October 2022, Vol. 2, No. 1

OMI Research Newsletter



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

Noteworthy contributions in this month's edition include two papers from the OMI. The first classifies trade flow based on a measure of order book imbalance, and investigates price impact of the decomposed trade flow to design profitable trading strategies. The second introduces a transfer ranking model applied to cryptocurrencies using a momentum strategy.

Another notable contribution from KCL, Liverpool, and Warwick provides a new open source repository for unified gym environments that allow one to benchmark the performance of RL algorithms for a range of model-based LOB trading problems.

This month's featured work is from MIT and Yale University. The authors extend learning with misspecified models to incorporate learning with memory biases. The work provides the theoretical foundations to understand the long-run implications of learning with selective memory where the data generating process is a consequence of the agent's own actions.

Machine Learning for Finance

Trade Co-occurrence, Trade Flow Decomposition, and Conditional Order Imbalance in Equity Markets

Y. Lu, G. Reinert, and M. Cucuringu
University of Oxford

The time proximity of high-frequency trades can contain a salient signal. In this paper, we propose a method to classify every trade, based on its proximity with other trades in the market within a short period of time, into five types. By means of a suitably defined normalized order imbalance associated to each type of trade, which we denote as conditional order imbalance (COI), we investigate the price impact of the decomposed trade flows. Our empirical findings indicate strong positive correlations between contemporaneous returns and COIs. In terms of predictability, we document that associations with future returns are positive for COIs of trades which are isolated from trades of stocks other than themselves, and negative otherwise. Furthermore, trading strategies which we develop using COIs achieve conspicuous returns and Sharpe ratios, in an extensive experimental setup on a universe of 457 stocks using daily data for a period of three years.

SEC Form 13F-HR: Statistical investigation of trading imbalances and profitability analysis

D. Miori and M. Cucuringu
University of Oxford

US Institutions with more than $100 million assets under management must disclose part of their long positions into the SEC Form 13F-HR on a quarterly basis. We consider the number of variations in holdings between consecutive reporting periods, and compute imbalances in buying versus selling behaviour for the assets under consideration. A significant opportunity for profit arises if an external investor is willing to trade contrarian to the 13F filings imbalances. Indeed, imbalances capture the amount of information already consumed in the market and the related trades tend to be inflated by crowding and herding. Betting on a relatively short-term movement of prices against the sign of imbalances results in a profitable strategy especially when using a time horizon between 21 and 42 trading days (corresponding to 1-2 calendar months) after each financial quarter ends.

Neural variance reduction for stochastic differential equations

P.D. Hinds and M.V. Tretyakov

University of Nottingham

Variance reduction techniques are of crucial importance for the efficiency of Monte Carlo simulations in finance applications. We propose the use of neural SDEs, with control variates parameterized by neural networks, in order to learn approximately optimal control variates and hence reduce variance as trajectories of the SDEs are being simulated. We consider SDEs driven by Brownian motion and, more generally, by Lévy processes including those with infinite activity. For the latter case, we prove optimality conditions for the variance reduction. Several numerical examples from option pricing are presented.

Transfer Ranking in Finance - Applications to Cross-Sectional Momentum with Data Scarcity

D. Poh, S. Roberts, and S. Zohren

University of Oxford – Oxford-Man Institute

Cross-sectional strategies are a classical and popular trading style, with recent high performing variants incorporating sophisticated neural architectures. While these strategies have been applied successfully to data-rich settings involving mature assets with long histories, deploying them on instruments with limited samples generally produce over-fitted models with degraded performance. In this paper, we introduce Fused Encoder Networks – a novel and hybrid parameter-sharing transfer ranking model. The model fuses information extracted using an encoder-attention module operated on a source dataset with a similar but separate module focused on a smaller target dataset of interest. This mitigates the issue of models with poor generalisability that are a consequence of training on scarce target data. Additionally, the self-attention mechanism enables interactions among instruments to be accounted for, not just at the loss level during model training, but also at inference time. Focusing on momentum applied to the top ten cryptocurrencies by market capitalisation as a demonstrative use-case, the Fused Encoder Networks outperforms the reference benchmarks on most performance measures, delivering a three-fold boost in the Sharpe ratio over classical momentum as well as an improvement of approximately 50% against the best benchmark model without transaction costs. It continues outperforming baselines even after accounting for the high transaction costs associated with trading cryptocurrencies.

Learning Value-at-Risk and Expected Shortfall

D. Barrera, S. Crepey, E. Gobet, Hoang-Dung Nguyen, B. Saadeddine – Universidad de los Andes, Université Paris Cité, Ecole Polytechnique, and Paris-Saclay

We propose a non-asymptotic convergence analysis of a two-step approach to learn a conditional value-at-risk (VaR) and expected shortfall (ES) in a nonparametric setting using Rademacher and Vapnik-Chervonenkis bounds. Our approach for the VaR is extended to the problem of learning at once multiple VaRs corresponding to different quantile levels. This results in efficient learning schemes based on neural network quantile and least-squares regressions. An a posteriori Monte Carlo (non-nested) procedure is introduced to estimate distances to the ground-truth VaR and ES without access to the latter. This is illustrated using numerical experiments in a Gaussian toy-model and a financial case-study where the objective is to learn a dynamic initial margin.

Time-Series Forecasting using Dynamic Graphs: Case Studies with Dyn-STGCN and Dyn-GWN on Finance and Traffic Datasets

S. Ibrahim, M. R. Tell, and R. Mazumder

Massachusetts Institute of Technology

Spatio-temporal modeling is an essential lens to understand many real-world phenomena from traffic to epidemiology. Although forecasting time-series is an exceptionally well-studied problem, recent years have seen impressive gains in the performance of graph learning as a paradigm for spatial learning problems. Some recent work has explored the intersection of these two fields, but often assumes that the spatial structure is static. We propose a new framework for spatio-temporal learning from dynamic graphs. The two main components of our models are: (i) Temporal convolutions on the time-varying adjacency space, (ii) Tensor Graph Convolutional Layer (TGCL) which aggregates latent temporal representations of time-varying node features and time-varying graphs. We generalize previous models to leverage both dynamic and static graphs while being attractive in terms of computational efficiency. We demonstrate our proposals with two new time-varying graph-based methods Dyn-STGCN and Dyn-GWN for time-series forecasting. Experiments demonstrate the efficacy of these model across datasets from different domains. Interestingly, our Dyn-STGCN and Dyn-GWN models are superior at handling dynamic graphs than existing state-of-the-art time-varying graph-based methods e.g., EvolveGCN and TM-GCN in terms of strong generalization while providing both efficient training and inference.

CT-GCN: a Phishing Identification Model for Blockchain Cryptocurrency Transactions

B. Fu, X. Yu, T. Feng
Yunnan University of Finance and Economics

With the widespread application of blockchain technology, the cyberspace security issue of phishing has also appeared in the emerging blockchain cryptocurrency ecosystem. Because phishing fraud in cryptocurrency transactions has its own unique characteristics compared to traditional phishing, many existing phishing detection algorithms are not usable. Therefore, based on graph convolutional networks, we have researched and built a high-performance model for detecting blockchain cryptocurrency phishing fraud. Our model divides the constructed blockchain cryptocurrency transaction graph into “Sender” and “Receiver” graphs, according to the sending and receiving directions. Then, the edge features in the graph are transferred. Finally, we use a double-layer graph convolution network for feature learning and send it to the classifier for fraud detection. After completing training on the actual dataset collected from Ethereum, the model achieved an accuracy of 88.02% and an F1 score of 88.14% on the test data, which had a better performance than that of the other models. Our model provides a new concept for the detection of phishing scams in blockchain cryptocurrency networks.

Mathematical Finance

Multivariate Hawkes-based Models in LOB: European, Spread and Basket Option Pricing

Q. Guo, A. Swishchuk, and B. Rémillard
University of Calgary and HEC Montréal

In this paper, we consider pricing of European options and spread options for Hawkes-based model for the limit order book. We introduce multivariate Hawkes process and the multivariable general compound Hawkes process. Exponential multivariate general compound Hawkes processes and limit theorems for them, namely, LLN and FCLT, are considered then. We also consider a special case of one-dimensional EMGCHP and its limit theorems. Option pricing with 1D EGCHP in LOB, hedging strategies, and numerical example are presented. We also introduce greeks calculations for those models. Margrabe’s spread options valuations with Hawkes-based models for two assets and numerical example are presented. Also, Margrabe’s spread option pricing with two 2D EMGCHP and numerical example are included. Basket options valuations with numerical example are included. We finally discuss the implied volatility and implied order flow. It reveals the relationship between stock volatility and the order flow in the limit order book system. In this way, the Hawkes-based model can provide more market forecast information than the classical Black-Scholes model.

Data-Driven Nonparametric Robust Control under Dependence Uncertainty

E. Bayraktar and T. Chen
University of Michigan

We consider a multi-period stochastic control problem where the multivariate driving stochastic factor of the system has known marginal distributions but uncertain dependence structure. To solve the problem, we propose to implement the nonparametric adaptive robust control framework. We aim to find the optimal control against the worst-case copulae in a sequence of shrinking uncertainty sets which are generated from continuously observing the data. Then, we use a stochastic gradient descent ascent algorithm to numerically handle the corresponding high dimensional dynamic inf-sup optimization problem. We present the numerical results in the context of utility maximization and show that the controller benefits from knowing more information about the uncertain model.

Optimal Bubble Riding: A Mean Field Game with Varying Entry Times

L. Tangpi and S. Wang
Princeton University

Recent financial bubbles such as the emergence of cryptocurrencies and "meme stocks" have gained increasing attention from both retail and institutional investors. In this paper, we propose a game-theoretic model on optimal liquidation in the presence of an asset bubble. Our setup allows the influx of players to fuel the price of the asset. Moreover, traders will enter the market at possibly different times and take advantage of the uptrend at the risk of an inevitable crash. In particular, we consider two types of crashes: an endogenous burst which results from excessive selling, and an exogenous burst which cannot be anticipated and is independent from the actions of the traders.

The popularity of asset bubbles suggests a large-population setting, which naturally leads to a mean field game (MFG) formulation. We introduce a class of MFGs with varying entry times. In particular, an equilibrium will depend on the entry-weighted average of conditional optimal strategies. To incorporate the exogenous burst time, we adopt the method of progressive enlargement of filtrations. We prove existence of MFG equilibria using the weak formulation in a generalized setup, and we show that the equilibrium strategy can be decomposed into before-and-after-burst segments, each part containing only the market information. We also perform numerical simulations of the solution, which allow us to provide

some intriguing results on the relationship between the bubble burst and equilibrium strategies.

Path Integrals for Fractional Brownian Motion and Fractional Gaussian Noise

B. Meerson, O. Bénichou, and G. Oshanin
Hebrew University of Jerusalem, Sorbonne Université, and Politecnico di Torino

The Wiener's path integral played a major role in the studies of Brownian motion (Bm). Here we derive exact path-integral representations for the more general *fractional* Brownian motion (fBm) and for its time derivative process – the fractional Gaussian noise (fGn). These paradigmatic non-Markovian stochastic processes found numerous applications across the disciplines, ranging from anomalous diffusion in cellular environments to mathematical finance. Still, exact path-integral representations for the fBm and fGn were previously unknown. Our formalism exploits the Gaussianity of the fBm and fGn, relies on theory of singular integral equations and overcomes some technical difficulties by representing the action functional for the fBm in terms of the fGn for the sub-diffusive fBm, and in terms of the first derivative of the fGn for the super-diffusive fBm. We also extend the formalism to include external forcing. The path-integral representations open a new avenue in the studies of the fBm and fGn,just as the Wiener's path integral did for the standard Bm.

Economics

Smooth Fictitious Play in Stochastic Games with Perturbed Payoffs and Unknown Transitions

L. Baudin and R. Laraki
Université Paris-Dauphine - PSL

Recent extensions to dynamic games of the well-known fictitious play learning procedure in static games were proved to globally converge to stationary Nash equilibria in two important classes of dynamic games (zero-sum and identical-interest discounted stochastic games). However, those decentralized algorithms need the players to know exactly the model (the transition probabilities and their payoffs at every stage). To overcome these strong assumptions, our paper introduces regularizations of the systems in (Leslie 2020; Baudin 2022) to construct a family of new decentralized learning algorithms which are model-free (players don't know the transitions and their payoffs are perturbed at every stage). Our procedures can be seen as extensions to stochastic games of the classical smooth fictitious play learning procedures in static games (where the players best responses are regularized, thanks to a smooth strictly concave perturbation of their payoff functions). We prove the convergence of our family of procedures to stationary regularized Nash equilibria in zero-sum and identical-interest discounted stochastic games. The proof uses the continuous smooth best-response dynamics counterparts, and stochastic approximation methods. When there is only one player, our problem is an instance of Reinforcement Learning and our procedures are proved to globally converge to the optimal stationary policy of the regularized MDP. In that sense, they can be seen as an alternative to the well known Q-learning procedure.

Selective Memory Equilibrium

D. Fudenberg, G. Lanzani, and Philipp Strack
Massachusetts Institute of Technology and Yale University

We study agents who are more likely to remember some xperiences than others, but update beliefs as if the experiences they remember are the only ones that occurred. If the agent's behavior converges, their limit strategy is a selective memory equilibrium. We illustrate how selective memory equilibrium can be used to understand the long-run effects of several well-documented memory biases, such as ego-boosting bias, associativeness, interference, and extreme experience bias. We then extend our analysis to cases where the expected number of recalled experiences is bounded and experiences that are recalled once are more likely to be recalled again. Here the frequency of recalled experiences does not converge, but we characterize the long-run action frequencies that can arise.

Others

Model-based gym environments for limit order book trading

J. Jerome, L. Sánchez-Betancourt, R. Savani, and M. Herdegen – University of Liverpool, King's College London, and University of Warwick

Within the mathematical finance literature there is a rich catalogue of mathematical models for studying algorithmic trading problems – such as market-making and optimal execution – in limit order books. This paper introduces **mbt–gym**, a Python module that provides a suite of gym environments for training reinforcement learning (RL) agents to solve such model-based trading problems. The module is set up in an extensible way to allow the combination of different aspects of different models. It supports highly efficient implementations of vectorized environments to allow faster training of RL agents. In this paper, we motivate the challenge of using RL to solve such model-based limit order book problems in mathematical finance, we explain the design

of our gym environment, and then demonstrate its use in solving standard and non-standard problems from the literature. Finally, we lay out a roadmap for further development of our module, which we provide as an open source repository on GitHub so that it can serve as a focal point for RL research in model-based algorithmic trading.

Systematization of Knowledge: Synthetic Assets, Derivatives, and On-Chain Portfolio Management

A. Rahman, V. Shi, M. Ding, and E. Choi
University of California, Berkeley

Synthetic assets are decentralized finance (DeFi) analogues of derivatives in the traditional finance (TradFi) world - financial arrangements which derive value from and are directly pegged to fluctuations in the value of an underlying asset (ex: futures and options). Synthetic assets occupy a unique niche, serving to facilitate currency exchange, giving traders a means to speculate on the value of crypto assets without directly holding them, and powering more complex financial tools such as yield optimizers and portfolio management suites. Unfortunately, the academic literature on this topic is highly disparate and struggles to keep up with rapid changes in the space. We present the first Systematization of Knowledge (SoK) in this area, focusing on presenting the key mechanisms, protocols, and issues in an accessible fashion to highlight risks for participants as well as areas of research interest. This paper takes a broad perspective in establishing a general framework for synthetic assets, from the ideological origins of crypto to legal barriers for firms in this space, encapsulating the basic mechanisms underpinning derivatives markets as well as presenting data-driven analyses of major protocols.

Adaptive Gaussian Process Change Point Detection

E. Caldarelli, P. Wenk, S. Bauer, and A. Krause
– ETH Zurich, KTH, and CSIC-UPC

Detecting change points in time series, i.e., points in time at which some observed process suddenly changes, is a fundamental task that arises in many real-world applications, with consequences for safety and reliability. In this work, we propose ADAGA, a novel Gaussian process-based solution to this problem, that leverages a powerful heuristics we developed based on statistical hypothesis testing. In contrast to prior approaches, ADAGA adapts to changes both in mean and covariance structure of the temporal process. In extensive experiments, we show its versatility and applicability to different classes of change points, demonstrating that it is significantly more accurate than current state-of-the-art alternatives.

Acknowledgement

The OMI Research Newsletter was curated by Prof. Álvaro Cartea, Director of the Oxford-Man Institute, and Fayçal Drissi, D.Phil. Candidate. Paper suggestions are from fellow students at the institute: Álvaro Arroyo, Patrick Chang, Dragos Gorduza, and Marcello Monga.