May 2022, No 3

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 minimalist model by NBER, which isolates overlooked economic forces impacting the Russian Ruble during the Ukraine invasion; multi-step forecasting of FX rates using Bayesian optimisation methods by authors from UCSB and Northwestern; extension of the Fundamental Asset Pricing Theory by authors from Zurich; work from Ecole Polytechnique incorporating climate risk into credit risk models; and the work linking unconventional monetary policy with equity convexity by authors from Amundi.

Machine Learning for Finance

The Virtue of Complexity in Machine Learning Portfolios.

G. Pinter, C. Wang, and J. Zou
AQR, Ecole Polytechnique & Yale University

The work theoretically characterises the behavior of machine learning portfolios in the high complexity regime, i.e. when the number of parameters exceeds the number of observations. They demonstrate a surprising "virtue of complexity", and empirically document it in the US equity market timing strategies. High complexity models deliver economically large and statistically significant out-of-sample portfolio gains relative to simpler models, due in large part to their remarkable ability to predict recessions.

Convergence and Implicit Regularization Properties of Gradient Descent for Deep Residual Networks.

R. Cont, A. Rossier, and R. Xu
University of Oxford & Southern California

The work proves linear convergence of gradient descent to a global minimum for the training of deep residual networks with constant layer width and smooth activation function. They further show that the trained weights, as a function of the layer index, admits a scaling limit which is Hölder continuous as the depth of the network tends to infinity. The proofs are based on non-asymptotic estimates of the loss function and of norms of the network weights along the gradient descent path. We illustrate the relevance of their theoretical results to practical settings using detailed numerical experiments on supervised learning problems.

Minimax-Regret Climate Policy with Deep Uncertainty in Climate Modeling and Intergenerational Discounting.

S. DeCanio, C. Manski, and A. Sanstad
UCSB, Northwestern Uni. & NBER

Integrated assessment models have become the primary tools for comparing climate policies that seek to reduce greenhouse gas emissions. Policy comparisons have often been performed by considering a planner who seeks to make optimal trade-offs between the costs of carbon

abatement and the economic damages from climate change. Studying climate policy as an optimal control problem presumes that a planner knows enough to make optimisation feasible, but physical and economic uncertainties abound. Manski, Sanstad, and DeCanio (2021) studied use of the minimax-regret (MMR) decision criterion to account for deep uncertainty in climate modeling. The work studies the choice of climate policy that minimizes maximum regret with deep uncertainty regarding both the correct climate model and the appropriate time discount rate to use in intergenerational assessment of policy consequences. The analysis specifies a range of discount rates to express both empirical and normative uncertainty about the appropriate rate. The findings regarding climate policy are novel and informative. The MMR analysis points to use of a relatively low discount rate of 0.02 for climate policy. The MMR decision rule keeps the maximum future temperature increase below 2°C above the 1900-10 level for most of the parameter values used to weight costs and damages.

Forecasting Foreign Exchange Rates With Parameter-Free Regression Networks Tuned By Bayesian Optimisation.

L. Li, P. Matt, and C. Heumann
UCSB, Northwestern University & NBER

The article is concerned with the problem of multi-step financial time series forecasting of Foreign Exchange (FX) rates. To address this problem, they introduce a parameter-free regression network termed RegPred Net. The exchange rate to forecast is treated as a stochastic process. It is assumed to follow a generalization of Brownian motion and the mean-reverting process referred to as the generalized Ornstein-Uhlenbeck (OU) process, with time-dependent coefficients. Using past observed values of the input time series, these coefficients can be regressed online by the cells of the first half of the network (Reg). The regressed coefficients depend only on - but are very sensitive to - a small number of hyperparameters required to be set by a global optimisation procedure for which, Bayesian optimisation is an adequate heuristic. Thanks to its multi-layered architecture, the second half of the regression network (Pred) can project time-dependent values for the OU process coefficients and generate realistic trajectories of the time series. Predictions can be easily derived in the form of expected values estimated by averaging values obtained by Monte Carlo simulation. The forecasting accuracy on a 100 days horizon is evaluated for several of the most important FX rates such as EUR/USD, EUR/CNY, and EUR/GBP. Their experimental results show that the RegPred Net significantly outperforms ARMA, ARIMA, LSTMs, and Autoencoder-LSTM models in this task.

Policy Gradient Stock GAN for Realistic Discrete Order Data Generation in Financial Markets.

M. Hirano, H. Sakaji, and K. Izumi
University of Tokyo

This study proposes a new generative adversarial network (GAN) for generating realistic orders in financial markets. In some previous works, GANs for financial markets generated fake orders in continuous spaces because of GAN architectures' learning limitations. However, in reality, the orders are discrete, such as order prices, which has minimum order price unit, or order types. Thus, they change the generation method to place the generated fake orders into discrete spaces in this study. Because this change disabled the ordinary GAN learning algorithm, this study employed a policy gradient, frequently used in reinforcement learning, for the learning algorithm. Through their experiments, they show that their proposed model outperforms previous models in generated order distribution. As an additional benefit of introducing the policy gradient, the entropy of the generated policy can be used to check GAN's learning status.

Risk Management

Banking with an Overreaction.

A.K. Parija, and V. Mathur
Indian Institute of Management

The work analyses the expectations of U.S. banks using the construct of "expected loss" from loan loss provisions and find evidence of departure from rational expectations, in particular overreaction to losses observed in the recent past. Their findings are explained by a model of belief distortion rather than any strategic accrual manipulation. To capture the unjustified belief (or sentiment) of banks, they construct a "Credit Risk Sentiment" (CRS) metric. Using CRS, they find the following evidence of neglect of risks by banks: improvement in sentiment is associated with a contemporaneous rise in credit growth and poor financial performance (as measured by an increase in bad loans and decline in profitability) in the future. But despite the neglect of risks by banks, shareholders fail to discipline them; CRS predicts lower equity returns in subsequent one to three years.

Relevance of Wrong-Way Risk in Funding Valuation Adjustments.

T. Zwaard, L.A. Grzelak, and C.W. Oosterlee
Utrecht University

March 2020, the world was thrown into a period of financial distress. This manifested through increased

uncertainty in the financial markets. Many interest rates collapsed and funding spreads surged significantly, which increased due to the market turmoil. In light of these events, it is key to understand and model Wrong-Way Risk (WWR) in a Funding Valuation Adjustment (FVA) context. WWR might currently not be incorporated in FVA calculations in banks' Valuation Adjustment (xVA) engines. However, they demonstrate that WWR effects are non-negligible in FVA modeling from a risk-management perspective. They look at the impact of various modeling choices such as including the default times of the relevant parties and they consider different choices of funding spread. A case study is presented for interest rate derivatives.

A Minimalist Model for the Ruble During the Russian Invasion of Ukraine.

G. Lorenzoni, and I. Werning
NBER

This note isolates an overlooked economic force for the Ruble to appreciate in response to international sanctions limiting exports to Russia. The economic intuition is that when Russians are unable to buy the mix of foreign goods they wish, then foreign goods becomes less attractive, increasing the demand for domestic goods; to reestablish an equilibrium a real appreciation is needed to raise the relative price of domestic goods and incentivizing the accumulation of foreign assets and the import from non-sanctioning countries. They also review well-known forces for a depreciation (e.g. Russian export reduction). Their analysis emphasizes that the exchange rate is an inadequate signal of welfare impacts and the effectiveness of sanctions.

The Climate Extended Risk Model (CERM).

J. Garnier, J. Gaudemet, and A. Gruz
Ecole polytechnique & Green RWA

This paper addresses estimates of climate risk embedded within a bank credit portfolio. The proposed Climate Extended Risk Model (CERM) adapts well known credit risk models and makes it possible to calculate incremental credit losses on a loan portfolio that are rooted into physical and transition risks. The paper provides detailed description of the model hypotheses and steps.

Fundamental Theorem of Asset Pricing with Acceptable Risk in Markets with Frictions.

M. Arduca, and C. Munari
LUISS University & University of Zurich

They study the range of prices at which a rational agent should contemplate transacting a financial contract outside a given securities market. Trading is subject to nonproportional transaction costs and portfolio constraints and full replication by way of market instruments is not always possible. Rationality is defined in terms of consistency with market prices and acceptable risk thresholds. They obtain a direct and a dual description of market-consistent prices with acceptable risk. The dual characterization requires an appropriate extension of the classical Fundamental Theorem of Asset Pricing where the role of arbitrage opportunities is played by good deals, i.e., costless investment opportunities with acceptable risk-reward tradeoff. In particular, they highlight the importance of scalable good deals, i.e., investment opportunities that are good deals regardless of their volume.

The Safe Asset Frontier.

M. Arduca, and C. Munari
Deutsche Bundesbank & University of Auckland

They identify the frontier between safe and unsafe assets and show how the growth rate of the economy and its fiscal capacity interact with differences of opinion amongst investors to determine the safe asset equilibrium. Multiple equilibria emerge in their set-up due to strategic complementarities across counterparties, and the safety of the bond depends on the extent to which investors' opinions diverge from the credit rating of the asset.

Derivatives & Volatility

Improved Robust Price Bounds for Multi-Asset Derivatives under Market-Implied Dependence Information.

J. Ansari, E. Lütkebohmert, A. Neufeld, and J. Sester
University of Freiburg & Nanyang Technological University

The work shows how inter-asset dependence information derived from observed market prices of liquidly traded options can lead to improved model-free price bounds for multi-asset derivatives. Depending on the type of the observed liquidly traded option, they either extract correlation information, or they derive restrictions on the set of admissible copulas that capture the inter-asset dependencies. To compute the resultant price bounds for some multi-asset options of interest, they apply a modified martingale optimal transport approach. In particular, they derive an adjusted pricing-hedging

duality. Several examples based on simulated and real market data illustrate the improvement of the obtained price bounds and thus provide evidence for the relevance and tractability of their approach.

On Model Robustness of the Regime Switching Approach for Pegged FX Markets.

Y. Zhang, and S. Drapeau
Jiao Tong University

The work test the robustness of the regime switching model for pegged markets introduced by Drapeau et al. In particular, there are two disputable underlying assumptions: (1) a Black and Scholes model with low volatility for the pre-depegging regime, and (2) a thin tail distribution—Exponential type—for the time of the depegging. For the pre-depegging regime, they consider a bounded model within the peg—from Ingersoll and Rady. For the depegging time, they consider fat tail distributions more in line with catastrophic events—Pareto/Fréchet. They derive the option prices formula for each combination of these models. They then calibrate to option data from USD-HKD as well as EUR-CHF. In comparison to the benchmark model in Drapeau et al., it turns out that the relevant resulting characteristics—probability of a depegging before maturity, appreciation/depreciation at the depegging time as well as post-depegging volatility—are strongly robust in terms of model choice for this regime switching approach. However, from a term structure perspective, fat tail distributions fit the data significantly better and provide more rational depegging probabilities for short and long maturities.

Cross Impact in Derivative Markets.

M. Tomas, U. Mastromatteo, and M. Benzaquen
Ecole Polytechnique & CFM

They test the robustness of the regime switching model for pegged markets introduced by Drapeau et al. In particular, there are two disputable underlying assumptions: (1) a Black and Scholes model with low volatility for the pre-depegging regime and (2) a thin tail distribution—Exponential type—for the time of the depegging. For the pre-depegging regime, they consider a bounded model within the peg—from Ingersoll and Rady. For the depegging time, they consider fat tail distributions more in line with catastrophic events—Pareto/Fréchet. They derive the option prices formula for each combination of these models. They then calibrate to option data from USD-HKD as well as EUR-CHF. In comparison to the benchmark model in Drapeau et al., it turns out that the relevant resulting characteristics—probability of a depegging before maturity, appreciation/depreciation at the depegging time as well as post-depegging volatility—are strongly robust in terms of model choice for this regime switching approach. However, from a term structure perspective, fat tail distributions fit the data significantly better and provide more rational depegging probabilities for short and long maturities..

Local Volatility Under Rough Volatility.

F. Bourgey, F. Marco, P. Friz, P. Pigato
Ecole Polytechnique & TUB

Several asymptotic results for the implied volatility generated by a rough volatility model have been obtained in recent years (notably in the small-maturity regime), providing a better understanding of the shapes of the volatility surface induced by rough volatility models, and supporting their calibration power to SP500 option data. Rough volatility models also generate a local volatility surface, via the so-called Markovian projection of the stochastic volatility. The work complements existing results on the implied volatility by studying the asymptotic behavior of the local volatility surface generated by a class of rough stochastic volatility models, encompassing the rough Bergomi model. Notably, they observe that the celebrated "1/2 skew rule" linking the short-term at-the-money skew of the implied volatility to the short-term at-the-money skew of the local volatility, a consequence of the celebrated "harmonic mean formula" of Berestycki et al., is replaced by a new rule: the ratio of the at-the-money implied and local volatility skews tends to the constant $1/(H + 3/2)$ (as opposed to the constant $1/2$), where H is the regularity index of the underlying instantaneous volatility process.

Quantitative Finance & Economics

Equity Convexity and Unconventional Monetary Policy.

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

This work is aimed at 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: they 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 capitalization-weighted benchmark, especially in turbulent markets. Convexity exposure appears particularly relevant in a context of monetary policy normalization.

Scale Dependencies and Self-Similarity Through Wavelet Scattering Covariance.

R. Morel, G. Rochette, R. Leonarduzzi, J.P. Bouchaud, and S. Mallat

CFM

They introduce a scattering covariance matrix which provides non-Gaussian models of time-series having stationary increments. A complex wavelet transform computes signal variations at each scale. Dependencies across scales are captured by the joint covariance across time and scales of complex wavelet coefficients and their modulus. This covariance is nearly diagonalized by a second wavelet transform, which defines the scattering covariance. They show that this set of moments characterizes a wide range of non-Gaussian properties of multi-scale processes. This is analyzed for a variety of processes, including fractional Brownian motions, Poisson, multifractal random walks and Hawkes processes. They prove that self-similar processes have a scattering covariance matrix which is scale invariant. This property can be estimated numerically and defines a class of wide-sense self-similar processes. They build maximum entropy models conditioned by scattering covariance coefficients, and generate new time-series with a microcanonical sampling algorithm. Applications are shown for highly non-Gaussian financial and turbulence time-series.

Bridging the Gap: Decoding the Intrinsic Nature of Time in Market Data.

J. Glattfelder1 and A. Golub

Flov

Intrinsic time is an example of an event-based conception of time, used to analyze financial time series. Here, for the first time, they reveal the connection between intrinsic time and physical time. In detail, they present an analytic relationship which links the two different time paradigms. Central to this discovery are the emergence of scaling laws. Indeed, a novel empirical scaling law is presented, relating to the variability of what is know as overshoots in the intrinsic time framework. To evaluate the validity of the theoretically derived expressions, three time series are analyzed; in detail, Brownian motion and two tick-by-tick empirical currency market data sets (one crypto and one fiat). Finally, the time series analyzed in physical time can be decomposed into their liquidity and volatility components, both only visible in intrinsic time, further highlighting the utility of this temporal kinship.

Information Chasing Versus Adverse Selection.

G. Pinter, C. Wang, and J. Zou

Bank of England

Contrary to the prediction of the classic adverse selection theory, a more informed trader could receive better pricing relative to a less informed trader in over-the-counter financial markets. Dealers chase informed orders to better position their future quotes and avoid winner's curse in subsequent trades. When dealers are perfectly competitive and risk averse, their incentive of information chasing dominates their fear of adverse selection. In a more general setting, information chasing can dominate adverse selection when dealers face differentially informed speculators, while adverse selection dominates when dealers face differentially informed trades from a given speculator. These two seemingly contrasting predictions are supported by empirical evidence from the UK government bond market.

Does Non-linear Factorization of Financial Returns Help Build Better and Stabler Portfolios?

B. Spilak, and W. Hardle

Humboldt University of Berlin & Singapore Management University

A portfolio allocation method based on linear and non-linear latent constrained conditional factors is presented. The factor loadings are constrained to always be positive in order to obtain long-only portfolios, which is not guaranteed by classical factor analysis or PCA. In addition, the factors are to be uncorrelated among clusters in order to build long-only portfolios. Their approach is based on modern machine learning tools: convex Non-negative Matrix Factorization (NMF) and autoencoder neural networks, designed in a specific manner to enforce the learning of useful hidden data structure such as correlation between the assets' returns. Their technique finds lowly correlated linear and nonlinear conditional latent factors which are used to build outperforming global portfolios consisting of cryptocurrencies and traditional assets, similar to hierarchical clustering method. They study the dynamics of the derived non-linear factors in order to forecast tail losses of the portfolios and thus build more stable ones.

Limits of Semistatic Trading Strategies

M. Nutz, J. Wiesel, and L. Zhao

Columbia University

The work show that pointwise limits of semistatic trading

strategies in discrete time are again semistatic strategies. The analysis is carried out in full generality for a two-period model, and under a probabilistic condition for multi-period, multi-stock models. Their result contrasts with a counterexample of Acciaio, Larsson and Schachermayer, and shows that their observation is due to a failure of integrability rather than instability of the semistatic form. Mathematically, their results relate to the decomposability of functions as studied in the context of Schrödinger bridges.

Cryptocurrencies & Blockchain

Optimal Routing for Constant Function Market Makers.

G. Angeris, T. Chitra, A. Evans, and S. Boyd
University of Stanford, Bain & Gauntlet

They consider the problem of optimally executing an order involving multiple cryptoassets, sometimes called tokens, on a network of multiple constant function market makers (CFMMs). When they ignore the fixed cost associated with executing an order on a CFMM, this optimal routing problem can be cast as a convex optimisation problem, which is computationally tractable. When they include the fixed costs, the optimal routing problem is a mixed-integer convex problem, which can be solved using (sometimes slow) global optimisation methods, or approximately solved using various heuristics based on convex optimisation. The optimal routing problem includes as a special case the problem of identifying an arbitrage present in a network of CFMMs, or certifying that none exists.

Graphs-based Learning

Regulating systemic crises: Stemming the Contagion Risk in Networked-Loans Through Deep Graph Learning.

D. Cheng, Z. Niu, J. Li, and C. Jiang
Tongji University

In networked-loans, guarantor enterprises have a legal duty to repay debt to the commercial bank when the guaranteed borrower enterprise defaults (fail to repay). During an economic recession, the risk of defaults may spread like wildfire – the loan network structure could amplify both reach and impact; thus leading to a large-scale corporation defaults even systemic financial crises. The Central Bank urges advanced regulation technology to recognize and act on the contagion risk in order to avoid the "gray rhino". Therefore, they present a novel approach to help the regulators quantify the systemic risk and provide stemming clues. In particular, they report a state-of-the-art graph neural network architecture (iConReg) for detecting and isolating of contagion risk in China's national-wide networked-loans. The overall accuracy of their model reaches over 91% of AUC, which considerably outperforms the compared benchmark methods. By isolating the top 1% of predicted high-risky nodes in the contagion chains, iConReg reports a significant shrink (averaged 25.8%) of loan default rates. Their presented approach opens up a new direction of using deep graph learning techniques to regulate the contagion risk of networked-loans, which enables the authorities to design more prompt prevention measures against systemic financial crises.

Sparsification and Filtering for Spatial-temporal GNN in Multivariate Time-series.

Y. Wang, and T. Aste
UCL

They propose an end-to-end architecture for multivariate time-series prediction that integrates a spatial-temporal graph neural network with a matrix filtering module. This module generates filtered (inverse) correlation graphs from multivariate time series before inputting them into a GNN. In contrast with existing sparsification methods adopted in graph neural network, their model explicitly leverage time-series filtering to overcome the low signal-to-noise ratio typical of complex systems data. They present a set of experiments, where they predict future sales from a synthetic time-series sales dataset. The proposed spatial-temporal graph neural network displays superior performances with respect to baseline approaches, with no graphical information, and with fully connected, disconnected graphs and unfiltered graphs.

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, Patrick Chang, and Marcello Monga.