

Machine Learning in Finance

Themed seminar series, 22/23 Autumn

Wednesday, 12 October hour via [Zoom](#)

Speaker Stéphane Crépey, Department of Mathematics, Université Paris Cité (Paris 7)/LPSM
Stéphane Crépey graduated from ENSAE ParisTech and holds a PhD in differential games and mathematical finance from Ecole Polytechnique and INRIA Sophia Antipolis. He is Distinguished Professor of Mathematics at the Université de Paris, Laboratoire de Probabilités, Statistique et Modélisation (LPSM), Team Mathematical Finance and Numerical Probability (MathFiProNum).

Stephane's research interests are financial modeling; counterparty credit risk, XVA analysis, risk measures; risk management for central counterparties; simulation, calibration, training, and machine learning techniques; uncertainty quantification, model risk; and the related mathematical topics in the fields of backward stochastic differential equations, random times modeling, enlargement of filtration, and numerical probability. He is the author of numerous research papers and two books: "Financial Modeling: A Backward Stochastic Differential Equations Perspective" (S. Crépey, Springer Finance Textbook Series, 2013) and "Counterparty Risk and Funding, a Tale of Two Puzzles" (S. Crépey, T. Bielecki and D. Brigo, Chapman & Hall/CRC Financial Mathematics Series, 2014). He is an associate editor of SIAM Journal on Financial Mathematics, Journal of Computational Finance, International Journal of Theoretical and Applied Finance, Journal of Dynamics and Games, and a member of the scientific council of the French financial markets authority (AMF).

Talk Learning Value-at-Risk and Expected Shortfall

Abstract: We propose a non-asymptotic convergence analysis of a two-step approach to learn a conditional value-at-risk (VaR) and expected shortfall (ES) in a nonparametric setting using Rademacher and Vapnik-Chervonenkis bounds. Our approach for the VaR is extended to the problem of learning at once multiple VaRs corresponding to different quantile levels. This results in efficient learning schemes based on neural network quantile and least-squares regressions. An a posteriori Monte Carlo (non-nested) procedure is introduced to estimate distances to the ground-truth VaR and ES without access to the latter. This is illustrated using numerical experiments in a Gaussian toy-model and a financial case-study where the objective is to learn a dynamic initial margin.

Wednesday, 19 October hour via [Zoom](#)

Speaker Phillip Murray, JP Morgan Chase and Co., London

Phillip Murray is a final-year PhD student at Imperial College London and a Quantitative Research Associate at JP Morgan within the Equity Derivatives Group. His research focusses on applications of machine learning to the pricing and hedging of derivatives portfolios.

Talk Deep Bellman Hedging

Abstract: We present an actor-critic-type reinforcement learning algorithm for solving the problem of hedging a portfolio of financial instruments such as securities and over-the-counter derivatives using purely historic data. The key characteristics of our approach are: the ability to hedge with derivatives such as forwards, swaps, futures, options; incorporation of trading frictions such as trading cost and liquidity constraints; applicability for any reasonable portfolio of financial instruments; realistic, continuous state and action spaces;

and formal risk-adjusted return objectives. Most importantly, the trained model provides an optimal hedge for arbitrary initial portfolios and market states without the need for re-training. We also prove existence of finite solutions to our Bellman equation, and show the relation to our vanilla Deep Hedging approach [BGTW19]

Wednesday, 26 October 5-6PM via [Zoom](#)

Speaker John Hull, Rotman School of Business, University of Toronto

John Hull is the Maple Financial Professor of Derivatives and Risk Management at the Joseph L. Rotman School of Management, University of Toronto. He was in 2016 awarded the title of University Professor (an honor granted to only 2% of faculty at University of Toronto.) He is an internationally recognized authority on derivatives and risk management and has many publications in this area. His work has an applied focus. His areas of research have included the impact of stochastic volatility on the pricing and hedging of options, the valuation of interest rate derivatives and credit derivatives, the calculation of value at risk, the evaluation of model risk, and the regulation of financial institutions. Recently, machine learning has been the main focus of his research and teaching. He has used machine learning to understand volatility surface movements, hedge derivatives portfolios, and value exotic options.

He was, with Alan White, one of the winners of the Nikko-LOR research competition for his work on the Hull-White interest rate model, which is widely used by practitioners. In 1999 he was voted Financial Engineer of the Year by the International Association of Financial Engineers. He has acted as consultant to many financial institutions throughout the world and has won many teaching awards, including University of Toronto's prestigious Northrop Frye award. His current research interests are concerned with the application of machine learning to finance. He is well known for his four books: "Risk Management and Financial Institutions" (now in its 5th edition); "Options, Futures, and Other Derivatives" (now in its 11th edition); "Fundamentals of Futures and Options Markets" (now in its 9th edition); and "Machine Learning in Business: An Introduction to the World of Data Science" (now in its 3rd edition). The books have been translated into many languages and are widely used by practicing managers as well as in the classroom.

Dr. Hull is academic director of FinHub (Rotman's Financial Innovation Lab). He is a Senior Research Fellow at the Global Risk Institute and a Senior Advisor to the Global Association of Risk Professionals. In addition to the University of Toronto, Dr. Hull has taught at York University, University of British Columbia, New York University, Cranfield University, and London Business School.

Talk Gamma and Vega Hedging Using Deep Distributional Reinforcement Learning

Abstract: We use deep distributional reinforcement learning (RL) to develop a hedging strategy for a trader responsible for derivatives that arrive stochastically and depend on a single underlying asset. The transaction costs associated with trading the underlying asset are usually quite small. The trader therefore normally carries out delta hedging daily, or even more frequently, to ensure that the current portfolio is almost completely insensitive to small movements in the asset's price. Hedging the portfolio's exposure to large asset price movements and volatility changes (gamma and vega hedging) is more expensive because this requires trades in derivatives, for which transaction costs are quite large. Our analysis takes account of these transaction cost differences. It shows how RL can be used to develop a strategy for using options to manage gamma and vega risk with three different objective functions. These objective functions involve a mean-variance trade-off, value at risk, and conditional value at risk. We illustrate how the optimal hedging strategy depends on the asset price process, the trader's objective function, the level of transaction costs when options are traded, and the maturity of the options used for hedging. We also investigate the robustness of the hedging strategy to the process assumed for the underlying asset.

Wednesday, 16 November hour via [Zoom](#)

Speaker Sanjiv Das, Leavey School of Business, Santa Clara University

Sanjiv Das is the William and Janice Terry Professor of Finance at Santa Clara University's Leavey School of Business.

He previously held faculty appointments as Associate Professor at Harvard Business School and UC Berkeley. He holds post-graduate degrees in Finance (M.Phil and Ph.D. from New York University), Computer Science (M.S. from UC Berkeley), an MBA from the Indian Institute of Management, Ahmedabad, B.Com in Accounting and Economics (University of Bombay, Sydenham College), and is also a qualified Cost and Works Accountant. He is a senior editor of The Journal of Investment Management, co-editor of The Journal of Derivatives, and Associate Editor of other academic journals.

Prior to being an academic, he worked in the derivatives business in the Asia-Pacific region as a Vice-President at Citibank. His current research interests include: the modeling of default risk, machine learning, social networks, derivatives pricing models, portfolio theory, and venture capital. He has published over eighty articles in academic journals, and has won numerous awards for research and teaching. His recent book "Derivatives: Principles and Practice" was published in May 2010. He currently also serves as a Senior Fellow at the FDIC Center for Financial Research.

Talk Multimodal Machine Learning for Finance

Abstract: Credit ratings are traditionally generated using models that use financial statement data and market data, which is tabular (numeric and categorical). Practitioner and academic models do not include text data. Using an automated approach to combine long-form text from SEC filings with the tabular data, we show how multimodal machine learning using stack ensembling and bagging can generate more accurate rating predictions. This paper demonstrates a methodology to use big data to extend tabular data models, which have been used by the ratings industry for decades, to the class of multimodal machine learning models.

Wednesday, 23 November 5-6PM via [Zoom](#)

Speaker Arthur Böök, Recover X, San Francisco

Arthur Böök works as a machine learning researcher (w/ focus on Natural Language Processing) and has previous experience as a derivatives structurer at a top tier European bank. Fascinated by the information content of markets and the potential of modern computing, he conducts research in option pricing and AI.

Talk Smiling in Action

This talk is on Arthur's latest paper published in Wilmott July issue.

Abstract: The implied volatility surface (IVS) contains market participants' expectations of the dynamics of an underlying asset's price w.r.t time and its spot price. The IVS is a cornerstone for pricing and risk management of non-linear derivatives.

We present a single framework for multi-step ahead prediction of the entire IVS by letting the parameters of a parametric volatility model be learned and evolved with a stack of convolutional LSTM (ConvLSTM) layers in an encoding-forecasting structure. Our framework allows for an arbitrary choice of parametric volatility models and thus accommodates constraints such as arbitrage-freeness and can thus used for pricing and hedging options, performing risk analysis, as well as for volatility trading.

In this presentation, we will explore the performance of our model on the S&P 500 index option prices several steps ahead by computing some measures of accuracy and comparing them against a benchmark. On average, our model systematically outperforms the naive approach at predicting long term forecasts for short to mid-range maturities. This indicates that the dynamics of the IVS are dominated by trend and

mean reversion, and can be captured with a suitable model.

30 November hour via [Zoom](#)

Speaker Matthew Dixon, Applied Math Department and Stuart School of Business, Illinois Tech

Matthew Dixon is a British applied mathematician working in the area of algorithmic finance. His research focuses on applying concepts in computational and applied mathematics to financial modeling, especially in the area of algorithmic trading and derivatives. Matthew's research is currently funded by Intel Corporation and he develops codes for high performance architectures. His work in deep learning with Diego Klabjan (NWU) has brought wide recognition and he is a frequently invited speaker at quant and fintech events around the world in addition to be referenced as a computational finance expert in multiple reputed media outlets including the Financial Times and Bloomberg Markets.

Talk Deep Partial Least Squares for Empirical Asset Pricing

Abstract: We use deep partial least squares (DPLS) to estimate an asset pricing model for individual stock returns that exploits conditioning information in a flexible and dynamic way while attributing excess returns to a small set of statistical risk factors. The novel contribution is to resolve the non-linear factor structure, thus advancing the current paradigm of deep learning in empirical asset pricing which uses linear stochastic discount factors under an assumption of Gaussian asset returns and factors. This non-linear factor structure is extracted by using projected least squares to jointly project firm characteristics and asset returns on to a subspace of latent factors and using deep learning to learn the non-linear map from the factor loadings to the asset returns. The result of capturing this non-linear risk factor structure is to characterize anomalies in asset returns by both linear risk factor exposure and interaction effects. Thus the well known ability of deep learning to capture outliers, shed lights on the role of convexity and higher order terms in the latent factor structure on the factor risk premia. On the empirical side, we implement our DPLS factor models and exhibit superior performance to LASSO and plain vanilla deep learning models. Furthermore, our network training times are significantly reduced due to the more parsimonious architecture of DPLS. Specifically, using 3290 assets in the Russell 1000 index over a period of December 1989 to January 2018, we assess our DPLS factor model and generate information ratios that are approximately 1.2x greater than deep learning. DPLS explains variation and pricing errors and identifies the most prominent latent factors and firm characteristics.