VOL. 2, NO. 10

OMI Research Newsletter – July 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, École Polytechnique, Imperial College London, University of Oxford, University of Edinburgh, and Queen Mary University. In this work, the authors address uncertainty in the data generating process with an adversarial approach to robustify deep hedging. Here, the authors study the cost of misspecifying the market impact model for the dynamics of the price in execution problems. The authors of this work propose a new method to study the generalisation error of algorithms for empirical risk minimisation. Finally, here the authors propose an efficient online inference algorithm for nonlinear functions using a non-stationary data stream.

Machine Learning & Statistics

Robust Hedging GANs

Y. LIMMER, B. HORVATH
OXFORD-MAN INSTITUTE

The availability of deep hedging has opened new horizons for solving hedging problems under a large variety of realistic market conditions. At the same time, any model – be it a traditional stochastic model or a market generator – is at best an approximation of market reality, prone to model-misspecification and estimation errors. This raises the question, how to furnish a modelling setup with tools that can address the risk of discrepancy between anticipated distribution and market reality, in an automated way. Automated robustification is currently attracting increased attention in numerous investment problems, but it is a delicate task due to its imminent implications on risk management. Hence, it is beyond doubt that more activity can be anticipated on this topic to converge towards a consensus on best practices.

This paper presents a natural extension of the original deep hedging framework to address uncertainty in the data generating process via an adversarial approach inspired by GANs to automate robustification in our hedging objective. This is achieved through an interplay of three modular components: (i) a (deep) hedging engine, (ii) a data-generating process (that is model agnostic permitting a large variety of classical models as well as machine learning-based market generators), and (iii) a notion of distance on model space to measure deviations between our market prognosis and reality. We do not restrict the ambiguity set to a region around a reference model, but instead penalise deviations from the anticipated distribution. Our suggested choice for each component is motivated by our aims for: Adaptability to generic path-dependent and exotic payoffs without major modifications of the setup, and applicability to highly realistic data structures and market environments, but other choices for each of the components are also possible. We demonstrate this in numerical experiments to benchmark our framework against other existing results. Since all individual components are already used in practice, we believe that our framework is easily adaptable to existing functional settings.

Lifelong Bandit optimisation: No Prior and No Regret

F. SCHUR, P. KASSRAIE, J. ROTHFUSS, A. KRAUSE
ETH ZÜRICH

Machine learning algorithms are often repeatedly applied to problems with similar structure over and over again. We focus on solving a sequence of bandit optimisation tasks and develop LIBO, an algorithm which adapts to

the environment by learning from past experience and becomes more sample-efficient in the process. We assume a kernelised structure where the kernel is unknown but shared across all tasks. LIBO sequentially meta-learns a kernel that approximates the true kernel and solves the incoming tasks with the latest kernel estimate. Our algorithm can be paired with any kernelised or linear bandit algorithm and guarantees oracle optimal performance, meaning that as more tasks are solved, the regret of LIBO on each task converges to the regret of the bandit algorithm with oracle knowledge of the true kernel. Naturally, if paired with a sublinear bandit algorithm, LIBO yields a sublinear lifelong regret. We also show that direct access to the data from each task is not necessary for attaining sublinear regret. We propose F-LIBO, which solves the lifelong problem in a federated manner.

Copula-Based Deep Survival Models for Dependent Censoring

A. H. G. Foomani, M. Cooper, R. Greiner, R. G. Krishnan

University of Alberta, University of Toronto

A survival dataset describes a set of instances (e.g., patients) and provides, for each, either the time until an event (e.g., death), or the censoring time (e.g., when lost to follow-up – which is a lower bound on the time until the event). We consider the challenge of survival prediction: learning, from such data, a predictive model that can produce an individual survival distribution for a novel instance. Many contemporary methods of survival prediction implicitly assume that the event and censoring distributions are independent conditional on the instance's covariates — a strong assumption that is difficult to verify (as we observe only one outcome for each instance) and which can induce significant bias when it does not hold. This paper presents a parametric model of survival that extends modern non-linear survival analysis by relaxing the assumption of conditional independence. On synthetic and semi-synthetic data, our approach significantly improves estimates of survival distributions compared to the standard that assumes conditional independence in the data.

Scaling MLPs: A Tale of Inductive Bias

G. Bachmann, S. Anagnostidis, T. Hofmann

ETH Zürich

In this work we revisit the most fundamental building block in deep learning, the multi-layer perceptron (MLP), and study the limits of its performance on vision tasks. Empirical insights into MLPs are important for multiple reasons. (1) Given the recent narrative "less inductive bias is better", popularised due to transformers eclipsing convolutional models, it is natural to explore the limits of this hypothesis. To that end, MLPs offer an ideal test bed, being completely free of any inductive bias. (2) MLPs have almost exclusively been the main protagonist in the deep learning theory literature due to their mathematical simplicity, serving as a proxy to explain empirical phenomena observed for more complex architectures. Surprisingly, experimental datapoints for MLPs are very difficult to find in the literature, especially when coupled with large pre-training protocols. This discrepancy between practice and theory is worrying: Do MLPs reflect the empirical advances exhibited by practical models? Or do theorists need to rethink the role of MLPs as a proxy? We provide insights into both these aspects. We show that the performance of MLPs drastically improves with scale (93% on CIFAR10, 79% on CIFAR100, 69% on TinyImageNet), highlighting that lack of inductive bias can indeed be compensated. We observe that MLPs mimic the behaviour of their modern counterparts faithfully, with some components in the learning setting however surprisingly exhibiting stronger or unexpected behaviours. Due to their inherent computational efficiency, large pre-training experiments become more accessible for academic researchers. All of our experiments were run on a single GPU.

K-Nearest-Neighbor Resampling for Off-Policy Evaluation in Stochastic Control

M. Giegricha, R. Oomen, C. Reisinger

University of Oxford, London School of Economics

In this paper, we propose a novel K-nearest neighbor resampling procedure for estimating the performance of a policy from historical data containing realised episodes of a decision process generated under a different policy. We focus on feedback policies that depend deterministically on the current state in environments with continuous state-action spaces and system-inherent stochasticity effected by chosen actions. Such settings are common in a wide range of high-stake applications in diverse fields such as finance, engineering and operations research and are actively investigated in the context of stochastic control. Assuming continuity in the rewards and the state transitions, our procedure exploits that similar state/action pairs (in a metric sense) are associated with similar rewards and state transitions. This enables our resampling procedure to tackle the counterfactual estimation problem underlying off-policy evaluation (OPE) by simulating trajectories similarly to Monte Carlo methods. Compared to other OPE methods, our algorithm does not require optimisation, can be efficiently implemented via tree-based nearest neighbor search and parallelisation and does not explicitly assume a parametric model for the environment's dynamics. These properties make the proposed resam-

pling algorithm particularly useful for stochastic control environments. We prove that our method is statistically consistent in estimating the performance of a policy in the OPE setting under weak assumptions and for data sets containing entire episodes rather than independent transitions. To establish the consistency, we generalise Stone's Theorem, a well-known result in nonparametric statistics on local averaging, to include episodic data and the counterfactual estimation underlying OPE. Numerical experiments demonstrate the effectiveness of the algorithm in a variety of stochastic control settings including a linear quadratic regulator, trade execution in limit order books and online stochastic bin packing. We show that our method has a vanishing mean-squared error for increasing amounts of data and is often superior to existing baselines.

Conditional Generators for Limit Order Book Environments: Explainability, Challenges, and Robustness

A. Coletta, J. Jerome, R. Savani, S. Vyetrenko
University of Liverpool, Alan Turing Institute

Limit order books are a fundamental and widespread market mechanism. This paper investigates the use of conditional generative models for order book simulation. For developing a trading agent, this approach has drawn recent attention as an alternative to traditional backtesting due to its ability to react to the presence of the trading agent. Using a state-of-the-art CGAN (from Coletta et al. (2022)), we explore its dependence upon input features, which highlights both strengths and weaknesses. To do this, we use "adversarial attacks" on the model's features and its mechanism. We then show how these insights can be used to improve the CGAN, both in terms of its realism and robustness. We finish by laying out a roadmap for future work.

Online Bootstrap Inference with Nonconvex Stochastic Gradient Descent Estimator

Y. Zhong, T. Kuffner, S. Lahiri
Washington University,

In this paper, we investigate the theoretical properties of stochastic gradient descent (SGD) for statistical inference in the context of nonconvex optimisation problems, which have been relatively unexplored compared to convex settings. Our study is the first to establish provable inferential procedures using the SGD estimator for general nonconvex objective functions, which may contain multiple local minima. We propose two novel online inferential procedures that combine SGD and the multiplier bootstrap technique. The first procedure employs a consistent covariance matrix estimator, and we establish its error convergence rate. The second procedure approximates the limit distribution using bootstrap SGD estimators, yielding asymptotically valid bootstrap confidence intervals. We validate the effectiveness of both approaches through numerical experiments. Furthermore, our analysis yields an intermediate result: the in-expectation error convergence rate for the original SGD estimator in nonconvex settings, which is comparable to existing results for convex problems. We believe this novel finding holds independent interest and enriches the literature on optimisation and statistical inference.

Agent market orders representation through a contrastive learning approach

R. Ruan, E. Bacry, J-F. Muzy
Université Paris-Dauphine PSL, Université de Corse

Due to the access to the labeled orders on the CAC40 data from Euronext, we are able to analyze agents' behaviors in the market based on their placed orders. In this study, we construct a self-supervised learning model using triplet loss to effectively learn the representation of agent market orders. By acquiring this learned representation, various downstream tasks become feasible. In this work, we utilize the K-means clustering algorithm on the learned representation vectors of agent orders to identify distinct behavior types within each cluster.

Low-rank Extended Kalman Filtering for Online Learning of Neural Networks from Streaming Data

P. G. Chang, G. Durán-Martín, A. Shestopaloff, M. Jones, K. Murphy
University of Chicago, Queen Mary University, University of Colorado

We propose an efficient online approximate Bayesian inference algorithm for estimating the parameters of a nonlinear function from a potentially non-stationary data stream. The method is based on the extended Kalman filter (EKF), but uses a novel low-rank plus diagonal decomposition of the posterior precision matrix, which gives a cost per step which is linear in the number of model parameters. In contrast to methods based on stochastic variational inference, our method is fully deterministic, and does not require step-size tuning. We show experimentally that this results in much faster (more sample efficient) learning, which results in more rapid adaptation to changing distributions, and faster accumulation of reward when used as part of a contextual bandit algorithm.

Machine Learning architectures for price formation models with common noise

D. Gomes, J. Gutierrez, and M. Laurière
King Abdullah University, NYU Shanghai

We propose a machine learning method to solve a mean-field game price formation model with common noise. This involves determining the price of a commodity traded among rational agents subject to a market clearing condition imposed by random supply, which presents additional challenges compared to the deterministic counterpart. Our approach uses a dual recurrent neural network architecture encoding noise dependence and a particle approximation of the mean-field model with a single loss function optimised by adversarial training. We provide a posteriori estimates for convergence and illustrate our method through numerical experiments.

Wasserstein distributional robustness of neural networks

X. Bai, G. He, Y. Jiang, J. Obłój
University of Oxford

Deep neural networks are known to be vulnerable to adversarial attacks (AA). For an image recognition task, this means that a small perturbation of the original can result in the image being misclassified. Design of such attacks as well as methods of adversarial training against them are subject of intense research. We re-cast the problem using techniques of Wasserstein distributionally robust optimisation (DRO) and obtain novel contributions leveraging recent insights from DRO sensitivity analysis. We consider a set of distributional threat models. Unlike the traditional pointwise attacks, which assume a uniform bound on perturbation of each input data point, distributional threat models allow attackers to perturb inputs in a non-uniform way. We link these more general attacks with questions of out-of-sample performance and Knightian uncertainty. To evaluate the distributional robustness of neural networks, we propose a first-order AA algorithm and its multi-step version. Our attack algorithms include Fast Gradient Sign Method (FGSM) and Projected Gradient Descent (PGD) as special cases. Furthermore, we provide a new asymptotic estimate of the adversarial accuracy against distributional threat models. The bound is fast to compute and first-order accurate, offering new in- sights even for the pointwise AA. It also naturally yields out-of-sample performance guarantees. We conduct numerical experiments on the CIFAR-10 dataset using DNNs on RobustBench to illustrate our theoretical results. Our code is available at https://github.com/JanObloj/W-DRO-Adversarial-Methods

Deep Learning for Population-Dependent Controls in Mean Field Control Problems

G. Dayanikli, M. Laurière, J. Zhang
Columbia University, University of California

In this paper, we propose several approaches to learn optimal population-dependent controls, in order to solve mean field control problems (MFC). Such policies enable us to solve MFC problems with generic common noise. We analyze the convergence of the proposed approximation algorithms, particularly the N-particle approximation. The effectiveness of our algorithms is supported by three different experiments, including systemic risk, price impact and crowd motion. We first show that our algorithms converge to the correct solution in an explicitly solvable MFC problem. Then, we conclude by showing that population-dependent controls outperform state-dependent controls. Along the way, we show that specific neural network architectures can improve the learning further.

Mathematical Finance

Equilibrium in Functional Stochastic Games with Mean-Field Interaction

E. Abi Jaber, E. Neuman, M. Voÿ
École Polytechnique, Imperial College London, University of California Los Angeles

We consider a general class of finite-player stochastic games with mean-field interaction, in which the linear-quadratic cost functional includes linear operators acting on controls in L^2. We propose a novel approach for deriving the Nash equilibrium of the game explicitly in terms of operator resolvents, by reducing the associated first order conditions to a system of stochastic Fredholm equations of the second kind and deriving their closed form solution. Furthermore, by proving stability results for the system of stochastic Fredholm equations we derive the convergence of the equilibrium of the N-player game to the corresponding mean-field equilibrium. As a by-product we also derive an ϵ-Nash equilibrium for the mean-field game, which is valuable in this setting as we show that the conditions for existence of an equilibrium in the mean-field limit are less restrictive than in the finite-player game. Finally we apply our general framework to solve various examples, such as stochastic Volterra linear-quadratic games, models of systemic risk and advertising with delay, and optimal liquidation games with transient price impact.

The Cost of Misspecifying Price Impact

N. Hey, J-P. Bouchaud, I. Mastromatteo, J. Muhle-Karbe, Kevin Webste
École Polytechnique, Imperial College London

Portfolio managers' orders trade off return and trading cost predictions. Return predictions rely on alpha models, whereas price impact models quantify trading costs. This paper studies what happens when trades are based on an incorrect price impact model, so that the portfolio either over- or under-trades its alpha signal. We derive tractable formulas for these misspecification costs and illustrate them on proprietary trading data. The misspecification costs are naturally asymmetric: underestimating impact concavity or impact decay shrinks profits, but overestimating concavity or impact decay can even turn profits into losses.

Randomness and early termination: what makes a game exciting?

G. Guo, D. Possamaï, C. Reisinger
Université Paris-Saclay, ETH Zürich,

In this short paper, we revisit an open problem posed by Aldous on the max-entropy win-probability martingale: given two players of equal strength, such that the win-probability is a martingale diffusion, which of these processes has maximum entropy and hence gives the most excitement for the spectators? From a stochastic control perspective, the corresponding value function can be characterised by a nonlinear parabolic PDE, for which we show existence and uniqueness of classical solutions. We establish key qualitative properties of the solution including concavity, monotonicity, and convergence to a steady state. Moreover, we construct convergent numerical approximations, which allow us to highlight the behaviour of the win-probability process in the present case where the match may end early, in contrast to recent work by Backhoff-Veraguas and Beiglböck where the match always runs the full length.

The Uniform Diversification Strategy is Optimal for Expected Utility Maximisation Under High Model Ambiguity

L. Carassus, J. Wiesel
Columbia University, ESILV

We investigate an expected utility maximisation problem under model uncertainty in a one-period financial market. We formalize the model uncertainty by taking a Wasserstein ball of radius k around P and consider the corresponding Wasserstein distributionally robust optimisation problem. We show that optimal solutions converge to the uniform diversification strategy when uncertainty is increasingly large, i.e. when the radius k tends to infinity.

Economics

Mean-field Analysis of generalisation Errors

G. Aminian, S. N. Cohen, Ł. Szpruch
University of Oxford, University of Edinburgh

We propose a novel framework for exploring weak and L2 generalisation errors of algorithms through the lens of differential calculus on the space of probability measures. Specifically, we consider the KL-regularised empirical risk minimisation problem and establish generic conditions under which the generalisation error convergence rate, when training on a sample of size n, is $\mathcal{O}(1/n)$. In the context of supervised learning with a one-hidden layer neural network in the mean-field regime, these conditions are reflected in suitable integrability and regularity assumptions on the loss and activation functions.

Opinion dynamics in communities with major influencers and implicit social influence via mean-field approximation

D. Coculescu, M. Motte, H. Pham
University of Zurich, Université Paris Cité

We study binary opinion formation in a large population where individuals are influenced by the opinions of other individuals. The population is characterised by the existence of (i) communities where individuals share some similar features, (ii) opinion leaders that may trigger unpredictable opinion shifts in the short term (iii) some degree of incomplete information in the observation of the individual or public opinion processes. In this setting, we study three different approximate mechanisms: common sampling approximation, independent sampling approximation, and, what will be our main focus in this paper, McKean-Vlasov (or mean-field) approximation. We show that all three approximations perform well in terms of different metrics that we introduce for measuring population level and individual level errors. In the presence of a common noise represented by the major influencers opinions processes, and despite the absence of idiosyncratic noises, we derive a propagation of chaos type result. For the particular case of a linear model and particular specifications of the major influencers opinion dynamics, we provide additional analysis, including long term behavior and fluctuations of the public opinion. The theoretical results are complemented by some concrete examples and numerical analysis, illustrating the formation of echo-chambers, the propagation of chaos, and phenomena such as snowball effect and social inertia.

Endogenous Barriers to Learning

O. Compte
Paris School of Economics

Motivated by the idea that lack of experience is a source of errors but that experience should reduce them, we model agents' behavior using a stochastic choice model, leaving endogenous the accuracy of their choice. In some games, increased accuracy is conducive to unstable best-response dynamics. We define the barrier to learning as the minimum level of noise which keeps the best-response dynamic stable. Using logit Quantal Response, this defines a limitQR Equilibrium. We apply the concept to centipede, travelers' dilemma, and 11-20 money-request games and to first-price and all-pay auctions, and discuss the role of strategy restrictions in reducing or amplifying barriers to learning.

Adaptive parsimony as an evolutionary solution to the equilibrium selection problem

J-B. André
Institut Jean Nicod

Many games, especially repeated games, have multiple Nash equilibria, which limits the predictive power of game theory for understanding animal behavior. In this article, I propose a solution to this problem inspired by the notion of stability by convergence from adaptive dynamics. The multiplicity of equilibria is due to the possibility of strategies that are arbitrary in the sense that they are individually adaptive only because others use them. While arbitrariness is a possibility in game theory, since it is always possible to design patterns of behavior as complex as one wishes, it cannot be gradually shaped by biological evolution. Evolution can stay in an arbitrary equilibrium if it starts there, but it cannot converge to an arbitrary equilibrium if it starts from a different initial state. I propose an equilibrium refinement, the concept of evolutionarily parsimonious equilibrium, that captures this convergence constraint by relying on the premise that a complex strategy cannot evolve as a by-product of selection for a simple strategy. Using examples, I show that this refinement supports the selection of biologically reasonable equilibria in many of the most important games in the literature. In particular, it eliminates the vast majority of equilibria in repeated games and provides a better understanding of the conditions necessary for the evolution of reciprocal cooperation.

Acknowledgement

The OMI Research Newsletter was curated by Álvaro Cartea, Director of the Oxford-Man Institute, and Fayçal Drissi, postdoctoral researcher at the institute. Paper suggestions are from fellow Ph.D. students at the institute: Álvaro Arroyo, Patrick Chang, Gerardo Durán-Martín, and Marcello Monga.

References

Abi Jaber, Eduardo, Eyal Neuman, and Moritz Voss. 2023. "Equilibrium in Functional Stochastic Games with Mean-Field Interaction." *Available at SSRN*.

Aminian, Gholamali, Samuel N Cohen, and Łukasz Szpruch. 2023. "Mean-field Analysis of Generalization Errors." *arXiv preprint arXiv:2306.11623*.

André, Jean-Baptiste. 2023. "Adaptive parsimony as an evolutionary solution to the equilibrium selection problem."

Bachmann, Gregor, Sotiris Anagnostidis, and Thomas Hofmann. 2023. "Scaling MLPs: A Tale of Inductive Bias." *arXiv preprint arXiv:2306.13575*.

Bai, Xingjian, Guangyi He, Yifan Jiang, and Jan Obloj. 2023. "Wasserstein distributional robustness of neural networks." *arXiv preprint arXiv:2306.09844*.

Carassus, Laurence, and Johannes Wiesel. 2023. "The uniform diversification strategy is optimal for expected utility maximization under high model ambiguity." *arXiv preprint arXiv:2306.01503*.

Chang, Peter, Gerardo Duràn-Martín, Alexander Y Shestopaloff, Matt Jones, and Kevin Murphy. 2023. "Low-rank extended Kalman filtering for online learning of neural networks from streaming data." *arXiv preprint arXiv:2305.19535*.

Coculescu, Delia, Médéric Motte, and Huyên Pham. 2023. "Opinion dynamics in communities with major influencers and implicit social influence via mean-field approximation." *arXiv preprint arXiv:2306.16553*.

Coletta, Andrea, Joseph Jerome, Rahul Savani, and Svitlana Vyetrenko. 2023. "Conditional Generators for Limit Order Book Environments: Explainability, Challenges, and Robustness." *arXiv preprint arXiv:2306.12806*.

Compte, Olivier. 2023. "Endogenous Barriers to Learning." *arXiv preprint arXiv:2306.16904*.

Dayanikli, Gokce, Mathieu Lauriere, and Jiacheng Zhang. 2023. "Deep Learning for Population-Dependent Controls in Mean Field Control Problems." *arXiv preprint arXiv:2306.04788*.

Foomani, Ali Hossein Gharari, Michael Cooper, Russell Greiner, and Rahul G Krishnan. 2023. "Copula-Based Deep Survival Models for Dependent Censoring." *arXiv preprint arXiv:2306.11912*.

Giegrich, Michael, Roel Oomen, and Christoph Reisinger. 2023. "K-Nearest-Neighbor Resampling for Off-Policy Evaluation in Stochastic Control." *arXiv preprint arXiv:2306.04836*.

Gomes, Diogo, Julian Gutierrez, and Mathieu Laurière. 2023. "Machine Learning architectures for price formation models with common noise." *arXiv preprint arXiv:2305.17618*.

Guo, Gaoyue, Dylan Possamaï, and Christoph Reisinger. 2023. "Randomness and early termination: what makes a game exciting?" *arXiv preprint arXiv:2306.07133*.

Hey, Natascha, Jean-Philippe Bouchaud, Iacopo Mastromatteo, Johannes Muhle-Karbe, and Kevin Webster. 2023. "The Cost of Misspecifying Price Impact." *arXiv preprint arXiv:2306.00599*.

Limmer, Yannick, and Blanka Horvath. 2023. "Robust Hedging GANs." *Available at SSRN 4489029*.

Ruan, Ruihua, Emmanuel Bacry, and Jean-François Muzy. 2023. "Agent market orders representation through a contrastive learning approach." *arXiv preprint arXiv:2306.05987*.

Schur, Felix, Parnian Kassraie, Jonas Rothfuss, and Andreas Krause. 2022. "Lifelong Bandit Optimization: No Prior and No Regret." *arXiv preprint arXiv:2210.15513*.

Zhong, Yanjie, Todd Kuffner, and Soumendra Lahiri. 2023. "Online Bootstrap Inference with Nonconvex Stochastic Gradient Descent Estimator." *arXiv preprint arXiv:2306.02205*.