November 2022, Vol. 2, No. 2

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 members of the OMI. The first uses Causal Convolutions to forecast volatility with high-frequency data. The second won the Best Paper Award at the ICAIF'22 conference and uses generative adversarial networks to learn and generate realistic order flow dynamics based on market features.

Another notable contribution from Princeton University explores order attribution between signed and anonymous trading flow and uses Mean Field Game tools to analyse its effects on optimal execution strategies.

Machine Learning for Finance

DeepVol: Volatility Forecasting from High-Frequency Data with Dilated Causal Convolutions

F. Moreno-Pinoa and S. Zohren

Oxford-Man Institute, Universidad Carlos III de Madrid

Volatility forecasts play a central role among equity risk measures. Besides traditional statistical models, modern forecasting techniques, based on machine learning, can readily be employed when treating volatility as a univariate, daily time-series. However, econometric studies have shown that increasing the number of daily observations with high-frequency intraday data helps to improve predictions. In this work, we propose DeepVol, a model based on Dilated Causal Convolutions to forecast day-ahead volatility by using high-frequency data. We show that the dilated convolutional filters are ideally suited to extract relevant information from intraday financial data, thereby naturally mimicking (via a data-driven approach) the econometric models which incorporate realised measures of volatility into the forecast. This allows us to take advantage of the abundance of intraday observations, helping us to avoid the limitations of models that use daily data, such as model misspecification or manually designed handcrafted features, whose devise involves optimising the trade-off between accuracy and computational efficiency and makes models prone to lack of adaptation into changing circumstances. In our analysis, we use two years of intraday data from NASDAQ-100 to evaluate DeepVol's performance. The reported empirical results suggest that the proposed deep learning-based approach learns global features from high-frequency data, achieving more accurate predictions than traditional methodologies, yielding to more appropriate risk measures.

Detecting asset price bubbles using deep learning

F. Biagini, L. Gonon, A. Mazzon, and T. Meyer-Brandis

Imperial College Londo, Ludwig-Maximilians Universitat

In this paper we employ deep learning techniques to detect financial asset bubbles by using observed call option prices. The proposed algorithm is widely applicable and model-independent. We test the accuracy of our methodology in numerical experiments within a wide range of models and apply it to market data of tech stocks in order to assess if asset price bubbles are present. In addition, we provide a theoretical foundation of our approach in the framework of local volatility models. To this purpose, we give a new necessary and sufficient condition for a process with time-dependent local volatility function to be a strict local martingale.

Deep neural network expressivity for optimal stopping problems

L. Gonon

Imperial College London

This article studies deep neural network expression rates for optimal stopping problems of discrete-time Markov processes on high-dimensional state spaces. A general framework is established in which the value function and continuation value of an optimal stopping problem can be approximated with error at most ε by a deep ReLU neural network of size at most $\kappa d^{\mathfrak{q}}\varepsilon^{-\mathfrak{r}}$. The constants $\kappa, \mathfrak{q}, \mathfrak{r} \geq 0$ do not depend on the dimension d of the state space or the approximation accuracy ε. This proves that deep neural networks do not suffer from the curse of dimensionality when employed to solve optimal stopping problems. The framework covers, for example, exponential Lévy models, discrete diffusion processes and their running minima and maxima. These results mathematically justify the use of deep neural networks for numerically solving optimal stopping problems and pricing American options in high dimensions.

Dynamic Calibration of Order Flow Models with Generative Adversarial Networks

F. Prenzel, Rama Cont, Mihai Cucuringu, and Jonathan Kochems

University of Oxford, JP Morgan

Classical models for order flow dynamics based on point processes, such as Poisson or Hawkes processes, have been studied intensively. Often, several days of limit border book (LOB) data is used to calibrate such models, thereby averaging over different dynamics - such as intraday effects or different trading volumes.

This work uses generative adversarial networks (GANs) to learn the distribution of calibrations – obtained by many calibrations based on short time frames. The trained GAN can then be used to generate synthetic, realistic calibrations based on external conditions such as time of the day or volatility. Results show that GANs easily reproduce patterns of the order arrival intensities and can fit the distribution well without heavy parameter tuning.

The synthetic calibrations can then be used to simulate order streams which contain new dynamics such as temporary drifts, different volatility regimes, but also intra-day patterns such as the commonly observed U-shape that reflects stylized behaviour around open and close of market hours.

Optimal Stopping with Gaussian Processes

K. Dwarakanath, D. Dervovic, P. Tavallali, S. S Vyetrenko, and T. Balch

J.P.Morgan AI Research

We propose a novel group of Gaussian Process based algorithms for fast approximate optimal stopping of time series with specific applications to financial markets. We show that structural properties commonly exhibited by financial time series (e.g., the tendency to mean-revert) allow the use of Gaussian and Deep Gaussian Process models that further enable us to analytically evaluate optimal stopping value functions and policies. We additionally quantify uncertainty in the value function by propagating the price model through the optimal stopping analysis. We compare and contrast our proposed methods against a sampling-based method, as well as a deep learning based benchmark that is currently considered the state-of-the-art in the literature. We show that our family of algorithms outperforms benchmarks on three historical time series datasets that include intra-day and end-of-day equity stock prices as well as the daily US treasury yield curve rates.

Efficient and Near-Optimal Online Portfolio Selection

R. Jézéquel, D. Ostrovskii, P. Gaillard

Ecole Normale Supérieure, University of Southern California, Inria

In the problem of online portfolio selection as formulated by Cover (1991), the trader repeatedly distributes her capital over d assets in each of $T > 1$ rounds, with the goal of maximizing the total return. Cover proposed an algorithm, termed Universal Portfolios, that performs nearly as well as the best (in hindsight) static assignment of a portfolio, with an $\mathcal{O}(d\log(T))$ regret in terms of the logarithmic return. Without imposing any restrictions on the market this guarantee is known to be worst-case optimal, and no other algorithm attaining it has been discovered so far. Unfortunately, Cover's algorithm crucially relies on computing certain d-dimensional integral which must be approximated in any implementation; this results in a prohibitive $\mathcal{O}(d^4(T+d)^{14})$ per-round runtime for the fastest known implementation due to Kalai and Vempala (2002). We propose an algorithm for online portfolio selection that admits essentially the same regret guarantee as Universal Portfolios – up to a constant factor and replacement of $log(T)$ with $log(T+d)$ – yet has a drastically reduced runtime of $\mathcal{O}(d^2(T+d))$ per round. The selected portfolio minimizes the current logarithmic loss regularized by the log-determinant of its Hessian – equivalently, the hybrid logarithmic-volumetric barrier of the polytope specified by the asset return vectors. As such, our work reveals surprising connections of online portfolio selection with two classical topics in optimization theory: cutting-plane and interior-point algorithms.

Mathematical Finance

Optimal Execution with Identity Optionality

R. Carmona and C. Zeng
Princeton University

This paper investigates the impact of anonymous trading on the agents' strategy in an optimal execution framework. It mainly explores the specificity of order attribution on the Toronto Stock Exchange, where brokers can choose to either trade with their own identity or under a generic anonymous code that is common to all the brokers. We formulate a stochastic differential game for the optimal execution problem of a population of N brokers and incorporate permanent and temporary price impacts for both the identity-revealed and anonymous trading processes. We then formulate the limiting mean-field game of controls with common noise and obtain a solution in closed-form via the probabilistic approach for the Almgren-Chris price impact framework. Finally, we perform a sensitivity analysis to explore the impact of the model parameters on the optimal strategy.

Quantitative Fundamental Theorem of Asset Pricing

B. Acciaio, J. Backhoff, and G. Pammer
ETH Zurich

In this paper we provide a quantitative analysis to the concept of arbitrage, that allows to deal with model uncertainty without imposing the no-arbitrage condition. In markets that admit "small arbitrage", we can still make sense of the problems of pricing and hedging. The pricing measures here will be such that asset price processes are close to being martingales, and the hedging strategies will need to cover some additional cost. We show a quantitative version of the Fundamental Theorem of Asset Pricing and of the Super-Replication Theorem. Finally, we study robustness of the amount of arbitrage and existence of respective pricing measures, showing stability of these concepts with respect to a strong adapted Wasserstein distance.

Fast and Slow Optimal Trading with Exogenous Information

A. Micheli and E. Neuman
Imperial College London

We consider a stochastic game between a slow institutional investor and a high-frequency trader who are trading a risky asset and their aggregated order-flow impacts the asset price. We model this system by means of two coupled stochastic control problems, in which the high-frequency trader exploits the available information on a price predicting signal more frequently, but is also subject to periodic "end of day" inventory constraints. We first derive the optimal strategy of the high-frequency trader given any admissible strategy of the institutional investor. Then, we solve the problem of the institutional investor given the optimal signal-adaptive strategy of the high-frequency trader, in terms of the resolvent of a Fredholm integral equation, thus establishing the unique multi-period Stackelberg equilibrium of the game. Our results provide an explicit solution to the game, which shows that the high-frequency trader can adopt either predatory or cooperative strategies in each period, depending on the tradeoff between the order-flow and the trading signal. We also show that the institutional investor's strategy is considerably more profitable when the order-flow of the high-frequency trader is taken into account in her trading strategy.

The rough Hawkes Heston stochastic volatility model

A. Bondi, S. Pulido, and S. Scotti
Université Paris-Saclay, Università di Pisa

We study an extension of the Heston stochastic volatility model that incorporates rough volatility and jump clustering phenomena. In our model, named the rough Hawkes Heston stochastic volatility model, the spot variance is a rough Hawkes-type process proportional to the intensity process of the jump component appearing in the dynamics of the spot variance itself and the log returns. The model belongs to the class of affine Volterra models. In particular, the Fourier-Laplace transform of the log returns and the square of the volatility index can be computed explicitly in terms of solutions of deterministic Riccati-Volterra equations, which can be efficiently approximated using a multi-factor approximation technique. We calibrate a parsimonious specification of our model characterized by a power kernel and an exponential law for the jumps. We show that our parsimonious setup is able to simultaneously capture, with a high precision, the behavior of the implied volatility smile for both S&P 500 and VIX options. In particular, we observe that in our setting the usual shift in the implied volatility of VIX options is explained by a very low value of the power in the kernel. Our findings demonstrate the relevance, under an affine framework, of rough volatility and self-exciting jumps in order to capture the joint evolution of the S&P 500 and VIX.

Statistical inference for rough volatility: Minimax Theory

C. Chong, M. Hoffmann, Y. Liu, M. Rosenbaum, and G. Szymanski
Columbia University, Ecole Polytechnique

Rough volatility models have gained considerable interest in the quantitative finance community in recent years. In

this paradigm, the volatility of the asset price is driven by a fractional Brownian motion with a small value for the Hurst parameter H. In this work, we provide a rigorous statistical analysis of these models. To do so, we establish minimax lower bounds for parameter estimation and design procedures based on wavelets attaining them. We notably obtain an optimal speed of convergence of $n^{1/(4H+2)}$ for estimating H based on n sampled data, extending results known only for the easier case $H > 1/2$ so far. We therefore establish that the parameters of rough volatility models can be inferred with optimal accuracy in all regimes.

High order approximations of the Cox-Ingersoll-Ross process semigroup using random grids

A. Alfonsi and E. Lombardo
Ecole des Ponts ParisTech

We present new high order approximations schemes for the Cox-Ingersoll-Ross (CIR) process that are obtained by using a recent technique developed by [Alfonsi and Bally (2021)] for the approximation of semigroups. The idea consists in using a suitable combination of discretization schemes calculated on different random grids to increase the order of convergence. This technique coupled with the second order scheme proposed by [Alfonsi (2010)] for the CIR leads to weak approximations of order $2k$, for all $k \in \mathbb{N}_\star$. Despite the singularity of the square-root volatility coefficient, we show rigorously this order of convergence under some restrictions on the volatility parameters. We illustrate numerically the convergence of these approximations for the CIR process and for the Heston stochastic volatility model and show the computational time gain they give.

Economics

Can maker-taker fees prevent algorithmic cooperation in market making?

B. Han
University of Michigan

In a semi-realistic market simulator, independent reinforcement learning algorithms may facilitate market makers to maintain wide spreads even without communication. This unexpected outcome challenges the current antitrust law framework. We study the effectiveness of maker-taker fee models in preventing cooperation via algorithms. After modeling market making as a repeated general-sum game, we experimentally show that the relation between net transaction costs and maker rebates is not necessarily monotone. Besides an upper bound on taker fees, we may also need a lower bound on maker rebates to destabilize the cooperation. We also consider the taker-maker model and the effects of mid-price volatility, inventory risk, and the number of agents.

Artificial Collusion: Examining Supracompetitive Pricing by Q-Learning Algorithms

A. den Boer, J. Meylahn, M. Schinkel
University of Amsterdam

We examine recent claims that a particular Q-learning algorithm used by competitors 'autonomously' and systematically learns to collude, resulting in supracompetitive prices and extra profits for the firms sustained by collusive equilibria. A detailed analysis of the inner workings of this algorithm reveals that there is no immediate reason for alarm. We set out what is needed to demonstrate the existence of a colluding price algorithm that does form a threat to competition.

Towards Multi-Agent Reinforcement Learning driven Over-The-Counter Market Simulations

N. Vadori, L. Ardon, S. Ganesh, T. Spooner, S. Amrouni, J. Vann, M. Xu, Z. Zheng, T. Balch, and M. Veloso
J.P. Morgan AI Research

We study a game between liquidity provider and liquidity taker agents interacting in an over-the-counter market, for which the typical example is foreign exchange. We show how a suitable design of parameterized families of reward functions coupled with associated shared policy learning constitutes an efficient solution to this problem. Precisely, we show that our deep-reinforcement-learning-driven agents learn emergent behaviors relative to a wide spectrum of incentives encompassing profit-and-loss, optimal execution and market share, by playing against each other. In particular, we find that liquidity providers naturally learn to balance hedging and skewing as a function of their incentives, where the latter refers to setting their buy and sell prices asymmetrically as a function of their inventory. We further introduce a novel RL-based calibration algorithm which we found performed well at imposing constraints on the game equilibrium, both on toy and real market data.

Others

Optimization-Informed Neural Networks

D. Wu and A. Lisser
Université Paris-Saclay, CNRS

Solving constrained nonlinear optimization problems (CNLPs) is a longstanding problem that arises in various fields, e.g., economics, computer science, and engineering. We propose optimization-informed neural networks (OINN), a deep learning approach to solve CNLPs. By neurodynamic optimization methods, a CNLP is first

reformulated as an initial value problem (IVP) involving an ordinary differential equation (ODE) system. A neural network model is then used as an approximate solution for this IVP, with the endpoint being the prediction to the CNLP. We propose a novel training algorithm that directs the model to hold the best prediction during training. In a nutshell, OINN transforms a CNLP into a neural network training problem. By doing so, we can solve CNLPs based on deep learning infrastructure only, without using standard optimization solvers or numerical integration solvers. The effectiveness of the proposed approach is demonstrated through a collection of classical problems, e.g., variational inequalities, nonlinear complementary problems, and standard CNLPs.

Rhino: Deep Causal Temporal Relationship Learning With History-dependent Noise

W. Gong, J. Jennings, C. Zhang, and N. Pawlowski

Microsoft Research, Cambridge

Discovering causal relationships between different variables from time series data has been a long-standing challenge for many domains such as climate science, finance, and healthcare. Given the complexity of real-world relationships and the nature of observations in discrete time, causal discovery methods need to consider non-linear relations between variables, instantaneous effects and history-dependent noise (the change of noise distribution due to past actions). However, previous works do not offer a solution addressing all these problems together. In this paper, we propose a novel causal relationship learning framework for time-series data, called Rhino, which combines vector auto-regression, deep learning and variational inference to model non-linear relationships with instantaneous effects while allowing the noise distribution to be modulated by historical observations. Theoretically, we prove the structural identifiability of Rhino. Our empirical results from extensive synthetic experiments and two real-world benchmarks demonstrate better discovery performance compared to relevant baselines, with ablation studies revealing its robustness under model misspecification.

Factor Augmented Sparse Throughput Deep ReLU Neural Networks for High Dimensional Regression

J. Fan and Y. Gu

Princeton University

This paper introduces a Factor Augmented Sparse Throughput (FAST) model that utilizes both latent factors and sparse idiosyncratic components for nonparametric regression. The FAST model bridges factor models on one end and sparse nonparametric models on the other end. It encompasses structured nonparametric models such as factor augmented additive models and sparse low-dimensional nonparametric interaction models and covers the cases where the covariates do not admit factor structures. Via diversified projections as estimation of latent factor space, we employ truncated deep ReLU networks to nonparametric factor regression without regularization and to a more general FAST model using nonconvex regularization, resulting in factor augmented regression using neural network (FAR-NN) and FAST-NN estimators respectively. We show that FAR-NN and FAST-NN estimators adapt to the unknown low-dimensional structure using hierarchical composition models in nonasymptotic minimax rates. We also study statistical learning for the factor augmented sparse additive model using a more specific neural network architecture. Our results are applicable to the weak dependent cases without factor structures. In proving the main technical result for FAST-NN, we establish a new deep ReLU network approximation result that contributes to the foundation of neural network theory. Our theory and methods are further supported by simulation studies and an application to macroeconomic data.

Axioms for Automated Market Makers: A Mathematical Framework in FinTech and Decentralized Finance

M. Bichuch, and Z. Feinstein

Stevens Institute of Technology, University at Buffalo

Within this work we consider an axiomatic framework for Automated Market Makers (AMMs). By imposing reasonable axioms on the underlying utility function, we are able to characterize the properties of the swap size of the assets and of the resulting pricing oracle. We have analyzed many existing AMMs and shown that the vast majority of them satisfy our axioms. We have also considered the question of fees and divergence loss. In doing so, we have proposed a new fee structure so as to make the AMM indifferent to transaction splitting. Finally, we have proposed a novel AMM that has nice analytical properties and provides a large range over which there is no divergence loss.

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.