VOL. 2, NO. 5

OMI Research Newsletter – Feb 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

Noteworthy contributions in this month's edition include works from the OMI, University of Cambridge, Université Paris 1 Panthéon-Sorbonne, and Harvard University. In this work, the authors make a considerable contribution to game theory and prove that algorithms can learn to collude. Here, the authors propose a probabilistic scheme that uses deep learning techniques to solve high-dimensional semi-linear PDEs. Finally, in this work, the authors leverage Riemannian geometry and in particular product manifolds to improve the performance of graph neural networks.

Machine Learning

Fin-GAN: Forecasting and Classifying Financial Time Series via Generative Adversarial Networks

M. VULETIĆ, F. PRENZEL, AND M. CUCURINGU
OXFORD-MAN INSTITUTE

We investigate the use of Generative Adversarial Networks (GANs) for probabilistic forecasting of financial time series. To this end, we introduce a novel economics-driven loss function for the generator. This newly designed loss function renders GANs more suitable for a classification task, and places them into a supervised learning setting, whilst producing full conditional probability distributions of price returns given previous historical values. Our approach moves beyond the point estimates traditionally employed in the forecasting literature, and allows for uncertainty estimates. Numerical experiments on equity data showcase the effectiveness of our proposed methodology, which achieves higher Sharpe Ratios compared to classical supervised learning models, such as LSTMs and ARIMA.

Latent Graph Inference using Product Manifolds

H. S. DE OCÁRIZ BORDE, A. KAZI, F. BARBERO, AND P. LIÒ
OXFORD-MAN INSTITUTE, UNIVERSITY OF OXFORD, HARVARD UNIVERSITY, UNIVERSITY OF CAMBRIDGE

Graph Neural Networks usually rely on the assumption that the graph topology is available to the network as well as optimal for the downstream task. Latent graph inference allows models to dynamically learn the intrinsic graph structure of problems where the connectivity patterns of data may not be directly accessible. In this work, we generalize the discrete Differentiable Graph Module (dDGM) for latent graph learning. The original dDGM architecture used the Euclidean plane to encode latent features based on which the latent graphs were generated. By incorporating Riemannian geometry into the model and generating more complex embedding spaces, we can improve the performance of the latent graph inference system. In particular, we propose a computationally tractable approach to produce product manifolds of constant curvature model spaces that can encode latent features of varying structure. The latent representations mapped onto the inferred product manifold are used to compute richer similarity measures that are leveraged by the latent graph learning model to obtain optimized latent graphs. Moreover, the curvature of the product manifold is learned during training alongside the rest of the network parameters and based on the downstream task, rather than it being a static embedding space. Our novel approach is tested on a wide range of datasets, and outperforms the original dDGM model.

Sheaf Attention Networks

F. Barbero, C. Bodnar, H. S. de Ocariz Borde, and P. Liò

Oxford-Man Institute, University of Oxford, University of Cambridge

Attention has become a central inductive bias for deep learning models irrespective of domain. However, increasing theoretical and empirical evidence suggests that Graph Attention Networks (GATs) suffer from the same pathological issues affecting many other Graph Neural Networks (GNNs). First, GAT's features tend to become progressively smoother as more layers are stacked, and second, the model performs poorly in heterophilic graphs. Sheaf Neural Networks (SNNs), a new class of models inspired by algebraic topology and geometry, have shown much promise in tackling these two issues. Building upon the recent success of SNNs and the wide adoption of attention-based architectures, we propose Sheaf Attention Networks (SheafANs). By making use of a novel and more expressive attention mechanism equipped with geometric inductive biases, we show that this type of construction generalizes popular attention-based GNN models to cellular sheaves. We demonstrate that these models help tackle the oversmoothing and heterophily problems and show that, in practice, SheafANs consistently outperform GAT on synthetic and real-world benchmarks.

Scaleformer: Iterative Multi-scale Refining Transformers for Time Series Forecasting

M. A. Shabani, A. Abdi, L. Meng, and T. Sylvain

Simon Fraser University, Borealis AI

The performance of time series forecasting has recently been greatly improved by the introduction of transformers. In this paper, we propose a general multi-scale framework that can be applied to the state-of-the-art transformer-based time series forecasting models (FEDformer, Autoformer, etc.). By iteratively refining a forecasted time series at multiple scales with shared weights, introducing architecture adaptations, and a specially-designed normalization scheme, we are able to achieve significant performance improvements, from 5.5% to 38.5% across datasets and transformer architectures, with minimal additional computational overhead. Via detailed ablation studies, we demonstrate the effectiveness of each of our contributions across the architecture and methodology. Furthermore, our experiments on various public datasets demonstrate that the proposed improvements outperform their corresponding baseline counterparts. Our code is publicly available in this https URL.

Statistical Learning with Sublinear Regret of Propagator Models

E. Neuman and Y. Zhang

Imperial College, London School of Economics

We consider a class of learning problems in which an agent liquidates a risky asset while creating both transient price impact driven by an unknown convolution propagator and linear temporary price impact with an unknown parameter. We characterize the trader's performance as maximization of a revenue-risk functional, where the trader also exploits available information on a price predicting signal. We present a trading algorithm that alternates between exploration and exploitation phases and achieves sublinear regrets with high probability. For the exploration phase we propose a novel approach for non-parametric estimation of the price impact kernel by observing only the visible price process and derive sharp bounds on the convergence rate, which are characterised by the singularity of the propagator. These kernel estimation methods extend existing methods from the area of Tikhonov regularisation for inverse problems and are of independent interest. The bound on the regret in the exploitation phase is obtained by deriving stability results for the optimizer and value function of the associated class of infinite-dimensional stochastic control problems. As a complementary result we propose a regression-based algorithm to estimate the conditional expectation of non-Markovian signals and derive its convergence rate.

Robust Machine Learning Pipelines For Trading Market-Neutral Stock Portfolios

T. Wong and M. Barahona

Imperial College

The application of deep learning algorithms to financial data is difficult due to heavy non-stationarities which can lead to over-fitted models that underperform under regime changes. Using the Numerai tournament data set as a motivating example, we propose a machine learning pipeline for trading market-neutral stock portfolios based on tabular data which is robust under changes in market conditions. We evaluate various machine-learning models, including Gradient Boosting Decision Trees (GBDTs) and Neural Networks with and without simple feature engineering, as the building blocks for the pipeline. We find that GBDT models with dropout display high performance, robustness and generalisability with relatively low complexity and reduced computational cost. We then show that online learning techniques can be used in post-prediction processing to enhance the results. In particular, dynamic feature neutralisation, an efficient procedure that requires no retraining of models and can be applied post-prediction to any machine learning model, improves robustness by reducing drawdown in

volatile market conditions. Furthermore, we demonstrate that the creation of model ensembles through dynamic model selection based on recent model performance leads to improved performance over baseline by improving the Sharpe and Calmar ratios. We also evaluate the robustness of our pipeline across different data splits and random seeds with good reproducibility of results.

Asynchronous Deep Double Duelling Q-Learning for Trading-Signal Execution in Limit Order Book Markets

P. Nagy, J-P. Calliess, and S. Zohren
Oxford-Man Institute

We employ deep reinforcement learning (RL) to train an agent to successfully translate a high-frequency trading signal into a trading strategy that places individual limit orders. Based on the ABIDES limit order book simulator, we build a reinforcement learning OpenAI gym environment and utilise it to simulate a realistic trading environment for NASDAQ equities based on historic order book messages. To train a trading agent that learns to maximise its trading return in this environment, we use Deep Duelling Double Q-learning with the APEX (asynchronous prioritised experience replay) architecture. The agent observes the current limit order book state, its recent history, and a short-term directional forecast. To investigate the performance of RL for adaptive trading independently from a concrete forecasting algorithm, we study the performance of our approach utilising synthetic alpha signals obtained by perturbing forward-looking returns with varying levels of noise. Here, we find that the RL agent learns an effective trading strategy for inventory management and order placing that outperforms a heuristic benchmark trading strategy having access to the same signal.

Mathematical Finance

Efficient Risk Estimation for the Credit Valuation Adjustment

M. B. Giles, A-L. Haji-Ali, and J. Spence
University of Oxford, Heriot-Watt University

The valuation of over-the-counter derivatives is subject to a series of valuation adjustments known as xVA, which pose additional risks for financial institutions. Associated risk measures, such as the value-at-risk of an underlying valuation adjustment, play an important role in managing these risks. Monte Carlo methods are often regarded as inefficient for computing such measures. As an example, we consider the value-at-risk of the Credit Valuation Adjustment (CVA-VaR), which can be expressed using a triple nested expectation. Traditional Monte Carlo methods are often inefficient at handling several nested expectations. Utilising recent developments in multilevel nested simulation for probabilities, we construct a hierarchical estimator of the CVA-VaR which reduces the computational complexity by 3 orders of magnitude compared to standard Monte Carlo.

Decarbonization of financial markets: a mean-field game approach

P. Lavigne and P. Tankov
Institut Louis Bachelier, CREST-ENSAE

We build a model of a financial market where a large number of firms determine their dynamic emission strategies under climate transition risk in the presence of both green-minded and neutral investors. The firms aim to achieve a trade-off between financial and environmental performance, while interacting through the stochastic discount factor, determined in equilibrium by the investors' allocations. We formalize the problem in the setting of mean-field games and prove the existence and uniqueness of a Nash equilibrium for firms. We then present a convergent numerical algorithm for computing this equilibrium and illustrate the impact of climate transition risk and the presence of green-minded investors on the market decarbonization dynamics and share prices. We show that uncertainty about future climate risks and policies leads to higher overall emissions and higher spreads between share prices of green and brown companies. This effect is partially reversed in the presence of environmentally concerned investors, whose impact on the cost of capital spurs companies to reduce emissions.

A pure jump model for the valuation of options on a credit index

Y. Shirai
University of Maryland

A two-dimensional pure jump process is proposed to model the evolution of the risk-free rate and default intensities for the purpose of evaluating option contracts on a credit index. Time evolution in credit markets is assumed to follow a gamma process evaluated at calendar time in order to reflect different levels of business activity in the credit and Treasury markets, which ultimately reflect differences in preferences and incentives of credit products investors, as well as the structure of the credit market itself, with those of their respective counterparts in the Treasury market. Formulas for the characteristic function, zero coupon bonds and moments of the process are derived, and its parameters calibrated to market prices of options on a credit index. Model and market implied credit spreads moments are estimated and compared.

Deep Runge-Kutta schemes for BSDEs

J-F Chassagneux, J. Chen, and N. Frikha
Université Paris Cité, Université Paris 1 Panthéon-Sorbonne

We propose a new probabilistic scheme which combines deep learning techniques with high order schemes for backward stochastic differential equations belonging to the class of Runge-Kutta methods to solve high-dimensional semi-linear parabolic partial differential equations. Our approach notably extends the one introduced in [*Huré, C., Pham, H., Warin, X. (2020)*] for the implicit Euler scheme to schemes which are more efficient in terms of discrete-time error. We establish some convergence results for our implemented schemes under classical regularity assumptions. We also illustrate the efficiency of our method for different schemes of order one, two and three. Our numerical results indicate that the Crank-Nicolson schemes is a good compromise in terms of precision, computational cost and numerical implementation.

Measuring tail risk at high-frequency: An L_1-regularized extreme value regression approach with unit-root predictors

J. Hambuckers, L. Sun, and L. Trapin
University of Bologna, HEC Liège

We study tail risk dynamics in high-frequency financial markets and their connection with trading activity and market uncertainty. We introduce a dynamic extreme value regression model accommodating both stationary and local unit-root predictors to appropriately capture the time-varying behaviour of the distribution of high-frequency extreme losses. To characterize trading activity and market uncertainty, we consider several volatility and liquidity predictors, and propose a two-step adaptive L1-regularized maximum likelihood estimator to select the most appropriate ones. We establish the oracle property of the proposed estimator for selecting both stationary and local unit-root predictors, and show its good finite sample properties in an extensive simulation study. Studying the high-frequency extreme losses of nine large liquid U.S. stocks using 42 liquidity and volatility predictors, we find the severity of extreme losses to be well predicted by low levels of price impact in period of high volatility of liquidity and volatility.

Constrained monotone mean-variance problem with random coefficients

Y. Hu, X. Shi, and Z. Q. Xu
Hong Kong Polytechnic University, Shandong University, CNRS

This paper studies the monotone mean-variance (MMV) problem and the classical mean-variance (MV) with convex cone trading constraints in a market with random coefficients. We provide explicit optimal strategies and optimal values for both problems via certain backward stochastic differential equations (BSDEs). After noting the links between these BSDEs, we find that both the two problems share the same optimal portfolio and optimal value. This generalizes existing results from deterministic coefficients to the random case.

Fully-dynamic risk measures: horizon risk, time-consistency, and relations with BSDEs and BSVIEs

G. Di Nunno and E. R. Gianin
University of Oslo, Norwegian School of Economics, University of Milano Bicocca

In a dynamic framework, we identify a new concept associated with the risk of assessing the financial exposure by a measure that is not adequate to the actual time horizon of the position. This will be called horizon risk. We clarify that dynamic risk measures are subject to horizon risk, so we propose to use the fully-dynamic version. To quantify horizon risk, we introduce h-longevity as an indicator. We investigate these notions together with other properties of risk measures as normalization, restriction property, and different formulations of time-consistency. We also consider these concepts for fully-dynamic risk measures generated by backward stochastic differential equations (BSDEs), backward stochastic Volterra integral equations (BSVIEs), and families of these. Within this study, we provide new results for BSVIEs such as a converse comparison theorem and the dual representation of the associated risk measures.

Optimal stopping: Bermudan strategies meet non-linear evaluations

M. Grigorova, M-C Quenez, and P. Yuan
University of Oslo, Norwegian School of Economics, University of Milano Bicocca

We address an optimal stopping problem over the set of Bermudan-type strategies θ (which we understand in a more general sense than the stopping strategies for Bermudan options in finance) and with non-linear operators (non-linear evaluations) assessing the rewards, under general assumptions on the non-linear operators ρ. We provide a characterization of the value family V in terms of what we call the (ρ, θ)-Snell envelope of the payoff family. We establish a Dynamic Programming Principle. We provide an optimality criterion in terms of a (ρ, θ)-martingale property of V on a stochastic interval. We investigate the (ρ, θ)-martingale structure and we show that the first time when the value family coincides with the pay-off family is optimal. The reasoning simplifies in the case where there is a finite number n of predescribed stopping times, where n does

not depend on the scenario ω. We provide examples of non-linear operators entering our framework.

Economics

Learning to Collude: A Folk Theorem for Algorithms

Álvaro Cartea, P. Chang, José Penalva, and H. Waldon

Oxford-Man Institute, Universidad Carlos III de Madrid, The University of Texas

We introduce state-dependent smooth fictitious play and we use this algorithm to prove a Folk theorem for repeated potential games under one-period perfect monitoring. Our result proves that decentralized learning algorithms can learn to collude through repeated interactions and without communication. Specifically, the algorithms learn to sustain a non-Nash supra-competitive outcome of the stage game by using a credible threat of punishment to suppress competition.

Asymptotic Convergence and Performance of Multi-Agent Q-Learning Dynamics

A. Hussain, F. Belardinelli, and G. Piliouras

Imperial College, SUTD

Achieving convergence of multiple learning agents in general N-player games is imperative for the development of safe and reliable machine learning (ML) algorithms and their application to autonomous systems. Yet it is known that, outside the bounds of simple two-player games, convergence cannot be taken for granted. To make progress in resolving this problem, we study the dynamics of smooth Q-Learning, a popular reinforcement learning algorithm which quantifies the tendency for learning agents to explore their state space or exploit their payoffs. We show a sufficient condition on the rate of exploration such that the Q-Learning dynamics is guaranteed to converge to a unique equilibrium in any game. We connect this result to games for which Q-Learning is known to converge with arbitrary exploration rates, including weighted Potential games and weighted zero sum polymatrix games. Finally, we examine the performance of the Q-Learning dynamic as measured by the Time Averaged Social Welfare, and comparing this with the Social Welfare achieved by the equilibrium. We provide a sufficient condition whereby the Q-Learning dynamic will outperform the equilibrium even if the dynamics do not converge.

The Effect of Outsourcing Pricing Algorithms on Market Competition

J. Harrington

University of Pennsylvania

A third party developer designs and sells a pricing algorithm that enhances a firm's ability to tailor prices to a source of demand variation, whether high-frequency demand shocks or market segmentation. The equilibrium pricing algorithm is characterized that maximizes the third party's profit given firms' optimal adoption decisions. Outsourcing the pricing algorithm does not reduce competition but does make prices more sensitive to the demand variation, and this is shown to decrease consumer welfare and increase industry profit. This effect is larger when products are more substitutable.

Acknowledgement

The OMI Research Newsletter was curated by 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.

References

Barbero, Federico, Cristian Bodnar, Haitz Sáez de Ocáriz Borde, and Pietro Lio. 2022. "Sheaf attention networks."

Borde, Haitz Sáez de Ocáriz, Anees Kazi, Federico Barbero, and Pietro Liò. 2022. "Latent Graph Inference using Product Manifolds." *arXiv preprint arXiv:2211.16199*.

Cartea, Álvaro, Patrick Chang, José Penalva, and Harrison Waldon. 2022. "Learning to Collude: A Folk Theorem for Algorithms." *Available at SSRN 4293831*.

Chassagneux, Jean-François, Junchao Chen, and Noufel Frikha. 2022. "Deep Runge-Kutta schemes for BSDEs." *arXiv preprint arXiv:2212.14372*.

Di Nunno, Giulia, and Emanuela Rosazza Gianin. 2023. "Fully-dynamic risk measures: horizon risk, time-consistency, and relations with BSDEs and BSVIEs." *arXiv preprint arXiv:2301.04971*.

Giles, Michael B, Abdul-Lateef Haji-Ali, and Jonathan Spence. 2023. "Efficient Risk Estimation for the Credit Valuation Adjustment." *arXiv preprint arXiv:2301.05886*.

Grigorova, Miryana, Marie-Claire Quenez, and Peng Yuan. 2023. "Optimal stopping: Bermudan strategies meet non-linear evaluations." *arXiv preprint arXiv:2301.11102*.

Hambuckers, Julien, Li Sun, and Luca Trapin. 2023. "Measuring tail risk at high-frequency: An L_1-regularized extreme value regression approach with unit-root predictors." *arXiv preprint arXiv:2301.01362*.

Harrington Jr, Joseph E. 2022. "The effect of outsourcing pricing algorithms on market competition." *Management Science*, 68(9): 6889–6906.

Hussain, Aamal Abbas, Francesco Belardinelli, and Georgios Piliouras. 2023. "Asymptotic Convergence and Performance of Multi-Agent Q-Learning Dynamics." *arXiv preprint arXiv:2301.09619*.

Hu, Ying, Xiaomin Shi, and Zuo Quan Xu. 2022. "Constrained monotone mean-variance problem with random coefficients." *arXiv preprint arXiv:2212.14188*.

Lavigne, Pierre, and Peter Tankov. 2023. "Decarbonization of financial markets: a mean-field game approach." *arXiv preprint arXiv:2301.09163*.

Nagy, Peer, Jan-Peter Calliess, and Stefan Zohren. 2023. "Asynchronous Deep Double Duelling Q-Learning for Trading-Signal Execution in Limit Order Book Markets." *arXiv preprint arXiv:2301.08688*.

Neuman, Eyal, and Yufei Zhang. 2023. "Statistical Learning with Sublinear Regret of Propagator Models." *arXiv preprint arXiv:2301.05157*.

Shabani, Amin, Amir Abdi, Lili Meng, and Tristan Sylvain. 2022. "Scaleformer: iterative multi-scale refining transformers for time series forecasting." *arXiv preprint arXiv:2206.04038*.

Shirai, Yoshihiro. 2023. "A pure jump model for the valuation of options on a credit index." *arXiv preprint arXiv:2301.05332*.

Vuletić, Milena, Felix Prenzel, and Mihai Cucuringu. 2023. "Fin-GAN: Forecasting and Classifying Financial Time Series via Generative Adversarial Networks." *Available at SSRN 4328302*.

Wong, Thomas, and Mauricio Barahona. 2022. "Robust machine learning pipelines for trading market-neutral stock portfolios." *arXiv preprint arXiv:2301.00790*.