OMI Research Newsletter – Mar 2023



The Oxford-Man Institute of Quantitative Finance (OMI) is a world-leading centre for interdisciplinary research in financial markets. The OMI addresses fundamental problems in quantitative finance with a strong focus on machine learning and data driven models.

*The **OMI Research Newsletter** is a curated summary that covers the most recent research, insights, and tools related to financial decision making.*

Highlights

Noteworthy contributions in this month's edition include works from the OMI, University of Oxford, TU Berlin, and Université Paris Cité. In this work, the authors develop new policy gradient methods for stochastic control and show that it outperforms PDE or neural networks techniques in a model-based setting. Here, the authors propose a method to construct dynamic networks of stocks using spectral clustering. Here, the authors obtain closed-form solutions for optimal investments when an agent is dynamically informed about the future evolution of one of the independent Brownian motions driving a stock's price fluctuations. Finally, in this work, the authors present a framework to model LOBs which captures the dynamics of order flow and market clearing.

Machine Learning

Policy gradient learning methods for stochastic control with exit time and applications to share repurchase pricing

M. HAMDOUCHE, P. HENRY-LABORDERE, AND H. PHAM

UNIVERSITÉ PARIS CITÉ

We develop policy gradients methods for stochastic control with exit time in a model-free setting. We propose two types of algorithms for learning either directly the optimal policy or by learning alternately the value function (critic) and the optimal control (actor). The use of randomized policies is crucial for overcoming notably the issue related to the exit time in the gradient computation. We demonstrate the effectiveness of our approach by implementing our numerical schemes in the application to the problem of share repurchase pricing. Our results show that the proposed policy gradient methods outperform PDE or other neural networks techniques in a model-based setting. Furthermore, our algorithms are flexible enough to incorporate realistic market conditions like e.g. price impact or transaction costs.

Co-trading networks for modeling dynamic interdependency structures and estimating high-dimensional covariances in US equity markets

Y. LU, G. REINERT, AND M. CUCURINGU

OXFORD-MAN INSTITUTE, UNIVERSITY OF OXFORD

The time proximity of trades across stocks reveals interesting topological structures of the equity market in the United States. In this article, we investigate how such concurrent cross-stock trading behaviors, which we denote as co-trading, shape the market structures and affect stock price co-movements. By leveraging a co-trading-based pairwise similarity measure, we propose a novel method to construct dynamic networks of stocks. Our empirical studies employ high-frequency limit order book data from 2017-01-03 to 2019-12-09. By applying spectral clustering on co-trading networks, we uncover economically meaningful clusters of stocks. Beyond the static Global Industry Classification Standard (GICS) sectors, our data-driven clusters capture the time evolution of the dependency among stocks. Furthermore, we demonstrate statistically significant positive relations between low-latency co-trading and return covariance. With the aid of co-trading networks, we develop a robust estimator for high-dimensional covariance matrix, which yields superior economic value on portfolio allocation. The mean-variance portfolios based on our covariance estimates achieve both lower volatility and higher Sharpe ratios than standard benchmarks.

Spatio-Temporal Momentum: Jointly Learning Time-Series and Cross-Sectional Strategies

W. Ling Tan, S. Roberts, and S. Zohren
Oxford-Man Institute, University of Oxford

We introduce Spatio-Temporal Momentum strategies, a class of models that unify both time-series and cross-sectional momentum strategies by trading assets based on their cross-sectional momentum features over time. While both time-series and cross-sectional momentum strategies are designed to systematically capture momentum risk premia, these strategies are regarded as distinct implementations and do not consider the concurrent relationship and predictability between temporal and cross-sectional momentum features of different assets. We model spatio-temporal momentum with neural networks of varying complexities and demonstrate that a simple neural network with only a single fully connected layer learns to simultaneously generate trading signals for all assets in a portfolio by incorporating both their time-series and cross-sectional momentum features. Backtesting on portfolios of 46 actively-traded US equities and 12 equity index futures contracts, we demonstrate that the model is able to retain its performance over benchmarks in the presence of high transaction costs of up to 5-10 basis points. In particular, we find that the model when coupled with least absolute shrinkage and turnover regularization results in the best performance over various transaction cost scenarios.

SurvivalGAN: Generating Time-to-Event Data for Survival Analysis

A. Norcliffe, B. Cebere, F. Imrie, P. Liò, and M. van der Schaar
University of Cambridge

Synthetic data is becoming an increasingly promising technology, and successful applications can improve privacy, fairness, and data democratization. While there are many methods for generating synthetic tabular data, the task remains non-trivial and unexplored for specific scenarios. One such scenario is survival data. Here, the key difficulty is censoring: for some instances, we are not aware of the time of event, or if one even occurred. Imbalances in censoring and time horizons cause generative models to experience three new failure modes specific to survival analysis: (1) generating too few at-risk members; (2) generating too many at-risk members; and (3) censoring too early. We formalize these failure modes and provide three new generative metrics to quantify them. Following this, we propose SurvivalGAN, a generative model that handles survival data firstly by addressing the imbalance in the censoring and event horizons, and secondly by using a dedicated mechanism for approximating time-to-event/censoring. We evaluate this method via extensive experiments on medical datasets. SurvivalGAN outperforms multiple baselines at generating survival data, and in particular addresses the failure modes as measured by the new metrics, in addition to improving downstream performance of survival models trained on the synthetic data.

Exploring the Advantages of Transformers for High-Frequency Trading

F. Barez, P. Bilokon, A. Gervais, and N. Lisitsyn
University of Edinburgh, Imperial College London

This paper explores the novel deep learning Transformers architectures for high-frequency Bitcoin-USDT log-return forecasting and compares them to the traditional Long Short-Term Memory models. A hybrid Transformer model, called **HFformer**, is then introduced for time series forecasting which incorporates a Transformer encoder, linear decoder, spiking activations, and quantile loss function, and does not use position encoding. Furthermore, possible high-frequency trading strategies for use with the HFformer model are discussed, including trade sizing, trading signal aggregation, and minimal trading threshold. Ultimately, the performance of the HFformer and Long Short-Term Memory models are assessed and results indicate that the HFformer achieves a higher cumulative PnL than the LSTM when trading with multiple signals during backtesting.

View fusion vis-à-vis a Bayesian interpretation of Black-Litterman for portfolio allocation

T. Spears, S. Zohren, and S. Roberts
Oxford Man Institute

The Black-Litterman model extends the framework of the Markowitz Modern Portfolio Theory to incorporate investor views. We consider a case where multiple view estimates, including uncertainties, are given for the same underlying subset of assets at a point in time. This motivates our consideration of data fusion techniques for combining information from multiple sources. In particular, we consider consistency-based methods that yield fused view and uncertainty pairs; such methods are not common to the quantitative finance literature. We show a relevant, modern case of incorporating machine learning model-derived view and uncertainty estimates, and the impact on portfolio allocation, with an example subsuming Arbitrage Pricing Theory. Hence we show the value of the Black-Litterman model in combination with information fusion and artificial intelligence-grounded prediction methods.

Mathematical Finance

Optimal investment with a noisy signal of future stock prices

P. Bank and Yan Dolinsky
TU Berlin, Hebrew University

We consider an investor who is dynamically informed about the future evolution of one of the independent Brownian motions driving a stock's price fluctuations. With linear temporary price impact the resulting optimal investment problem with exponential utility turns out to be not only well posed, but it even allows for a closed-form solution. We describe this solution and the resulting problem value for this stochastic control problem with partial observation by solving its convex-analytic dual problem.

A mathematical framework for modelling order book dynamics

R. Cont, P. Degond, and L. Xuan
University of Oxford

We present a general framework for modelling the dynamics of limit order books, built on the combination of two modelling ingredients: the order flow, modelled as a general spatial point process, and market clearing, modelled via a deterministic mass transport operator acting on distributions of buy and sell orders. At the mathematical level, this corresponds to a natural decomposition of the infinitesimal generator describing the evolution of the limit order book into two operators: the generator of the order flow and the clearing operator. Our model provides a flexible framework for modelling and simulating order book dynamics and studying various scaling limits of discrete order book models. We show that our framework includes previous models as special cases and yields insights into the interplay between order flow and price dynamics.

Joint calibration to SPX and VIX options with signature-based models

C. Cuchiero, G. Gazzani, J. Möller, and S. Svaluto-Ferro
Vienna University, University of Verona

We consider a stochastic volatility model where the dynamics of the volatility are described by linear functions of the (time extended) signature of a primary underlying process, which is supposed to be some multidimensional continuous semimartingale. Under the additional assumption that this primary process is of polynomial type, we obtain closed form expressions for the VIX squared, exploiting the fact that the truncated signature of a polynomial process is again a polynomial process. Adding to such a primary process the Brownian motion driving the stock price, allows then to express both the log-price and the VIX squared as linear functions of the signature of the corresponding augmented process. This feature can then be efficiently used for pricing and calibration purposes. Indeed, as the signature samples can be easily precomputed, the calibration task can be split into an offline sampling and a standard optimization. For both the SPX and VIX options we obtain highly accurate calibration results, showing that this model class allows to solve the joint calibration problem without adding jumps or rough volatility.

Convergence of the Euler–Maruyama particle scheme for a regularised McKean–Vlasov equation arising from the calibration of local-stochastic volatility models

C. Reisinger and M. O. Tsianni
University of Oxford

In this paper, we study the Euler–Maruyama scheme for a particle method to approximate the McKean–Vlasov dynamics of calibrated local-stochastic volatility (LSV) models. Given the open question of well-posedness of the original problem, we work with regularised coefficients and prove that under certain assumptions on the inputs, the regularised model is well-posed. Using this result, we prove the strong convergence of the Euler–Maruyama scheme to the particle system with rate 1/2 in the step-size and obtain an explicit dependence of the error on the regularisation parameters. Finally, we implement the particle method for the calibration of a Heston-type LSV model to illustrate the convergence in practice and to investigate how the choice of regularisation parameters affects the accuracy of the calibration.

Signature SDEs from an affine and polynomial perspective

C. Cuchiero, S. Svaluto-Ferro, and J. Teichmann
ETH Zurich, Vienna University, University of Verona

Signature stochastic differential equations (SDEs) constitute a large class of stochastic processes, here driven by Brownian motions, whose characteristics are entire or real-analytic functions of their own signature, i.e. of iterated integrals of the process with itself, and allow therefore for a generic path dependence. We show that their prolongation with the corresponding signature is an affine and polynomial process taking values in subsets of group-like elements of the extended tensor algebra. By relying on the duality theory for affine and polynomial processes we obtain explicit formulas in terms of novel and proper notions of converging power series for the Fourier-Laplace transform and the expected value of

entire functions of the signature process. The coefficients of these power series are solutions of extended tensor algebra valued Riccati and linear ordinary differential equations (ODEs), respectively, whose vector fields can be expressed in terms of the entire characteristics of the corresponding SDEs. In other words, we construct a class of stochastic processes, which is universal within Itô processes with path-dependent characteristics and which allows for a relatively explicit characterization of the Fourier-Laplace transform and hence the full law on path space. We also analyze the special case of one-dimensional signature SDEs, which correspond to classical SDEs with real-analytic characteristics. Finally, the practical feasibility of this affine and polynomial approach is illustrated by several numerical examples.

A parsimonious inverse Cox-Ingersoll-Ross process for financial price modeling

L. Lin, and D. Sornette

East China University of Technology and Science, University of Geneva

We propose a formulation to construct new classes of financial price processes based on the insight that the key variable driving prices P is the earning-over-price ratio $\gamma \approx 1/P$, which we refer to as the earning yield and is analogous to the yield-to-maturity of an equivalent perpetual bond. This modeling strategy is illustrated with the choice for real-time γ in the form of the Cox-Ingersoll-Ross (CIR) process, which allows us to derive analytically many stylised facts of financial prices and returns, such as the power law distribution of returns, transient super-exponential bubble behavior, and the fat-tailed distribution of prices before bubbles burst. Our model sheds new light on rationalizing the excess volatility and the equity premium puzzles. The model is calibrated to five well-known historical bubbles in the US and China stock markets via a quasi-maximum likelihood method with the L-BFGS-B optimization algorithm. Using ϕ-divergence statistics adapted to models prescribed in terms of stochastic differential equations, we show the superiority of the CIR process for γt against three alternative models.

Economics

A rational measure of irrationality

D. Carpentiere, A. Giarlotta, and S. Watson

York University, University of Catania

All possible types of deterministic choice behavior are classified by their degree of irrationality. This classification is performed in three steps: (1) select a benchmark of rationality, for which this degree is zero; (2) endow the set of choices with a metric to measure deviations from rationality; and (3) compute the distance of any choice behavior from the selected benchmark. The natural candidate for step 1 is the family of all rationalizable behaviors. A possible candidate for step 2 is the metric described by Klamler (2008), who incorrectly claims this is the only one satisfying five intuitive properties. While proving a correct characterization of this metric, we determine the causes of its low discriminating power, and design a high-discerning variation of it. In step 3 we use this new metric to establish the minimum distance of any choice behavior from the benchmark of rationality. We conclude by describing a measure of stochastic irrationality, which employs the random utility model as a benchmark of rationality, and the Block-Marschak polynomials to measure deviations from it.

Ruin Probabilities for Risk Processes in Stochastic Networks

H. Amini, Z. Cao, A. Minca, and A. Sulem

University of Florida, INRIA Paris, Cornell University

We study multidimensional Cramér-Lundberg risk processes where agents, located on a large sparse network, receive losses form their neighbors. To reduce the dimensionality of the problem, we introduce classification of agents according to an arbitrary countable set of types. The ruin of any agent triggers losses for all of its neighbours. We consider the case when the loss arrival process induced by the ensemble of ruined agents follows a Poisson process with general intensity function that scales with the network size. When the size of the network goes to infinity, we provide explicit ruin probabilities at the end of the loss propagation process for agents of any type. These limiting probabilities depend, in addition to the agents' types and the network structure, on the loss distribution and the loss arrival process. For a more complex risk processes on open networks, when in addition to the internal networked risk processes the agents receive losses from external users, we provide bounds on ruin probabilities.

Coordination via Selling Information

A. Bonatti, M. Dahleh, T. Horel, and A. Nouripour

Massachusetts Institute of Technology

We consider games of incomplete information in which the players' payoffs depend both on a privately observed type and an unknown but common "state of nature". External to the game, a data provider knows the state of nature and sells information to the players, thus solving a joint information and mechanism design problem: deciding which information to sell while eliciting the player' types and collecting payments. We restrict ourselves to a general class of symmetric games with quadratic payoffs that includes games of both strategic substitutes (e.g.

Cournot competition) and strategic complements (e.g. Bertrand competition, Keynesian beauty contest). By to the Revelation Principle, the sellers' problem reduces to designing a mechanism that truthfully elicits the player' types and sends action recommendations that constitute a Bayes Correlated Equilibrium of the game. We fully characterize the class of all such Gaussian mechanisms (where the joint distribution of actions and private signals is a multivariate normal distribution) as well as the welfare- and revenue- optimal mechanisms within this class. For games of strategic complements, the optimal mechanisms maximally correlate the players' actions, and conversely maximally anticorrelate them for games of strategic substitutes. In both cases, for sufficiently large uncertainty over the players' types, the recommendations are deterministic (and linear) conditional on the state and the type reports, but they are not fully revealing.

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, and Marcello Monga.