VOL. 2, NO. 4

OMI Research Newsletter – January 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 a paper from the OMI and a paper from Princeton University and Bloomberg L.P. The first solves an indefinite Riccati equation that arises in various trading problems where the agent uses signals and past prices to enhance trading performance, and uses the result to derive optimal strategies for market making and execution. The second introduces functional Taylor expansion (FTE) as a pathwise tool combining functional Ito calculus and signatures to unravel path dependence, and applies the approach for the pricing and hedging of exotic claims.

Machine Learning for Finance

Langevin algorithms for Markovian Neural Networks and Deep Stochastic control

P. BRAS
SORBONNE UNIVERSITÉ

Stochastic Gradient Descent Langevin Dynamics (SGLD) algorithms, which add noise to the classic gradient descent, are known to improve the training of neural networks in some cases where the neural network is very deep. In this paper we study the possibilities of training acceleration for the numerical resolution of stochastic control problems through gradient descent, where the control is parametrized by a neural network. If the control is applied at many discretization times then solving the stochastic control problem reduces to minimizing the loss of a very deep neural network. We numerically show that Langevin algorithms improve the training on various stochastic control problems like hedging and resource management, and for different choices of gradient descent methods.

Multiarmed Bandits Problem Under the Mean-Variance Setting

H. HU, A. CHARPENTIER, M. GHOSSOUB, AND A. SCHIED
UNIVERSITY OF WATERLOO, UNIVERSITÉ DU QÉBEC

Stochastic Gradient Descent Langevin Dynamics (SGLD) algorithms, which add noise to the classic gradient descent, are known to improve the training of neural networks in some cases where the neural network is very deep. In this paper we study the possibilities of training acceleration for the numerical resolution of stochastic control problems through gradient descent, where the control is parametrized by a neural network. If the control is applied at many discretization times then solving the stochastic control problem reduces to minimizing the loss of a very deep neural network. We numerically show that Langevin algorithms improve the training on various stochastic control problems like hedging and resource management, and for different choices of gradient descent methods. The classical multi-armed bandit (MAB) problem involves a learner and a collection of K independent arms, each with its own ex ante unknown independent reward distribution. At each one of a finite number of rounds, the learner selects one arm and receives new information. The learner often faces an exploration-exploitation dilemma: exploiting the current information by playing the arm with the highest estimated reward versus exploring all arms to gather more reward information. The design objective aims to maximize the expected cumulative reward over all rounds. However, such an objective does not account for a risk-reward tradeoff, which is often a fundamental precept in many areas of applications, most notably in finance and economics. In this paper, we build upon Sani et al. (2012) and extend the classical MAB problem to a mean-variance setting. Specifically, we relax the

assumptions of independent arms and bounded rewards made in Sani et al. (2012) by considering sub-Gaussian arms. We introduce the Risk Aware Lower Confidence Bound (RALCB) algorithm to solve the problem, and study some of its properties. Finally, we perform a number of numerical simulations to demonstrate that, in both independent and dependent scenarios, our suggested approach performs better than the algorithm suggested by Sani et al. (2012).

Empirical Asset Pricing via Ensemble Gaussian Process Regression

D. Filipović and P. Pasricha
Ecole Polytechnique Fédérale de Lausanne

We introduce an ensemble learning method based on Gaussian Process Regression (GPR) for predicting conditional expected stock returns given stock-level and macro-economic information. Our ensemble learning approach significantly reduces the computational complexity inherent in GPR inference and lends itself to general online learning tasks. We conduct an empirical analysis on a large cross-section of US stocks from 1962 to 2016. We find that our method dominates existing machine learning models statistically and economically in terms of out-of-sample R-squared and Sharpe ratio of prediction-sorted portfolios. Exploiting the Bayesian nature of GPR, we introduce the mean-variance optimal portfolio with respect to the predictive uncertainty distribution of the expected stock returns. It appeals to an uncertainty averse investor and significantly dominates the equal- and value-weighted prediction-sorted portfolios, which outperform the S&P 500.

DeFi: data-driven characterisation of Uniswap v3 ecosystem & an ideal crypto law for liquidity pools

M. Cucuringu and D. Miori
Oxford-Man Institute

The Uniswap v3 ecosystem is built upon liquidity pools, where pairs of tokens are exchanged subject to a fee. We propose a systematic workflow to extract a meaningful but tractable sub-universe out of the current $> 6{,}000$ pools. We filter by imposing minimum levels on individual pool features, e.g. liquidity locked and agents' activity, but also maximising the interconnection between the chosen pools to support broader dynamics. Then, we investigate liquidity consumption behaviour on the most relevant pools for Jan-June 2022. We propose to describe each liquidity taker by a transaction graph, which is a complete graph where nodes are transactions on pools and edges have weights from the time elapsed between pairs of transactions. Each graph is embedded into a vector by our own variant of the NLP rooted graph2vec algorithm. Thus, we are able to investigate the structural equivalence of liquidity takers behaviour and extract seven clusters with interpretable features. Finally, we introduce an ideal crypto law inspired from the ideal gas law of thermodynamics. Our model tests a relationship between variables that govern the mechanisms of each pool, i.e. liquidity provision, consumption, and price variation. If the law is satisfied, we say the pool has high cryptoness and demonstrate that it constitutes a better venue for the activity of market participants. Our metric could be employed by regulators and practitioners for developing pool health monitoring tools and establishing minimum levels of requirements.

Axial-LOB: High-Frequency Trading with Axial Attention

D. Kisiel and D. Gorse
University College London

Previous attempts to predict stock price from limit order book (LOB) data are mostly based on deep convolutional neural networks. Although convolutions offer efficiency by restricting their operations to local interactions, it is at the cost of potentially missing out on the detection of long-range dependencies. Recent studies address this problem by employing additional recurrent or attention layers that increase computational complexity. In this work, we propose Axial-LOB, a novel fully-attentional deep learning architecture for predicting price movements of stocks from LOB data. By utilizing gated position-sensitive axial attention layers our architecture is able to construct feature maps that incorporate global interactions, while significantly reducing the size of the parameter space. Unlike previous works, Axial-LOB does not rely on hand-crafted convolutional kernels and hence has stable performance under input permutations and the capacity to incorporate additional LOB features. The effectiveness of Axial-LOB is demonstrated on a large benchmark dataset, containing time series representations of millions of high-frequency trading events, where our model establishes a new state of the art, achieving an excellent directional classification performance at all tested prediction horizons.

Understanding stock market instability via graph auto-encoders

D. Gorduza, X. Dong, and S. Zohren
Oxford-Man Institute

Understanding stock market instability is a key question in financial management as practitioners seek to forecast breakdowns in asset co-movements which expose portfolios to rapid and devastating collapses in value. The structure of these co-movements can be described as a graph where companies are represented by nodes and edges capture correlations between their price movements. Learning a timely indicator of co-movement

breakdowns (manifested as modifications in the graph structure) is central in understanding both financial stability and volatility forecasting. We propose to use the edge reconstruction accuracy of a graph auto-encoder (GAE) as an indicator for how spatially homogeneous connections between assets are, which, based on financial network literature, we use as a proxy to infer market volatility. Our experiments on the S&P 500 over the 2015-2022 period show that higher GAE reconstruction error values are correlated with higher volatility. We also show that out-of-sample autoregressive modeling of volatility is improved by the addition of the proposed measure. Our paper contributes to the literature of machine learning in finance particularly in the context of understanding stock market instability.

Mathematical Finance

Solvability of Differential Riccati Equations and Applications to Algorithmic Trading with Signals

F. Drissi
Oxford-Man Institute, Université Paris 1 Panthéon-Sorbonne

We study a differential Riccati equation (DRE) with indefinite matrix coefficients, which arises in a wide class of practical problems. We show that the DRE solves an associated control problem, which is key to provide existence and uniqueness of a solution. As an application, we solve two algorithmic trading problems in which the agent adopts a constant absolute risk-aversion (CARA) utility function, and where the optimal strategies use signals and past observations of prices to improve their performance. First, we derive a multi-asset market making strategy in over-the-counter markets, where the market maker uses an external trading venue to hedge risk. Second, we derive an optimal trading strategy that uses prices and signals to learn the drift in the asset prices.

Multi-asset market making under the quadratic rough Heston

M. Rosenbaum and J. Zhang
Ecole polytechnique

Given the promising results on joint modeling of SPX/VIX smiles of the recently introduced quadratic rough Heston model, we consider a multi-asset market making problem on SPX and its derivatives, e.g. VIX futures, SPX and VIX options. The market maker tries to maximize its profit from spread capturing while controlling the portfolio's inventory risk, which can be fully explained by the value change of SPX under the particular setting of the quadratic rough Heston model. The high dimensionality of the resulting optimization problem is relaxed by several approximations. An asymptotic closed-form solution can be obtained. The accuracy and relevance of the approximations are illustrated through numerical experiments.

Functional Expansions

B. Dupire and V. Tissot-Daguette
Bloomberg L.P, Princeton University

Path dependence is omnipresent in social science, engineering, and finance. It reflects the influence of the past on the future, often expressed through functionals. However, non-Markovian problems are often infinite-dimensional, thus challenging from a conceptual and computational perspective. In this work, we shed light on expansions of functionals. First, we treat static expansions made around paths of fixed length and propose a generalization of the Wiener series–the intrinsic value expansion (IVE). In the dynamic case, we revisit the functional Taylor expansion (FTE). The latter connects the functional Itô calculus with the signature to quantify the effect in a functional when a "perturbation" path is concatenated with the source path. In particular, the FTE elegantly separates the functional from future trajectories. The notions of real analyticity and radius of convergence are also extended to the path space. We discuss other dynamic expansions arising from Hilbert projections and the Wiener chaos, and finally show financial applications of the FTE to the pricing and hedging of exotic contingent claims.

Measuring price impact and information content of trades in a time-varying setting

F. Campigli, G. Bormetti, and F. Lillo
Scuola Normale Superiore, Universit'a di Bologna

The estimation of market impact is crucial for measuring the information content of trades and for transaction cost analysis. Hasbrouck's (1991) seminal paper proposed a Structural-VAR (S-VAR) to jointly model mid-quote changes and trade signs. Recent literature has highlighted some pitfalls of this approach: S-VAR models can be misspecified when the impact function has a non-linear relationship with the trade sign, and they lack parsimony when they are designed to capture the long memory of the order flow. Finally, the instantaneous impact of a trade is constant, while market liquidity highly fluctuates in time. This paper fixes these limitations by extending Hasbrouck's approach in several directions. We consider a nonlinear model where we use a parsimonious parametrization allowing to consider hundreds of past lags. Moreover we adopt an observation driven approach to model the time-varying impact parameter, which adapts to market information flow and can be easily

estimated from market data. As a consequence of the non-linear specification of the dynamics, the trade information content is conditional both on the local level of liquidity, as modeled by the dynamic instantaneous impact coefficient, and on the state of the market. By analyzing NASDAQ data, we find that impact follows a clear intra-day pattern and quickly reacts to pre-scheduled announcements, such as those released by the FOMC. We show that this fact has relevant consequences for transaction cost analysis by deriving an expression for the permanent impact from the model parameters and connecting it with the standard regression procedure. Monte Carlo simulations and empirical analyses support the reliability of our approach, which exploits the complete information of tick-by-tick prices and trade signs without the need for aggregation on a macroscopic time scale.

Robust Bond Portfolio Construction via Convex-Concave Saddle Point Optimization

E. Luxenberg, P. Schiele, and S. Boyd
Stanford University,
Ludwig-Maximilians-Universitat Munchen

The minimum (worst case) value of a long-only portfolio of bonds, over a convex set of yield curves and spreads, can be estimated by its sensitivities to the points on the yield curve. We show that sensitivity based estimates are conservative, i.e., underestimate the worst case value, and that the exact worst case value can be found by solving a tractable convex optimization problem. We then show how to construct a long-only bond portfolio that includes the worst case value in its objective or as a constraint, using convex-concave saddle point optimization.

Joint SPX-VIX calibration with Gaussian polynomial volatility models: deep pricing with quantization hints

E. Abi Jaber, C. Illand, and S. Li
Ecole Polytechnique, Université Paris 1
Panthéon-Sorbonne

We consider the joint SPX-VIX calibration within a general class of Gaussian polynomial volatility models in which the volatility of the SPX is assumed to be a polynomial function of a Gaussian Volterra process defined as a stochastic convolution between a kernel and a Brownian motion. By performing joint calibration to daily SPX-VIX implied volatility surface data between 2012 and 2022, we compare the empirical performance of different kernels and their associated Markovian and non-Markovian models, such as rough and non-rough path-dependent volatility models. In order to ensure an efficient calibration and a fair comparison between the models, we develop a generic unified method in our class of models for fast and accurate pricing of SPX and VIX derivatives based on functional quantization and Neural Networks. For the first time, we identify a *conventional one-factor Markovian continuous stochastic volatility model* that is able to achieve remarkable fits of the implied volatility surfaces of the SPX and VIX together with the term structure of VIX futures. What is even more remarkable is that our conventional one-factor Markovian continuous stochastic volatility model outperforms, in all market conditions, its rough and non-rough path-dependent counterparts with the same number of parameters.

Arbitrage theory in a market of stochastic dimension

E. Bayraktar, D. Kim, and A. Tilva
University of Michigan, Columbia University

This paper studies an equity market of stochastic dimension, where the number of assets fluctuates over time. In such a market, we develop the fundamental theorem of asset pricing, which provides the equivalence of the following statements: (i) there exists a supermartingale numéraire portfolio; (ii) each dissected market, which is of a fixed dimension between dimensional jumps, has locally finite growth; (iii) there is no arbitrage of the first kind; (iv) there exists a local martingale deflator; (v) the market is viable. We also present the Optional decomposition theorem, which characterizes a given nonnegative process as the wealth process of some investment-consumption strategy. These results are developed in an equity market model where the price process is given by a piecewise continuous semimartingale of stochastic dimension. Without the continuity assumption on the price process, we present similar results but without explicit characterization of the numéraire portfolio.

Mean-field neural networks-based algorithms for McKean-Vlasov control problems

H. Pham and X. Warin
Université Paris Cité

This paper is devoted to the numerical resolution of McKean-Vlasov control problems via the class of mean-field neural networks introduced in our companion paper [Pham and Warin, 2022] in order to learn the solution on the Wasserstein space. We propose several algorithms either based on dynamic programming with control learning by policy or value iteration, or backward SDE from stochastic maximum principle with global or local loss functions. Extensive numerical results on different examples are presented to illustrate the accuracy of each of our eight algorithms. We discuss and compare the pros and cons of all the tested methods.

Explicit Caplet Implied Volatilities for Quadratic Term-Structure Models

M. Lorig and N. Suaysom
University of Washington

We derive an explicit asymptotic approximation for implied volatilities of caplets under the assumption that the short-rate is described by a generic quadratic term-structure model. In addition to providing an asymptotic accuracy result, we perform experiments in order to gauge the numerical accuracy of our approximation.

Economics

Regulating Collusion

S. Chassang and J. Ortner
Princeton University, Boston University

We attempt to provide a systemic view of the process of regulating collusion, including detection and prosecution, but also bargaining between firms and regulators via consent orders, the production of evidence, and containment measures that may be taken if collusion cannot be addressed in more direct means. In addition, we try to do justice to the peculiarities of the legal system: modeling the courts as they are, rather than as economists think they should be, is essential for economic analysis to improve the way collusion is regulated.

Strategic Observational Learning

D. Migrow
University of Edinburgh

This paper characterizes equilibrium behavior of forward-looking agents in a simple repeated-action setting of social learning with two privately informed players. If players have access to a symmetric signal structure (in a sense to be made precise), then myopic strategies always constitute an equilibrium. Myopic equilibrium is unique under symmetric threshold strategies, and under the simplest symmetric non-monotonic strategies. Under arbitrary threshold strategies, only myopic equilibrium can fully aggregate dispersed private information. If the players are asymmetric and the game is infinite, there is no equilibrium in myopic strategies, for any positive degree of patience.

Anticipatory Fictitious Play

A. Cloud, A. Wang, and W. Kerr
North Carolina State University

Fictitious play is an algorithm for computing Nash equilibria of matrix games. Recently, machine learning variants of fictitious play have been successfully applied to complicated real-world games. This paper presents a simple modification of fictitious play which is a strict improvement over the original: it has the same theoretical worst-case convergence rate, is equally applicable in a machine learning context, and enjoys superior empirical performance. We conduct an extensive comparison of our algorithm with fictitious play, proving an optimal convergence rate for certain classes of games, demonstrating superior performance numerically across a variety of games, and concluding with experiments that extend these algorithms to the setting of deep multiagent reinforcement learning.

Multilayer structure enhances the optimal outcome of coordination games

T. Raducha and M. San Miguel
Universidad Carlos III de Madrid

We study mechanisms of synchronisation, coordination, and equilibrium selection in two-player coordination games on multilayer networks. We apply the approach from evolutionary game theory with three possible update rules: the replicator dynamics (RD), the best response (BR), and the unconditional imitation (UI). Players interact on a two-layer random regular network. The population on each layer plays a different game, with layer I preferring the opposite strategy to layer II. We measure the difference between the two games played on the layers by a difference in payoffs ΔS while the inter-connectedness is measured by a node overlap parameter q. We discover a critical value $q_c(\Delta S)$ below which layers do not synchronise. For $q > q_c$ in general both layers coordinate on the same strategy. Surprisingly, there is a symmetry breaking in the selection of equilibrium – for RD and UI there is a phase where only the payoff-dominant equilibrium is selected. Our work is an example of previously observed differences between the update rules on a single network. However, we took a novel approach with the game being played on two inter-connected layers. As we show, the multilayer structure enhances the abundance of the Pareto-optimal equilibrium in coordination games with imitative update rules.

Machine Learning

WaveBound: Dynamic Error Bounds for Stable Time Series Forecasting

Y. Cho, D. Kim, D. Kim, M. Azam Khan, and J. Choo
KAIST AI

Time series forecasting has become a critical task due to its high practicality in real-world applications such as traffic, energy consumption, economics and finance, and

disease analysis. Recent deep-learning-based approaches have shown remarkable success in time series forecasting. Nonetheless, due to the dynamics of time series data, deep networks still suffer from unstable training and overfitting. Inconsistent patterns appearing in real-world data lead the model to be biased to a particular pattern, thus limiting the generalization. In this work, we introduce the dynamic error bounds on training loss to address the overfitting issue in time series forecasting. Consequently, we propose a regularization method called WaveBound which estimates the adequate error bounds of training loss for each time step and feature at each iteration. By allowing the model to focus less on unpredictable data, WaveBound stabilizes the training process, thus significantly improving generalization. With the extensive experiments, we show that WaveBound consistently improves upon the existing models in large margins, including the state-of-the-art model.

Generative Time Series Forecasting with Diffusion, Denoise, and Disentanglement

Y. Li, X. Lu, Y. Wang, and D. Dou

Baidu Research, Zhejiang University

Time series forecasting has been a widely explored task that is of great importance in many applications. However, it is common that real-world time series data are recorded in a short time period, which results in a big gap between the deep model and the limited and noisy time series. In this work, we propose to address the time series forecasting problem with generative modeling and propose a bidirectional variational auto-encoder (BVAE) equipped with diffusion, denoise, and disentanglement, namely D3VAE. Specifically, a coupled diffusion probabilistic model is proposed to augment the time series data without increasing the aleatoric uncertainty and implement a more tractable inference process with BVAE. To ensure the generated series move toward the true target, we further propose to adapt and integrate the multiscale denoising score matching into the diffusion process for time series forecasting. In addition, to enhance the interpretability and stability of the prediction, we treat the latent variable in a multivariate manner and disentangle them on top of minimizing total correlation. Extensive experiments on both synthetic and real-world data show that D3VAE outperforms competitive algorithms with remarkable margins. Our implementation is available here.

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

Abi Jaber, Eduardo, Camille Illand, and Shaun Li. 2024. "Joint SPX & VIX calibration with Gaussian polynomial volatility models: Deep pricing with quantization hints." *Mathematical Finance*.

Bayraktar, Erhan, Donghan Kim, and Abhishek Tilva. 2024. "Arbitrage theory in a market of stochastic dimension." *Mathematical Finance*, 34(3): 847–895.

Bras, Pierre, and Gilles Pagès. 2023. "Langevin algorithms for markovian neural networks and deep stochastic control." 1–8, IEEE.

Campigli, Francesco, Giacomo Bormetti, and Fabrizio Lillo. 2022. "Measuring price impact and information content of trades in a time-varying setting." *arXiv preprint arXiv:2212.12687*.

Cho, Youngin, Daejin Kim, Dongmin Kim, Mohammad Azam Khan, and Jaegul Choo. 2022. "WaveBound: dynamic error bounds for stable time series forecasting." *Advances in Neural Information Processing Systems*, 35: 19579–19591.

Cloud, Alex, Albert Wang, and Wesley Kerr. 2022. "Anticipatory fictitious play." *arXiv preprint arXiv:2212.09941*.

Drissi, Fayçal. 2022. "Solvability of differential Riccati equations and applications to algorithmic trading with signals." *Applied Mathematical Finance*, 29(6): 457–493.

Drissi, Fayçal. 2023. "Models of market liquidity: Applications to traditional markets and automated market makers." *Available at SSRN 4424010*.

Dupire, Bruno, and Valentin Tissot-Daguette. 2022. "Functional expansions." *arXiv preprint arXiv:2212.13628*.

Filipović, Damir, and Puneet Pasricha. 2022. "Empirical Asset Pricing via Ensemble Gaussian Process Regression." *arXiv preprint arXiv:2212.01048*.

Gorduza, Dragos, Xiaowen Dong, and Stefan Zohren. 2022. "Understanding stock market instability via graph auto-encoders." *arXiv preprint arXiv:2212.04974*.

Kisiel, Damian, and Denise Gorse. 2022. "Axial-lob: High-frequency trading with axial attention." 1327–1333, IEEE.

Li, Yan, Xinjiang Lu, Yaqing Wang, and Dejing Dou. 2022. "Generative time series forecasting with diffusion, denoise, and disentanglement." *Advances in Neural Information Processing Systems*, 35: 23009–23022.

Lorig, Matthew, and Natchanon Suaysom. 2024. "Explicit caplet implied volatilities for quadratic term-structure models." *International Journal of Financial Engineering*, 11(01): 2350041.

Luxenberg, Eric, Philipp Schiele, and Stephen Boyd. 2024. "Robust Bond Portfolio Construction via Convex–Concave Saddle Point Optimization." *Journal of Optimization Theory and Applications*, 1–27.

Migrow, Dimitri. 2022. "Strategic observational learning." *arXiv preprint arXiv:2212.09889*.

Miori, Deborah, and Mihai Cucuringu. 2022. "DeFi: data-driven characterisation of Uniswap v3 ecosystem & an ideal crypto law for liquidity pools." *arXiv preprint arXiv:2301.13009*.

Pham, Huyên, and Xavier Warin. 2022. "Mean-field neural networks-based algorithms for McKean-Vlasov control problems." *arXiv preprint arXiv:2212.11518*.

Raducha, Tomasz, and Maxi San Miguel. 2023. "Evolutionary games on multilayer networks: coordination and equilibrium selection." *Scientific Reports*, 13(1): 11818.

Rosenbaum, Mathieu, and Jianfei Zhang. 2022. "Multi-asset market making under the quadratic rough Heston." *arXiv preprint arXiv:2212.10164*.