VOL. 4, NO. 9

OMI Research Newsletter – June 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.*

Highlights

Key contributions in this month's edition: Cattaneo, Crump and Wang (2024) introduce a two-step nonparametric method for beta-sorted portfolio estimation with valid inference. Fortin, Gagliardini and Scaillet (2025) develop a pseudo-maximum likelihood framework for latent factor analysis in short panels, and derive an asymptotically uniformly most powerful test for the number of factors. Aït-Sahalia and Jacod (2025) propose a high-frequency specification test to determine whether volatility can be modelled as a local function of the underlying variable or requires a latent stochastic process. It empirically rejects local volatility for stock prices, interest rates, exchange rates, and volatility indices. Fan, Yan and Zheng (2025) develop a novel asymptotic framework to verify that the principal component analysis remains consistent and asymptotically normal under weak factor models. Mo and Ouyang (2025) synthesize the burgeoning literature on the intersection of (generative) AI and financial economics. Li et al. (2025) demonstrate that while machine-learning strategies built on a real-time "universe" of signals generate economically meaningful return forecasts, their performance is weaker than models using curated signal sets, highlighting the crucial role of feature engineering and inductive biases in investment strategies. Chen et al. (2025) propose a diffusion factor model that embeds an unknown low-dimensional factor structure to efficiently simulate high-dimensional asset returns, and demonstrate its effectiveness in portfolio construction.

Statistics, Econometrics & Machine Learning

Beta-Sorted Portfolios

M. D. CATTANEO, R. K. CRUMP, W. WANG

PRINCETON UNIVERSITY, FEDERAL RESERVE BANK OF NEW YORK, UNIVERSITY OF GRONINGEN

Beta-sorted portfolios—portfolios comprised of assets with similar covariation to selected risk factors—are a popular tool in empirical finance to analyze models of (conditional) expected returns. Despite their widespread use, little is known of their econometric properties in contrast to comparable procedures such as two-pass regressions. We formally investigate the properties of beta-sorted portfolio returns by casting the procedure as a two-step nonparametric estimator with a nonparametric first step and a beta-adaptive portfolios construction. Our framework rationalizes the well-known estimation algorithm with precise economic and statistical assumptions on the general data generating process. We provide conditions which ensure valid estimation and inference allowing for a range of hypotheses of interest in financial applications. We show that the rate of convergence of the estimator changes depending on the value of beta. We demonstrate that valid inference depends critically on the object of interest and discuss shortcomings of the widely-used Fama-MacBeth variance estimator. To address these limitations, we propose a new variance estimator. In an empirical application, we introduce a novel risk factor—a measure of the business credit cycle—and show that it is strongly predictive of both the cross-section and time-series behavior of U.S. stock returns.

Latent Factor Analysis in Short Panels

A.-P. FORTIN, P. GAGLIARDINI, O. SCAILLET

UNIVERSITY OF MONTREAL, UNIVERSITÀ DELLA SVIZZERA ITALIANA, SWISS FINANCE INSTITUTE, UNIVERSITY OF GENEVA

We develop inferential tools for latent factor analysis in

short panels. The pseudo maximum likelihood setting under a large cross-sectional dimension n and a fixed time series dimension T relies on a diagonal $T \times T$ covariance matrix of the errors without imposing sphericity or Gaussianity. We outline the asymptotic distributions of the latent factor and error covariance estimates as well as of an asymptotically uniformly most powerful invariant (AUMPI) test based on the likelihood ratio statistic for tests of the number of factors. We derive the AUMPI characterization from inequalities ensuring the monotone likelihood ratio property for positive definite quadratic forms in normal variables. An empirical application to a large panel of monthly U.S. stock returns separates date after date systematic and idiosyncratic risks in short subperiods of bear vs. bull market based on the selected number of factors. We observe an uptrend in idiosyncratic volatility while the systematic risk explains a large part of the cross-sectional total variance in bear markets but is not driven by a single factor. Rank tests reveal that observed factors struggle spanning latent factors with a discrepancy between the dimensions of the two factor spaces decreasing over time.

Testing Whether Volatility is Local or Stochastic

Y. Aït-Sahalia, J. Jacod
Princeton University, Sorbonne Université

There is broad empirical agreement that the volatility of most economic and financial variables is time-varying. One possible way to capture this in continuoustime models is to make the volatility its own latent process, leading to a stochastic volatility model. Another approach consists in making the volatility of a given variable a function of the variable itself, leading to a local volatility model. Local models are more restrictive, but also simpler and lower dimensional, and are often used in applications. This paper tests whether the restriction imposed by a local volatility model is empirically justified. The proposed test is based on high frequency asymptotics, is robust to the potential presence of jumps and/or noise in the data, and has power against a wide range of alternative non-local volatility models. The test is applied to stock price, interest rate, currency and volatility data and suggests that volatility is not local in all of them.

When Can Weak Latent Factors Be Statistically Inferred?

J. Fan, Y. Yan, Y. Zheng
Princeton University, University of Wisconsin-Madison

This article establishes a new and comprehensive estimation and inference theory for principal component analysis (PCA) under the weak factor model that allow for cross-sectional dependent idiosyncratic components under the nearly minimal factor strength relative to the noise level or signal-to-noise ratio. Our theory is applicable regardless of the relative growth rate between the cross-sectional dimension N and temporal dimension T. This more realistic assumption and noticeable result require completely new technical device, as the commonly-used leave-one-out trick is no longer applicable to the case with cross-sectional dependence. Another notable advancement of our theory is on PCA inference — for example, under the regime where $N \asymp T$, we show that the asymptotic normality for the PCA-based estimator holds as long as the signal-to-noise ratio (SNR) grows faster than a polynomial rate of $\log N$. This finding significantly surpasses prior work that required a polynomial rate of N. Our theory is entirely non-asymptotic, offering finite-sample characterizations for both the estimation error and the uncertainty level of statistical inference. A notable technical innovation is our closed-form first-order approximation of PCA-based estimator, which paves the way for various statistical tests. Furthermore, we apply our theories to design easy-to-implement statistics for validating whether given factors fall in the linear spans of unknown latent factors, testing structural breaks in the factor loadings for an individual unit, checking whether two units have the same risk exposures, and constructing confidence intervals for systematic risks. Our empirical studies uncover insightful correlations between our test results and economic cycles.

An Efficient Multi-Scale Leverage Effect Estimator Under Dependent Microstructure Noise

Z. Xiong, Z. Chen, C. D. Wang
Fudan University, NYU Shanghai

Estimating the leverage effect from high-frequency data is vital but challenged by complex, dependent microstructure noise, often exhibiting non-Gaussian higher-order moments. This paper introduces a novel multi-scale framework for efficient and robust leverage effect estimation under such flexible noise structures. We develop two new estimators, the Subsampling-and-Averaging Leverage Effect (SALE) and the Multi-Scale Leverage Effect (MSLE), which adapt subsampling and multi-scale approaches holistically using a unique shifted window technique. This design simplifies the multi-scale estimation procedure and enhances noise robustness without requiring the pre-averaging approach. We establish central limit theorems and stable convergence, with MSLE achieving convergence rates of an optimal $n^{-1/4}$ and a near-optimal $n^{-1/9}$ for the noise-free and noisy settings, respectively. A cornerstone of our framework's efficiency is a specifically designed MSLE weighting strategy that leverages covariance structures across scales. This significantly reduces asymptotic variance and, critically, yields substantially smaller finite-sample errors than existing methods under both noise-

free and realistic noisy settings. Extensive simulations and empirical analyses confirm the superior efficiency, robustness, and practical advantages of our approach.

Recurrent Self-Attention Dynamics: An Energy-Agnostic Perspective from Jacobians

A. Tomihari, R. Karakida
AIST Japan, University of Tokyo

The theoretical understanding of self-attention (SA) has been steadily progressing. A prominent line of work studies a class of SA layers that admit an energy function decreased by state updates. While it provides valuable insights into inherent biases in signal propagation, it often relies on idealized assumptions or additional constraints not necessarily present in standard SA. Thus, to broaden our understanding, this work aims to relax these energy constraints and provide an energy-agnostic characterization of inference dynamics by dynamical systems analysis. In more detail, we first consider relaxing the symmetry and single-head constraints traditionally required in energy-based formulations. Next, to investigate more general SA architectures capable of oscillatory dynamics without necessarily admitting an energy function, we analyze the Jacobian matrix of the state. We reveal that normalization layers effectively normalize the Jacobian's complex eigenvalues, forcing the dynamics close to a critical state. This significantly enhances inference performance. Furthermore, we utilize the Jacobian perspective to develop regularization methods for training and a pseudo-energy for monitoring inference dynamics.

Forecasting Intraday Volume in Equity Markets with Machine Learning

M. Cucuringu, K. Li, C. Zhang
University of California Los Angeles, University of Oxford, HKUST Guangzhou

This study focuses on forecasting intraday trading volumes, a crucial component for portfolio implementation, especially in high-frequency (HF) trading environments. Given the current scarcity of flexible methods in this area, we employ a suite of machine learning (ML) models enriched with numerous HF predictors to enhance the predictability of intraday trading volumes. Our findings reveal that intraday stock trading volume is highly predictable, especially with ML and considering commonality. Additionally, we assess the economic benefits of accurate volume forecasting through Volume Weighted Average Price (VWAP) strategies. The results demonstrate that precise intraday forecasting offers substantial advantages, providing valuable insights for traders to optimize their strategies.

Finance and Economics

(Generative) AI in Financial Economics

H. Mo, S. Ouyang
Saïd Business School - University of Oxford

This review article synthesizes the burgeoning literature on the intersection of (generative) artificial intelligence (AI) and financial economics. We organize our review around six key areas: (1) the emergent role of generative AI as analytic tools, external shocks to the economy, and autonomous economic agents; (2) corporate finance, focusing on how firms respond to and benefit from AI; (3) asset pricing, examining how AI brings novel methodologies for return predictability, stochastic discount factor estimation, and investment; (4) household finance, investigating how AI promotes financial inclusion and improves financial services; (5) labor economics, analyzing AI's impact on labor market dynamics; and (6) the risks and challenges associated with AI in financial markets. We conclude by identifying unanswered questions and discussing promising avenues for future research.

Machine Learning from a "Universe" of Signals: The Role of Feature Engineering

B. Li, A. G. Rossi, X. Yan, L. Zhang
Wuhan University, Georgetown University, Lehigh University, Renmin University of China

We construct real-time machine learning strategies based on a "universe" of fundamental signals. The out-of-sample performance of these strategies is economically meaningful and statistically significant, but considerably weaker than those documented by prior studies that use curated sets of signals as predictors. Strategies based on a simple recursive ranking of each signal's past performance also yield substantially better out-of-sample performance. We find qualitatively similar results when examining past-return-based signals. Our results underscore the key role of feature engineering and, more broadly, inductive biases in enhancing the economic benefits of machine learning investment strategies.

Renewable Asset Price Volatility and Its Implications for Decarbonization

H. Hong, J. D. Kubik, E. P. Shore
Columbia University, NBER, Syracuse University, University of New South Wales

We find that the price volatility of renewable assets is significantly greater than that of brown assets. Our causal estimates leverage the response of electricity and credit markets to US state-level renewable portfolio standards that require some utilities to use renewables while exempting others. This extra risk is related to

more volatile electricity prices and revenues, consistent with uncertainties including renewables intermittency. Using a growth model where the share of green capital balances climate damages and diversification benefits, we find that greater green-asset volatility is a more important determinant of economy-wide decarbonization than productivity differences of green versus brown capital.

Short-Term Basis Reversal

A. G. Rossi, Y. Zhang, Y. Zhu
Georgetown University, Peking University

We identify a previously undocumented form of return predictability in commodity futures markets, which we refer to as short-term basis reversal: The return spread between adjacent maturity contracts exhibits systematic negative autocorrelation. Basis reversal is independent of the short-term reversal of the individual contracts. Instead, it stems from the differential price sensitivity to news across the futures curve and is stronger among futures contracts with higher return volatility, more autocorrelated returns, and less correlated returns across maturities. Consistent with a preferred-habitat and limits to arbitrage interpretation, we show that basis reversal is also present in other assets characterized by a term structure, such as stock index futures, corporate bonds, and treasury bonds.

How Much Should You Pay for Alpha? Measuring the Value of Active Management with Utility Calculations

A. Ang, D. Basu
Columbia University, BlackRock

Not as much as you might think! We compute the amount an investor is willing to pay to access an active fund starting from an optimal equity-bond portfolio. We find that an investor optimally holding a 60/40 equity-bond portfolio would be willing to pay 19 bps, on average, to access active US equity funds over January 2016 to December 2024–with only 12% of these funds increasing the investor's utility. This compares with the average alpha of 108 bps of these funds conditional on their utility benefits being positive. In the case of active bond funds, the 60/40 investor is willing to pay only 10 bps, on average, conditional on them positively contributing to utility. The utility benefits do not increase significantly when we construct synthetic alpha funds which hedge out market exposures or increase the magnitude of alphas.

Cryptoeconomics

K. John, F. Saleh
New York University, University of Florida

Cryptoeconomics refers to the economic analysis of blockchain technology. This research area can be divided into two broad sub-fields. In that context, one sub-field is the study of economic mechanisms that seek to ensure successful transaction settlement on a blockchain. Such mechanisms are known as consensus protocols because they achieve transaction settlement by enabling all blockchain validators to form a consensus regarding which transactions have been added to the blockchain. More explicitly, if all validators agree that a transaction has been included in the blockchain and, further, if the validators will not deviate from this agreement in the future, then the transaction is deemed to have settled on the blockchain. The consensus protocol dictates both the incentives to generate the initial agreement regarding whether a transaction has been added to the blockchain and also the incentives for the initial agreement to persist indefinitely. The second broad sub-field within cryptoeconomics is the study of business applications that can be deployed on blockchain. Notably, when a blockchain is able to successfully generate transaction settlement, then the blockchain can serve as a commitment device for business policy. In particular, computer programs known as smart contracts can be deployed on a blockchain and a properly-functioning consensus protocol implies that the computer code within a given smart contract cannot be altered. To provide more detail, just as settlement of a peer-to-peer transaction implies that all validators indefinitely agree regarding the associated value transfer, settlement of a smart contract deployment implies that all validators indefinitely agree on the code within the smart contract. In turn, when the blockchain's consensus protocol is functioning properly, the deployment of a smart contract entails a commitment to the underlying computer code because the validators achieve consensus on that code. Importantly, within a business context, smart contracts can be used to commit to particular policies (e.g., pricing policy) by specifying those policies within smart contract code and then deploying to a blockchain. Prominent business applications deployed on blockchains include decentralized exchanges and decentralized lending protocols. Decentralized exchanges enable trading of assets via a blockchain without an intermediary. Decentralized lending protocols enable borrowing and lending of assets via a blockchain without an intermediary.

Reputational Cheap Talk Versus Prior Information

A. Vong
National University of Singapore

I study a setting where a sender communicates her information about a hidden state with a receiver who wants to match his action with that state. The sender strives to appear informed at the receiver's expense. I

characterize the receiver's tradeoff between relying on the sender or on prior information about the state when picking his action in all informative equilibria. I show that the receiver relies only on prior information, irrespective of how informed the sender may be, if and only if prior information is neither sufficiently noisy nor sufficiently informative. My results highlight a complementarity that informative prior information facilitates high-quality communication under reputation concern.

Endogenous Clustering and Analogy-Based Expectation Equilibrium

P. Jehiel, G. Weber
Paris School of Economics, University College London, University of Bologna

Normal-form two-player games are categorized by players into K analogy classes so as to minimize the prediction error about the behavior of the opponent. This results in Clustered Analogy-Based Expectation Equilibria in which strategies are analogy-based expectation equilibria given the analogy partitions and analogy partitions minimize the prediction errors given the strategies. We distinguish between environments with self-repelling analogy partitions in which some mixing over partitions is required and environments with self-attractive partitions in which several analogy partitions can arise, thereby suggesting new channels of belief heterogeneity and equilibrium multiplicity. Various economic applications are discussed.

Berk-Nash Rationalizability

I. Esponda, D. Pouzo
University of California, Santa Barbara, University of California, Berkeley

We introduce Berk–Nash rationalizability, a new solution concept for misspecified learning environments. It parallels rationalizability in games and captures all actions that are optimal given beliefs formed using the model that best fits the data in the agent's misspecified model class. Our main result shows that, with probability one, every *limit action* – any action played or approached infinitely often – is Berk–Nash rationalizable. This holds regardless of whether behavior converges. We apply the concept to known examples and identify classes of environments where it is easily characterized. The framework provides a general tool for bounding long-run behavior without assuming convergence to a Berk–Nash equilibrium.

Mathematical and Computational Finance

Diffusion Factor Models: Generating High-Dimensional Returns with Factor Structure

M. Chen, R. Xu, Y. Xu, R. Zhang
Princeton University, Sorbonne Université

Financial scenario simulation is essential for risk management and portfolio optimization, yet it remains challenging especially in high-dimensional and small data settings common in finance. We propose a diffusion factor model that integrates latent factor structure into generative diffusion processes, bridging econometrics with modern generative AI to address the challenges of the curse of dimensionality and data scarcity in financial simulation. By exploiting the low-dimensional factor structure inherent in asset returns, we decompose the score function–a key component in diffusion models–using time-varying orthogonal projections, and this decomposition is incorporated into the design of neural network architectures. We derive rigorous statistical guarantees, establishing nonasymptotic error bounds for both score estimation at $O(d^{5/2}n^{-2/(k+5)})$ and generated distribution at $O(d^{5/4}n^{-1/2(k+5)})$, primarily driven by the intrinsic factor dimension k rather than the number of assets d, surpassing the dimension-dependent limits in the classical nonparametric statistics literature and making the framework viable for markets with thousands of assets. Numerical studies confirm superior performance in latent subspace recovery under small data regimes. Empirical analysis demonstrates the economic significance of our framework in constructing mean-variance optimal portfolios and factor portfolios. This work presents the first theoretical integration of factor structure with diffusion models, offering a principled approach for high-dimensional financial simulation with limited data.

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: Álvaro Arroyo, Patrick Chang, Fayçal Drissi, Gerardo Durán-Martín, and Qi Jin.

References

Ang, Andrew, and Debarshi Basu. 2025. "How Much Should You Pay for Alpha? Measuring the Value of Active Management with Utility Calculations." *SSRN Working Paper*.

Aït-Sahalia, Yacine, and Jean Jacod. 2025. "Testing Whether Volatility is Local or Stochastic." *SSRN Working Paper*.

Cattaneo, Matias D., Richard K. Crump, and Weining Wang. 2024. "Beta-Sorted Portfolios." *SSRN Working Paper*.

Chen, Minshuo, Renyuan Xu, Yumin Xu, and Ruixun Zhang. 2025. "Diffusion Factor Models: Generating High-Dimensional Returns with Factor Structure." *arXiv preprint arXiv:2504.06566*.

Cucuringu, Mihai, Kang Li, and Chao Zhang. 2025. "Forecasting Intraday Volume in Equity Markets with Machine Learning." *arXiv preprint arXiv:2505.08180*.

Esponda, Ignacio, and Demian Pouzo. 2025. "Berk-Nash Rationalizability." *arXiv preprint arXiv:2505.20708*.

Fan, Jianqing, Yuling Yan, and Yuheng Zheng. 2025. "When Can Weak Latent Factors Be Statistically Inferred?" *SSRN Working Paper*.

Fortin, Alain-Philippe, Patrick Gagliardini, and Olivier Scaillet. 2025. "Latent Factor Analysis in Short Panels." *SSRN Working Paper*.

Hong, Harrison G., Jeffrey D. Kubik, and Edward P. Shore. 2025. "Renewable Asset Price Volatility and Its Implications for Decarbonization." *SSRN Working Paper*.

Jehiel, Philippe, and Giacomo Weber. 2025. "Endogenous clustering and analogy-based expectation equilibrium." *arXiv preprint arXiv:2505.13022*.

John, Kose, and Fahad Saleh. 2025. "Cryptoeconomics." *SSRN Working Paper*.

Li, Bin, Alberto G. Rossi, Xuemin Yan, and Lingling Zheng. 2025. "Machine Learning from a "Universe" of Signals: The Role of Feature Engineering." *SSRN Working Paper*.

Mo, Hongwei, and Shumiao Ouyang. 2025. "(Generative) AI in Financial Economics." *SSRN Working Paper*.

Rossi, Alberto G., Yingguang Zhang, and Yandi Zhu. 2025. "Short-Term Basis Reversal." *SSRN Working Paper*.

Tomihari, Akiyoshi, and Ryo Karakida. 2025. "Recurrent Self-Attention Dynamics: An Energy-Agnostic Perspective from Jacobians." *arXiv preprint arXiv:2505.19458*.

Vong, Allen. 2025. "Reputational Cheap Talk Versus Prior Information." *arXiv preprint arXiv:2505.11877*.

Xiong, Ziyang, Zhao Chen, and Christina Dan Wang. 2025. "An Efficient Multi-Scale Leverage Effect Estimator Under Dependent Microstructure Noise." *arXiv preprint arXiv:2505.08654*.