VOL. 3, NO. 8

OMI Research Newsletter – July 2022

The Oxford-Man Institute of Quantitative Finance (OMI) is a world-leading centre for interdisciplinary research in financial markets. The OMI addresses fundamental problems in quantitative finance with a strong focus on machine learning and data driven models.

*The **OMI Research Newsletter** is a curated summary that covers the most recent research, insights, and tools related to financial decision making.*

Highlights

Research focus. Noteworthy works in this month's edition include: two papers from the OMI, one proposing a novel framework to solve risk-sensitive reinforcement learning problems, where the agent optimises dynamic spectral risk measures; and another studying optimal execution problem in automated market makers. Consistent with the news trend, a meaningfully larger number of works focused on inflation, of which two are noteworthy - a note by the Bank of England, studying the inflation-uncertainty relationship; and the work from University of California and LSE, developing methods to estimate complementary tail probabilities, for persistently very high or very low inflation, using the price of inflation options. Another interesting work is from the Polish Academy of Sciences, which shows that the inter-transaction times in cryptocurrencies show long-range power-law autocorrelations, that lead to multifractility.

Machine Learning for Finance

Conditionally Elicitable Dynamic Risk Measures for Deep Reinforcement Learning

A. COACHE, S. JAIMUNGAL, AND A. CARTEA
UNIVERSITY OF OXFORD UNIVERSITY OF TORONTO

The work proposes a novel framework to solve risk-sensitive reinforcement learning (RL) problems where the agent optimises time-consistent dynamic spectral risk measures. Based on the notion of conditional elicitability, heir methodology constructs (strictly consistent) scoring functions that are used as penalizers in the estimation procedure. Their contribution is threefold: they (i) devise an efficient approach to estimate a class of dynamic spectral risk measures with deep neural networks, (ii) prove that these dynamic spectral risk measures may be approximated to any arbitrary accuracy using deep neural networks, and (iii) develop a risk-sensitive actor-critic algorithm that uses full episodes and does not require any additional nested transitions. They compare their conceptually improved reinforcement learning algorithm with the nested simulation approach and illustrate its performance in two settings: statistical arbitrage and portfolio allocation on both simulated and real data.

Deep Partial Least Squares for Empirical Asset Pricing

M. DIXON, N. POLSON, AND K. GOICOECHEA
ILLINOIS INSTITUTE OF TECH., AND UNI. OF CHICAGO

The work uses deep partial least squares (DPLS) to estimate an asset pricing model for individual stock returns that exploits conditioning information in a flexible and dynamic way while attributing excess returns to a small set of statistical risk factors. The novel contribution is to resolve the non-linear factor structure, thus advancing the current paradigm of deep learning in empirical asset pricing which uses linear stochastic discount factors under an assumption of Gaussian asset returns and factors. This non-linear factor structure is

extracted by using projected least squares to jointly project firm characteristics and asset returns on to a subspace of latent factors and using deep learning to learn the non-linear map from the factor loadings to the asset returns. The result of capturing this non-linear risk factor structure is to characterize anomalies in asset returns by both linear risk factor exposure and interaction effects. Thus the well known ability of deep learning to capture outliers, shed lights on the role of convexity and higher order terms in the latent factor structure on the factor risk premia. On the empirical side, they implement their DPLS factor models and exhibit superior performance to LASSO and plain vanilla deep learning models. Furthermore, their network training times are significantly reduced due to the more parsimonious architecture of DPLS. Specifically, using 3290 assets in the Russell 1000 index over a period of December 1989 to January 2018, they assess their DPLS factor model and generate information ratios that are approximately 1.2x greater than deep learning. DPLS explains variation and pricing errors and identifies the most prominent latent factors and firm characteristics.

Cryptocurrency Bubble Detection: A New Stock Market Dataset, Financial Task Hyperbolic Models.

R. Sawhney. S. Agarwal, V. Mittal, P. Rosso, V. Nanda, and S. Chava

Georgia Tech., Uni. of Illinois & Uni. of Texas

The rapid spread of information over social media influences quantitative trading and investments. The growing popularity of speculative trading of highly volatile assets such as cryptocurrencies and meme stocks presents a fresh challenge in the financial realm. Investigating such "bubbles" - periods of sudden anomalous behavior of markets are critical in better understanding investor behavior and market dynamics. However, high volatility coupled with massive volumes of chaotic social media texts, especially for underexplored assets like cryptocoins pose a challenge to existing methods. Taking the first step towards NLP for cryptocoins, they present and publicly release CryptoBubbles, a novel multispan identification task for bubble detection, and a dataset of more than 400 cryptocoins from 9 exchanges over five years spanning over two million tweets. Further, they develop a set of sequence-to-sequence hyperbolic models suited to this multi-span identification task based on the power-law dynamics of cryptocurrencies and user behavior on social media. They further test the effectiveness of their models under zero-shot settings on a test set of Reddit posts pertaining to 29 meme stocks, which see an increase in trade volume due to social media hype. Through quantitative, qualitative, and zero-shot analyses on Reddit and Twitter spanning cryptocoins and meme-stocks, they show the practical applicability of CryptoBubbles and hyperbolic models.

Risk Management

Dynamic Co-Quantile Regression.

T. Dimitriadis, and Y. Hoga

Heidelberg Uni. & Uni. of Duisburg-Essen

The popular systemic risk measure CoVaR (conditional Value-at-Risk) is widely used in economics and finance. Formally, it is defined as an (extreme) quantile of one variable (e.g., losses in the financial system) conditional on some other variable (e.g., losses in a bank's shares) being in distress and, hence, measures the spillover of risks. This article proposes a dynamic "Co-Quantile Regression", which jointly models VaR and CoVaR semiparametrically. They propose a two-step M-estimator drawing on recently proposed bivariate scoring functions for the pair (VaR, CoVaR). Among others, this allows for the estimation of joint dynamic forecasting models for (VaR, CoVaR). They prove the asymptotic normality of the proposed estimator and simulations illustrate its good finite-sample properties. They apply their co-quantile regression to correct the statistical inference in the existing literature on CoVaR, and to generate CoVaR forecasts for real financial data, which are shown to be superior to existing methods.

Excess Out-of-Sample Risk and Fleeting Modes.

J. Bouchaud, I. Mastromatteo, M. Potters, and K. Tikhonov

CFM

Using Random Matrix Theory, they propose a universal and versatile tool to reveal the existence of "fleeting modes", i.e. portfolios that carry statistically significant excess risk, signalling ex-post a change in the correlation structure in the underlying asset space. Their proposed test is furthermore independent of the "true" (but unknown) underlying correlation structure. They show empirically that such fleeting modes exist both in futures markets and in equity markets. They proposed a metric to quantify the alignment between known factors and fleeting modes and identify momentum as a source of excess risk in the equity space.

Multivariate Quadratic Hawkes Processes – Part I: Theoretical Analysis.

C. Aubrun, M. Benzaquen, and J. Bouchaud
CFM & Ecole Polytechnique

Quadratic Hawkes (QHawkes) processes have proved effective at reproducing the statistics of price changes, capturing many of the stylised facts of financial markets. Motivated by the recently reported strong occurrence of endogenous co-jumps (simultaneous price jumps of several assets) they extend QHawkes to a multivariate framework (MQHawkes), that is considering several financial assets and their interactions. Assuming that quadratic kernels write as the sum of a time-diagonal component and a rank one (trend) contribution, they investigate endogeneity ratios and the resulting stationarity conditions. They then derive the so-called Yule-Walker equations relating covariances and feedback kernels, which are essential to calibrate the MQHawkes process on empirical data. Finally, they investigate the volatility distribution of the process and find that, as in the univariate case, it exhibits power-law behavior, with an exponent that can be exactly computed in some limiting cases.

Jumps, Leverage and Risk Premiums.

T. Bollerslev
Duke University

Jumps in asset prices are ubiquitous, yet the apparent high price of jump risk observed empirically is widely viewed as puzzling. Importantly however, in addition to direct price risks, jumps may also trigger simultaneous changes in other distributional features of asset returns. The author develops a model-free short-time expansions for two alternative risk-neutral variance measures, allowing him to more clearly delineate the role played by jumps in generating both price and variance risks and the pricing of "jump leverage" in particular. Relying on high-frequency SP 500 index and options data, he finds that the jump leverage effect commands a sizeable risk premium.

Diversification Quotients: Quantifying Diversification via Risk Measures.

X. Han, L. Lin, and R. Wang
University of Waterloo

To overcome several limitations of existing diversification indices, they introduce the diversification quotient (DQ). Defined through a parametric family of risk measures, DQ satisfies three natural properties, namely, non-negativity, location invariance and scale invariance, which are shown to be conflicting for any traditional diversification index based on a single risk measure. They pay special attention to the two most important classes of risk measures in banking and insurance, the Value-at-Risk (VaR) and the Expected Shortfall (ES, also called CVaR). DQs based on VaR and ES enjoy many convenient technical properties, and they are efficient to optimize in portfolio selection. By analyzing the popular multivariate models of elliptical and regular varying distributions, they find that DQ can properly capture tail heaviness and common shocks which are neglected by traditional diversification indices. When illustrated with financial data, DQ is intuitive to interpret, and its performance is competitive when contrasted with other diversification methods in portfolio optimisation.

Derivatives & Volatility

How Likely Is an Inflation Disaster?

J. Hilscher, A. Raviv, and R. Reis
Uni. of California, Bar-Ilan Uni. & NBER

The prices of long-dated inflation swap contracts provide a much-used estimate of expected inflation at far horizons. This paper develops the methods to estimate complementary tail probabilities for persistently very high or very low inflation using the prices of inflation options. For the object of interest—inflation disasters at long horizons—they show that three adjustments to conventional measures are crucial: for real payoffs, risk, and horizon. Applying the method to the United States (US) and the Eurozone (EZ) they find that (i) the probability of US deflation in 2011-14 was not very high, (ii) the tail probability of a deflation trap in the EZ post 2015 has been high throughout in spite of varying policies meant to address this issue, as well as shocks, and (iii) there was a significant steady rise in 2021 in the risk of persistent high US inflation, and a sharp rise in 2022 in the EZ.

On the Universality of the Volatility Formation Process: When Machine Learning and Rough Volatility Agree

M. Rosenbaum, and J. Zhang
Ecole Polytechnique

The work presents an LSTM network based on a pooled dataset made of hundreds of liquid stocks aiming to forecast the next daily realized volatility for all stocks. Showing the consistent outperformance of this universal LSTM relative to other asset-specific parametric models, they uncover nonparametric evidences of a universal volatility formation mechanism across assets relating past market realizations, including daily returns and volatilities, to current volatilities. A parsimonious parametric forecasting device combining the rough fractional stochastic volatility and quadratic rough Heston models with fixed parameters results in the same level of performance as the universal LSTM, which confirms the

universality of the volatility formation process from a parametric perspective.

On the Directional Destabilizing Feedback Effects of Option Hedging.

D. Sornette, F. Ulmann, and A. Wehrli
ETH Zurich & Swiss National Bank

The work investigates the feedback effect of option hedging activity on the stability of the price of the underlying. While previous literature has focused on the effect of hedging activity on the volatility of the underlying, this paper focuses on directional instabilities arising from feedback effects. They propose a model in which the drift of the underlying is affected by delta hedging and couple the predictions of this model with an approach to identify short-lived, locally explosive trends in the expected price evolution. They show that such directional instabilities indeed occur with higher or lower intensity (depending on the option parameters and its delta) when option hedging is present, in line with the predictions of their model. Analytical results and synthetic experiments furthermore suggest that the effect of hedging on price stability is strongly asymmetric, with a significantly more pronounced effect on the spot price when the option market maker is short compared to long the option. Since the regime where the market maker is short is the predominant regime empirically, this suggests that option hedging indeed can be expected to impact and even destabilize prices. Using the example of the GameStop stock in early 2021, they investigate the predictions of their model on empirical option positions and document that there are indeed instances where explosive trends due to hedging were more likely. In line with other studies of the event, they however also find only limited evidence that option hedging was a main contributor to the observed large-scale price dislocations in this particular case.

Long Memory in Option Pricing: A Fractional Discrete-time Approach.

M. Augustyniak, A. Badescu, J. Begin, and S. Jayaraman

Uni. of Montreal, SFU & Uni. of Calgary

This article studies the impact of long memory on volatility modelling and option pricing. They propose a general discrete-time pricing framework based on affine multi-component volatility models that admit ARCH(∞) representations. This not only nests a large variety of option pricing models from the literature, but also allows for the introduction of novel covariance-stationary long-memory affine GARCH models. Using an infinite sum characterization of the log-asset price's cumulant generating function, they derive semi-explicit expressions for the valuation of European-style derivatives under a general variance-dependent stochastic discount factor. Moreover, they carry out an extensive empirical analysis using returns and cross sections of SP 500 options over the period 1996 to 2019. Overall, they find that once the informational content from options is incorporated into the parameter estimation process, the inclusion of long-memory dynamics in volatility is beneficial for improving the out-of-sample option pricing performance. The largest improvements in the implied volatility root-mean-square errors generated by their long-memory models occur for options with maturities longer than one year, reaching 33% and 13% when compared to one- and two-component short-memory models, respectively.

Quantitative Finance & Economics

Inflation and Uncertainty: A Note.

G. Pinter
Bank of England

This note studies the inflation-uncertainty relationship in a New Keynesian framework. Inflation in these models can be expressed as the discounted sum of current and expected future real marginal costs. The main point of this note is to highlight that real marginal costs in general equilibrium tend to be a convex function of output. This, ceteris paribus, makes higher uncertainty about future output increase current inflation, which can quantitatively off-set the deflationary effect of uncertainty via the precautionary savings channel.

Agenda Control Under Policy Uncertainty.

S. Callander, and N. McCarty
Stanford University & Princeton University

Models of agenda-setting are central to the positive analysis of political institutions. Elaborations of the classical agenda setting model of Romer and Rosenthal (1978) have long been used to make predictions about policy outcomes and the distribution of influence among different political actors. While the canonical version is based on complete and perfect information about preferences and policy outcomes, some extensions have relaxed these assumptions to include uncertainty about preferences and reversion points. In this paper, they consider a different type of uncertainty: incomplete knowledge of the mapping between policies and outcomes. In describing the optimal agenda setting under this form of uncertainty, they show that it amends substantively the properties that emerge in the Romer-Rosenthal model. They then extend the model dynamically and show the rich dynamics that emerge under policy uncertainty, which also contrast with predictions of the classic model. Over a longer horizon,

they find that agenda control suppresses the incentive of legislators to experiment with policy, leading to less policy learning and worse outcomes than are socially efficient.

Fear (No More) of Floating: Asset Purchases and Exchange Rate Dynamics.

Y. Mimir, and E. Sunel
OECD & ESM

The work provides a theory on currency dynamics, capital flows and conditions for central bank asset purchases to leave room for manoeuvre on conventional monetary policy. Asset purchases ease financial conditions and boost banks' foreign borrowing capacity. The ensuing capital inflows lead to an exchange rate appreciation, which passes through to inflation and facilitates a policy rate easing. They find asset purchases curb the financial amplification of capital flow quantity shocks (but not price shocks) through this transmission mechanism. Their framework sheds light on stable currency dynamics observed after the unprecedented asset purchase announcements in emerging markets during the COVID-19 crisis.

Macroeconomic Drivers and the Pricing of Uncertainty, Inflation, and Bonds.

B. Bok, T. Mertens, and J. Williams
Federal Reserve Bank of NY & San Francisco

This paper analyzes a new stylised fact: According to financial market prices, the correlation between uncertainty shocks, as measured by changes in the VIX, and changes in break-even inflation rates has declined and turned negative over the past quarter century. It rationalises this uncertainty-inflation correlation within a standard New Keynesian model with a lower bound on interest rates combined with a decline in the natural rate of interest. With a lower natural rate, the likelihood of the lower bound binding increased and the effects of uncertainty on the economy became more pronounced. In such an environment, increases in uncertainty raise the possibility that the central bank will be unable to eliminate inflation shortfalls following negative demand shocks. As a result, the observed decline in the correlation between uncertainty and inflation expectations emerges. Average-inflation targeting policies can mitigate the longer-run effects of increases in uncertainty on the real economy.

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

F. Bianchi, S. Ludvigson, and S. Ma
NYU, NBER, Duke & Federal Reserve Board

The work investigates if an alpha-weight angle resulting from unconstrained quadratic portfolio optimisations has an upper bound dependent on the condition number of the covariance matrix. This implies that better conditioned covariance matrices produce weights from unconstrained mean-variance optimisations that are better aligned with each assets expected return. They provide further clarity on the mathematical insights that relate the inequality between the α-weight angle and the condition number and extend the result to include portfolio optimisations with gearing constraints. They provide an extended family of robust optimisations that include the gearing constraints, and discuss their interpretation.

Graph-based Learning

An Entity-weights-based Convolutional Neural Network for Large-sale Complex Knowledge Embedding.

Z. Lei, S. Yang, A. Haq, D. Zhang, A. Francis
Zhengzhou University

Knowledge graph (KG) has increasingly been seen as a significant resource in financial applications (e.g., risk control, auditing and anti-fraud). However, there are few prior studies that focus on multi-relational circles, extracting additional information under the completed KG and selecting similarity measures for knowledge representation. In this paper, they introduce multi-relational circles and propose a novel embedding model, which considers entity weights calculated by PageRank algorithm to improve TransE method. In order to extract additional information, they use entity weights to convert embeddings into an on-map mining problem, and propose a model called CNNe based on entity weights and a convolutional neural network with three hidden layers, which converts vectors of entities, entity weights and relationships into matrices to perform link prediction in the same way as image processing. With the help of ten different similarity measures, it is demonstrated that the choice of distance measure greatly effect the results of the translation embedding models. Moreover, they propose two embedding methods, sMFE and tMFE, to enhance the results using matrix factorization. The complete incidence matrix is first applied to knowledge embedding, which contains the most comprehensive topological properties of the graph. Experimental results on standard benchmark datasets demonstrate that the proposed models are effective. In particular, CNNe achieves a mean rank of 166 less than the baseline

method and an improvement of 2.1% on the proportion of correct entities ranked in the top ten on YAGO3-10 dataset.

Learning to Rank Complex Network Node Based on the Self-supervised Graph Convolution Model.

L. Chen, C. Tingting, and Z. Lixin
Uni. of Shanghai for Science & Technology

Node ranking in complex networks is of great value in many fields, such as identifying opinion leaders in social networks and risky institutions in financial networks. However, traditional ranking methods are built on heuristic rules and may be effective in certain datasets and fail in others. To address the issues, they formulate node ranking on complex networks as a learning to rank problem, and propose a novel model based on self-supervised learning and graph convolution model to rank nodes based on integrated information from node features and network structure.

Network Diversification for a Robust Portfolio Allocation.

M. Jaeger, and D. Marinelli
Munich Reinsurance

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.

Natural Language Processing

Deciphering Federal Reserve Communication via Text Analysis of Alternative FOMC Statements.

T. Doh, D. Song, and S. Yang
Federal Reserve Bank of Kansas City

The work constructs a new measure of monetary policy surprise based on a natural language processing algorithm designed to capture contextual nuances in FOMC statements. Specifically, they exploit cross-sectional variations across alternative FOMC statements to identify the statement's tone, and compare current and previous FOMC statements to obtain the novelty. They use high-frequency bond price movements around FOMC announcements to compute the surprise component of the monetary policy announcement. According to their measure, the stock market declines after unexpected policy tightening. Their text-based approach allows them to assess the counterfactual effects of an altered FOMC statement on the stock market.

Gaining Insights on U.S. Senate Speeches Using a Time Varying Text Based Ideal Point Model.

P. Hofmarcher, and S. Adhikari
University of Salzburg

Estimating political positions of lawmakers has a long tradition in political science. The work presents the time varying text based ideal point model to study the political positions of lawmakers based on text data. In addition to identifying political positions, their model also provides insights into topical contents and their change over time. They use their model to analyse speeches given in the U.S. Senate between 1981 and 2017 and demonstrate how the results allow to conclude that partisanship between Republicans and Democrats increased in recent years. Further they investigate the political positions of speakers over time as well as at a specific point in time to identify speakers which are positioned at the extremes of their political party based on their speeches. The topics extracted are inspected to assess how their term compositions differ in dependence of the political position as well as how these term compositions change over time.

Digital Assets

The Economics of Automated Market Makers.

R. Fritsch, S. Kaser, and R. Wattenhofer
ETH Zurich

This paper studies the question whether automated market maker protocols, such as Uniswap, can sustainably retain a portion of their trading fees for the

protocol. They approach the problem by modelling how to optimally choose a pool's take rate, i.e the fraction of fee revenue that remains with the protocol, in order to maximise the protocol's revenue. The model suggests that if AMMs have a portion of loyal trade volume, they can sustainably set a non-zero take rate, even without losing liquidity to competitors with a zero take rate. Furthermore, they determine the optimal take rate depending on a number of model parameters, including how much loyal trade volume pools have, and how high the competitors' take rates are.

Analysis of inter-transaction time fluctuations in the cryptocurrency Market

J. KWAPIEN, M. WATOREK, M. BEZBRADICA, M. CRANE, T. MAI, AND S. DROZDZ

PAS, CRACOW UNI. OF TECH. DUBLIN CITY UNI.

The work analyses tick-by-tick data representing major cryptocurrencies traded on some different cryptocurrency trading platforms. They focon such quantities like the inter-transaction times, the number of transactions in time unit, the traded volume, and volatility. They show that the inter-transaction times show long-range power-law autocorrelations that lead to multifractality expressed by the right-side asymmetry of the singularity spectra $f(\alpha)$ indicating that the periods of increased market activity are characterised by richer multifractality compared to the periods of quiet market. They show that the inter-transaction times show long-range power-law autocorrelations. These lead to multifractality expressed by the right-side asymmetry of the singularity spectra $f(\alpha)$ indicating that the periods of increased market activity are characterised by richer multifractality compared to the periods of quiet market. They also show that neither the stretched exponential distribution nor the power-law-tail distribution are able to model universally the cumulative distribution functions of the quantities considered in this work. For each quantity, some data sets can be modeled by the former, some data sets by the latter, while both fail in other cases. An interesting, yet difficult to account for, observation is that parallel data sets from different trading platforms can show disparate statistical properties.

Decentralised Finance and Automated Market Making: Execution and Speculation

A. CARTEA, F. DRISSI, AND M. MONGA

UNIVERSITY OF OXFORD

Automated market makers (AMMs) are a new prototype of decentralised exchanges which are revolutionising market interactions. The majority of AMMs are constant product markets (CPMs) where exchange rates are set by a trading function. This work studies optimal trading and statistical arbitrage in CPMs where balancing exchange rate risk and execution costs is key. Empirical evidence shows that execution costs are accurately estimated by the convexity of the trading function. These convexity costs are linear in the trade size and are nonlinear in the depth of liquidity and in the exchange rate. We develop models for when exchange rates form in a competing centralised exchange, in a CPM, or in both venues. Finally, we derive computationally efficient strategies that account for stochastic convexity costs and we showcase their out-of-sample performance.

Acknowledgement

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

References

Aubrun, Cecilia, Michael Benzaquen, and Jean-Philippe Bouchaud. 2023. "Multivariate quadratic Hawkes processes—part I: theoretical analysis." *Quantitative Finance*, 23(5): 741–758.

Bouchaud, Jean-Philippe, Iacopo Mastromatteo, Marc Potters, and Konstantin Tikhonov. 2022. "Excess Out-of-Sample Risk and Fleeting Modes." *arXiv preprint arXiv:2205.01012*.

Callander, Steven, and Nolan McCarty. 2024. "Agenda Control under Policy Uncertainty." *American Journal of Political Science*, 68(1): 210–226.

Cartea, Álvaro, Fayçal Drissi, and Marcello Monga. 2023. "Decentralised finance and automated market making: Execution and speculation." *arXiv preprint arXiv:2307.03499*.

Cartea, Álvaro, Fayçal Drissi, and Marcello Monga. 2024. "Decentralized Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision." *SIAM Journal on Financial Mathematics*, 15(3): 931–959.

Dixon, Matthew F, Nicholas G Polson, and Kemen Goicoechea. 2022. "Deep partial least squares for empirical asset pricing." *arXiv preprint arXiv:2206.10014*.