VOL. 2, NO. 12

OMI Research Newsletter – September 2023

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 contributions in this month's edition include works from the OMI, École Normale Superieure, Collège de France, and University of Toronto. Cartea, Drissi and Monga (2023) study the predictable losses of liquidity providers in constant function markets and concentrated liquidity markets. Zohren et al. (2023) model the interconnections among assets using networks and exploits them to predict future returns. Morel, Mallat and Bouchaud (2023) introduce a new Monte-Carlo method to predict future paths for any generative model. Cheng, Coache and Jaimungal (2023) identify the risk aversion of agents using interactive questioning. Finally, Futter, Horvath and Wiese (2023) study portfolio optimisation using rough path signatures to incorporate path-dependencies in the joint-asset dynamics. Finally, Frey et al. (2023) introduce the first GPU-enabled LOB simulator with implementation in JAX to process thousands of books in parallel.

Machine Learning & Statistics

Network Momentum across Asset Classes

S. PU, S. ROBERTS, X. DONG, S. ZOHREN
OXFORD-MAN INSTITUTE

We investigate the concept of network momentum, a novel trading signal derived from momentum spillover across assets. Initially observed within the confines of pairwise economic and fundamental ties, such as the stock-bond connection of the same company and stocks linked through supply-demand chains, momentum spillover implies a propagation of momentum risk premium from one asset to another. The similarity of momentum risk premium, exemplified by co-movement patterns, has been spotted across multiple asset classes including commodities, equities, bonds and currencies. However, studying the network effect of momentum spillover across these classes has been challenging due to a lack of readily available common characteristics or economic ties beyond the company level. In this paper, we explore the interconnections of momentum features across a diverse range of 64 continuous future contracts spanning these four classes. We utilise a linear and interpretable graph learning model with minimal assumptions to reveal the intricacies of the momentum spillover network. By leveraging the learned networks, we construct a network momentum strategy that exhibits a Sharpe ratio of 1.5 and an annual return of 22%, after volatility scaling, from 2000 to 2022. This paper pioneers the examination of momentum spillover across multiple asset classes using only pricing data, presents a multi-asset investment strategy based on network momentum, and underscores the effectiveness of this strategy through robust empirical analysis.

Path Shadowing Monte-Carlo

R. MORELA, S. MALLAT, J-P BOUCHAUD
ÉCOLE NORMALE SUPERIEURE, COLLÈGE DE FRANCE, CAPITAL FUND MANAGEMENT

We introduce a Path Shadowing Monte-Carlo method, which provides prediction of future paths, given any generative model. At any given date, it averages future quantities over generated price paths whose past history matches, or 'shadows', the actual (observed)

history. We test our approach using paths generated from a maximum entropy model of financial prices, based on a recently proposed multi-scale analogue of the standard skewness and kurtosis called 'Scattering Spectra'. This model promotes diversity of generated paths while reproducing the main statistical properties of financial prices, including stylised facts on volatility roughness. Our method yields state-of-the-art predictions for future realised volatility and allows one to determine conditional option smiles for the S&P500 that outperform both the current version of the Path-Dependent Volatility model and the option market itself

Wisdom of the Crowds or Ignorance of the Masses? A data-driven guide to WSB

V. Semenova, D. Gorduza, W. Wildi, X. Dong, S. Zohren

Oxford-Man Institute

A trite yet fundamental question in economics is: What causes large asset price fluctuations? A tenfold rise in the price of GameStop equity, between the 22nd and 28th of January 2021, demonstrated that herding behaviour among retail investors is an important contributing factor. This paper presents a data-driven guide to the forum that started the hype – WallStreetBets (WSB). Our initial experiments decompose the forum using a large language topic model and network tools. The topic model describes the evolution of the forum over time and shows the persistence of certain topics (such as the market / SP500 discussion), and the sporadic interest in others, such as COVID or crude oil. Network analysis allows us to decompose the landscape of retail investors into clusters based on their posting and discussion habits; several large, correlated asset discussion clusters emerge, surrounded by smaller, niche ones. A second set of experiments assesses the impact of WSB on the market. We show that forum activity has a Granger-causal relationship with the returns of several assets, some of which are now commonly classified as 'meme stocks', while others have gone under the radar. The paper extracts a set of short-term trade signals from posts and long-term (monthly and weekly) trade signals from forum dynamics, and considers their predictive power at different time horizons. In addition to the analysis, the paper presents the dataset, as well as an interactive dashboard, in order to promote further research.

Learning to Learn Financial Networks for Optimising Momentum Strategies

S. Pu, S. Roberts, X. Dong, S. Zohren

Oxford-Man Institute

Network momentum provides a novel type of risk premium, which exploits the interconnections among assets in a financial network to predict future returns. However, the current process of constructing financial networks relies heavily on expensive databases and financial expertise, limiting accessibility for small-sized and academic institutions. Furthermore, the traditional approach treats network construction and portfolio optimisation as separate tasks, potentially hindering optimal portfolio performance. To address these challenges, we propose L2GMOM, an end-to-end machine learning framework that simultaneously learns financial networks and optimises trading signals for network momentum strategies. The model of L2GMOM is a neural network with a highly interpretable forward propagation architecture, which is derived from algorithm unrolling. The L2GMOM is flexible and can be trained with diverse loss functions for portfolio performance, e.g. the negative Sharpe ratio. Backtesting on 64 continuous future contracts demonstrates a significant improvement in portfolio profitability and risk control, with a Sharpe ratio of 1.74 across a 20-year period.

Multitask Learning with No Regret: from Improved Confidence Bounds to Active Learning

P. G. Sessa, P. Laforgue, N. Cesa-Bianchi, A. Krause

ETH Zürich, Università degli Studi di Milano

Multitask learning is a powerful framework that enables one to simultaneously learn multiple related tasks by sharing information between them. Quantifying uncertainty in the estimated tasks is of pivotal importance for many downstream applications, such as online or active learning. In this work, we provide novel multitask confidence intervals in the challenging agnostic setting, i.e., when neither the similarity between tasks nor the tasks' features are available to the learner. The obtained intervals do not require i.i.d. data and can be directly applied to bound the regret in online learning. Through a refined analysis of the multitask information gain, we obtain new regret guarantees that, depending on a task similarity parameter, can significantly improve over treating tasks independently. We further propose a novel online learning algorithm that achieves such improved regret without knowing this parameter in advance, i.e., automatically adapting to task similarity. As a second key application of our results, we introduce a novel multitask active learning setup where several tasks must be simultaneously optimised, but only one of them can be queried for feedback by the learner at each round. For this problem, we design a no-regret algorithm that uses our confidence intervals to decide which task should be queried. Finally, we empirically validate our bounds and algorithms on synthetic and real-world (drug discovery) data.

Eliciting Risk Aversion with Inverse Reinforcement Learning via Interactive Questioning

Z. Cheng, A. Coache, S. Jaimungal
University of Toronto

This paper proposes a novel framework for identifying an agent's risk aversion using interactive questioning. Our study is conducted in two scenarios: a one-period case and an infinite horizon case. In the one-period case, we assume that the agent's risk aversion is characterized by a cost function of the state and a distortion risk measure. In the infinite horizon case, we model risk aversion with an additional component, a discount factor. Assuming the access to a finite set of candidates containing the agent's true risk aversion, we show that asking the agent to demonstrate her optimal policies in various environment, which may depend on their previous answers, is an effective means of identifying the agent's risk aversion. Specifically, we prove that the agent's risk aversion can be identified as the number of questions tends to infinity, and the questions are randomly designed. We also develop an algorithm for designing optimal questions and provide empirical evidence that our method learns risk aversion significantly faster than randomly designed questions in simulations. Our framework has important applications in robo-advising and provides a new approach for identifying an agent's risk preferences.

JAX-LOB: A GPU-Accelerated limit order book simulator to unlock large scale reinforcement learning for trading

S. Frey, K. Li, P. Nagy, S. Sapora, C. Lu, S. Zohren, J. Foerster, A. Calinescu
Oxford-Man Institute, University of Oxford

Financial exchanges across the world use limit order books (LOBs) to process orders and match trades. For research purposes it is important to have large scale efficient simulators of LOB dynamics. LOB simulators have previously been implemented in the context of agent-based models (ABMs), reinforcement learning (RL) environments, and generative models, processing order flows from historical data sets and hand-crafted agents alike. For many applications, there is a requirement for processing multiple books, either for the calibration of ABMs or for the training of RL agents. We showcase the first GPU-enabled LOB simulator designed to process thousands of books in parallel, with a notably reduced per-message processing time. The implementation of our simulator – JAX-LOB – is based on design choices that aim to best exploit the powers of JAX without compromising on the realism of LOB-related mechanisms. We integrate JAX-LOB with other JAX packages, to provide an example of how one may address an optimal execution problem with reinforcement learning, and to share some preliminary results from end-to-end RL training on GPUs.

Mathematical Finance

Predictable Losses of Liquidity Provision in Constant Function Markets and Concentrated Liquidity Markets

Á. Cartea, F. Drissi, M. Monga
Oxford-Man Institute

We introduce a new comprehensive and model-free measure for the unhedgeable and predictable loss (PL) incurred by liquidity providers in constant function markets (CFMs) and in concentrated liquidity markets. PL compares the value of the LP's holdings in the CFM liquidity pool (assuming no fee revenue) with that of a self-financing portfolio that (i) continuously replicates the dynamic holdings of the LP in the pool to offset the market risk of the LP's position, and (ii) invests in a risk-free account. We provide closed-form formulae for PL in CFMs with and without concentrated liquidity, and show that the losses stem from two sources: convexity cost, which depends on liquidity taking activity and the convexity of the pool's trading function; and opportunity cost, which is due to locking the LP's assets in the pool. For liquidity providers, PL is the appropriate measure to assess the cost of liquidity provision in CFMs, so that fees and compensation to LPs provide the right incentives for a well-functioning market. When prices form outside of the pool, we show that PL is reduced when liquidity taking is costly, i.e., when the convexity of the pool's trading function is high. On the other hand, when prices form in the pool, PL is reduced when liquidity taking is cheap, i.e., when the convexity of the trading function is low. Finally, we use Uniswap v3 and Binance transaction data to compute PL and fees collected by LPs and show that, at present, liquidity provision in CFMs is a loss-leading activity.

Passport Option Machine Learning-powered Pricing of the Multidimensional Passport Option

J. Teichmann, H. Wutte
ETH Zurich

Introduced in the late 90s, the passport option gives its holder the right to trade in a market and receive any positive gain in the resulting traded account at maturity. Pricing the option amounts to solving a stochastic control problem that for $d > 1$ risky assets remains an open problem. Even in a correlated Black-Scholes (BS) market with $d = 2$ risky assets, no optimal trading strategy has been derived in closed form. In this paper, we derive a discrete-time solution for multi-dimensional BS markets with uncorrelated assets. Moreover, inspired

by the success of deep reinforcement learning in, e.g., board games, we propose two machine learning-powered approaches to pricing general options on a portfolio value in general markets. These approaches prove to be successful for pricing the passport option in one-dimensional and multi-dimensional uncorrelated BS markets.

Understanding the least well-kept secret of high-frequency trading

S. Pulido, M. Rosenbaum, E. Sfendourakis
Universite Paris-Saclay, Ecole polytechnique

Volume imbalance in a limit order book is often considered as a reliable indicator for predicting future price moves. In this work, we seek to analyse the nuances of the relationship between prices and volume imbalance. To this end, we study a market-making problem which allows us to view the imbalance as an optimal response to price moves. In our model, there is an underlying efficient price driving the mid-price, which follows the model with uncertainty zones. A single market maker knows the underlying efficient price and consequently the probability of a mid-price jump in the future. She controls the volumes she quotes at the best bid and ask prices. Solving her optimisation problem allows us to understand endogenously the price-imbalance connection and to confirm in particular that it is optimal to quote a predictive imbalance. The value function of the market maker's control problem can be viewed as a family of functions, indexed by the level of the market maker's inventory, solving a coupled system of PDEs. We show existence and uniqueness of classical solutions to this coupled system of equations. In the case of a continuous inventory, we also prove uniqueness of the market maker's optimal control policy.

Equilibria and incentives for illiquid auction markets

J. Derchu, D. Kavvathas, T. Mastrolia, M. Rosenbaum
École Polytechnique, UC Berkeley

We study a toy two-player game for periodic double auction markets to generate liquidity. The game has imperfect information, which allows us to link market spreads with signal strength. We characterize Nash equilibria in cases with or without incentives from the exchange. This enables us to derive new insights about price formation and incentives design. We show in particular that without any incentives, the market is inefficient and does not lead to any trade between market participants. We however prove that quadratic fees indexed on each players half spread leads to a transaction and we propose a quantitative value for the optimal fees that the exchange has to propose in this model to generate liquidity.

The Paradox of Adversarial Liquidation in Decentralised Lending

S. N. Cohen, M. Sabate Vidales, Ł. Szpruch, M. G. Delaunay
University of Oxford, University of Edinburgh

We present a holistic framework for risk and reward management for decentralised lending protocols, such as AAVE or Compound. This framework highlights tensions and trade-offs in protocol design and the choice of various risks and incentives parameters. Using this, we identify the fundamental paradox of adversarial liquidation in decentralised lending: If the rewards to liquidators are chosen such that there is an incentive for liquidators to act when a loan is in distress (as is needed for well functioning of the protocol), then it is also the case that liquidators are incentivised to manipulate prices (through front-running the liquidation process), leading to suboptimal outcomes for borrowers.

Detecting Financial Market Manipulation with Statistical Physics Tools

H. Li, M. Polukarov, C. Ventre
King's College London

We take inspiration from statistical physics to develop a novel conceptual framework for the analysis of financial markets. We model the order book dynamics as a motion of particles and define the momentum measure of the system as a way to summarise and assess the state of the market. Our approach proves useful in capturing salient financial market phenomena: in particular, it helps detect the market manipulation activities called spoofing and layering. We apply our method to identify pathological order book behaviours during the flash crash of the LUNA cryptocurrency, uncovering widespread instances of spoofing and layering in the market. Furthermore, we establish that our technique outperforms the conventional Z-score- based anomaly detection method in identifying market manipulations across both LUNA and Bitcoin cryptocurrency markets.

Optimal Robust Reinsurance with Multiple Insurers

E. KROELL, S. JAIMUNGAL, S. M. PESENTI
UNIVERSITY OF TORONTO

We study a reinsurer who faces multiple sources of model uncertainty. The reinsurer offers contracts to n insurers whose claims follow different compound Poisson processes. As the reinsurer is uncertain about the insurers' claim severity distributions and frequencies, they design reinsurance contracts that maximise their expected wealth subject to an entropy penalty. Insurers meanwhile seek to maximise their expected utility without ambiguity. We solve this continuous-time Stackelberg game for general reinsurance contracts and find that the reinsurer prices under a distortion of the barycentre of the insurers' models. We apply our results to proportional reinsurance and excess-of-loss reinsurance contracts, and illustrate the solutions numerically. Furthermore, we solve the related problem where the reinsurer maximises, still under ambiguity, their expected utility and compare the solutions.

Signature Trading: A Path-Ddependent Extension of the Mean-Variance Framework with Exogenous Signals

O. FUTTER, B. HORVATH, M. WIESE
OXFORD-MAN INSTITUTE, IMPERIAL COLLEGE LONDON

In this article we introduce a portfolio optimisation framework, in which the use of rough path signatures (Lyons, 1998) provides a novel method of incorporating path-dependencies in the joint signal-asset dynamics, naturally extending traditional factor models, while keeping the resulting formulas lightweight, tractable and easily interpretable. Specifically, we achieve this by representing a trading strategy as a linear functional applied to the signature of a path (which we refer to as "Signature Trading" or "Sig-Trading"). This allows the modeller to efficiently encode the evolution of past time-series observations into the optimisation problem. In particular, we derive a concise formulation of the dynamic mean-variance criterion alongside an explicit solution in our setting, which naturally incorporates a drawdown control in the optimal strategy over a finite time horizon. Secondly, we draw parallels between classical portfolio strategies and Sig-Trading strategies and explain how the latter leads to a pathwise extension of the classical setting via the "Signature Efficient Frontier". Finally, we give explicit examples when trading under an exogenous signal as well as examples for momentum and pair-trading strategies, demonstrated both on synthetic and market data. Our framework combines the best of both worlds between classical theory (whose appeal lies in clear and concise formulae) and between modern, flexible data-driven methods (usually represented by ML approaches) that can handle more realistic datasets. The advantage of the added flexibility of the latter is that one can bypass common issues such as the accumulation of heteroskedastic and asymmetric residuals during the optimisation phase. Overall, Sig-Trading combines the flexibility of data-driven methods without compromising on the clarity of the classical theory and our presented results provide a compelling toolbox that yields superior results for a large class of trading strategies.

Joint Calibration of Local Volatility Models with Stochastic Interest Rates using Semimartingale Optimal Transport

B. JOSEPH, G. LOEPER, J, OBŁÓJ
UNIVERSITY OF OXFORD

We develop and implement a non-parametric method for joint exact calibration of a local volatility model and a correlated stochastic short rate model using semimartingale optimal transport. The method relies on the duality results established in Joseph, Loeper, and Obłój, 2023 and jointly calibrates the whole equity-rate dynamics. It uses an iterative approach which starts with a parametric model and tries to stay close to it, until a perfect calibration is obtained. We demonstrate the performance of our approach on market data using European SPX options and European cap interest rate options. Finally, we compare the joint calibration approach with the sequential calibration, in which the short rate model is calibrated first and frozen.

Economics

Adapting to Game Trees in Zero-Sum Imperfect Information Games

C. FIEGEL, P. MENARD, T. KOZUNO, R M.UNOS, V. PERCHET, M. VALKO
ENSAE, ENS LYON, DEEPMIND

Imperfect information games (IIG) are games in which each player only partially observes the current game state. We study how to learn - optimal strategies in a zero-sum IIG through self-play with trajectory feedback. We give a problem-independent lower bound $\tilde{\mathcal{O}}(H(A_{\mathcal{X}} + B_{\mathcal{Y}})/\epsilon^2)$ on the required number of realizations to learn these strategies with high probability, where H is the length of the game, AX and BY are the total number of actions for the two players. We also propose two Follow the Regularised leader (FTRL) algorithms for this set- ting: Balanced FTRL which matches this lower bound, but requires the knowledge of the informa- tion set structure beforehand to define the regularisation; and Adaptive FTRL which needs $\tilde{\mathcal{O}}(H^2(A_{\mathcal{X}} + B_{\mathcal{Y}})/\epsilon^2)$

realizations without this requirement by progressively adapting the regularisation to the observations.

Revealed Preferences for Dynamically Inconcsistent Models

D. Echenique, G. Tserenjigmid
UC Berkeley

We study the testable implications of models of dynamically inconsistent choices when planned choices are unobservable, and thus only "on path" data is available. First, we discuss the approach in Blow, Browning, and Crawford (2021), who characterize first-order rationalizability of the model of quasi-hyperbolic discounting. We show that the first-order approach does not guarantee rationalizability by means of the quasi-hyperbolic model. This motivates consideration of an abstract model of intertemporal choice, under which we provide a characterization of different behavioral models – including the naive and sophisticated paradigms of dynamically inconsistent choice.

Multiplayer War of Attrition with Asymmetric Private Information

H. Li, J. Zhang
Yale University, Columbia University

This paper analyzes a multiplayer war of attrition game with asymmetric type distributions in the setting of private provision of an indivisible public good. In the unique equilibrium, asymmetry leads to a stratified behavior pattern where one agent provides the good instantly with a positive probability, while each of the others has no probability of provision before a certain moment and this moment is idiosyncratic for different agents. Comparative statics show that an agent with less patience, lower cost of provision, and higher reputation in valuation provides the good uniformly faster. The cost of delay is mainly determined by the strongest type of the instant-exit agent. This paper investigates two types of variations of asymmetry: raising the strongest type tends to improve efficiency, whereas controlling the strongest type associates the effect of asymmetry with the sign of an intuitive measure of "the cost of symmetry".

Acknowledgement

The OMI Research Newsletter was curated by Álvaro Cartea, Director of the Oxford-Man Institute, and Fayçal Drissi, postdoctoral researcher at the institute. Paper suggestions are from fellow Ph.D. students at the institute: Álvaro Arroyo, Patrick Chang, Gerardo Durán-Martín, and Marcello Monga.

References

Arroyo, Alvaro, Alvaro Cartea, Fernando Moreno-Pino, and Stefan Zohren. 2023. "Deep Attentive Survival Analysis in Limit Order Books: Estimating Fill Probabilities with Convolutional-Transformers." *arXiv preprint arXiv:2306.05479*.

Cartea, Álvaro, Fayçal Drissi, and Marcello Monga. 2022. "Decentralised Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision." *Available at SSRN 4273989*.

Cartea, Álvaro, Fayçal Drissi, and Marcello Monga. 2023. "Predictable Losses of Liquidity Provision in Constant Function Markets and Concentrated Liquidity Markets." *Available at SSRN 4541034*.

Cheng, Ziteng, Anthony Coache, and Sebastian Jaimungal. 2023. "Eliciting Risk Aversion with Inverse Reinforcement Learning via Interactive Questioning." *arXiv preprint arXiv:2308.08427*.

Cohen, Samuel N, Marc Sabate-Vidales, Lukasz Szpruch, and Mathis Gontier Delaunay. 2023. "The Paradox of Adversarial Liquidation in Decentralised Lending." *Available at SSRN 4540333*.

Derchu, Joffrey, Dimitrios Kavvathas, Thibaut Mastrolia, and Mathieu Rosenbaum. 2023. "Equilibria and incentives for illiquid auction markets." *arXiv preprint arXiv:2307.15805*.

Echenique, Federico, and Gerelt Tserenjigmid. 2023. "Revealed preferences for dynamically inconsistent models." *arXiv preprint arXiv:2305.14125*.

Fiegel, Côme, Pierre Ménard, Tadashi Kozuno, Rémi Munos, Vianney Perchet, and Michal Valko. 2023. "Adapting to game trees in zero-sum imperfect information games." 10093–10135, PMLR.

Frey, Sascha, Kang Li, Peer Nagy, Silvia Sapora, Chris Lu, Stefan Zohren, Jakob Foerster, and Anisoara Calinescu. 2023. "JAX-LOB: A GPU-Accelerated limit order book simulator to unlock large scale reinforcement learning for trading." *arXiv preprint arXiv:2308.13289*.

Futter, Owen, Blanka Horvath, and Magnus Wiese. 2023. "Signature Trading: A Path-Dependent Extension of the Mean-Variance Framework with Exogenous Signals." *Available at SSRN 4541830*.

Kroell, Emma, Sebastian Jaimungal, and Silvana M Pesenti. 2023. "Optimal Robust Reinsurance with Multiple Insurers." *arXiv preprint arXiv:2308.11828*.

Li, Haochen, Maria Polukarova, and Carmine Ventre. 2023. "Detecting Financial Market Manipulation with Statistical Physics Tools." *arXiv preprint arXiv:2308.08683*.

Li, Hongcheng. 2023. "Multiplayer War of Attrition with Asymmetric Private Information." *arXiv preprint arXiv:2302.09427*.

Loeper, Gregoire, Jan Obloj, and Benjamin Joseph. 2023. "Calibration of Local Volatility Models with Stochastic Interest Rates using Optimal Transport." *arXiv preprint arXiv:2305.00200*.

Morel, Rudy, Stéphane Mallat, and Jean-Philippe Bouchaud. 2023. "Path Shadowing Monte-Carlo." *arXiv preprint arXiv:2308.01486*.

Pulido, Sergio, Mathieu Rosenbaum, and Emmanouil Sfendourakis. 2023. "Understanding the least well-kept secret of high-frequency trading." *arXiv preprint arXiv:2307.15599*.

Roberts, Stephen, Xiaowen Dong, Stefan Zohren, et al. 2023. "Network Momentum across Asset Classes." *arXiv preprint arXiv:2308.11294*.

Semenova, Valentina, Dragos Gorduza, William Wildi, Xiaowen Dong, and Stefan Zohren. 2023. "Wisdom of the Crowds or Ignorance of the Masses? A data-driven guide to WSB." *arXiv preprint arXiv:2308.09485*.

Sessa, Pier Giuseppe, Pierre Laforgue, Nicolò Cesa-Bianchi, and Andreas Krause. 2023. "Multitask Learning with No Regret: from Improved Confidence Bounds to Active Learning." *arXiv preprint arXiv:2308.01744*.

Teichmann, Josef, and Hanna Wutte. 2023. "Machine Learning-powered Pricing of the Multidimensional Passport Option." *arXiv preprint arXiv:2307.14887*.

Zohren, Stefan, Stephen Roberts, Xiaowen Dong, et al. 2023. "Learning to Learn Financial Networks for Optimising Momentum Strategies." *arXiv preprint arXiv:2308.12212*.