Quantitative Finance Research Newsletter
Oxford-Man Institute
OMIReNew



The Oxford-Man Institute of Quantitative Finance is a world-leading centre for interdisciplinary research in financial markets. This newsletter is an eclectic summary that covers the most recent research of relevance to quantitative finance, from a wide range of disciplines in academia and industry.

***This issue**: Machine Learning for Finance, Graph-based Learning, Derivatives & Volatility, Quantitative Finance & Economics, and Digital Assets.*

Highlights

Noteworthy works in this month's edition include: two works from University of Oxford - one considering double-execution strategies using path signatures, and another considering transfer ranking applied to cross-sectional momentum with data scarcity; a work from erstwhile governors of the Reserve Bank of India, which presents a quantitative framework examining sovereign debt dynamics; and a work from MIT constructing optimal impact portfolios. Finally, readers might also be interested in an intriguingly-idiosyncratic work from Oakland University, linking the human brain - presented as a Bayesian prediction engine and pricing of out-of-the-money put options.

Machine Learning for Finance

Transfer Ranking in Finance: Applications to Cross-Sectional Momentum with Data Scarcity

D. POH, S. ROBERTS, AND S. ZOHREN
UNIVERSITY OF OXFORD

Cross-sectional strategies are a classical and popular trading style, with recent high performing variants incorporating sophisticated neural architectures. While these strategies have been applied successfully to data-rich settings involving mature assets with long histories, deploying them on instruments with limited samples generally produce over-fitted models with degraded performance. In this paper, they introduce Fused Encoder Networks – a novel and hybrid parameter-sharing transfer ranking model. The model fuses information extracted using an encoder-attention module operated on a source dataset with a similar but separate module focused on a smaller target dataset of interest. This mitigates the issue of models with poor generalisability that are a consequence of training on scarce target data. Additionally, the self-attention mechanism enables interactions among instruments to be accounted for, not just at the loss level during model training, but also at inference time. Focusing on momentum applied to the top ten cryptocurrencies by market capitalisation as a demonstrative use-case, the Fused Encoder Networks outperforms the reference benchmarks on most performance measures, delivering a three-fold boost in the Sharpe ratio over classical momentum as well as an improvement of approximately 50% against the best benchmark model without transaction costs. It continues outperforming baselines even after accounting for the high transaction costs associated with trading cryptocurrencies.

A Simple Learning Agent Interacting with an Agent-based Market Model

M. DICKS, AND T. GEBBIE
UNIVERSITY OF CAPE TOWN

This work considers the learning dynamics of a single reinforcement learning optimal execution trading agent when it interacts with an event driven agent-based financial market model. Trading takes place asynchronously through a matching engine in event time. The optimal

execution agent is considered at different levels of initial order-sizes and differently sized state spaces. The resulting impact on the agent-based model and market are considered using a calibration approach that explores changes in the empirical stylised facts and price impact curves. Convergence, volume trajectory and action trace plots are used to visualise the learning dynamics. Here the smaller state space agents had the number of states they visited converge much faster than the larger state space agents, and they were able to start learning to trade intuitively using the spread and volume states. They find that the moments of the model are robust to the impact of the learning agents except for the Hurst exponent, which was lowered by the introduction of strategic order-splitting. The introduction of the learning agent preserves the shape of the price impact curves but can reduce the trade-sign auto-correlations when their trading volumes increase.

Limited Attention and Dynamically Distorted Beliefs

N. Schweizer, and M. Vries
Tilburg School of Economics and Management & Vrije Uni. Amsterdam

This work studies the effects of limited attention on belief formation in a dynamic environment. They propose a simple theoretical model of inattentive learning and derive closed-form expressions for the limiting distortion of the true underlying distribution. When attention is drawn towards extreme observations, the limiting probability distortion takes the form of an S-shape. These belief distortions yield behavior that is similar to that implied by the preference-based distortions of prospect theory. They thus find that part of the empirical success of prospect theory might be due to abstracting away the impact of limited attention on learning. They show that their model of inattentive learning explains several aspects of investors behavior, for instance, a preference for skewed assets; overextrapolation of past returns; stock market momentum; and the simultaneous demand for lotteries and insurance. When attention has cognitive costs, limited attention thus provides a rational unification of several biases in investor behavior.

Graph-based Learning

Stock Trend Forecasting Based on Multi-Information Fusion with Deep Learning

Y. Yao, Z. Zhang, Y. Zhao, and C. Li
Henan University

Stock forecasting is the practice of using past data to predict future stock prices. Stock forecasting is effective in assisting financial regulators and investors to make informed decisions. Most previous studies have used the time series of a single stock to forecast. In the capital market, a change in other stock prices also affects a change in the price of a single stock. Therefore, information about a single stock may lead to inaccurate stock price trend prediction. This study proposes a stock price trend forecasting model based on multi-information fusion. This model uses the Graph Convolutional Network (GCN) to extract the spatial features of information affecting the price change of the target stock and uses the Convolutional Neural Network (CNN) to extract the spatial features of the target stock information. Then, the model combines the extracted two spatial features and the Gated Recurrent Unit (GRU) to extract the temporal features and forecast stock price. Finally, the annualized rate of return, classification accuracy, and other related indicators are used for evaluation in the experimental analysis. The results show that the proposed Graph Convolutional Incorporating GRU (GCN-CNN-GRU) hybrid neural network forecasting model offers benefits for estimating the stock price trend.

The Network of Commodity Risk

B. Foroni, G. Morelli, and L. Petrealla
Sapienza University of Rome

This paper investigates the interconnections among and within the Energy, Agricultural, and Metal commodities, operating in a risk management framework with a twofold goal. First, they estimate the Value-at-Risk (VaR) employing GARCH and Markov-switching GARCH models with different error term distributions. The use of such models allows us to take into account well-known stylized facts shown in the time series of commodities as well as possible regime changes in their conditional variance dynamics. They rely on backtesting procedures to select the best model for each commodity. Second, they estimate the sparse Gaussian Graphical model of commodities exploiting the Graphical LASSO (GLASSO) methodology to detect the most relevant conditional dependence structure among and within the sectors. A novel feature of their framework is that GLASSO estimation is achieved exploring the precision matrix of the multivariate Gaussian distribution obtained using a Gaussian copula with marginals given by the residuals of the aforementioned selected models. They apply their approach to the sample of twenty-four series of commodity futures prices over the years 2005–2022. They find that Soybean Oil, Cotton, and Coffee represent the major sources of propagation of financial distress in commodity markets while Gold, Natural Gas UK, and Heating Oil are depicted as safe-haven commodities. They further find that connectedness does not always increase in response to crisis events.

Multilayer Financial Complex Networks and Their Applications

X. Li, X. Xu, J. Liu, J. Dong, and J. Lu

Chinese Academy of Sciences

In the context of global integration, the theory of financial complex networks has made significant contributions to the establishment of stable financial systems and effective regulatory systems. This survey paper presents a systematic methodology for the multi-layer financial complex networks and their applications. Several typical financial networks in existing research are first summarized: the interbank network, the credit network, the international trade network, and the dealer network. The main methods used in the existing literature to analyze the structure of financial networks include maximum entropy methods, network characterization, community detection, and dynamic multi-layer network. Financial risk contagion as well as spillover effects are the core issues that most of the literature focuses on. Finally, this paper reviews the shortcomings of the existing literature and suggests future research directions in this area.

Derivatives & Volatility

Robust Estimation of Integrated and Spot Volatility

Z. Li, and O. Linton

Chinese Uni. of Hong Kong & Uni. of Cambridge

This work introduces a new method to estimate the integrated volatility (IV) and the spot volatility (SV) based on noisy high-frequency data. Their method employs the ReMeDI approach introduced by Li and Linton (2022, Econometrica) to estimate the moments of microstructure noise and thereby eliminate their influence, and the pre-averaging method to target the volatility parameter. The method is robust: it can be applied when the efficient price exhibits stochastic volatility and jumps, the observation times are random, and the noise process is nonstationary, autocorrelated, asymptotically vanishing and dependent on the efficient price. They derive the limit distributions for the proposed estimators under the infill asymptotics in a general setting. Their extensive simulation studies demonstrate the robustness, accuracy and computational efficiency of their estimators compared to several alternative estimators recently proposed in the literature. Empirically, they show that neglecting the complexities of noise and the random observation times yields substantial biases in volatility estimation and may lead to a different intraday volatility pattern.

Regime-Based Implied Stochastic Volatility Model for Crypto Option Pricing

D. Saef, Y. Wang, and T. Aste

University College London

The increasing adoption of Digital Assets (DAs), such as Bitcoin (BTC), rises the need for accurate option pricing models. Yet, existing methodologies fail to cope with the volatile nature of the emerging DAs. Many models have been proposed to address the unorthodox market dynamics and frequent disruptions in the microstructure caused by the non-stationarity, and peculiar statistics, in DA markets. However, they are either prone to the curse of dimensionality, as additional complexity is required to employ traditional theories, or they overfit historical patterns that may never repeat. Instead, they leverage recent advances in market regime (MR) clustering with the Implied Stochastic Volatility Model (ISVM). Time-regime clustering is a temporal clustering method, that clusters the historic evolution of a market into different volatility periods accounting for non-stationarity. ISVM can incorporate investor expectations in each of the sentiment-driven periods by using implied volatility (IV) data. In this paper, they applied this integrated time-regime clustering and ISVM method (termed MR-ISVM) to high-frequency data on BTC options at the popular trading platform Deribit. They demonstrate that MR-ISVM contributes to overcome the burden of complex adaption to jumps in higher order characteristics of option pricing models. This allows us to price the market based on the expectations of its participants in an adaptive fashion.

Quantitative Finance & Economics

Inflation as a Fiscal Limit

F. Bianchi, and L. Melosi

John Hopkins Uni. & FRB of Chicago

Low and stable inflation requires an appropriate fiscal framework aimed at stabilizing government debt. Historically, trend inflation is critically influenced by actual or perceived changes to this framework, while cost-push shocks only account for short-lasting movements in inflation. Before the pandemic, a moderate level of fiscal inflation has counteracted deflationary pressures, helping the central bank to avoid deflation. The recent fiscal interventions in response to the COVID pandemic have altered the private sector's beliefs about the fiscal framework, accelerating the recovery, but also determining an increase in fiscal inflation. This increase in inflation could not have been averted by simply tightening monetary policy. The conquest of post-pandemic inflation requires mutually consistent monetary and fiscal policies to avoid fiscal stagflation.

Risk allocation through Shapley decompositions with applications to variable annuities

F. Godin, E. Hamel, P. Gaillardetz, and E. Ng
Concordia University & Aviva

This paper introduces a flexible risk decomposition method for life insurance contracts embedding several risk factors. Hedging can be naturally embedded in the framework. Although the method is applied to variable annuities in this work, it is also applicable in general to other insurance or financial contracts. The approach relies on applying an allocation principle to components of a Shapley decomposition of the gain and loss. The implementation of the allocation method requires a stochastic on stochastic algorithm. Numerical examples studying the relative impact of equity, interest rate and mortality risk for guaranteed minimal maturity benefit (GMMB) policies conclude out analysis.

Risk and State-Dependent Financial Frictions

M. Harding, and R. Wouters
Bank of Canada

This work augments a standard New Keynesian model with a financial accelerator mechanism and show that financial frictions generate large state-dependent amplification effects. They fit the model to US data and show that, when shocks drive the model far away from the steady state, the nonlinear model produces much stronger propagation of shocks than the linearized model. They document that these amplification effects are due to endogenous variation in financial conditions and not due to other nonlinearities in the model. Motivated by these findings, they propose a regime-switching dynamic stochastic general equilibrium framework where financial frictions endogenously fluctuate between moderate (low risk) and severe (high risk), depending on the state of the economy. This framework allows for efficient estimation with many state variables and improves fit with respect to the linear model.

Shrinking the Term Structure

D. Filipovic, M. Pelger, and Y. Ye
Swiss Finance Institute & Stanford University

This work develops a conditional factor model for the term structure of treasury bonds, which unifies non-parametric curve estimation with cross-sectional asset pricing. Their factors correspond to the optimal non-parametric basis functions spanning the discount curve. They are investable portfolios estimated with cross-sectional ridge regressions and derived from economic first principles. Empirically, they show that four factors explain the discount bond excess return curve and term structure premium. Cash flows are covariances, as cash flows of coupon bonds fully explain the factor exposure. The term structure premium depends on the market complexity measured by the time-varying importance of higher order factors.

Double-Execution Strategies using Path Signatures

A. Cartea, I. Arribas, and L. Sanchez-Betancourt
University of Oxford & Kings College London

This work employs the expected signature of equity and foreign exchange markets to derive an optimal double-execution trading strategy. The signature of a path of a stochastic process is a sequence of real numbers that provides a full description of the evolution of the process. The double-execution strategy maximises the wealth (in units of the domestic currency) of an investor who liquidates a block of shares in a foreign stock market. Their approach is model agnostic because they do not specify the dynamics of the market. They prove that the optimal strategy is a linear combination of the terms in the expected signature of the market and employ high-frequency data from Nasdaq and for various currencies to compute the signature of the market. Data for ten stocks and four currency pairs are employed to implement the strategy. Their results show that the performance of the signature-based double-execution strategy is superior than the performance of the benchmarks. In most cases, the outperformance increases further when the signature of the market is enhanced with the price dynamics of the SPY; a tracker of the Standard Poor's 500 index.

When is Sovereign Debt Odious? A Theory of Government Repression, Growth Traps, and Growth Boosts

V. Acharya, R. Rajan, and J. Shim
New York Uni., NBER & Uni. of Chicago

This work examines the dynamics of a country's growth, consumption, and sovereign debt, assuming that the government is myopic and wants to maximize short-term, self-interested spending. Surprisingly, government myopia can increase a country's access to external borrowing. In turn, access to borrowing can extend the government's effective horizon; the government's ability to borrow hinges on it convincing creditors they will be repaid, which gives it a stake in generating future revenues. In a high-saving country, the lengthening of the government's effective horizon can incentivize it to tax less, resulting in a 'growth boost', with higher steady-state household consumption than if it could not borrow. However, in a country that saves little, the government may engage in more repressive policies to enhance its debt capacity. This could lead to a 'growth trap' where household steady-state consumption is lower than if

the government had no access to debt. They discuss the effectiveness of alternative debt policies, including declaring debt odious, debt forgiveness, and debt ceilings. They also analyse the impact of unanticipated shocks on the country's welfare.

Optimal Impact Portfolios with General Dependence and Marginals

A. Lo, L. Wu, R. Zhang, and C. Zhao
MIT & Peking University

Impact investing typically involves ranking and selecting assets based on a non-financial impact factor, such as the environmental, social, and governance (ESG) score and the prospect of developing a disease-curing drug. They develop a framework for constructing optimal impact portfolios and quantifying their financial performances. Under general bivariate distributions of the impact factor and residual returns from a multi-factor asset-pricing model, the construction and performance of optimal impact portfolios depend critically on the dependence structure (copula) between the two, which reduces to a correlation under normality assumptions. More generally, they explicitly derive the optimal portfolio weights under two widely-used copulas—the Gaussian copula and the Archimedean copula family, and find that the optimal weights depend on the tail characteristics of the copula. In addition, when the marginal distribution of residual returns is skewed or heavy-tailed, assets with the most extreme impact factors have lower weights than non-extreme assets due to their high risk. Their framework requires the estimation of only a constant number of parameters as the number of assets grow, an advantage over traditional Markowitz portfolios. Overall, these results provide a recipe for constructing and quantifying the performance of optimal impact portfolios with arbitrary dependence structures and return distributions.

Digital Assets

Automated Market Making and Loss-Versus-Rebalancing

J. Milionis, C. Moallemi, T. Roughgarden, A. Zhang
Columbia University & University of Chicago

This work considers the market microstructure of automated market making and, specifically, constant function market makers (CFMMs), from the economic perspective of passive liquidity providers (LPs). In a frictionless, continuous-time Black-Scholes setting and in the absence of trading fees, they decompose the return of an LP into a instantaneous market risk component and a non-negative, non-decreasing, and predictable component which they call "loss-versus-rebalancing" (LVR, pronounced "lever"). Market risk can be fully hedged, but once eliminated, LVR remains as a running cost that must be offset by trading fee income in order for liquidity provision to be profitable. They show how LVR can be interpreted in many ways: as the cost of pre-commitment, as the time value for giving up future optionality, as the compensator in a Doob-Meyer decomposition, as an adverse selection cost in the form the profits of arbitrageurs trading against the pool, and as an information cost because the pool does not have access to accurate market prices. LVR is distinct from the more commonly known metric of "impermanent loss" or "divergence loss"; this latter metric is more fundamentally described as "loss-versus-holding" and is not a true running cost. They express LVR simply and in closed-form: instantaneously, it is is the scaled product of the variance of prices and the marginal liquidity available in the pool, i.e., LVR the floating leg of a generalized variance swap. As such, LVR is easily calibrated to market data and specific CFMM structure. LVR provides tradeable insight in both the ex ante and ex post assessment of CFMM LP investment decisions, and can also inform the design of CFMM protocols.

Acknowledgement

This newsletter was curated by Utkarsh Sharma, D.Phil. Candidate, and Prof. Álvaro Cartea, Director, at the Oxford-Man Institute, with paper suggestions from fellow students at the institute - Dragos Gorduza, Marcello Monga, and Patrick Chang.