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, Risk Management, Quantitative Finance & Economics, Graph-based Learning, Derivatives & Volatility, Risk Management, Cryptocurrencies & Blockchain.*

Highlights

Noteworthy works in this month's edition include: a work from Luasanne, Athens and Sorborne proposing geometric tools to detect low-volatility anomalies in the stock market; another from UCL pricing options on overnight futures, which required convexity adjustment with an Asian flavour; works from OMI researchers studying the impact of tick size on algorithmic collusion in electronic markets and the provision of AI driven liquidity in OTC markets; and the work from Amundi discussing novel metrics to measure carbon emissions.

Machine Learning for Finance

A Machine Learning Approach to Portfolio Pricing and Risk Management for High-dimensional Problems.

L. Fernandez-Arjona, and D. Filipovic
University of Zurich

They present a general framework for portfolio risk management in discrete time, based on a replicating martingale. This martingale is learned from a finite sample in a supervised setting. Their method learns the features necessary for an effective low-dimensional representation, overcoming the curse of dimensionality common to function approximation in high-dimensional spaces, and applies for a wide range of model distributions. They show numerical results based on polynomial and neural network bases applied to high-dimensional Gaussian models. In these examples, both bases offer superior results to naive Monte Carlo methods and regress-now least squares Monte Carlo.

Dynamic and Context-Dependent Stock Price Prediction Using Attention Modules and News Sentiment.

N. Königstein
Impactvise

The growth of machine-readable data in finance, such as alternative data, requires new modeling techniques that can handle non-stationary and non-parametric data. Due to the underlying causal dependence and the size and complexity of the data, they propose a new modeling approach for financial time series data, the t-RIM (recurrent independent mechanism). This architecture makes use of key-value attention to integrate top-down and bottom-up information in a context-dependent and dynamic way. To model the data in such a dynamic manner, the α_t-RIM utilises an exponentially smoothed recurrent neural network, which can model non-stationary times series data, combined with a modular and independent recurrent structure. They apply their approach to the closing prices of three selected stocks of the SP 500 universe as well as their news sentiment score. The results suggest that the α_t-RIM is capable of reflecting the causal structure between stock prices and news sentiment, as well as the seasonality and

trends. Consequently, this modeling approach markedly improves the generalization performance, that is, the prediction of unseen data, and outperforms state-of-the-art networks such as long short-term memory models.

QuantNet: Transferring Learning Across Trading Strategies.

A. Koshiyama, S. Blumberg, N. Firoozye, P. Treleaven, and S. Flennerhag

UCL & University of Manchester

Systematic financial trading strategies account for over 80% of trade volume in equities and a large chunk of the foreign exchange market. In spite of the availability of data from multiple markets, current approaches in trading rely mainly on learning trading strategies per individual market. In this paper, they take a step towards developing fully end-to-end global trading strategies that leverage systematic trends to produce superior market-specific trading strategies. They introduce QuantNet: an architecture that learns market-agnostic trends and use these to learn superior market-specific trading strategies. Each market-specific model is composed of an encoder-decoder pair. The encoder transforms market-specific data into an abstract latent representation that is processed by a global model shared by all markets, while the decoder learns a market-specific trading strategy based on both local and global information from the market-specific encoder and the global model. QuantNet uses recent advances in transfer and meta-learning, where market-specific parameters are free to specialize on the problem at hand, whilst market-agnostic parameters are driven to capture signals from all markets. By integrating over idiosyncratic market data they can learn general transferable dynamics, avoiding the problem of overfitting to produce strategies with superior returns. They evaluate QuantNet on historical data across 3103 assets in 58 global equity markets. Against the top performing baseline, QuantNet yielded 51% higher Sharpe and 69% Calmar ratios. In addition, they show the benefits of their approach over the non-transfer learning variant, with improvements of 15% and 41% in Sharpe and Calmar ratios.

Are Missing Values Important for Earnings Forecasts? A Machine Learning Perspective.

A. Uddin, X. Tao, C. Chou, and D. Yu

New Jersey Institute of Technology

Analysts' forecasts are one of the most common and important estimators for firms' future earnings. However, they are challenging to fully utilize because of missing values. This study applies machine learning techniques to estimate missing values in individual analysts' forecasts and subsequently to predict firms' future earnings based on both estimated and observed forecasts. After estimating missing values, forecast error is reduced by 41% compared to the mean forecast, suggesting that missing values after estimating are indeed useful for earnings forecasts. They analyse multiple estimation methods and show that the out-performance of matrix factorization (MF) is consistent using different evaluation measures and across firms. Finally, they propose a stochastic gradient descent based coupled matrix factorization (CMF) to augment the estimation quality of missing values with multiple datasets. CMF further reduces the error of earnings forecasts by 19% compared to MF with a single dataset.

What is the Value of the Cross-sectional Approach to Deep Reinforcement Learning?

A. Aboussalah, Z. Cu, and C. Lee

University of Toronto

Reinforcement learning (RL) for dynamic asset allocation is an emerging field of study. Total return, the common performance metric, is useful for comparing algorithms but does not help determine how close an RL algorithm is to an optimal solution. In real-world financial applications, a bad decision could prove to be fatal. One of the key ideas of their work is to combine the two paradigms of the mean-variance optimization approach (Markowitz criteria) and the optimal capital growth approach (Kelly criteria) via the actor-critic approach. By using an actor-critic approach, they can balance optimization of risk and growth by configuring the actor to optimize the mean-variance while the critic is configured to maximize growth. They propose a Geometric Policy Score used by the critic to assess the quality of the actions taken by the actor. This could allow portfolio manager practitioners to better understand the investment RL policy. They present an extensive and in-depth study of RL algorithms for use in portfolio management (PM). They studied eight published policy-based RL algorithms which are preferred over value-based RL because they are better suited for continuous action spaces and are considered to be state of the art, Deterministic Policy Gradient (DPG), Stochastic Policy Gradients (SPG), Deep Deterministic Policy Gradient (DDPG), Trust Region Policy Optimization (TRPO), Proximal Policy Optimization (PPO), Twin Delayed Deep Deterministic Policy Gradient (TD3), Soft Actor Critic (SAC), and Evolution Strategies (ES) for Policy Optimization. They implemented all eight and they were able to modify all of them for PM but their initial testing determined that their performance was not satisfactory. Most algorithms showed difficulty converging during the training process due to the non-stationary and noisy nature of financial environments, along with other challenges. They selected the four most promising algorithms DPG, SPG, DDPG, PPO for further improvements. The modification of RL algorithms to finance required unconventional changes.

They have developed a novel approach for encoding multi-type financial data in a way that is compatible with RL. They use a multi-channel convolutional neural network (CNN-RL) framework, where each channel corresponds to a specific type of data such as high-low-open-close prices and volumes. They also designed a reward function based on concepts such as alpha, beta, and diversification that are financially meaningful while still being learnable by RL. In addition, portfolio managers will typically use a blend of time series analysis and cross-sectional analysis before making a decision. They extend their approach to incorporate, for the first time, cross-sectional deep RL in addition to time series RL. Finally, they demonstrate the performance of the RL agents and benchmark them against commonly used passive and active trading strategies, such as the uniform buy-and-hold (UBAH) index and the dynamical multi-period Mean-Variance-Optimization (MVO) model.

Machine Learning Techniques in Joint Default Assessment

M. Doria, E. Luciano, and P. Semeraro
Credit Suisse & University of Torino

This paper studies the consequences of capturing non-linear dependence among the covariates that drive the default of different obligors and the overall riskiness of their credit portfolio. Joint default modeling is, without loss of generality, the classical Bernoulli mixture model. Using an application to a credit card dataset they show that, even when Machine Learning techniques perform only slightly better than Logistic Regression in classifying individual defaults as a function of the covariates, they do outperform it at the portfolio level. This happens because they capture linear and non-linear dependence among the covariates, whereas Logistic Regression only captures linear dependence. The ability of Machine Learning methods to capture non-linear dependence among the covariates produces higher default correlation compared with Logistic Regression. As a consequence, on their data, Logistic Regression underestimates the riskiness of the credit portfolio.

Risk Management

Randomized Geometric Tools for Anomaly Detection in Stock Markets.

C. Bachelard, A. Chalkis, V. Fisikopoulos, and E. Tsigaridas
University of Luasanne, University of Athens & Sorbonne University and Paris University

They propose novel randomized geometric tools to detect low-volatility anomalies in stock markets; a principal problem in financial economics. Their modeling of the (detection) problem results in sampling and estimating the (relative) volume of geodesically non-convex and non-connected spherical patches that arise by intersecting a non-standard simplex with a sphere. To sample, they introduce two novel Markov Chain Monte Carlo (MCMC) algorithms that exploit the geometry of the problem and employ state-of-the-art continuous geometric random walks (such as Billiard walk and Hit-and-Run) adapted on spherical patches. To their knowledge, this is the first geometric formulation and MCMC-based analysis of the volatility puzzle in stock markets. Ther analyses provide accurate detection and new insights into the distribution of portfolios' performance characteristics. Moreover, they use their tools to show that classical methods for low-volatility anomaly detection in finance form bad proxies that could lead to misleading or inaccurate results.

Market Impact: Empirical Evidence, Theory and Practice.

E.Said
University of Montreal & Abu Dhabhi Investment Authority

They propose a theory of the market impact of metaorders based on a coarse-grained approach where the microscopic details of supply and demand is replaced by a single parameter $\in [0, +\infty]$ shaping the supply-demand equilibrium and the market impact process during the execution of the metaorder. Their model provides an unified explanation of most of the empirical observations that have been reported and establishes a strong connection between the excess volatility puzzle and the order-driven view of the markets through the square-root law.

Volatility Sensitive Bayesian Estimation of Portfolio VaR and CVaR.

T. Bodnar, V. Niklasson, and E. Thorsén
Stockholm University

In this paper, a new way to integrate volatility information for estimating value at risk (VaR) and conditional value at risk (CVaR) of a portfolio is suggested. The new method is developed from the perspective of Bayesian statistics and it is based on the idea of volatility clustering. By specifying the hyperparameters in a conjugate prior based on two different rolling window sizes, it is possible to quickly adapt to changes in volatility and automatically specify the degree of certainty in the prior. This constitutes an advantage in comparison to existing Bayesian methods that are less sensitive to such changes in volatilities and also usually lack standardized ways of expressing the degree of belief. They illustrate their new approach using both simulated and empirical data. Compared to some other well known homoscedastic and heteroscedastic models, the new method provides

a good alternative for risk estimation, especially during turbulent periods where it can quickly adapt to changing market conditions.

Derivatives & Volatility

Evaluation of Deep Learning Algorithms for Quadratic Hedging.

Z. Dai, L. Li, and G. Zhang

The Chinese University of Hong Kong

They solve the quadratic hedging problem by deep learning in discrete time. They consider three deep learning algorithms corresponding to three architectures of neural network approximation: approximating controls of different periods by different feedforward neural networks (FNNs) as proposed by Han(2016), using a single FNN with decision time as an input to approximate controls of different periods, and using a recursive neural network (RNN) to utilize historical information. They evaluate these algorithms under the discrete-time Black-Scholes model and the DCC-GARCH model for hedging basket options on portfolios of up to 100 assets with time to maturity up to one year. They compare them in terms of their hedging error on the test data, the extent of overlearning, the learned hedging strategy, training speed and scalability. Their results favour the single FNN and RNN approximations overall while the multiple FNN approximation can fail for a large portfolio and a long maturity. They also evaluate the performance of the single FNN and RNN algorithms in a data-driven framework, where data is generated by resampling without assuming any parametric model.

Options on Overnight Futures.

M. Henrad

UCL

With the transitions to overnight benchmarks as the main benchmarks in some currencies, futures based on overnight rates are becoming more common. The most traded futures on overnight rates settle gains compounded rates. The pricing of those futures requires some convexity adjustments with an Asian flavour due to the composition. Together with the greater liquidity of those futures came the market for options on futures. The options are traded with the usual futures daily margin mechanisms and can be standard options or mid-curve options; the options themselves are European or American. The pricing of those options, require different adjustments for the margin and composition features. In this note they propose the pricing of those options in the Gaussian HJM framework.

On the Skew and Curvature of Implied and Local Volatilities.

E. Alos, D. Garcia-Lorite, and M. Pravosud

Barcelona School of Economics & Caixabank

This paper studies the relationship between the short-end of the local and the implied volatility surfaces. Their results, based on Malliavin calculus techniques, recover the recent $\frac{1}{H+3/2}$ rule (where H denotes the Hurst parameter of the volatility process) for rough volatilitites (see Bourgey, De Marco, Friz, and Pigato (2022)), that states that the short-time skew slope of the at-the-money implied volatility is $\frac{1}{H+3/2}$ the corresponding slope for local volatilities. Moreover, they see that the at-the-money shortend curvature of the implied volatility can be written in terms of the short-end skew and curvature of the local volatility and viceversa, and that this relationship depends on H.

Market Impact: A Systematic Study of the High Frequency Options Market.

E. Said, A. Ayed, D. Thillou, J. Rabeyrin, and F. Abergel

BNP Paribas

This paper deals with a fundamental subject that has seldom been addressed in recent years, that of market impact in the options market. Their analysis is based on a proprietary database of metaorders - large orders that are split into smaller pieces before being sent to the market - on one of the main Asian markets. They propose an algorithmic approach to identify metaorders, based on some implied volatility parameters, the at the money forward volatility and at the money forward skew. In both cases, they obtain results similar to the now well understood equity market: Square-Root Law, Fair Pricing Condition and Market Impact Dynamics.

Deep Generators on Commodity Markets Application to Deep Hedging.

N. Boursin, C. Remlinger, J. Mikael, and C. Hargreaves

EDF Lab & NUS Singapore

Driven by the good results obtained in computer vision, deep generative methods for time series have been the subject of particular attention in recent years, particularly from the financial industry. In this article, they focus on commodity markets and test four state-of-the-art generative methods, namely Time Series Generative Adversarial Network (GAN), Causal Optimal Transport GAN, Signature GAN and the conditional Euler generator, are adapted and tested on commodity time series. A first series of experiments deals with the joint generation of historical time series on commodities. A second set deals with deep hedging of commodity options trained on the

generated time series. This use case illustrates a purely data-driven approach to risk hedging.

Quantitative Finance & Economics

Algorithmic Collusion in Electronic Markets: The Impact of Tick Size.

Á. Cartea, P. Chang, and J. Penalva
University of Oxford & Universidad Carlos III

They characterise the stochastic interaction of independent learning algorithms as a deterministic system of ordinary differential equations and use it to understand the long-term behaviour of the algorithms in a repeated game. In a symmetric bimatrix repeated game, they prove that the dynamics of many learning algorithms converge to the outcomes of pure strategy Nash equilibria of the stage game. In contrast, they prove that competition between Q-learning algorithms does not always lead to outcomes of Nash equilibria of the stage game. They apply these results to study how the size of the tick in a limit order book facilitates or obstructs tacit collusion among algorithms that compete to provide liquidity. They characterise the set of pure strategy Nash equilibria of the market making stage game with a discrete action space (e.g., the price grid of a limit order book) and its relation to the Bertrand–Nash equilibrium of the game with a continuous action space (i.e., an idealised limit order book with a zero tick size). They derive the bounds that define the set of Nash equilibria of the stage game when the action space is discrete and show that the bounds converge to the Bertrand–Nash equilibrium as the tick size of the limit order book tends to zero. In the repeated game played by all the algorithms they consider, their findings show that a large tick size obstructs competition; a smaller tick size lowers trading costs for liquidity takers, but slows down the speed of convergence to a rest point. For the algorithms with theoretical guarantees to reach a Nash equilibrium outcome, there is no assurance that the outcome reached is the outcome of the most competitive equilibrium. Indeed, tacit collusion can and often occurs when the algorithms tacitly coordinate on a more profitable Nash equilibrium outcome. However, the excess profits are bounded by the range of those of the possible Nash equilibria of the stage game, which shrinks as the tick size decreases. Finally, for Q-learning, many of the outcomes are sub-optimal for both the market makers and the liquidity takers.

AI Drive Liquidity Provision in OTC Financial Markets.

Á. Cartea, P. Chang, M. Mroczka, and R. Oomen
University of Oxford & Deutsche Bank

Providing liquidity in over-the-counter markets is a challenging under-taking, in large part because a market maker does not observe where their competitors quote, nor do they typically know how many rivals they compete with or what the trader's overall liquidity demand is. Optimal pricing strategies can be derived in theory assuming full knowledge of the competitive environment, but these results do not translate into practice where information is incomplete and asymmetric. This paper studies whether artificial intelligence, in the form of multi-armed bandit reinforcement learning algorithms, can be used by liquidity providers to dynamically set spreads using only information that is commonly available to them. They also investigate whether collusive effects can arise when competing liquidity providers all employ such algorithms. Their findings are as follows. In a single-agent setup where only one liquidity provider is optimising pricing in an otherwise static environment, all the algorithms considered are able to locate the theoretically optimal pricing policy, albeit they do so quite inefficiently when compared to a model-based approach. In a multi-agent setting where competing liquidity providers simultaneously and independently use algorithms to optimise pricing, they demonstrate that for one class of algorithms (pseudo) collusion can not arise, while for another it can theoretically arise in certain circumstances and they provide examples where it does. The scenarios where collusive effects appear, however, are fragile and sensitive to the specific configuration and exceedingly unlikely to occur in practice. Moreover, with a modest number of competitors, collusive effects that might otherwise arise in some of the most contrived scenarios are largely or entirely eliminated.

Monetary-based Asset Pricing: A Mixed-Frequency Structural Approach.

F. Bianchi, S. Ludvigson, and S. Ma

Federal Reserve, Duke Uni. & NYU

They integrate a high-frequency monetary event study into a mixed-frequency macro-finance model and structural estimation. The model and estimation allow for jumps at Fed announcements in investor beliefs, providing granular detail on why markets react to central bank communications. They find that the reasons involve a mix of revisions in investor beliefs about the economic state and/or future regime change in the conduct of monetary policy, and subjective reassessments of financial market risk. However, the structural estimation also finds that much of the causal impact of monetary policy on markets occurs outside of tight windows around policy announcements.

An Explosion Time Characterization of Asset Price Bubbles.

R. Jarrow, and S. Kwok
University of Sydney

In a standard continuous time asset pricing model, this paper provides an explosion time characterization of asset price bubbles that extends the existing characterization theorems in the literature from diffusion processes to general semi-martingales (which can include jumps). This characterization has a nice economic interpretation, not emphasized in the existing literature.

Portfolio Construction with Gaussian Mixture Returns and Exponential Utility via Convex Optimization.

E. Luxenberg, and S. Boyd
University of Sydney

They consider the problem of choosing an optimal portfolio, assuming the asset returns have a Gaussian mixture (GM) distribution, with the objective of maximizing expected exponential utility. In this paper they show that this problem is convex, and readily solved exactly using domain-specific languages for convex optimisation, without the need for sampling or scenarios. They then show how the closely related problem of minimizing entropic value at risk can also be formulated as a convex optimisation problem.

Inflation and GDP Dynamics in Production Networks: A Sufficient Statistics Approach.

H. Afrouzi, and S. Bhattarai
Columbia University & University of Texas Austin

They provide closed-form solutions for inflation and GDP dynamics in multi-sector New Keynesian economies with arbitrary input-output linkages. In particular, they show that a sufficient statistic for dynamics of sectoral prices is the principal square root of the Leontief matrix that is appropriately adjusted for the sectoral frequencies of price adjustments. They measure this sufficient statistic for the U.S. and find a significant quantitative role for production networks in the propagation of monetary and sectoral supply shocks. In response to monetary shocks, input-output linkages significantly increase the persistence of inflation and lead to a GDP response that is more than three times the response in an economy with a horizontal production network. In response to a negative supply shock in the "computers and electronics industry," input-output linkages propagate this shock by increasing the cost of downstream sectors, which acts like a markup shock to these sectors. These spillover effects generate a bigger and more persistent increase in aggregate inflation relative to an economy with a horizontal production network, which leads to about a three times bigger contraction in aggregate GDP. They show that a monetary policy response which fully offsets the inflationary effect of this negative sectoral supply shock on impact generates an aggregate GDP contraction that is more than twice as large as the case with no such policy response.

Machine Learning the Gravity Equation for International Trade.

S. Verstyuk, and M. Douglas
Harvard University

Machine learning (ML) is becoming more and more important throughout the mathematical and theoretical sciences. They apply modern ML methods to gravity models of pairwise interactions in international economics. They explain the formulation of graphical neural networks (GNNs), models for graph-structured data that respect the properties of exchangeability and locality. GNNs are a natural and theoretically appealing class of models for international trade, which they demonstrate empirically by fitting them to a large panel of annual-frequency country-level data. They then use a symbolic regression algorithm to turn their fits into interpretable models with performance comparable to state of the art hand-crafted models motivated by economic theory. The resulting symbolic models contain objects resembling market access functions, which were developed in modern structural literature, but in their analysis arise ab initio without being explicitly postulated. Along the way, they also produce several model-consistent and model-agnostic ML-based measures of bilateral trade accessibility.

Quantifying "Quantitative Tightening" (QT): How Many Rate Hikes is QT Equivalent To?

Bin Wei
Federal Reserve Bank of Atlanta

How many interest rate hikes is "quantitative tightening" (QT) equivalent to? In this paper, they examine this question based on the preferred-habitat model. The author defines the equivalence between rate hikes and QT such that they both have the same impact on the 10-year yield. Based on the model calibrated to fit the nominal Treasury data between 1999 and 2022, they show that a passive roll-off of $2.2 trillion over 3 years is equivalent to an increase of 29 basis points in the current federal funds rate at normal times. However, during a crisis period with risk aversion being doubled, it is equivalent to a 90-basis-point increase. They also quantify the effect of QT implemented by active sales. Lastly, based on their model-based estimates, they show that if the Treasury were to issue bills to offset maturing securities, the resulting equivalent rate hikes in the

current federal funds rate would decrease dramatically to 7.4 (12.6) basis points during normal (crisis) periods.

Net Zero Carbon Metrics

T. Guenedal, F. Lombard, T. Roncalli, and T. Sekine

Amundi

They propose a carbon budget approach that incorporates novel metrics for measuring the carbon emissions reduction targets and the relative positioning with respect to the net zero emissions (NZE) scenario. Indeed, current carbon emissions data are not sufficient to build portfolio alignment. The purpose of this paper is then to define net zero carbon metrics, which are necessary to enhance the disclosure and the debate on corporates' emissions. These carbon metrics can be divided into two families. The static measures are NZE duration, NZE gap, NZE slope and NZE budget. They can be computed using a target scenario or the linear trend model. The dynamic NZE measures incorporate the past trajectory and the future scenarios of carbon emissions. For instance, they break down the carbon budget by error and revision time contributions. They also propose a velocity measure of the carbon emissions trend and two main dynamic NZE measures that are necessary to assess the performance of an issuer compared to the NZE scenario: the zero-velocity scenario and the burn-out scenario. These different measures can then be used to define the PAC framework, that analyzes the participation, ambition and credibility of issuers' NZE policies. Finally, they apply this framework to the CDP database. Empirical results show that net zero carbon emissions are challenging for many issuers for two reasons. The first is that some issuers have a lack of ambition concerning their NZE scenario. The second is that some targets are not compatible with past trends.

Cryptocurrencies & Blockchain

The Coming Battle of Digital Currencies.

L. Cong, and S. Mayer

University of Chicago & Cornell University

They model the dynamic competition among national fiat currencies, cryptocurrencies, and Central Bank Digital Currencies (CBDCs), whereby a country's fiscal strength and currency strength are mutually reinforcing. The rise of cryptocurrencies hurts stronger fiat currencies, but can benefit weaker fiat currencies by reducing competition from stronger ones. Countries strategically implement CBDCs in response to competition from emerging cryptocurrencies and other currencies. Their model reveals the following pecking order: countries with strong but non-dominant currencies have the highest incentives to launch CBDCs to gain technological first-mover advantage; countries with the strongest currencies are the next in line to benefit from developing CBDC early on to nip cryptocurrency growth in the bud and to counteract competitors' CBDCs; nations with the weakest currencies forgo implementing CBDCs and adopt cryptocurrencies instead. They further extend the framework to understand the role of stablecoins in currency competition, and study the effects of currency competition and cryptocurrencies on financial innovation. Their findings help rationalize recent developments in currency and payment digitization, while providing insights into the battle of currencies and the future of money.

Graphs-based Learning

Network Diversification for a Robust Portfolio Allocation.

M. Jaeger, and D. Marinelli

Munich Re

Portfolio allocation strategies often seek risk budgeting and diversification by relying only on correlation matrices to model relationships between assets. Although this approach can capture, in normal times, most of the dependencies between asset prices, it faces several challenges in terms of noise resistance, capturing non-linear relations that can naturally appear in the market and extreme allocations in long-short portfolio strategies. This paper presents novel network-based strategies that combine equal volatility allocation with network centrality measures to construct efficiently diversified portfolios and deliver stable strategies also suitable for long-short investments. Networks can encode linear and non-linear relationships between asset prices. To encode several layers of information simultaneously multiplex networks - a particular form of a multilayer network - can be deployed. Associated centrality measures can agnostically account for each asset's (ir)relevance in diversifying the risks of the portfolios. The results show that network-based portfolios can outperform several competing alternatives, maintaining a favourable risk characteristic.

Can Abnormality be Detected by Graph Neural Networks?

Z. Chai, S. You, Y. Yang, S. Pu, J. Xu, H. Cai, and W. Jiang

Zhejiang University, Carnegie Mellon University & Fudan University

Anomaly detection in graphs has attracted considerable interests in both academia and industry due to its wide applications in numerous domains ranging from finance to biology. Meanwhile, graph neural networks (GNNs) is emerging as a powerful tool for modeling graph data. A natural and fundamental question that arises here is:

can abnormality be detected by graph neural networks? In this paper, they aim to answer this question, which is nontrivial. As many existing works have explored, graph neural networks can be seen as filters for graph signals, with the favor of low frequency in graphs. In other words, GNN will smooth the signals of adjacent nodes. However, abnormality in a graph intuitively has the characteristic that it tends to be dissimilar to its neighbors, which are mostly normal samples. It thereby conflicts with the general assumption with traditional GNNs. To solve this, they propose a novel Adaptive Multi-frequency Graph Neural Network (AMNet), aiming to capture both low-frequency and high-frequency signals, and adaptively combine signals of different frequencies. Experimental results on real-world datasets demonstrate that their model achieves a significant improvement comparing with several state-of-the-art baseline methods.

An Adaptive Hierarchical Network Model for Studying the Structure of Economic Network.

X. Yang, Z. Wu, and S. Javaid
Shaanxi Normal University & Tongji University

The interdependence of financial institutions is primarily responsible for creating a systemic hierarchy in the industry. In this paper, an Adaptive Hierarchical Network Model is proposed to study the problem of hierarchical relationships arising from different individuals in the economic domain. In the presented dynamically evolving network model, new directed edges are generated depending on the existing nodes and the hierarchical structures among the network, and these edges decay over time. When the preference of nodes in the network for higher ranks exceeds a certain threshold value, the equality state in the network becomes unstable and rank states emerge. Meanwhile, they select four real data sets for model evaluation and observe the resilience in the network hierarchy evolution and the differences formed by different patterns of hierarchy preference mechanisms, which helped them better understand data science and network dynamics evolution.

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, Álvaro Arroyo, and Patrick Chang.