VOL. 4, NO. 12

OMI Research Newsletter – September 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 Big Data and Finance Workshop will be held at the Oxford-Man Institute (Eagle House) on 16–17 October 2025. Please find the schedule and register here.

Highlights

Key contributions in this month's edition: Medeiros and Pinro (2025) introduce a reinforcement-learning framework for dynamic forecast combination. Andrew, Chen and Tecchio (2025) explore why Hansen's J-statistic is especially informative under model misspecification. Cavaliere, Fanelli and Georgiev (2025) propose an easy-to-implement bootstrap-based specification test that keeps valid against a wide range of alternatives and free from pre-testing bias. Lu et al. (2025) build overnight–daytime lead–lag stock networks and introduce a novel clustering-based framework for portfolio construction, which delivers significant factor-adjusted alphas. Bali et al. (2025) find that option returns show momentum within the same half-day (intraday or overnight) and reversal across periods, and attribute this predictability mainly to supply-side market-maker capacity constraints. Freire et al. (2025) build cross-country, theme-specific sentiment linkages from 520 million news articles, and shows that out-of-sample market-timing strategies based on sentiment measures predict equity returns and yield sizable, cost-robust alphas. Hansen and Tong (2025*b*) introduce an option-pricing framework with time-varying volatility risk aversion and provide analytical pricing approximations. Cohen et al. (2025) analyse and test a neural actor-critic algorithm for high-dimensional Hamilton-Jacobi-Bellman (HJB) partial differential equations from stochastic control theory.

Statistics, Econometrics & Machine Learning

Time Series Embedding and Combination of Forecasts: A Reinforcement Learning Approach

M. C. MEDEIROS, J. M. PINRO
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, MINISTRY OF ECONOMY (BRAZIL)

The forecasting combination puzzle is a well-known phenomenon in forecasting literature, stressing the challenge of outperforming the simple average when aggregating forecasts from diverse methods. This study proposes a Reinforcement Learning-based framework as a dynamic model selection approach to address this puzzle. Our framework is evaluated through extensive forecasting exercises using simulated and real data. Specifically, we analyze the M4 Competition dataset and the Survey of Professional Forecasters (SPF). This research introduces an adaptable methodology for selecting and combining forecasts under uncertainty, offering a promising advancement in resolving the forecasting combination puzzle.

The Purpose of an Estimator Is What It Does: Misspecification, Estimands, and Over-Identification

I. ANDREWS, J. CHEN, O. TECCHIO
MASSACHUSETTS INSTITUTE OF TECHNOLOGY, NBER, STANFORD UNIVERSITY

In over-identified models, misspecification – the norm rather than exception – fundamentally changes what estimators estimate. Different estimators imply different estimands rather than different efficiency for the same

target. A review of recent applications of generalized method of moments in the American Economic Review suggests widespread acceptance of this fact: There is little formal specification testing and widespread use of estimators that would be inefficient were the model correct, including the use of "hand-selected" moments and weighting matrices. Motivated by these observations, we review and synthesize recent results on estimation under model misspecification, providing guidelines for transparent and robust empirical research. We also provide a new theoretical result, showing that Hansen's J-statistic measures, asymptotically, the range of estimates achievable at a given standard error. Given the widespread use of inefficient estimators and the resulting researcher degrees of freedom, we thus particularly recommend the broader reporting of J-statistics.

Bootstrap Diagnostic Tests

G. Cavaliere, L. Fanelli, I. Georgiev
University of Bologna, University of Exeter

Violation of the assumptions underlying classical (Gaussian) limit theory frequently leads to unreliable statistical inference. This paper shows the novel result that the bootstrap can detect such violation by means of simple and powerful tests which (a) induce no pre-testing bias, (b) can be performed using the same critical values in a broad range of applications, and (c) are consistent against deviations from asymptotic normality. By focusing on the discrepancy between the conditional distribution of a bootstrap statistic and the (limiting) Gaussian distribution which obtains under valid specification, we show how to assess whether this discrepancy is large enough to indicate specification invalidity. The method, which is computationally straightforward, only requires to measure the discrepancy between the bootstrap and the Gaussian distributions based on a sample of i.i.d. draws of the bootstrap statistic. We derive sufficient conditions for the randomness in the data to mix with the randomness in the bootstrap repetitions in a way such that (a), (b) and (c) above hold. To demonstrate the practical relevance and broad applicability of our diagnostic procedure, we discuss five scenarios where the asymptotic Gaussian approximation may fail: (i) weak instruments in instrumental variable regression; (ii) non-stationarity in autoregressive time series; (iii) parameters near or at the boundary of the parameter space; (iv) infinite variance innovations in a location model for i.i.d. data; (v) invalidity of the delta method due to (near-)rank deficiency in the implied Jacobian matrix. An illustration drawn from the empirical macroeconomic literature concludes.

Moments by Integrating the Moment-Generating Function

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

We introduce a novel method for obtaining a wide variety of moments of a random variable with a well-defined moment-generating function (MGF). We derive new expressions for fractional moments and fractional absolute moments, both central and non-central moments. The new moment expressions are relatively simple integrals that involve the MGF, but do not require its derivatives. We label the new method CMGF because it uses a complex extension of the MGF and can be used to obtain complex moments. We illustrate the new method with three applications where the MGF is available in closed-form, while the corresponding densities and the derivatives of the MGF are either unavailable or very difficult to obtain.

Adaptive Robust Confidence Intervals

Y. Luo, C. Gao
University of Chicago

This paper studies the construction of adaptive confidence intervals under Huber's contamination model when the contamination proportion is unknown. For the robust confidence interval of a Gaussian mean, we show that the optimal length of an adaptive interval must be exponentially wider than that of a non-adaptive one. An optimal construction is achieved through simultaneous uncertainty quantification of quantiles at all levels. The results are further extended beyond the Gaussian location model by addressing a general family of robust hypothesis testing. In contrast to adaptive robust estimation, our findings reveal that the optimal length of an adaptive robust confidence interval critically depends on the distribution's shape.

Finance and Economics

A Tug of War across the Market: Overnight-vs.-Daytime Lead-Lag Networks and Clustering-Based Portfolio Strategies

Y. Lu, N. Zhang, G. Reinert, M. Cucuringu
University of Oxford, University of California, Los Angeles

In this research, we show that the tug of war is not only at individual stock level but also a networked trading behaviour across the entire market. By decomposing daily returns into overnight and daytime components, we construct directed networks to capture the lead-lag relations between overnight speculations and daytime price corrections, and vice versa, across stocks. We originate a novel clustering-based framework to construct portfolios to capture the cross-stock tug of war. In order to

identify disjoint leader and lagger groups in directed lead-lag networks, we develop a specialized spectral clustering algorithm. By generating trading signals exclusively from the leader stocks to predict and trade lagger stocks, we isolate pure cross-stock interactions from autocorrelation within individual stocks. Our empirical results support the conclusion that both noise traders and arbitrageurs trade at the portfolio level and disseminate the tug of war across stocks. With backtests spanning from 2000-01-03 to 2024-12-31, the cross-stock lead-lag portfolios generate remarkable returns and significant alphas on top of portfolios representing firm-level tug-of-war reversals and other the pricing factors. Moreover, the performance of cross-stock lead-lag portfolios grow in recent years, while the stock-specific reversals decay.

In Search of Seasonality in Intraday and Overnight Option

T. G. Bali, A. Goyal, M. Moerke, F. Weigert

Georgetown University, University of Lausanne, Swiss Finance Institute, University of St. Gallen, University of Neuchatel

We uncover momentum and reversal patterns in half-day option returns that persist for up to at least 20 business days, with economic magnitudes of 0.22% to 0.45% per half-day. Specifically, returns show strong momentum within the same period (e.g., intraday-to-intraday) but reverse sharply across opposite periods (e.g., intraday-toovernight). These patterns increase over time, are robust to various delta-hedging schemes and option selection criteria, and persist across different subsamples. Momentum and reversal strengthen when market makers actively manage capacity constraints during intraday-overnight transitions, indicating supply-side constraints drive predictability.

Global News Networks and Return Predictability

G. Freire, A. Moin, A. Quaini, A. Soebhag

Erasmus University Rotterdam, Robeco

We employ a novel, high-dimensional global news corpus (GDELT) to examine the effects of news-based sentiment on international equity market returns. Leveraging over 520 million articles from 14 developed countries, spanning 260 major themes and classified by source and target country, we construct granular sentiment linkages that capture how one country's media reports on another. Using machine learning, we construct out-of-sample market-timing strategies based on local and global sentiment measures and document significant economic gains. Further, global sentiment measures contain distinct, incremental information compared to local measures, indicating market integration via sensitivity to shared news information and sentiment.

Institutional Granular Impact Is Benign on Asset Sales and Price Efficiency

Y. Fan, G. Feng, X. Qiao, S. Baronyan

City University of Hong Kong, Emerging Portfolio Fund Research

We construct two types of trading shocks and examine their effects on stock prices. Common shocks capture the shared trading activity across funds, whereas granular idiosyncratic shocks place emphasis on large players. Common shocks related to stock sales exhibit a significantly stronger price impact than those related to purchases, in contrast to symmetric effects of purchases and sales for granular idiosyncratic shocks. The initial price impact persists in the short run and partially reverses after six months, suggesting underreaction to institutional trading. Our results underscore the impact of the common component across various funds on asset prices and market efficiency.

Option Pricing with Time-Varying Volatility Risk Aversion

P. R. Hansen, C. Tong

University of North Carolina at Chapel Hill, Xiamen University

We introduce a pricing kernel with time-varying volatility risk aversion to explain the observed time variations in the shape of the pricing kernel. When combined with the Heston-Nandi GARCH model, this framework yields a tractable option pricing model in which the variance risk ratio (VRR) emerges as a key variable. We show that the VRR is closely linked to economic fundamentals, as well as sentiment and uncertainty measures. A novel approximation method provides analytical option pricing formulas, and we demonstrate substantial reductions in pricing errors through an empirical application to the SP 500 index, the CBOE VIX, and option prices.

Artificial or Human Intelligence?

Eric Gao

Massachusetts Institute of Technology

Artificial intelligence (AI) tools such as large language models (LLMs) are already altering student learning. Unlike previous technologies, LLMs can independently solve problems regardless of student understanding, yet are not always accurate (due to hallucination) and face sharp performance cutoffs (due to emergence). Access to these tools significantly alters a student's incentives to learn, potentially decreasing the sum knowledge of humans and AI. Additionally, the marginal benefit of learning changes depending on which side of the AI frontier a human is on, creating a discontinuous gap

between those that know more than or less than AI. This contrasts with downstream models of AI's impact on the labor force which assume continuous ability. Finally, increasing the portion of assignments where AI cannot be used can counteract student mis-specification about AI accuracy, preventing underinvestment. A better understanding of how AI impacts learning and student incentives is crucial for educators to adapt to this new technology.

False Cascades and the Cost of Truth

G. Lukyanov, D. Cheredina

Toulouse School of Economics, International College of Economics and Finance

We study sequential social learning when agents can sometimes pay to verify a claim and obtain hard, publicly checkable evidence. Each agent observes the public history, receives a private signal, may investigate at a cost (succeeding only when the claim is true), and can disclose or conceal any proof. Actions are binary or continuous, with a conformity pull toward the prevailing consensus. We characterize when false cascades persist and when societies self-correct. In the binary benchmark, we derive an investigation cutoff and show how its location relative to classic cascade bands governs breakability; a simple knife-edge condition guarantees that any wrong cascade at the boundary is overturned with positive probability. With continuous actions, coarse observation and conformity can recreate cascades, yet occasional disclosures collapse them. These forces yield a tractable "resilience frontier" with transparent comparative statics and policy levers.

Mathematical and Computational Finance

Neural Actor-Critic Methods for Hamilton-Jacobi-Bellman PDEs: Asymptotic Analysis and Numerical Studies

S. N. Cohen, J. Hebner, D. Jiang, J. Sirignano

University of Oxford

We mathematically analyze and numerically study an actor-critic machine learning algorithm for solving high-dimensional Hamilton-Jacobi-Bellman (HJB) partial differential equations from stochastic control theory. The architecture of the critic (the estimator for the value function) is structured so that the boundary condition is always perfectly satisfied (rather than being included in the training loss) and utilizes a biased gradient which reduces computational cost. The actor (the estimator for the optimal control) is trained by minimizing the integral of the Hamiltonian over the domain, where the Hamiltonian is estimated using the critic. We show that the training dynamics of the actor and critic neural networks converge in a Sobolev-type space to a certain infinite-dimensional ordinary differential equation (ODE) as the number of hidden units in the actor and critic $\to \infty$. Further, under a convexity-like assumption on the Hamiltonian, we prove that any fixed point of this limit ODE is a solution of the original stochastic control problem. This provides an important guarantee for the algorithm's performance in light of the fact that finite-width neural networks may only converge to a local minimizers (and not optimal solutions) due to the non-convexity of their loss functions. In our numerical studies, we demonstrate that the algorithm can solve stochastic control problems accurately in up to 200 dimensions. In particular, we construct a series of increasingly complex stochastic control problems with known analytic solutions and study the algorithm's numerical performance on them. These problems range from a linear-quadratic regulator equation to highly challenging equations with non-convex Hamiltonians, allowing us to identify and analyze the strengths and limitations of this neural actor-critic method for solving HJB equations.

Functional Itô Formula and Taylor Expansions for Non-Anticipative Maps of Càdlàg Rough Paths

C. Cuchiero, X. Guo, F. Primavera

University of Vienna, University of California, Berkeley

We compare two methodologies for calibrating implied volatility surfaces: a second-order asymptotic We derive a functional Itô-formula for non-anticipative maps of rough paths, based on the approximation properties of the signature of càdlàg rough paths. This result is a functional extension of the Itô-formula for càdlàg rough paths (by Friz and Zhang (2018)), which coincides with the change of variable formula formulated by Dupire (2009) whenever the functionals' representations, the notions of regularity, and the integration concepts can be matched. Unlike these previous works, we treat the vertical (jump) pertubation via the Marcus transformation, which allows for incorporating path functionals where the second order vertical derivatives do not commute, as is the case for typical signature functionals. As a byproduct, we show that sufficiently regular non-anticipative maps admit a functional Taylor expansion in terms of the path's signature, leading to an important generalization of the recent results by Dupire and Tissot-Daguette (2022).

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

Andrew, Isaiah, Jiafeng Chen, and Otavio Tecchio. 2025. "The Purpose of an Estimator Is What It Does: Misspecification, Estimands, and Over-Identification." *arXiv preprint arXiv:2508.13076*.

Bali, Turan G, Amit Goyal, Mathis Moerke, and Florian Weigert. 2025. "In Search of Seasonality in Intraday and Overnight Option Returns." *Available at SSRN 5386128*.

Cavaliere, Giuseppe, Luca Fanelli, and Iliyan Georgiev. 2025. "Bootstrap Diagnostic Tests." *arXiv preprint arXiv:2509.01351*.

Cohen, Samuel N, Jackson Hebner, Deqing Jiang, and Justin Sirignano. 2025. "Neural Actor-Critic Methods for Hamilton-Jacobi-Bellman PDEs: Asymptotic Analysis and Numerical Studies." *arXiv preprint arXiv:2507.06428*.

Cuchiero, Christa, Xin Guo, and Francesca Primavera. 2025. "Functional Itô Formula and Taylor Expansions for Non-Anticipative Maps of Càdlàg Rough Paths." *arXiv preprint arXiv:2504.06164*.

Fan, Yinghua, Guanhao Feng, Xiao Qiao, and Sayad Baronyan. 2025. "Institutional Granular Impact Is Benign on Asset Sales and Price Efficiency." *Journal of Financial Markets*, forthcoming.

Freire, Gustavo, Ali Moin, Ali Moin, Alberto Quaini, and Amar Soebhag. 2025. "Global News Networks and Return Predictability." *Available at SSRN 5421496*.

Gao, Eric. 2025. "Artificial or Human Intelligence?" *arXiv preprint arXiv:2509.02879*.

Hansen, Peter Reinhard, and Chen Tong. 2025*a*. "Moments by Integrating the Moment-Generating Function." *arXiv preprint arXiv:2410.23587*.

Hansen, Peter Reinhard, and Chen Tong. 2025*b*. "Option Pricing with Time-Varying Volatility Risk Aversion." *Review of Financial Studies*, forthcoming.

Lukyanov, Georgy, and Darina Cheredina. 2025. "False cascades and the cost of truth." *arXiv preprint arXiv:2508.20538*.

Luo, Yuetian, and Chao Gao. 2024. "Adaptive Robust Confidence Intervals." *arXiv preprint arXiv:2410.22647*.

Lu, Yutong, Ning Zhang, Gesine Reinert, and Mihai Cucuringu. 2025. "A Tug of War across the Market: Overnight-vs.-Daytime Lead-Lag Networks and Clustering-Based Portfolio Strategies." *Available at SSRN 5371952*.

Medeiros, Marcelo C, and Jeronymo M Pinro. 2025. "Time Series Embedding and Combination of Forecasts: A Reinforcement Learning Approach." *arXiv preprint arXiv:2508.20795*.