VOL. 2, NO. 3

OMI Research Newsletter – December 2022

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 four papers from members of the OMI. The first describes the predictable loss (PL) of liquidity providers in automated market makers and solves optimal liquidity provision in these venues. The second derives closed-form strategies for brokers interacting with informed traders (at a loss) and noise traders (at a profit). The third proposes a novel framework to decouple the learning of high-dimensional interactions from learning non-linear feature interactions in machine learning algorithms involving interactions between entities. The fourth forecasts realised covariance matrices by exploiting the predictive information of graphs that capture the interdependence in volatility and correlation.

Machine Learning for Finance

Graph-based Methods for Forecasting Realized Covariances

C. ZHANG, X. PU, M. CUCURINGU, AND X. DONG
UNIVERSITY OF OXFORD

We forecast the realized covariance matrix of asset returns in the U.S. equity market by exploiting the predictive information of graphs in volatility and correlation. Specifically, we augment the Heterogeneous Autoregressive (HAR) model via neighborhood aggregation on these graphs. Our proposed method allows for the modeling of interdependence in volatility (also known as spillover effect) and correlation, while maintaining parsimony and interpretability. We explore various graph construction methods, including sector membership and graphical LASSO (for modeling volatility), and line graph (for modeling correlation). The results generally suggest that the augmented model incorporating graph information yields both statistically and economically significant improvements for out-of-sample performance over the traditional models. Such improvements remain significant over horizons up to one month ahead, but decay in time. The robustness tests demonstrate that the forecast improvements are obtained consistently over the different out-of-sample sub-periods, and are insensitive to measurement errors of volatilities.

Symphony in the Latent Space: Provably Integrating High-dimensional Techniques with Non-linear Machine Learning Models

Q. WU, J. LI, Z. LIU, Y. LI, AND M. CUCURINGU
UNIVERSITY OF OXFORD, TSINGHUA UNIVERSITY, WORCESTER POLYTECHNIC INSTITUTE

This paper revisits building machine learning algorithms that involve interactions between entities, such as those between financial assets in an actively managed portfolio, or interactions between users in a social network. Our goal is to forecast the future evolution of ensembles of multivariate time series in such applications (e.g., the future return of a financial asset or the future popularity of a Twitter account). Designing ML algorithms for such systems requires addressing the challenges of high-dimensional interactions and non-linearity. Existing approaches usually adopt an ad-hoc approach to integrating high-dimensional techniques into non-linear models and recent studies have shown these approaches have questionable efficacy in time-evolving interacting systems. To this end, we propose a novel framework, which we dub as the additive influence model. Under our modeling assumption, we show that it is possible to decouple the learning of high-dimensional interactions from the learning of non-linear feature interactions. To learn the high-dimensional interactions,

we leverage kernel-based techniques, with provable guarantees, to embed the entities in a low-dimensional latent space. To learn the non-linear feature-response interactions, we generalize prominent machine learning techniques, including designing a new statistically sound non-parametric method and an ensemble learning algorithm optimized for vector regressions. Extensive experiments on two common applications demonstrate that our new algorithms deliver significantly stronger forecasting power compared to standard and recently proposed methods.

Optimal Estimation of Generic Dynamics by Path-dependent Neural Jump ODEs

F. Krach, M. Nübel, and J. Teichmann

ETH

This paper studies the problem of forecasting general stochastic processes using an extension of the Neural Jump ODE (NJ-ODE) framework. While NJ-ODE was the first framework to establish convergence guarantees for the prediction of irregularly observed time series, these results were limited to data stemming from Itô-diffusions with complete observations, in particular Markov processes where all coordinates are observed simultaneously. In this work, we generalise these results to generic, possibly non-Markovian or discontinuous, stochastic processes with incomplete observations, by utilising the reconstruction properties of the signature transform. These theoretical results are supported by empirical studies, where it is shown that the path-dependent NJ-ODE outperforms the original NJ-ODE framework in the case of non-Markovian data. Moreover, we show that PD-NJ-ODE can be applied successfully to limit order book (LOB) data.

FinRL-Meta: Market Environments and Benchmarks for Data-Driven Financial Reinforcement Learning

X. Liu, Z. Xia, J. Rui, J. Gao, H. Yang, M. Zhu, C. D. Wang, Z. Wang, J. Guo

Columbia; University of Hongkong, University of Virginia, New York University, Northwestern University

Finance is a particularly difficult playground for deep reinforcement learning. However, establishing high-quality market environments and benchmarks for financial reinforcement learning is challenging due to three major factors, namely, low signal-to-noise ratio of financial data, survivorship bias of historical data, and model overfitting in the backtesting stage. In this paper, we present an openly accessible FinRL-Meta library that has been actively maintained by the AI4Finance community. First, following a DataOps paradigm, we will provide hundreds of market environments through an automatic pipeline that collects dynamic datasets from real-world markets and processes them into gym-style market environments. Second, we reproduce popular papers as stepping stones for users to design new trading strategies. We also deploy the library on cloud platforms so that users can visualize their own results and assess the relative performance via community-wise competitions. Third, FinRL-Meta provides tens of Jupyter/Python demos organized into a curriculum and a documentation website to serve the rapidly growing community. FinRL-Meta is available here.

The Short-Term Predictability of Returns in Order Book Markets: a Deep Learning Perspective

L. Lucchese, M. Pakkanen, and A. Veraart

Imperial College London, University of Waterloo

In this paper, we conduct a systematic large-scale analysis of order book-driven predictability in high-frequency returns by leveraging deep learning techniques. First, we introduce a new and robust representation of the order book, the volume representation. Next, we carry out an extensive empirical experiment to address various questions regarding predictability. We investigate if and how far ahead there is predictability, the importance of a robust data representation, the advantages of multi-horizon modeling, and the presence of universal trading patterns. We use model confidence sets, which provide a formalized statistical inference framework particularly well suited to answer these questions. Our findings show that at high frequencies predictability in mid-price returns is not just present, but ubiquitous. The performance of the deep learning models is strongly dependent on the choice of order book representation, and in this respect, the volume representation appears to have multiple practical advantages.

Mathematical Finance

Brokers and Informed Traders: dealing with toxic flow and extracting trading signals

Á. Cartea and L. Sánchez-Betancourt

University of Oxford, King's College London

We derive closed-form strategies for a broker who provides liquidity to an informed trader and to a noise trader over a finite-time and infinite-time trading horizon. The flow of the noise trader is uninformative and the broker trades with the noise trader at a profit, on average. On the other hand, the informed trader has privileged information about the trend in the price of the asset, so the broker trades with the informed trader at a loss, on average. These losses are payment

for toxic flow from which the broker extracts the trend signal. The signal is one of the key ingredients in the broker's trading strategy to internalise (i.e., how much of the flow she keeps in her books), to externalise (i.e., how much she unwinds in a lit exchange), and to speculate in the lit market. The broker's dynamic strategy is a linear combination of four variables: the broker's inventory, the informed trader's inventory, the trend signal, and the uninformed trader's rate of trading with the broker. When the trading horizon is infinite (resp. finite), the coefficients of the four terms are constants (resp. deterministic functions of time). Finally, in the infinite horizon case, we show how the broker uses the flow of her clients to estimate the constant coefficients of the optimal strategy.

Decentralised Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision

Á. Cartea, F. Drissi, and M. Monga
University of Oxford, King's College London

We introduce a new comprehensive metric of predictable loss (PL) for liquidity providers in constant function automated market makers and derive an optimal liquidity provision strategy. PL compares the value of the LP's holdings in the liquidity pool (assuming no fee revenue) with that of a self-financing portfolio that replicates the LP's holdings and invests in a risk-free account. We provide closed-form formulae for PL, and show that the losses stem from two sources: the convexity cost, which depends on liquidity taking activity and the convexity of the pool's trading function; the opportunity cost, which is due to locking the LP's assets in the pool. For LPs in constant product market makers with concentrated liquidity, we derive a closed-form strategy that dynamically adjusts the range around the exchange rate as a function of market trend, volatility, and liquidity taking activity in the pool. We prove that the profitability of liquidity provision depends on the tradeoff between PL and fee income. Finally, we use Uniswap v3 data to show that LPs have traded at a significant loss, and to show that the out-of-sample performance of our strategy is considerably superior to the historical performance of LPs in the pool we consider.

Exploratory Control with Tsallis Entropy for Latent Factor Models

R. Donnelly and S. Jaimungal
King's College London, University of Toronto

We study optimal control in models with latent factors where the agent controls the distribution over actions, rather than actions themselves, in both discrete and continuous time. To encourage exploration of the state space, we reward exploration with Tsallis Entropy and derive the optimal distribution over states - which we prove is q-Gaussian distributed with location characterized through the solution of an FBSE and FBSDE in discrete and continuous time, respectively. We discuss the relation between the solutions of the optimal exploration problems and the standard dynamic optimal control solution. Finally, we develop the optimal policy in a model-agnostic setting along the lines of soft Q-learning. The approach may be applied in, e.g., developing more robust statistical arbitrage trading strategies.

Optimal Liquidation with Signals: the General Propagator Case

E. Abi Jaber and E. Neuman
Ecole Polytechnique, Imperial College London

We consider a class of optimal liquidation problems where the agent's transactions create transient price impact driven by a Volterra-type propagator along with temporary price impact. We formulate these problems as minimization of a revenue-risk functionals, where the agent also exploits available information on a progressively measurable price predicting signal. By using an infinite dimensional stochastic control approach, we characterize the value function in terms of a solution to a free-boundary L^2-valued backward stochastic differential equation and an operator-valued Riccati equation. We then derive analytic solutions to these equations which yields an explicit expression for the optimal trading strategy. We show that our formulas can be implemented in a straightforward and efficient way for a large class of price impact kernels with possible singularities such as the power-law kernel.

Dynamic Inventory Management with Mean-Field Competition

R. Donnelly and Z. Li
King's College London

Agents attempt to maximize expected profits earned by selling multiple units of a perishable product where their revenue streams are affected by the prices they quote as well as the distribution of other prices quoted in the market by other agents. We propose a model which captures this competitive effect and directly analyze the model in the mean-field limit as the number of agents is very large. We classify mean-field Nash equilibrium in terms of the solution to a Hamilton-Jacobi-Bellman equation and a consistency condition and use this to motivate an iterative numerical algorithm to compute equilibrium. Properties of the equilibrium pricing strategies and overall market dynamics are then investigated, in particular how they depend on the strength of the competitive interaction and the ability to oversell the product.

Stochastic Algorithms for Advanced Risk Budgeting

A. R. Cetingoz, J-D. Fermanian, and O. Guéant
Université Paris 1, ENSAE

Modern portfolio theory has provided for decades the main framework for optimizing portfolios. Because of its sensitivity to small changes in input parameters, especially expected returns, the mean-variance framework proposed by Markowitz (1952) has however been challenged by new construction methods that are purely based on risk. Among risk-based methods, the most popular ones are Minimum Variance, Maximum Diversification, and Risk Budgeting (especially Equal Risk Contribution) portfolios. Despite some drawbacks, Risk Budgeting is particularly attracting because of its versatility: based on Euler's homogeneous function theorem, it can indeed be used with a wide range of risk measures. This paper presents sound mathematical results regarding the existence and the uniqueness of Risk Budgeting portfolios for a very wide spectrum of risk measures and shows that, for many of them, computing the weights of Risk Budgeting portfolios only requires a standard stochastic algorithm.

Stressing Dynamic Loss Models

E. Kroell, S. M. Pesenti and S. Jaimungal
University of Toronto

Stress testing, and in particular, reverse stress testing, is a prominent exercise in risk management practice. Reverse stress testing, in contrast to (forward) stress testing, aims to find an alternative but plausible model such that under that alternative model, specific adverse stresses (i.e. constraints) are satisfied. Here, we propose a reverse stress testing framework for dynamic models. Specifically, we consider a compound Poisson process over a finite time horizon and stresses composed of expected values of functions applied to the process at the terminal time. We then define the stressed model as the probability measure under which the process satisfies the constraints and which minimizes the Kullback-Leibler divergence to the reference compound Poisson model. We solve this optimization problem, prove existence and uniqueness of the stressed probability measure, and provide a characterization of the Radon-Nikodym derivative from the reference model to the stressed model. We find that under the stressed measure, the intensity and the severity distribution of the process depend on time and the state space. We illustrate the dynamic stress testing by considering stresses on VaR and both VaR and CVaR jointly and provide illustrations of how the stochastic process is altered under these stresses. We generalize the framework to multivariate compound Poisson processes and stresses at times other than the terminal time. We illustrate the applicability of our framework by considering "what if" scenarios, where we answer the question: What is the severity of a stress on a portfolio component at an earlier time such that the aggregate portfolio exceeds a risk threshold at the terminal time? Moreover, for general constraints, we provide a simulation algorithm to simulate sample paths under the stressed measure.

A model-free approach to continuous-time finance

H. Chiu, R. Cont
Imperial College London, University of Oxford

We present a non-probabilistic, pathwise approach to continuous-time finance based on causal functional calculus. We introduce a definition of self-financing, free from any integration concept and show that the value of a self-financing portfolio is a pathwise integral (every self-financing strategy is a gradient) and that generic domain of functional calculus is inherently arbitrage-free. We then consider the problem of hedging a path-dependent payoff across a generic set of scenarios. We apply the transition principle of Isaacs in differential games and obtain a verification theorem for the optimal solution, which is characterised by a fully non-linear path-dependent equation. For the Asian option, we obtain explicit solution.

Stochastic Liquidity as a Proxy for Nonlinear Price Impact

J. Muhle-Karbe, Z. Wang, and K. Webster
Imperial College London

Optimal execution and trading algorithms rely on price impact models, like the propagator model, to quantify trading costs. Empirically, price impact is concave in trade sizes, leading to nonlinear models for which optimization problems are intractable and even qualitative properties such as price manipulation are poorly understood. However, we show that in the diffusion limit of small and frequent orders, the nonlinear model converges to a tractable linear model. In this high-frequency limit, a stochastic liquidity parameter approximates the original impact function's nonlinearity. We illustrate the approximation's practical performance using limit-order data.

Economics

On the Emergence of Cooperation in the Repeated Prisoner's Dilemma

M. Schaefer
Yale University

This article explores which parameters of the repeated Prisoner's Dilemma lead to cooperation. Using simula-

tions, I demonstrate that the potential function of the stochastic evolutionary dynamics of the Grim Trigger strategy is useful to predict cooperation between Q-learners. The frontier separating the parameter spaces that induce either cooperation or defection can be determined based on the kinetic energy exerted by the respective basins of attraction. When the incentive compatibility constraint of the Grim Trigger strategy is slack, a sudden increase in the observed cooperation rates occurs when the ratio of the kinetic energies approaches a critical value, which itself is a function of the discount factor, multiplied by a correction factor to account for the effect of the algorithms' exploration probability. Using metadata from laboratory experiments, I provide evidence that the insights obtained from the simulations are also useful to explain the emergence of cooperation between humans. The observed cooperation rates show a positive gradient at the frontier characterized by an exploration probability of approximately five percent. In the context of human-to-human interaction, the exploration probability can be viewed as the belief about the opponent's probability to deviate from the equilibrium action.

Zero-Sum Stochastic Stackelberg Games

D. Goktas, J. Zhao, and A. Greenwald
Brown University, Pomona College

Zero-sum stochastic games have found important applications in a variety of fields, from machine learning to economics. Work on this model has primarily focused on the computation of Nash equilibrium due to its effectiveness in solving adversarial board and video games. Unfortunately, a Nash equilibrium is not guaranteed to exist in zero-sum stochastic games when the payoffs at each state are not convex-concave in the players' actions. A Stackelberg equilibrium, however, is guaranteed to exist. Consequently, in this paper, we study zero-sum stochastic Stackelberg games. Going beyond known existence results for (non-stationary) Stackelberg equilibria, we prove the existence of recursive (i.e., Markov perfect) Stackelberg equilibria (recSE) in these games, provide necessary and sufficient conditions for a policy profile to be a recSE, and show that recSE can be computed in (weakly) polynomial time via value iteration. Finally, we show that zero-sum stochastic Stackelberg games can model the problem of pricing and allocating goods across agents and time. More specifically, we propose a zero-sum stochastic Stackelberg game whose recSE correspond to the recursive competitive equilibria of a large class of stochastic Fisher markets. We close with a series of experiments that showcase how our methodology can be used to solve the consumption-savings problem in stochastic Fisher markets.

Machine Learning

Non-stationary Transformers: Exploring the Stationarity in Time Series Forecasting

Y. Liu, H. Wu, J. Wang, and M. Long
Tsinghua University

Transformers have shown great power in time series forecasting due to their global-range modeling ability. However, their performance can degenerate terribly on non-stationary real-world data in which the joint distribution changes over time. Previous studies primarily adopt stationarization to attenuate the non- stationarity of original series for better predictability. But the stationarized series deprived of inherent non-stationarity can be less instructive for real-world bursty events forecasting. This problem, termed over-stationarization in this paper, leads Transformers to generate indistinguishable temporal attentions for different series and impedes the predictive capability of deep models. To tackle the dilemma between series predictability and model capability, we propose Non-stationary Transformers as a generic framework with two interdependent modules: Series Stationarization and Destationary Attention. Concretely, Series Stationarization unifies the statistics of each input and converts the output with restored statistics for better predictability. To address the over-stationarization problem, Destationary Attention is devised to recover the intrinsic non-stationary information into temporal dependencies by approximating distinguishable attentions learned from raw series. Our Non-stationary Transformers framework consistently boosts mainstream Transformers by a large margin, which reduces MSE by 49.43% on Transformer, 47.34% on Informer, and 46.89% on Reformer, making them the state-of-the-art in time series forecasting. Code is available at this repository.

Generating multivariate time series with COmmon Source CoordInated GAN (COSCI-GAN)

A. Seyfi, J-F Rajotte, and R. T. Ng
University of British Columbia

Generating multivariate time series is a promising approach for sharing sensitive data in many medical, financial, and IoT applications. A common type of multivariate time series originates from a single source such as the biometric measurements from a medical patient. This leads to complex dynamical patterns between individual time series that are hard to learn by typical generation models such as GANs. There is valuable information in those patterns that machine learning models can use to better classify, predict or perform other downstream tasks. We propose a novel framework that takes time series' common origin into account and favors channel/feature relationships preservation. The two key points of our method are: 1) the individual time series are generated from a common

point in latent space and 2) a central discriminator favors the preservation of inter-channel/feature dynamics. We demonstrate empirically that our method helps preserve channel/feature correlations and that our synthetic data performs very well in downstream tasks with medical and financial data.

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, and Eyal Neuman. 2022. "Optimal liquidation with signals: the general propagator case." *Available at SSRN 4264823*.

Cartea, Álvaro, Fayçal Drissi, and Marcello Monga. 2023. "Predictable losses of liquidity provision in constant function markets and concentrated liquidity markets." *Applied Mathematical Finance*, 30(2): 69–93.

Cartea, Álvaro, Fayçal Drissi, and Marcello Monga. 2024. "Decentralized Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision." *SIAM Journal on Financial Mathematics*, 15(3): 931–959.

Donnelly, Ryan, and Sebastian Jaimungal. 2024. "Exploratory control with Tsallis entropy for latent factor models." *SIAM Journal on Financial Mathematics*, 15(1): 26–53.

Donnelly, Ryan, and Zi Li. 2022. "Dynamic Inventory Management with Mean-Field Competition." *arXiv preprint arXiv:2210.17208*.

Goktas, Denizalp, Sadie Zhao, and Amy Greenwald. 2022. "Zero-sum stochastic stackelberg games." *Advances in Neural Information Processing Systems*, 35: 11658–11672.

Krach, Florian, Marc Nübel, and Josef Teichmann. 2022. "Optimal Estimation of Generic Dynamics by Path-Dependent Neural Jump ODEs." *arXiv preprint arXiv:2206.14284*.

Kroell, Emma, Silvana M Pesenti, and Sebastian Jaimungal. 2024. "Stressing dynamic loss models." *Insurance: Mathematics and Economics*, 114: 56–78.

Liu, Xiao-Yang, Ziyi Xia, Jingyang Rui, Jiechao Gao, Hongyang Yang, Ming Zhu, Christina Wang, Zhaoran Wang, and Jian Guo. 2022*a*. "FinRL-Meta: Market environments and benchmarks for data-driven financial reinforcement learning." *Advances in Neural Information Processing Systems*, 35: 1835–1849.

Liu, Yong, Haixu Wu, Jianmin Wang, and Mingsheng Long. 2022*b*. "Non-stationary transformers: Exploring the stationarity in time series forecasting." *Advances in Neural Information Processing Systems*, 35: 9881–9893.

Lucchese, Lorenzo, Mikko S Pakkanen, and Almut ED Veraart. 2024. "The short-term predictability of returns in order book markets: a deep learning perspective." *International Journal of Forecasting*.

Muhle-Karbe, Johannes, Zexin Wang, and Kevin Webster. 2024. "Stochastic liquidity as a proxy for nonlinear price impact." *Operations Research*, 72(2): 444–458.

Schaefer, Maximilian. 2022. "On the Emergence of Cooperation in the Repeated Prisoner's Dilemma." *arXiv preprint arXiv:2211.15331*.

Seyfi, Ali, Jean-Francois Rajotte, and Raymond Ng. 2022. "Generating multivariate time series with COmmon Source CoordInated GAN (COSCI-GAN)." *Advances in neural information processing systems*, 35: 32777–32788.

Wu, Qiong, Jian Li, Zhenming Liu, Yanhua Li, and Mihai Cucuringu. 2023. "Symphony in the latent space: Provably integrating high-dimensional techniques with non-linear machine learning models." Vol. 37, 10361–10369.

Zhang, Chao, Xingyue Pu, Mihai Cucuringu, and Xiaowen Dong. 2024. "Graph-based methods for forecasting realized covariances." *Journal of Financial Econometrics*, nbae026.