OMI Research Newsletter – October 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

News. The Conference on Learning in Games and Algorithmic Collusion Workshop will take place on 12 and 13 October 2023 at the University of Oxford. Theorists, empiricists, and legal experts will discuss topics related to learning in games, algorithmic pricing, regulation, and detecting collusion. Click here for a detailed schedule and to register.

The Blockchain@X-OMI Workshop on Blockchain and Decentralized Finance (DeFi) was held on September 21-22 2023 at the École Polytechnique campus. The workshop brought together 22 scholars from a variety of backgrounds, such as economics, computer science, and mathematics, who presented their working papers. Visit this link for more information.

Research focus. Noteworthy contributions in this month's edition include works from the OMI, MAN AHL, McMaster University, Illinois Institute of Technology. Cartea, Cucuringu and Jin (2023) construct statistical arbitrage portfolios based on clustering algorithms. Zhang et al. (2023) develop a cluster-driven methodology for robust detection of lead-lag relationships in multi-factor models. Luk (2023) provides a comprehensive overview of Generative AI models and their impact and applications in asset management. Yang et al. (2023) explore transfer learning between datasets. Finally, Bielecki, Cialenco and Liu (2023) study risk measures generated by distortion functions in a dynamic time setup

Machine Learning & Statistics

Correlation Matrix Clustering for Statistical Arbitrage Portfolios

Á. CARTEA, M. CUCURINGU, Q. JIN
OXFORD-MAN INSTITUTE OF QUANTITATIVE FINANCE

We propose a framework to construct statistical arbitrage portfolios with graph clustering algorithms. First, we use various clustering methods to partition the correlation matrix of market residual returns of stocks into clusters. Next, we construct and evaluate the performance of mean-reverting statistical arbitrage portfolios within each cluster. We explore five clustering algorithms and demonstrate that our proposed framework generates profitable trading strategies with over 10% annualized returns and statistically significant Sharpe ratios above one. The performance of our statistical arbitrage portfolios is neutral to the market and cannot be fully explained by intra-industry mean-reversion effects.

Generative AI for End-to-End Limit Order Book Modelling: A Token-Level Autoregressive Generative Model of Message Flow Using a Deep State Space Network

P. NAGY, S. FREY, S. SAPORA, K. LI, A. CALINESCU, S. ZOHREN, J. FOERSTER
OXFORD-MAN INSTITUTE, FOERSTER LAB FOR AI RESEARCH

Developing a generative model of realistic order flow in financial markets is a challenging open problem, with numerous applications for market participants. Addressing this, we propose the first end-to-end autoregressive generative model that generates tokenized limit order book (LOB) messages. These messages are interpreted by a Jax-LOB simulator, which updates the LOB state. To

handle long sequences efficiently, the model employs simplified structured state-space layers to process sequences of order book states and tokenized messages. Using LOBSTER data of NASDAQ equity LOBs, we develop a custom tokenizer for message data, converting groups of successive digits to tokens, similar to tokenization in large language models. Out-of-sample results show promising performance in approximating the data distribution, as evidenced by low model perplexity. Furthermore, the mid-price returns calculated from the generated order flow exhibit a significant correlation with the data, indicating impressive conditional forecast performance. Due to the granularity of generated data, and the accuracy of the model, it offers new application areas for future work beyond forecasting, e.g. acting as a world model in high-frequency financial reinforcement learning applications. Overall, our results invite the use and extension of the model in the direction of autoregressive large financial models for the generation of high-frequency financial data and we commit to open-sourcing our code to facilitate future research.

New News is Bad News

P. Glasserman, H. Mamaysky, J. Qin
Columbia Business School

An increase in the novelty of news predicts negative stock market returns and negative macroeconomic outcomes over the next year. We quantify news novelty - changes in the distribution of news text - through an entropy measure, calculated using a recurrent neural network applied to a large news corpus. Entropy is a better out-of-sample predictor of market returns than a collection of standard measures. Cross-sectional entropy exposure carries a negative risk premium, suggesting that assets that positively covary with entropy hedge the aggregate risk associated with shifting news language. Entropy risk cannot be explained by existing long-short factors.

An Offline Learning Approach to Propagator Models

E. Neuman, W. Stockinger, Y. Zhang
Imperial College London

We consider an offline learning problem for an agent who first estimates an unknown price impact kernel from a static dataset, and then designs strategies to liquidate a risky asset while creating transient price impact. We propose a novel approach for a nonparametric estimation of the propagator from a dataset containing correlated price trajectories, trading signals and metaorders. We quantify the accuracy of the estimated propagator using a metric which depends explicitly on the dataset. We show that a trader who tries to minimise her execution costs by using a greedy strategy purely based on the estimated propagator will encounter suboptimality due to so-called spurious correlation between the trading strategy and the estimator and due to intrinsic uncertainty resulting from a biased cost functional. By adopting an offline reinforcement learning approach, we introduce a pessimistic loss functional taking the uncertainty of the estimated propagator into account, with an optimiser which eliminates the spurious correlation, and derive an asymptotically optimal bound on the execution costs even without precise information on the true propagator. Numerical experiments are included to demonstrate the effectiveness of the proposed propagator estimator and the pessimistic trading strategy.

Dynamic Time Warping for Lead-Lag Relationships in Lagged Multi-Factor Models

Y. Zhang, M. Cucuringu, A. Y. Shestopaloff, S. Zohren
Oxford-Man Institute, Queen Mary University of London

In multivariate time series systems, lead-lag relationships reveal dependencies between time series when they are shifted in time relative to each other. Uncovering such relationships is valuable in downstream tasks, such as control, forecasting, and clustering. By understanding the temporal dependencies between different time series, one can better comprehend the complex interactions and patterns within the system. We develop a cluster-driven methodology based on dynamic time warping for robust detection of lead-lag relationships in lagged multi-factor models. We establish connections to the multireference alignment problem for both the homogeneous and heterogeneous settings. Since multivariate time series are ubiquitous in a wide range of domains, we demonstrate that our algorithm is able to robustly detect lead-lag relationships in financial markets, which can be subsequently leveraged in trading strategies with significant economic benefits.

Applying Deep Learning to Calibrate Stochastic Volatility Models

A. Sridi, P. Bilokon
Imperial College London

Stochastic volatility models, where the volatility is a stochastic process, can capture most of the essential stylized facts of implied volatility surfaces and give more realistic dynamics of the volatility smile/skew. However, they come with the significant issue that they take too long to calibrate. Alternative calibration methods based on Deep Learning (DL) techniques have been recently used to build fast and accurate solutions to the calibration problem. Huge and Savine developed a Differential Machine Learning (DML) approach, where Machine Learning models are trained on samples of not only features and labels but also differentials of

labels to features. The present work aims to apply the DML technique to price vanilla European options (i.e. the calibration instruments), more specifically, puts when the underlying asset follows a Heston model and then calibrate the model on the trained network. DML allows for fast training and accurate pricing. The trained neural network dramatically reduces Heston calibration's computation time. In this work, we also introduce different regularisation techniques, and we apply them notably in the case of the DML. We compare their performance in reducing overfitting and improving the generalisation error. The DML performance is also compared to the classical DL (without differentiation) one in the case of Feed-Forward Neural Networks. We show that the DML outperforms the DL. The complete code for our experiments is provided in the GitHub repository: this https URL.

Generative AI: Overview, Economic Impact, and Applications in Asset Management

M. Luk
Man AHL

This paper provides a comprehensive overview of the evolution and latest advancements in Generative AI models, alongside their economic impact and applications in asset management. The paper is divided into two sections. The first section outlines the key innovations and methodologies that underpin large language models like ChatGPT, while also covering image-based, multimodal, and tool-using Generative AI models. Transitioning from the technical to the practical, the second section reviews the impact of Generative AI on jobs, productivity, and various industries, ending with a focus on use-cases within investment management. This section also addresses the dangers and risks associated with the use of Generative AI, including the issue of hallucinations. Understanding how this technology works and how it can be effectively applied is crucial for assessing its future prospects.

Regret-Optimal Federated Transfer Learning for Kernel Regression with Applications in American Option Pricing

X. Yang, A. Kratsios, F. Krach, M. Grasselli, A. Lucchi
McMaster University, University of Basel, ETH Zurich

We propose an optimal iterative scheme for federated transfer learning, where a central planner has access to datasets $\mathcal{D}_1, \dots, \mathcal{D}_N$ for the same learning model f_θ. Our objective is to minimize the cumulative deviation of the generated parameters $\{\theta_i(t)\}_{t=0}^T$ across all T iterations from the specialized parameters $\theta_1^*, \dots, \theta_N^*$ obtained for each dataset, while respecting the loss function for the model $f_\theta(T)$ produced by the algorithm upon halting. We only allow for continual communication between each of the specialized models (nodes/agents) and the central planner (server), at each iteration (round). For the case where the model f_θ is a finite-rank kernel regression, we derive explicit updates for the regret-optimal algorithm. By leveraging symmetries within the regret-optimal algorithm, we further develop a nearly regret-optimal heuristic that runs with $\mathcal{O}(N\,p^2)$ fewer elementary operations, where p is the dimension of the parameter space. Additionally, we investigate the adversarial robustness of the regret-optimal algorithm showing that an adversary which perturbs q training pairs by at-most $\epsilon > 0$, across all training sets, cannot reduce the regret-optimal algorithm's regret by more than $\mathcal{O}(\epsilon\, q\, \overline{N}^{1/2})$, where $\overline{N}$ is the aggregate number of training pairs. To validate our theoretical findings, we conduct numerical experiments in the context of American option pricing, utilizing a randomly generated finite-rank kernel.

Mathematical Finance

Modeling liquidity in corporate bond markets: applications to price adjustments

P. Bergault, O. Guéant
Université Paris 1 Panthéon-Sorbonne, Université Paris Dauphine-PSL

To assign a value to a portfolio, it is common to use Mark-to-Market prices. But how to proceed when the securities are illiquid? When transaction prices are scarce, how to use other available real-time information? In this article dedicated to corporate bonds, we address these questions using an extension of the concept of micro-price recently introduced for assets exchanged on limit order books in the market microstructure literature and ideas coming from the recent literature on OTC market making. To account for liquidity imbalances in OTC markets, we use a novel approach based on Markov-modulated Poisson processes. Beyond an extension to corporate bonds of the concept of micro-price, we coin the new concept of Fair Transfer Price that can be used to value or transfer securities in a fair manner even when the market is illiquid and/or tends to be one-sided.

Thorough mathematical modelling and analysis of Uniswap v3

M. Echenim, E. Gobet, A-C. Maurice
Laboratoire d'Informatique de Grenoble, Ecole Polytechnique

Automated Market Markers emerged quite recently, and Uniswap is one of the most widely used platforms (it covers 96% of the available pools as of today). This protocol is challenging from a quantitative point of view,

as it allows participants to choose where they wish to to concentrate liquidity. There has been an increasing number of research papers on Uniswap v3 but often, these articles use heuristics or approximations that can be far from reality: for instance, the liquidity in the pool is assumed to be constant over time, which strongly contradicts the mechanism of the protocol. The objectives of this work are fourfold: first, to revisit Uniswap v3's mechanisms in detail (starting from the open source code) to build an unambiguous knowledge base. Second, to analyze the Impermanent Loss of a liquidity provider by detailing its evolution, in full generality on the swap trades and the liquidity events than can occur. Third, we introduce the notion of a liquidity curve. For each curve, we can deduce a payoff at a given maturity, net of fees. Conversely, we show how any concave payoff can be synthetized by an initial liquidity curve and some tokens outside the pool; this paves the way for using Uniswap v3 to create options. Fourth, we analyze the behavior of collected fees without any simplifying hypothesis (like a constant liquidity or zero Spot-Pool spread) under the mild assumption that the pool price follows a general Ito price dynamic. The value of the collected fees then coincides with an integral of call and put prices. Our derivations are supported by graphical illustrations and experiments.

Time consistency of dynamic risk measures and dynamic performance measures generated by distortion functions

T. R. Bielecki, I. Cialenco, H. Liu

Illinois Institute of Technology

The aim of this work is to study risk measures generated by distortion functions in a dynamic discrete time setup, and to investigate the corresponding dynamic coherent acceptability indices (DCAIs) generated by families of such risk measures. First we show that conditional version of Choquet integrals indeed are dynamic coherent risk measures (DCRMs), and also introduce the class of dynamic weighted value at risk measures. We prove that these two classes of risk measures coincides. In the spirit of robust representations theorem for DCAIs, we establish some relevant properties of families of DCRMs generated by distortion functions, and then define and study the corresponding DCAIs. Second, we study the time consistency of DCRMs and DCAIs generated by distortion functions. In particular, we prove that such DCRMs are sub-martingale time consistent, but they are not super-martingale time consistent. We also show that DCRMs generated by distortion functions are not weakly acceptance time consistent. We also present several widely used classes of distortion functions and derive some new representations of these distortions.

Optimal adaptive control with separable drift uncertainty

S. N. Cohen, C. Knochenhauer, A. Merkel

University of Oxford

We consider a problem of stochastic optimal control with separable drift uncertainty in strong formulation on a finite horizon. The drift coefficient of the state Y^u is multiplicatively influenced by an unknown random variable λ, while admissible controls u are required to be adapted to the observation filtration. Choosing a control actively influences the state and information acquisition simultaneously and comes with a learning effect. The problem, initially non-Markovian, is embedded into a higher-dimensional Markovian, full information control problem with control-dependent filtration and noise. To that problem, we apply the stochastic Perron method to characterize the value function as the unique viscosity solution to the HJB equation, explicitly construct ϵ-optimal controls and show that the values of strong and weak formulations agree. Numerical illustrations show a significant difference between the adaptive control and the certainty equivalence control.

Mean-field games of speedy information access with observation costs

D. Becherer, C. Reisinger, J. Tam

University of Oxford, Humboldt University of Berlin

We investigate a mean-field game (MFG) in which agents can exercise control actions that affect their speed of access to information. The agents can dynamically decide to receive observations with less delay by paying higher observation costs. Agents seek to exploit their active information gathering by making further decisions to influence their state dynamics to maximize rewards. In the mean field equilibrium, each generic agent solves individually a partially observed Markov decision problem in which the way partial observations are obtained is itself also subject of dynamic control actions by the agent. Based on a finite characterisation of the agents' belief states, we show how the mean field game with controlled costly information access can be formulated as an equivalent standard mean field game on a suitably augmented but finite state space.We prove that with sufficient entropy regularisation, a fixed point iteration converges to the unique MFG equilibrium and yields an approximate -Nash equilibrium for a large but finite population size. We illustrate our MFG by an example from epidemiology, where medical testing results at different speeds and costs can be chosen by the agents.

Tackling Nonlinear Price Impact with Linear Strategies

X. Brokmann, D. Itkin, J. Muhle-Karbe, P. Schmidt

Imperial College London, Qube Research

Empirical studies consistently find that the price impact of large trades approximately follows a nonlinear power law. Yet, tractable formulas for the portfolios that trade off predictive trading signals, risk, and trading costs in an optimal manner are only available for quadratic costs corresponding to linear price impact. In this paper, we show that the resulting linear strategies allow to achieve virtually optimal performance also for realistic nonlinear price impact, if the "effective" quadratic cost parameter is chosen appropriately. To wit, for a wide range of risk levels, this leads to performance losses below 2% compared to the numerical Viterbi algorithm of Kolm and Ritter (2014) run at very high accuracy. The effective quadratic cost depends on the portfolio risk, but can be computed without any sophisticated numerics by simply maximizing an explicit scalar function.

Economics

Game Connectivity and Adaptive Dynamics

T. Johnston, M. Savery, A. Scott, B. Tarbush

University of Oxford, University of Bristol

We analyse the typical structure of games in terms of the connectivity properties of their best-response graphs. In particular, we show that almost every 'large' generic game that has a pure Nash equilibrium is connected, meaning that every non-equilibrium action profile can reach every pure Nash equilibrium via best-response paths. This has implications for dynamics in games: many adaptive dynamics, such as the best-response dynamic with inertia, lead to equilibrium in connected games. It follows that there are simple, uncoupled, adaptive dynamics for which period-by-period play converges almost surely to a pure Nash equilibrium in almost every 'large' generic game that has one. We build on recent results in probabilistic combinatorics for our characterisation of game connectivity.

Measuring Higher-Order Rationality with Belief Control

W. J. Chen, M-J. Fong, P-H. Lin

California Institute of Technology, National Taiwan University

Determining an individual's strategic reasoning capability based solely on choice data is a complex task. This complexity arises because sophisticated players might have non-equilibrium beliefs about others, leading to non-equilibrium actions. In our study, we pair human participants with computer players known to be fully rational. This use of robot players allows us to disentangle limited reasoning capacity from belief formation and social biases. Our results show that, when paired with robots, subjects consistently demonstrate higher levels of rationality and maintain stable rationality levels across different games compared to when paired with humans. This suggests that strategic reasoning might indeed be a consistent trait in individuals. Furthermore, the identified rationality limits could serve as a measure for evaluating an individual's strategic capacity when their beliefs about others are adequately controlled.

Learning News Bias: Misspecifications and Consequences

L. Hu, M. Kovach, A. Li

Australian National University, Purdue University, University of Waterloo

We study how a decision maker (DM) learns about the bias of unfamiliar news sources. Absent any frictions, a rational DM uses known sources as a yardstick to discern the true bias of a source. If a DM has misspecified beliefs, this process fails. We derive long-run beliefs, behavior, welfare, and corresponding comparative statics, when the DM has dogmatic, incorrect beliefs about the bias of known sources. The distortion due to misspecified learning is succinctly captured by a single-dimensional metric we introduce. Our model generates the hostile media effect and false polarization, and has implications for fact-checking and misperception recalibration.

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

Becherer, Dirk, Christoph Reisinger, and Jonathan Tam. 2023. "Mean-field games of speedy information access with observation costs." *arXiv preprint arXiv:2309.07877*.

Bergault, Philippe, and Olivier Guéant. 2023. "Modeling liquidity in corporate bond markets: applications to price adjustments." *arXiv preprint arXiv:2309.04216*.

Bielecki, Tomasz R, Igor Cialenco, and Hao Liu. 2023. "Time consistency of dynamic risk measures

and dynamic performance measures generated by distortion functions." *arXiv preprint arXiv:2309.02570*.

Brokmann, Xavier, David Itkin, Johannes Muhle-Karbe, and Peter Schmidt. 2023. "Tackling Nonlinear Price Impact with Linear Strategies." *Available at SSRN*.

Cartea, Álvaro, Mihai Cucuringu, and Qi Jin. 2023. "Correlation Matrix Clustering for Statistical Arbitrage Portfolios." *Available at SSRN*.

Chen, Wei James, Meng-Jhang Fong, and Po-Hsuan Lin. 2023. "Measuring Higher-Order Rationality with Belief Control." *arXiv preprint arXiv:2309.07427*.

Cohen, Samuel N, Christoph Knochenhauer, and Alexander Merkel. 2023. "Optimal adaptive control with separable drift uncertainty." *arXiv preprint arXiv:2309.07091*.

Echenim, Mnacho, Emmanuel Gobet, and Anne-Claire Maurice. 2023. "Thorough mathematical modelling and analysis of Uniswap v3."

Glasserman, Paul, Harry Mamaysky, and Jimmy Qin. 2023. "New News is Bad News." *arXiv preprint arXiv:2309.05560*.

Hu, Lin, Matthew Kovach, and Anqi Li. 2023. "Learning News Bias: Misspecifications and Consequences." *arXiv preprint arXiv:2309.08740*.

Johnston, Tom, Michael Savery, Alex Scott, and Bassel Tarbush. 2023. "Game Connectivity and Adaptive Dynamics." *arXiv preprint arXiv:2309.10609*.

Luk, Martin. 2023. "Generative AI: Overview, Economic Impact, and Applications in Asset Management." *Economic Impact, and Applications in Asset Management (September 18, 2023)*.

Nagy, Peer, Sascha Frey, Silvia Sapora, Kang Li, Anisoara Calinescu, Stefan Zohren, and Jakob Foerster. 2023. "Generative AI for End-to-End Limit Order Book Modelling: A Token-Level Autoregressive Generative Model of Message Flow Using a Deep State Space Network." *arXiv preprint arXiv:2309.00638*.

Neuman, Eyal, Wolfgang Stockinger, and Yufei Zhang. 2023. "An Offline Learning Approach to Propagator Models." *arXiv preprint arXiv:2309.02994*.

Wei, Song, Andrea Coletta, Svitlana Vyetrenko, and Tucker Balch. 2023. "ATMS: Algorithmic Trading-Guided Market Simulation." *arXiv preprint arXiv:2309.01784*.

Yang, Xuwei, Anastasis Kratsios, Florian Krach, Matheus Grasselli, and Aurelien Lucchi. 2023. "Regret-Optimal Federated Transfer Learning for Kernel Regression with Applications in American Option Pricing." *arXiv preprint arXiv:2309.04557*.

Zhang, Yichi, Mihai Cucuringu, Alexander Y Shestopaloff, and Stefan Zohren. 2023. "Dynamic Time Warping for Lead-Lag Relationships in Lagged Multi-Factor Models." *arXiv preprint arXiv:2309.08800*.