The Decline of the Variance Risk Premium: Evidence from Traded and Synthetic Options

I. Dew-Becker, S. Giglio
Federal Reserve Bank of Chicago, Yale University, NBER

Equity index options historically displayed sharply negative returns and CAPM alphas. This could reflect investor risk preferences or intermediary frictions. We document that over the past 15 years, option alphas have become indistinguishable from zero. We also introduce synthetic options, that, under some conditions, reflect risk preferences of the average equity investor, independent of option-market frictions. Synthetic options never, over the

last 100 years, had negative alpha, indicating that equity investors never required high compensation for market downturns. An intermediarybased model explains the patterns in both synthetic and traded options, including the recent decline in the variance risk premium.

Schrödinger's Sparsity in the Cross Section of Stock Returns

D. Avramov, G. Feng, J. He, S. Xiao
Reichman University, City University of Hong Kong

This paper proposes Schrödinger's Sparsity, a Bayesian asset pricing framework that treats sparsity as a latent feature to be inferred rather than imposed. Departing from conventional models that assume either sparse or dense structures ex ante, our approach learns the probability of sparsity in both mispricing (alphas) and factor loadings (betas) directly from the data. Empirically, we find that alphas are considerably sparser than betas, and that sparsity varies systematically with asset complexity, pricing difficulty, and macroeconomic conditions. By allowing model complexity to emerge endogenously, our method yields interpretable, adaptive, and high-performing predictions, offering a principled bridge between sparse selection and dense shrinkage approaches.

Local Estimation for Option Pricing: Improving Forecasts with Market State Information

H. J. Kim, D. H. Oh
Federal Reserve System

We propose a novel estimation framework for option pricing models that incorporates local, state-dependent information to improve out-of-sample forecasting performance. Rather than modifying the underlying option pricing model, such as the Heston-Nandi GARCH or the Heston stochastic volatility framework, we introduce a local M-estimation approach that conditions on key state variables including VIX, realized volatility, and time. Our method reweights historical observations based on their relevance to current market conditions, using kernel functions with bandwidths selected via a validation procedure. This adaptive estimation improves the model's responsiveness to evolving dynamics while maintaining tractability. Empirically, we show that local estimators substantially outperform traditional non-local approaches in forecasting near-term option implied volatilities. The improvements are particularly pronounced in low-volatility environments and across the cross-section of options. The local estimators also outperform the non-local estimators in explaining future option returns. Our findings suggest that local information, when properly incorporated into the estimation process, can enhance the accuracy and robustness of option pricing models.

Volatility Forecasting Factors

M. Cinquetti, S. Y. Hong, I. Nolte, S. Nolte
Lancaster University, Nanyang Technological University

This paper examines the predictive role of high-frequency factor volatilities in modeling the volatility of individual stocks. We develop a dynamic forecasting framework that selects the most informative factor-specific realized volatility from a large cross-section of asset pricing anomalies. Embedded in a log-linear specification, the model integrates both market-wide and idiosyncratic components, allowing for a flexible representation of volatility dynamics. We prove the strong consistency of our selection procedure, and show that our selection rule asymptotically identifies the factor that truly drives volatility. We further show how measurement errors affect the adaptive selection. Empirical results based on a broad universe of U.S. equities demonstrate that the proposed method significantly outperforms standard benchmarks, both statistically and economically. The findings underscore the importance of incorporating high-frequency cross-sectional information in volatility modeling, offering a scalable and interpretable approach to understanding time-varying risks in equity markets.

AI Agents for Economic Research

A. Korinek
University of Virginia

The objective of this paper is to demystify AI agents - autonomous LLM-based systems that plan, use tools, and execute multi-step research tasks - and to provide hands-on instructions for economists to build their own, even if they do not have programming expertise. As AI has evolved from simple chatbots to reasoning models and now to autonomous agents, the main focus of this paper is to make these powerful tools accessible to all researchers. Through working examples and step-by-step code, it shows how economists can create agents that autonomously conduct literature reviews across myriads of sources, write and debug econometric code, fetch and analyze economic data, and coordinate complex research workflows. The paper demonstrates that by "vibe coding" (programming through natural language) and building on modern agentic frameworks like LangGraph, any economist can build sophisticated research assistants and other autonomous tools in minutes. By providing complete, working implementations alongside conceptual frameworks, this guide demonstrates how to employ AI agents in every stage of the research process, from initial investigation to final analysis.

A Research Agenda for the Economics of Transformative AI

E. Brynjolfsson, A. Korinek, A. K. Agrawal

Stanford University, University of Virginia, University of Toronto

As we approach Transformative Artificial Intelligence (TAI), there is an urgent need to advance our understanding of how it could reshape our economic models, institutions and policies. We propose a research agenda for the economics of TAI by identifying nine Grand Challenges: economic growth, innovation, income distribution, decision-making power, geoeconomics, information flows, safety risks, human well-being, and transition dynamics. By accelerating work in these areas, researchers can develop insights and tools to help fulfill the economic potential of TAI.

Mathematical and Computational Finance

Optimal Exit Time for Liquidity Providers in Automated Market Makers

P. Bergault, S. Bieber, L. Sánchez-Betancourt

University of Oxford

We study the problem of optimal liquidity withdrawal for a representative liquidity provider (LP) in an automated market maker (AMM). LPs earn fees from trading activity but are exposed to impermanent loss (IL) due to price fluctuations. While existing work has focused on static provision and exogenous exit strategies, we characterise the optimal exit time as the solution to a stochastic control problem with an endogenous stopping time. Mathematically, the LP's value function is shown to satisfy a Hamilton-Jacobi-Bellman quasi-variational inequality, for which we establish uniqueness in the viscosity sense. To solve the problem numerically, we develop two complementary approaches: a Euler scheme based on operator splitting and a Longstaff-Schwartz regression method. Calibrated simulations highlight how the LP's optimal exit strategy depends on the oracle price volatility, fee levels, and the behaviour of arbitrageurs and noise traders. Our results show that while arbitrage generates both fees and IL, the LP's optimal decision balances these opposing effects based on the pool state variables and price misalignments. This work contributes to a deeper understanding of dynamic liquidity provision in AMMs and provides insights into the sustainability of passive LP strategies under different market regimes.

Market Making with Competition: A Stackelberg Equilibrium

N. Chilenje, M. Daba, D. Feleppa, C. Fellner, L. Sánchez-Betancourt

University of Cape Town, Sapienza Università di Roma, University of Oxford

We study the Stackelberg equilibrium between a representative market maker (follower) and the aggregate competition (leader), representing the other liquidity providers in a given limit order book market. Our results extend the framework of Boyce et al. (SIAM J. Financ. Math. 16(2), 692-706) by going beyond exogenous competition. For a given class of time-dependent strategies for the leader and a linear-quadratic objective, we obtain an approximate closed-form solution for the follower's strategy. Using this solution, we characterize the leader's objective function as the solution to a differential Sylvester equation and show how the leader optimizes over their strategy class to approach the Stackelberg equilibrium. By studying this equilibrium as a function of model parameters, we identify testable hypotheses that can be used to assess the degree of competition or cooperation between a representative market maker and the rest of the liquidity providers.

The Price of Liquidity: Implied Volatility of Automated Market Maker Fees

M. Bichuch, Z. Feinstein

SUNY Buffalo, Stevens Institute of Technology

An automated market maker (AMM) provides a method for creating a decentralized exchange on the blockchain. For this purpose, individual investors lend liquidity to the AMM pool in exchange for a stream of fees earned from its operations as a market maker. Within this work, we reinterpret the loss-versus-rebalancing as the implied fee stream generated by an AMM so that a risk-neutral investor is indifferent in the decision of providing liquidity. With this implied fee structure, we propose a novel fixed-for-floating swap on the fees generated by an AMM in order to quote the implied volatilities and implied correlations of digital assets. We apply this theory to realized fees in different markets to empirically validate the relevance of the deduced fee-based volatility.

Pooling Liquidity Pools in AMMs

M. Bagnulo, A. Hernando-Veciana, E. Smyrniotis

Universidad Carlos III de Madrid, University of Bristol

Market fragmentation across multiple Automated Market Makers (AMMs) creates inefficiencies such as costly arbitrage, unnecessarily high slippage and delayed incorporation of new information into prices. These inefficiencies raise trading costs, reduce liquidity provider profits, and degrade overall market efficiency. To address these issues, we propose a modification of the Constant Product Market Maker (CPMM) pricing mechanism, called the Global Market Maker (GMM), which aggregates liquidity information from all AMMs to mitigate these inefficiencies. Through theoretical and numerical analyses, we demonstrate that the GMM enhances profits for both AMMs and traders by eliminating arbitrage opportunities.

Additionally, it reduces the profitability of sandwich attacks and minimizes impermanent losses.

Occupied Processes: Going with the Flow

VALENTIN TISSOT-DAGUETTE
BLOOMBERG

A stochastic process X becomes occupied when it is enlarged with its occupation flow $\mathcal{O}$, which tracks the time spent by the path at each level. Crucially, the occupied process $(\mathcal{O}, X)$ enjoys a Markov structure when X is Markov. We develop a novel Itô calculus for occupied processes that lies midway between Dupire's functional Itô calculus and the classical version. We derive a surprisingly simple Itô formula and, through Feynman-Kac, unveil a broad class of path-dependent PDEs where $\mathcal{O}$ plays the role of time. The space variable, given by the current value of X, remains finite-dimensional, thereby paving the way for standard elliptic PDE techniques and numerical methods.

In the financial applications, we demonstrate that occupation flows provide unified Markovian lifts for exotic options and variance instruments, allowing financial institutions to price and manage derivatives books with a single numerical solver. We then explore avenues in financial modeling where volatility is driven by the occupied process. In particular, we propose the local occupied volatility (LOV) model which not only calibrates to European vanilla options but also offers the flexibility to capture stylized facts of volatility and fit other instruments. We also present an extension of forward variance models that leverages the entire forward occupation surface.

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, and Kevin Schneider.

References

Avramov, Doron, Guanhao Feng, Jingyu He, and Shuhua Xiao. 2025. "Schrödinger's Sparsity in the Cross Section of Stock Returns." *Available at SSRN 5370960*.

Bagnulo, Marcelo, Angel Hernando-Veciana, and Efthymios Smyrniotis. 2025. "Pooling Liquidity Pools in AMMs." *arXiv preprint arXiv:2503.09765*.

Bergault, Philippe, Sébastien Bieber, and Leandro Sánchez-Betancourt. 2025. "Optimal Exit Time for Liquidity Providers in Automated Market Makers." *arXiv preprint arXiv:2509.06510*.

Bichuch, Maxim, and Zachary Feinstein. 2025. "The Price of Liquidity: Implied Volatility of Automated Market Maker Fees." *arXiv preprint arXiv:2509.23222*.

Bouamara, Nabil, Sébastien Laurent, and Shuping Shi. 2025. "A Stepwise Cauchy Combination Test for Multiple Testing Problems with Financial Applications." *Working Paper*.

Brynjolfsson, Erik, Anton Korinek, and Ajay K Agrawal. 2025. "A Research Agenda for the Economics of Transformative AI." *NBER Working Paper 34256*.

Chen, Zhanhui, Yi Ding, Yingying Li, and Xinghua Zheng. 2025. "Cross-Sectional Learning and Inference for the Stochastic Discount Factor." *Available at SSRN 5391303*.

Chilenje, Nehelo, Mandisa Daba, Davide Feleppa, Cameron Fellner, and Leandro Sánchez-Betancourt. 2025. "Market Making with Competition: A Stackelberg Equilibrium." *Available at SSRN 5505480*.

Chong, Carsten H, and Viktor Todorov. 2025. "Testing for Rough Volatility When Prices are Purely Discontinuous." *Available at SSRN 5344725*.

Cinquetti, Marco, Seok Young Hong, Ingmar Nolte, and Sandra Nolte. 2025. "Volatility Forecasting Factors." *Available at SSRN 5502438*.

Dew-Becker, Ian, and Stefano Giglio. 2025. "The Decline of the Variance Risk Premium: Evidence from Traded and Synthetic Options." *Available at SSRN 5525882*.

Fortin, Alain-Philippe, Patrick Gagliardini, and Olivier Scaillet. 2025. "Optimal Maximin GMM Tests for Sphericity in Latent Factor Analysis of Short Panels." *Available at SSRN 5169343*.

Jia, Peiwen, Haoyu Wei, and Xiaojun Song. 2025. "Smooth Tests for Normality in ANOVA." *arXiv preprint arXiv:2110.04849*.

Kim, Hyung Joo, and Dong Hwan Oh. 2025. "Local Estimation for Option Pricing: Improving Forecasts with Market State Information." *Available at SSRN 5535649*.

Korinek, Anton. 2025. "AI Agents for Economic Research, August 2025 Update to "Generative AI for Economic Research: Use Cases and Implications for Economists," Published in the *Journal of Economic Literature* 61(4)." *NBER Working Paper 34202*.

Pflueger, John Y Campbell Carolin, and Luis M Viceira. 2025. "Bond-Stock Comovements." *NBER Working Paper 34323*.

Tissot-Daguette, Valentin. 2025. "Occupied processes: Going with the flow." *arXiv preprint arXiv:2311.07936*.