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, Natural Language Processing (NLP) & Fundamental Data Analysis, Quantitative Finance, Graph-based Learning, Risk Optimisation & Analytics, Derivatives & Volatility*

Machine Learning for Finance

Asset Pricing with Attention Guided Deep Learning

P. Chatigny, R. Goyenko, and C. Zhang
University of Sherbrooke, McGill University

Trading costs approach to portfolio construction strongly recommends considering multiple firm characteristics to reduce transaction and re-balancing expenses on the portfolio level. Deep learning methods, which can accommodate wide ranges of various stock characteristics to identify optimal and tradable stochastic discount factor (SDF) have been criticised for losing their superior performance after trading costs. This work introduces attention-guided deep learning which, in a data driven way, allows to identify the most influential time-varying firm characteristics contributing to SDF. Attention dramatically improves SDF performance and reduces portfolio rebalancing costs. The attention guided SDF outperforms existing models after trading costs, excluding small and micro-cap stocks, avoids extreme portfolio weights, and unlike other models, exhibits the best performance during market regimes with the highest price efficiency.

Forecasting Natural Gas Price Trends Using Random Forest and Support Vector Machine Classifiers.

F. Castaneda, S. Niro, M. Schicks, and N. Hartmann
RWE Group

Price forecasting using statistical modeling methods and data mining has been a topic of great interest among data scientists around the world. In this paper, different machine learning approaches are applied to forecasting future yearly price trends in the natural gas Title Transfer Facility market in the Netherlands. The study compares two models: random forest and support vector classifiers. The identification of potential natural gas price drivers that improve the model's classification is crucial. The forecast horizon was set in a range from 10 to 60 trading days, considering that shorter time horizons have greater importance for trading. The results reflect values up to 85% of the area-under-the-curve score as a reaction of the models to the four different feature combinations used. This invites continued research on the multiple opportunities that these new technologies could create.

Recovering Missing Firm Characteristics with Attention-based Machine Learning

H.Beckmeyer, and T. Wiedemann
University of Muenster

Firm characteristics are often missing. This work sets up an attention-based machine learning model borrowing ideas from state-of-the-art research in natural language processing to understand how characteristics relate to the cross-section of other – observed – firm characteristics and their historical evolution. The work presents a

model which reconstructs firm characteristics with high accuracy and comfortably outperforms competing approaches. Revisiting the vast literature on risk factors in financial research reveals that disregarding the influence of missing observations likely overestimates the magnitude of factor premia.

Reinforcement Learning With Dynamic Convex Risk Measures

A. Coache, and S. Jaimungal
University of Toronto

The work develops an approach for solving time-consistent risk-sensitive stochastic optimisation problems using model-free reinforcement learning. It assumes that agents assess the risk of a sequence of random variables using dynamic convex risk measures. They employ a time-consistent dynamic programming principle to determine the value of a particular policy, and develop policy gradient update rules that aid in obtaining optimal policies. They further develop an actor-critic style algorithm using neural networks to optimise over policies. Finally, the work demonstrates the performance and flexibility of their approach by applying it to three optimisation problems: statistical arbitrage trading strategies, obstacle avoidance robot control, and financial hedging.

Enhancing Cross-Sectional Currency Strategies by Context-Aware Learning to Rank with Self-Attention

D. Poh, B. Lim, S. Zohren, and S. Roberts
University of Oxford

The performance of a cross-sectional currency strategy depends crucially on accurately ranking instruments prior to portfolio construction. While this ranking step is traditionally performed using heuristics, or by sorting the outputs produced by pointwise regression or classification techniques, strategies using Learning to Rank algorithms have recently presented themselves as competitive and viable alternatives. Although the rankers at the core of these strategies are learned globally and improve ranking accuracy on average, they ignore the differences between the distributions of asset features over the times when the portfolio is rebalanced. This flaw renders them susceptible to producing sub-optimal rankings, possibly at important periods when accuracy is actually needed the most. For example, this might happen during critical risk-off episodes, which consequently exposes the portfolio to substantial, unwanted drawdowns. The authors tackle this shortcoming with an analogous idea from information retrieval: that a query's top retrieved documents or the local ranking context provide vital information about the query's own characteristics, which can then be used to refine the initial ranked list. In this work, authors use a context-aware Learning-to-rank model that is based on the Transformer architecture to encode top/bottom ranked assets, learn the context and exploit this information to re-rank the initial results. Backtesting on a slate of 31 currencies, the proposed methodology increases the Sharpe ratio by around 30% and significantly enhances various performance metrics. Additionally, this approach also improves the Sharpe ratio when separately conditioning on normal and risk-off market states.

NLP & Fundamental Data Analysis

Federal Reserve Chair Communication Sentiments' Heterogeneity, Personal Characteristics, and their Impact on Uncertainty and Target Rate Discovery

J. Arismendi-Zambrano, E. Kypraios, and A. Paccagnini
Uni. College Dublin, Northwestern Uni.

The work estimates the personal communication risk-premium profile of the U.S. Federal Reserve Chair by measuring a new dataset of the sentiment revealed by their public statements during their tenure. They analyse the impact of such Fed communications' sentiment risk on the uncertainty of the monetary policy, and the market price discovery process of interest rates, in the aftermath of the Federal Open Market Committee (FOMC) meetings. After controlling for the evolving state of the economy surrounding the meetings, they find that there is a significant statistical and economic difference in the communications' sentiment that is heterogeneous across Chairs, depending on their personal traits. The sentiment in the Chairs' communications plays a role in moderating the potential surprises in the Fed announcements, and it can be effectively used as a tool for controlling and measuring monetary policy shocks.

Multimodal machine learning for credit modeling

C. Nguyen, S. Das, J. He, S. Yue, V. Hanumaiah, X. Ragot, and L. Zhang
Amazon

Credit ratings are traditionally generated using models that use financial statement data and market data, which is tabular (numeric and categorical). Practitioner and academic models do not include text data. Using an automated approach to combine long-form text from SEC filings with the tabular data, we show how multimodal machine learning using stack ensembling and bagging can generate more accurate rating predictions. This paper demonstrates a methodology to use big data to extend tabular data models, which have been used by the ratings industry for decades, to the class of multimodal machine learning models.

FinLex: An effective use of word embeddings for financial lexicon generation

S. Das, M. Donini, B. Zafar, J. He, and K. Kenthapadi

Amazon

We present a simple and effective methodology for the generation of lexicons (word lists) that may be used in natural language scoring applications. In particular, in the finance industry, word lists have become ubiquitous for sentiment scoring. These have been derived from dictionaries such as the Harvard Inquirer and require manual curation. Here, we present an automated approach to the curation of lexicons, which makes automatic preparation of any initial word list immediate, which can then be further curated. We show that our automated word lists deliver comparable performance to traditional lexicons on machine learning classification tasks. This new approach will enable finance academics and practitioners to create and deploy new word lists in addition to the few traditional ones in a facile manner.

Sustainable Alpha in Sovereign and Corporate Bonds

K. Kaul, K. Schwaiger, M. Si, and A. Ang

BlackRock

The work constructs fixed income portfolios for sovereign bonds and corporate bonds with sustainable insights. The climate methodology for sovereign bonds can be applied as an overlay on any benchmark and tilts towards sovereigns more prepared with the climate transition and away from those which are less prepared. The tilts reduce sovereign carbon emissions in line with the Paris Agreement. For corporate bonds, the work investigates three sustainable signals that predict excess returns: ESG scores of corporations scored across various rating and sector buckets, firm carbon emission intensities, and corporate commitments that signal reduced carbon emissions.

Quantifying the Impact of Impact Investing

A.W. Lo, and R. Zhang

Massachusetts Institute of Technology

The work proposes a quantitative framework for assessing the financial impact of any form of impact investing, including socially responsible investing, environmental, social, and governance (ESG) objectives, and other non-financial investment criteria. They derive conditions under which impact investing detracts from, improves on, or is neutral to the performance of traditional mean-variance optimal portfolios, which depends on whether the correlations between the impact factor and unobserved excess returns are negative, positive, or zero, respectively. Using Treynor-Black portfolios to maximise the risk-adjusted returns of impact portfolios, they propose a quantitative measure for the financial reward, or cost, of impact investing compared to passive index benchmarks. They illustrate their approach with applications to biotech venture philanthropy, divesting from "sin" stocks, investing in ESG, and "meme" stock rallies such as GameStop in 2021.

Beliefs and Portfolios: Causal Evidence

J. Beutel, and M. Weber

Bundesbank, NBER

The work causally tests alternative theories of expectation formation and asset pricing. Using a randomised information experiment they show: i) individuals form pro-cyclical beliefs, both about capital gains and about earnings growth, inconsistent with rational expectations models; ii) individuals are heterogeneous both at the information acquisition and at the information processing stage. Their reaction to stock-market news depends on their information preference; iii) beliefs and portfolio decisions are causally linked; iv) the sensitivity of portfolio shares with respect to expected returns is low, especially when accounting for individuals' low perceived variance of stock returns. Non-linear constraints on individuals' portfolio decisions resolve this puzzle. These results provide guidance on promising causal mechanisms for macro-finance models.

Assessing Macroeconomic Tail Risk

F. Loria, C. Matthes, and D. Zhang

Federal Reserve

Real GDP and industrial production in the US feature substantial tail risk. While this fact is well documented, several questions remain unanswered. Is this asymmetry driven by a specific structural shock? No. The work shows that the 10th percentile of the predictive growth distributions responds about three times more than the median to both monetary policy and financial shocks. What mechanism can generate this asymmetry in the data? They discuss nonlinear VAR models and a nonlinear equilibrium model that are capable of matching empirical findings in this work. Furthermore, they provide empirical evidence that allows us to differentiate between two competing theories.

Long-term macroeconomic effects of climate change: A cross-country analysis

M.E. Kahn, K. Mohaddesb, R. Ng, M. Pesaran, M. Raissi, and J. Yang

Uni. of Cambridge, Uni. of South. California

The work studies the long-term impact of climate change on economic activity across countries, using a stochastic growth model where productivity is affected by deviations of temperature and precipitation from their long-term moving average historical norms. Using a panel data set of 174 countries over the years 1960 to 2014, they find that the per-capita real output

growth is adversely affected by persistent changes in the temperature above or below its historical norm, but they do not obtain any statistically significant effects for changes in precipitation. They also show that the marginal effects of temperature shocks vary across climates and income groups. The counterfactual analysis suggests that a persistent increase in average global temperature by 0.04 °C per year, in the absence of mitigation policies, reduces world real GDP per capita by more than 7 percent by 2100. On the other hand, abiding by the Paris Agreement goals, thereby limiting the temperature increase to 0.01 °C per annum, reduces the loss substantially to about 1 percent. These effects vary significantly across countries depending on the pace of temperature increases and variability of climate conditions. The estimated losses would increase to 13 percent globally if country-specific variability of climate conditions were to rise commensurate with annual temperature increases of 0.04 °C

The Macroeconomic Impact of Social Unrest

M. Hadzi-Vaskov, S. Pienknagura, and L.A. Ricci
International Monetary Fund

This paper explores the macroeconomic impact of social unrest, using a novel index based on news reports. The findings are threefold. First, unrest has an adverse effect on economic activity, with GDP remaining on average 0.2 percentage points below the pre-shock baseline six quarters after a one-standard deviation increase in the unrest index. This is driven by sharp contractions in manufacturing and services (sectoral dimension), and consumption (demand dimension). Second, unrest lowers confidence and raises uncertainty; however, its adverse effect on GDP can be mitigated by strong institutions and by a country's policy space. Third, an unrest "event", which is captured by a large change in the unrest index, is associated with a 1 percentage point reduction in GDP six quarters after the event. Impacts differ by type of event: episodes motivated by socio-economic reasons result in sharper GDP contractions compared to those associated with politics/elections, and events triggered by a combination of both factors lead to sharpest contractions. Results are not driven by countries with adverse growth trajectories prior to unrest events or by fiscal consolidations, and are robust to instrumenting via regional unrest.

Interest Rate Skewness and Biased Beliefs

M. Bauer, and M. Chernov
University of Hamburg, UCLA Anderson

The conditional skewness of Treasury yields is an important indicator of the risks to the macroeconomic outlook. Positive skewness signals upside risk to interest rates during periods of accommodative monetary policy and an upward-sloping yield curve, and vice versa. Skewness has substantial predictive power for future bond excess returns, high-frequency interest rate changes around FOMC announcements, and survey forecast errors for interest rates. The estimated expectational errors, or biases in beliefs, are quantitatively important for statistical bond risk premia. These findings are consistent with a heterogeneous-beliefs model where one of the agents is wrong about consumption growth.

Flexible Demand Estimation with Search Data

T. Amano, A. Rhodes, and S. Seiler
Harvard University, Imperial College

Traditional methods for estimating demand are not always well-suited to online markets, where individual products are sold infrequently, unobserved factors such as webpage layout drive substitution, and often only a limited set of product characteristics is observed. In this work, authors propose a demand model where browsing data—which is abundant in many online settings—is used to infer individual consumers' consideration sets. In the proposed model, the underlying variables which drive consideration can be correlated arbitrarily across products. Authors estimate the model through a constraint maximization approach, based on the insight that these correlations should rationalize the product-pair co-search frequencies that are observed in the data. In turn, these correlations make it possible to estimate more flexible substitution patterns. Authors apply the model to data from an online retailer, recover the elasticity matrix, and solve for optimal prices.

Quantitative Finance

Harnessing the Overconfidence of the Crowd: A Theory of SPACs

S. Banerjee, and M. Szydlowski
UCSD, Uni. of Minnesota

In a SPAC transaction, a sponsor raises financing from investors using redeemable shares and rights. When investors are rational, these features dilute the sponsor's stake and lead to underinvestment in eficient projects. However, when investors are overconfident about their ability to respond to interim news, the optionality in such features is overpriced, and SPACs can lead to over-investment in inefficient projects. The model predicts different returns for short-term and long-term investors and overall under- performance, consistent with empirical evidence. They evaluate the impact of policy interventions, including eliminating redemption rights, mandating greater disclosure and transparency, limiting investor access, and restricting the rights offered.

Market making by an FX dealer: tiers, pricing ladders and hedging rates for optimal risk control

A. Barzykin, P. Bergault, and O. Guéant

HSBC, Ecole Polytechnique, Université Paris

Dealers make money by providing liquidity to clients but face flow uncertainty and thus price risk. They can efficiently skew their prices and wait for clients to mitigate risk (internalisation), or trade with other dealers in the open market to hedge their position and reduce their inventory (externalisation). Of course, the better control associated with externalisation comes with transaction costs and market impact. The internalisation vs. externalisation dilemma has been a topic of recent active discussion within the foreign exchange (FX) community. This paper offers an optimal control framework for market making tackling both pricing and hedging, thus answering a question well known to dealers: 'to hedge, or not to hedge?

How to Sell Hard Information

S.N. Ali, N. Haghpanah, X. Lin, and R. Siegel
Pennsylvania State University

The seller of an asset has the option to buy hard information about the value of the asset from an intermediary. The seller can then disclose this information before selling the asset in a competitive market. This work studies how the intermediary designs and sells hard information to robustly maximise the intermediary's revenue across all equilibria. Even though the intermediary could use an accurate test that reveals the asset's value, authors show that robust revenue maximisation leads to a noisy test with a continuum of possible scores. In addition, the intermediary always charges the seller for disclosing the test score to the market, but not necessarily for running the test. This enables the intermediary to robustly appropriate a significant share of the surplus resulting from the asset sale.

Long Term Risk: A Time Change Approach

P. Carr, and U. Cherubini
NYU, Uni. of Bologna

The work studies an asset pricing model in which a Stochastic Discount Factor (SDF) and a growth process have independent increments and are affected by a common stochastic clock. The stochastic clock is a strictly increasing process that may include a persistent component, that they call comonotonic, and represents sources of long run uncertainty, such as those induced by climate change issues. They find that long term log-returns are lower the higher the variance of the clock, and the more relevant its comonotonic component. The comonotonic component in the clock induces a separation of the SDF dynamics between persistent and transitory shocks similar to that assumed in the long term risk literature, with the difference that the system is non Markovian. They also provide an axiomatic definition of the compounding operation leading to an associative algebraic structure, with generator represented by the Laplace transform of the stochastic clock increments. The same generator characterises the Archimedean copula function linking the SDF and the growth process.

Backtesting Systemic Risk Forecasts using Multi-Objective Elicitability

T. Fissler, and Y. Hoga
Vienna University, University of Duisburg-Essen

Systemic risk measures such as CoVaR, CoES and MES are widely-used in finance, macroeconomics and by regulatory bodies. Despite their importance, they show that they fail to be elicitable and identifiable. This renders forecast comparison and validation, commonly summarised as 'backtesting', impossible. The novel notion of multi-objective elicitability solves this problem. Specifically, they propose Diebold– Mariano type tests utilising two-dimensional scores equipped with the lexicographic order. They illustrate the test decisions by an easy-to-apply traffic-light approach. The authors apply their traffic-light approach to DAX 30 and SP 500 returns, and infer some recommendations for regulators.

Risk-Neutral Market Simulation

M. Wiese, and P. Murray
University of Kaiserslautern, J.P. Morgan

The work develops a risk-neutral spot and equity option market simulator for a single underlying, under which the joint market process is a martingale. They leverage an efficient low-dimensional representation of the market which preserves no static arbitrage, and employ neural spline flows to simulate samples which are free from conditional drifts, but also highly realistic in the sense that among all possible risk-neutral simulators, they are a minimal Kullback-Leibler divergence to the historical data. The work demonstrates the effectiveness with numerical experiments which highlight both the drift removal and the fidelity of the calibrated simulator.

Exploring Risk Premia, Pricing Kernels, and No-Arbitrage Restrictions in Option Pricing Models

S.L. Heston, K. Jacobs, and H.J. Kim
Uni. of Maryland, Uni. of Houston

The literature on dynamic option valuation typically does not explicitly specify a pricing kernel. Instead it characterises the kernel indirectly by specifying prices of risk, or defines it implicitly as the ratio of the risk-neutral and physical probabilities. The authors propose explicit pricing kernels for stochastic volatility option pricing models that satisfy absence of arbitrage. Different affine price-of-risk specifications correspond to pricing kernels with radically different, and sometimes implausible, economic implications. It can be difficult to statistically distinguish between pricing kernels with widely different

economic implications and risk premia. The authors attribute this to the inherent statistical problem with the estimation of the equity and variance risk premia. This finding extends Merton's (1980) observations on the estimation of the market equity premium to joint estimation of the equity and variance risk premia using the cross-section of options and the underlying returns.

Price Impact of Order Flow Imbalances: Multi-level, Cross-asset and Forecasting

R. Cont, M. Cucuringu, and C. Zhang
University of Oxford

The work investigates the impact of order flow imbalance (OFI) on price movements in equity markets in a multiasset setting. First, authors show that taking into account multiple levels of the order book when defining order book imbalance leads to higher explanatory power for the contemporaneous price impact of OFI. Using a principal component analysis of OFI across order book levels, they define a notion of integrated OFI which shows superior explanatory power for market impact both in-sample and out-of-sample. Second, they examine the notion of cross-impact and show that, once the information from multiple levels is included in OFI, multi-asset models with cross-impact do not provide additional explanatory power for contemporaneous impact compared to a sparse model without cross-impact terms. However, they find evidence that cross-impact terms provide additional information for intraday forecasting of future returns

Optimal Turnover, Liquidity, and Autocorrelation

G. Ritter, B. Baldacci, and E. Benveniste
New York University

The steady-state turnover of a trading strategy is of clear interest to practitioners and portfolio managers, as is the steady-state Sharpe ratio. In this work, authors show that in a convenient Gaussian process model, the steady-state turnover can be computed explicitly, and obeys a clear relation to the liquidity of the asset and to the autocorrelation of the alpha forecast signals. They find that steady-state optimal turnover is given by $\gamma\sqrt{n+1}$ where γ is a liquidity-adjusted notion of risk-aversion, and n is the ratio of mean-reversion speed to .

Schrödinger Risk Diversification Portfolio

Y. Uchiyama, and K. Nakagawa
Mazin Inc.,Nomura A.M.

The mean-variance portfolio that considers the trade-off between expected return and risk has been widely used in the problem of asset allocation for multi-asset portfolios. However, since it is difficult to estimate the expected return and the out-of-sample performance of mean-variance portfolio is poor, risk-based portfolio construction methods focusing only on risk have been proposed, and are attracting attentions mainly in practice. In terms of risk, asset fluctuations that make up the portfolio are thought to have common factors behind them, and principal component analysis, which is a dimension reduction method, is applied to extract the factors. In this work, authors propose the Schrödinger risk diversification portfolio as a factor risk diversifying portfolio using Schrödinger principal component analysis that applies the Schrödinger equation in quantum mechanics. The Schrödinger principal component analysis can accurately estimate the factors even if the sample points are unequally spaced or in a small number, thus we can make efficient risk diversification. The proposed method was verified to outperform the conventional risk parity and other risk diversification portfolio constructions.

Systemic Implications of the Bail-In Design

J.D. Farmer, C. Goodhart, and A. Kleinnijenhuis
University of Oxford; University of Stanford

The 2007-2008 financial crisis forced governments to choose between the unattractive alternatives of either bailing out a systemically important bank (SIB) or allowing it to fail in a disorderly manner. Bail-in has been put forward as an alternative that potentially addresses the too-big-to-fail problem and contagion risk simultaneously. Though its efficacy has been demonstrated for smaller idiosyncratic SIB failures, its ability to maintain stability in cases of large SIB failures and system-wide crises remains untested. This paper's novelty is to assess the financial-stability implications of bail-in design, explicitly accounting for the multi-layered networked nature of the financial system. The authors present a model of the European financial system that captures five prevailing contagion channels. They demonstrate that it is essential to understand the interaction of multiple contagion mechanisms and that financial institutions other than banks play an important role. Their results indicate that stability hinges on the bank-specific and structural bail-in design. On the one hand, a well-designed bail-in buttresses financial resilience, but on the other hand, an ill-designed bail-in tends to exacerbate financial distress, especially in system-wide crises and when there are large SIB failures. Their analysis suggests that the current bail-in design may be in the region of instability. While policy makers can fix this, the political economy incentives make this unlikely.

Graph-based Learning

Beyond permutation equivariance in graph networks

E. Slade, and F. Farina
GlaxoSmithKline

Permutation Equivariance is a desirable property for training GNNs to solve problems, E(n)quivalent GNNs proposed by the authors are a new methodology designed to learn better deep representations that are invariant to transformations in their embedded representations to improve learning ability. - Graph Neural Networks with convolutional ARMA filters : Most existing filters for GNNs are graph convolutional filters operating in the spatial or spectral space with polynomial approximations of the filters, Bianchi et al. proposes a new filter based on the Autoregressigve Moving Average (ARMA) to give the model more flexibility as well as better predictive power.

Graph Attention Recurrent Neural Networks for Correlated Time Series Forecasting

R. Cristea, C. Guo, and B. Yang
Aalborg University

Dynamic Graph Learning is a difficult task with wide applicability from traffic to financial time series forecasting. One innovative strategy proposed by Cristea et al is to use attention to generate dynamic adjacency matrices to learn a better representation of links in a dynamic graph, the authors subsequently also propose a reccurent Graph Neural network to solve the problem of spatially and temporally evolving graphs that is portable to other time series analysis such as financial networks.

Time series analysis via network science: Concepts and algorithms

V.F. Silva, M.E. Silva, P. Ribeiro, and F. Silva
Universidade do Porto

Authors provide a good overview of the work done in using networks to forecast and analyse time series. Their work covers both the different application scenarios, data processing steps and adequate models to target both multivariate and univariate time series from a graph learning perspective.

Novel Features for Time Series Analysis: A Complex Networks Approach

V.F. Silva, M.E. Silva, P. Ribeiro, and F. Silva
Universidade do Porto

Authors propose in a followup paper to define a different similarity metric for time series basing themselves on the similarity of their graph mappings instead of only using the similarity of the time series to build similarity graphs.

Cross-Domain Few-Shot Graph Classification

K. Hassani
Autodesk AI Lab

Graph Neural Networks have not yet seen a period of flourishment in transfer learning methods in the vein of the NLP domains' enthusiastic adoption of BERT-like General Language Models. The paper by Hassani explores the possibility of transfer learning for graphs by proposing a model adept at few-shot learning across a variety of benchmarks with little data pre-training necessary.

Using Network-based Causal Inference to Detect the Sources of Contagion in the Currency Market

K. Hassani
Swiss Finance Institute

Financial contagion processes and spillover effects studied with the help of networks play a considerable role in the network science litterature for finance as grpah objects are uniquely suited to studying propagations and state-change effects as a result of changes in link patterns. The paper by Rigana et al. proposes to continue this inquiry by studying the spillover effects of crisis on currency markets.

Risk Optimisation & Analytics

One Step Preference Elicitation in Multi-Objective Bayesian Optimization

J. Ungredda, J. Branke, M. Marchi, and T. Montrone
University of Warwick

The work considers a multi-objective optimisation problem with objective functions that are expensive to evaluate. The decision maker (DM) has unknown preferences, and so the standard approach is to generate an approximation of the Pareto front and let the DM choose from the generated non-dominated designs. However, especially for expensive to evaluate problems where the number of designs that can be evaluated is very limited, the true best solution according to the DM's unknown preferences is unlikely to be among the small set of non-dominated solutions found, even if these solutions are truly Pareto optimal. The authors address this issue by using a multiobjective Bayesian optimisation algorithm and allowing the DM to select a preferred solution from a predicted continuous Pareto front just once before the end of the algorithm rather than selecting a solution after the end. This allows the algorithm to understand the DM's preferences and make a final attempt to identify a more preferred solution. The authors demonstrate the idea using ParEGO, and show empirically that the found solutions are significantly better in terms of true DM preferences than if the DM would simply pick a solution at the end.

Covariance estimation for risk-based portfolio optimization: an integrated approach

A.Butler, and R.H. Kwon
ReSolve Asset Management, University of Toronto

Many problems in quantitative finance involve both predictive forecasting and decision-based optimisation. Traditionally, covariance forecasting models are optimised with unique prediction-based objectives and constraints, and users are therefore unaware of how these predictions would ultimately be used in the context of the models' final decision-based optimisation. They present a stochastic optimisation framework for integrating time-varying factor covariance models in a risk-based portfolio optimisation setting. They make use of recent advances in neural-network architecture for efficient optimisation of batch convex programs, and consider three risk-based portfolio optimisations: minimum variance, maximum diversification and equal risk contribution. They provide a first-order method for performing the integrated optimisation, whilst presenting several historical simulations using US industry and stock data, and demonstrating the benefits of the integrated approach in comparison with the decoupled alternative.

Applications of Signature Methods to Market Anomaly Detection

E. Akyildiri, M. Gambara, J. Teichmann, and S. Zhou

ETH Zurich, University of Zurich

Anomaly detection is the process of identifying abnormal instances or events in data sets which deviate from the norm significantly. This work proposes a signatures based machine learning algorithm to detect rare or unexpected items in a given data set of time series type. They present applications of signature or randomised signature as feature extractors for anomaly detection algorithms; additionally authors provide an easy, representation theoretic justification for the construction of randomised signatures. Their first application is based on synthetic data and aims at distinguishing between real and fake trajectories of stock prices, which are indistinguishable by visual inspection. They also show a real life application by using transaction data from the cryptocurrency market. For this application, they are able to identify pump and dump attempts organised on social networks with F1 scores up to 88% by means of their unsupervised learning algorithm, thus achieving results that are close to the state-of-the-art in the field based on supervised learning.

Evaluating conditional covariance estimates via a new targeting approach and a networks-based analysis

c. Drago, and A. Scozzari

CESDE

Modeling and forecasting of dynamically varying covariances have received much attention in the literature. The two most widely used conditional covariances and correlations models are BEKK and DCC. In this paper, authors advance a new method to introduce targeting in both models to estimate matrices associated with financial time series. Their approach is based on specific groups of highly correlated assets in a financial market, and these relationships remain unaltered over time. Based on the estimated parameters, authors evaluate their targeting method on simulated series by referring to two well-known loss functions introduced in the literature and Network analysis. The authors find all the maximal cliques in correlation graphs to evaluate the effectiveness of their method. Results from an empirical case study are encouraging, mainly when the number of assets is not large

Black-box Bayesian inference for economic agent-based models

J. Dyer, P. Cannon, J.D. Farmer, and S. Schmon

University of Oxford

Simulation models, in particular agent-based models, are gaining popularity in economics. The considerable flexibility they offer, as well as their capacity to reproduce a variety of empirically observed behaviors of complex systems, give them broad appeal, and the increasing availability of cheap computing power has made their use feasible. Yet a widespread adoption in real-world modelling and decision-making scenarios has been hindered by the difficulty of performing parameter estimation for such models. In general, simulation models lack a tractable likelihood function, which precludes a straightforward application of standard statistical inference techniques. A number of recent works (Grazzini et al., 2017, Platt, 2020, 2021) have sought to address this problem through the application of likelihood-free inference techniques, in which parameter estimates are determined by performing some form of comparison between the observed data and simulation output. However, these approaches are (a) founded on restrictive assumptions, and/or (b) typically require many hundreds of thousands of simulations. These qualities make them unsuitable for large-scale simulations in economics and can cast doubt on the validity of these inference methods in such scenarios. This work investigates the efficacy of two classes of simulation-efficient black-box approximate Bayesian inference methods that have recently drawn significant attention within the probabilistic machine learning community: neural posterior estimation and neural density ratio estimation. They present a number of benchmarking experiments in which they demonstrate that neural network based black-box methods provide state of the art parameter inference for economic simulation models, and crucially are compatible with generic multivariate time-series data. In addition, they suggest appropriate assessment criteria for use in future benchmarking of approximate Bayesian inference procedures for economic simulation models.

Bayesian Imputation with Optimal Look-Ahead-Bias and Variance Tradeof

J. Blanchet, F. Hernandez, V.A. Nguyyen, M. Pelger, and X. Zhang

Stanford University

Missing time-series data is a prevalent problem in finance. Imputation methods for timeseries data are usually applied to the full panel data with the purpose of training a model for a downstream out-of-sample task. For example, the imputation of missing returns may be applied prior to estimating a portfolio optimisation model. However, this practice can result in a lookahead-bias in the future performance of the downstream task. There is an inherent trade-off between the look-ahead-bias of using the full data set for imputation and the larger variance in the imputation from using only the training data. By connecting layers of information revealed in time, they propose a Bayesian consensus posterior that fuses an arbitrary number of posteriors to optimally control the variance and look-ahead-bias trade-off in the imputation. They derive tractable two-step optimisation procedures for finding the optimal consensus posterior, with Kullback-Leibler divergence and Wasserstein distance as the measure of dissimilarity between posterior distributions. They demonstrate in simulations and an empirical study the benefit of their imputation mechanism for portfolio optimisation with missing returns.

Derivatives & Volatility

Multi-Asset Spot and Option Market Simulation

M. Wiese, B.Wood, A. Pachoud, R. Korn, H. Buehler, P. Murray, and L. Bai

University of Kaiserslautern, J.P. Morgan

The work constructs a realistic spot and equity option market simulators for a single underlying on the basis of normalizing flows. They address the high-dimensionality of market observed call prices through an arbitrage-free autoencoder that approximates efficient low-dimensional representations of the prices while maintaining no static arbitrage in the reconstructed surface. Given a multi-asset universe, this work leverages the conditional invertibility property of normalizing flows and introduce a scalable method to calibrate the joint distribution of a set of independent simulators while preserving the dynamics of each simulator. Empirical results highlight the goodness of the calibrated simulators and their fidelity.

Delta-gamma component VAR: nonlinear risk decomposition for any type of fund

M. Dixon, and J. Goldcamp

Illinois Institute of Technology, HedgeFacts

The authors develop an analytical methodology for decomposing non-linear portfolio risk not only by instrument, but also by fund managers or sub-portfolios for one single manager. Furthermore, the approach may be used by quantitative portfolio managers for risk decomposition by factors under a factor investing strategy.

Dynamically controlled kernel estimation

G. Lee

UBS

The work introduce a data-driven and model-agnostic approach for computing conditional expectations. The new method combines classical techniques with machine learning methods, in particular Gaussian process regression and kernel density estimation stabilised with a control variate. The method is applied to pricing and hedging of (multidimensional) exotic options for different models, including the rough Bergomi model

Forecasting volatility and market returns using the CBOE Volatility Index and its options

S.T. Stanley, and W.J. Trainor Jr.

ARGI Financial Investment Board

This study examines the Chicago Board Option Exchange Volatility Index (VIX) and its options. The VIX is the implied volatility calculated from short-term option prices on the Standard Poor's 500 (SP 500) stock index. Findings presented in this work demonstrate that VIX overestimates average volatility by approximately 3% but explains 55% of the SP 500's following month's volatility and 20% of its return. The smirks calculated from the VIX options' implied volatility add additional explanatory power for the SP 500 returns. None of the variables help predict skewness or kurtosis values (a measure of tail risk). A simple trading rule that buys the SP 500 when the VIX, the implied volatility from the options on the VIX and the VIX options' volatility smirk all decline is associated with an in average monthly return of 1.96% in the SP 500 relative to its 0.84% average. This only occurs approximately 10% of the time and would not beat a buy-and-hold strategy, but it could be used to adjust asset allocations at the margin. Buying equities only when the VIX decreases, which occurs approximately 50% of the time, outperforms a similar 50/50 stock/bond risk portfolio.

Option Volume Imbalance as a predictor for equity market returns

N. Michael, M. Cucuringu, and S. Howison

University of Oxford

The work investigates the use of the normalised imbalance between option volumes corresponding to positive and negative market views, as a predictor for directional price movements in the spot market. Via a nonlinear analysis, and using a decomposition of aggregated volumes into five distinct market participant classes, they find strong signs of predictability of excess

market overnight returns. The strongest signals come from Market-Maker volumes. Among other findings, they demonstrate that most of the predictability stems from high-implied-volatility option contracts, and that the informational content of put option volumes is greater than that of call options.

Multivariate Realised Volatility Forecasting with Graph Neural Network

Q. Chen, and C.Robert
Ecole Polytechnique, ENSAE

The existing publications demonstrate that the limit order book data is useful in predicting short-term volatility in stock markets. Since stocks are not independent, changes on one stock can also impact other related stocks. In this paper, the authors are interested in forecasting short-term realised volatility in a multivariate approach based on limit order book data and relational data. To achieve this goal, the authors introduce Graph Transformer Network for Volatility Forecasting. The model allows to combine limit order book features and an unlimited number of temporal and cross-sectional relations from different sources. Through experiments based on about 500 stocks from SP 500 index, authors find a better performance for their model than for other benchmarks.

Equity Convexity and Unconventional Monetary Policy

L. Stagnol, M. Abdallah, and P. Herfroy
Amundi

This works attempts to gain an understanding of the drivers of stock convexity, also known as gamma. First, using a bottom-up – firm-level – approach, they showcase that stock fundamentals, in particular metrics related to value (captured by the price-to-book ratio) and historical volatility, allow us to efficiently discriminate between convex and concave stocks. Building on this result, they investigate the ties between the gamma premium and traditional risk factors. Second, they adopt a top-down – macroeconomic driven – framework, to understand which economic environment is the most favorable to convexity: authors highlight the importance of the short-term interest rate, the VIX, but also oil price dynamics in a univariate cointegrating vector. These variables share long-term relationships. They then evaluate the ability of different models to forecast future convexity premium dynamics. Finally, they seek to employ these signals in the design of a systematic long convexity strategy and show that it leads to significantly improved risk-adjusted returns compared to a capitalisation-weighted benchmark, especially in turbulent markets. Convexity exposure appears particularly relevant in a context of monetary policy normalisation.

Volatility (Dis)Connect in International Markets

R. Colacito, M. M. Croce, Y. Liu, and I. Shaliastovich
University of North Carolina, CEPR

Lack of co-movement between consumption growth differentials and real exchange rates is a traditional indicator of a disconnect of foreign exchange markets from economic fundamentals (Backus-Smith 1993 anomaly). The work presents novel evidence for the (dis)connect between the volatilities, as opposed to their levels. The volatility correlations are below one, but they are larger than the level correlations. In the cross-section of countries, the volatility disconnect weakens for countries with low amount of expected growth risk and high amount of volatility risk. The authors provide an explanation of their empirical findings based on international risk-sharing of both expected growth and volatility news shocks.

Cheap Options Are Expensive

A. Boulatov, A. Eisdorfer, A. Goyal, and A. Zhdanov
Swiss Finance Institute

The work shows that (partial) inattention to the underlying stock prices generates a demand pressure for options on low-priced stocks, resulting in overpricing of such options. Empirically, they find that delta-hedged options on low-priced stocks underperform those on high-priced stocks by 0.63% per week for calls and 0.36% for puts. Natural experiments corroborate this finding; options tend to become relatively more expensive following stock splits, and options on mini-indices are overpriced relative to options written on otherwise identical regular indices. Skewness preference does not explain their results.

Bridging P-Q Modeling Divide with Factor HJM Modeling Framework

A. Lyashenko, and Y. Goncharov
Quantitative Risk Management Inc.

In this work, authors show how the factor modeling approach widely used to model yield curve evolution in real-world applications can be adapted to pricing applications using the Musiela HJM modeling framework. The resulting risk-neutral modeling framework combines the intuitiveness and computational efficiency of the factor modeling approach with the rigor of risk-neutral term structure pricing models.

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, Maxmillian Arhens, Marcello Monga, and Patrick Chang.