

Quantitative Finance Workshop
Abstracts

8–9 June 2026

June 8, 2026

Partially Ordered Peacocks

Anna Aksamit

The University of Sydney, Australia

Martingale optimal transport problem of transporting one measure to another under a martingale constraint emerges in the robust approach to pricing and hedging. The aim is to find optimal bounds on the prices of exotic derivatives in terms of the (market) prices of call options. No explicit assumptions are made about the dynamics of the price process of the underlying asset and information about the distribution of asset prices is derived from the call prices.

Kellerer's theorem shows that any peacock — that is, a family of probability measures increasing in convex order — can be embedded in a Markovian martingale whose marginals match those measures. This is a cornerstone result closely linked to martingale optimal transport. In our work, we look at what happens when the measures are only partially ordered in convex order — partial peacock — and study when an appropriate embedding is possible. This concept may be applied to verify compatibility of vanilla and barrier product prices.

In general, such embeddings do not always exist, and even simple cases can exhibit non-trivial interactions. By focusing on the subclass of min-stable embeddings, where the embedding extends to the lower semilattice completion of the measures, we will provide necessary and sufficient conditions for the existence of these min-stable embeddings. We will also present a convex-concave decomposition of measure pairs, as well as geometric characterisations of such embeddings.

Short-lived Gases, Carbon Markets, and Climate Risk Mitigation

Sara Biagini

LUISS Rome, Italy

We study the problem of optimal climate risk mitigation with short-term emission reduction targets and long-run temperature stabilization goals in the presence of firms generating greenhouse gases with different temporal persistency and warming potential. We investigate how the pervasive notion of carbon equivalence may undermine climate risk mitigation efforts when carbon markets can be used to trade short-lived gasses against long-lived ones. The findings are used to demonstrate the vulnerability of certain emission metrics and carbon accounting standards to greenwashing and to support the reporting of emissions in disaggregated form and native units of measure. Joint with E. Biffis and K. Selazadeh Nobari

Market Segmentation and Arbitrage

Umut Cetin

London School of Economics, UK

We study the optimal strategies for a high frequency arbitrageur who takes advantage of his superior trading speed to create arbitrage opportunities by investing in a stock traded in two venues. Arbitrage opportunities arise due to a combination of liquidity shocks and asynchronous price adjustment to news. We study the problem in a variant of the Kyle model. In a striking deviation from the Kyle mode, it is the market makers, as opposed to the strategic arbitrageur, that drive the prices to the fundamental value to limit the speculations of the arbitrageur towards the end of the trading period. Joint work with Eduardo Ferioli Gomes (Universidade Federal Fluminense).

Deep Learning for Continuous-time Stochastic Control with Jumps

Patrick Cheridito

ETH Zürich, Switzerland

In this paper, we introduce a model-based deep-learning approach to solve finite-horizon continuous-time stochastic control problems with jumps. We iteratively train two neural networks: one to represent the optimal policy and the other to approximate the value function. Leveraging a continuous-time version of the dynamic programming principle, we derive two different training objectives based on the Hamilton-Jacobi-Bellman equation, ensuring that the networks capture the underlying stochastic dynamics. Empirical evaluations on different problems illustrate the accuracy and scalability of our approach, demonstrating its effectiveness in solving complex high-dimensional stochastic control tasks.

Dynamic Hedging under Model Uncertainty

Rama Cont

University of Oxford, UK

It is commonly assumed that a detailed and accurate description of the joint dynamics of risk factors and asset prices is a pre-requisite for the design of hedging and risk management strategies for derivatives portfolios. This has motivated the development of increasingly complex stochastic models with many risk factors and parameters, which are challenging to estimate and implement, but whose benefit in terms of risk management is not clear.

We revisit the dynamic hedging problem as a feedback control problem under model uncertainty. We argue that the widely used approach of continuously recalibrating an auxiliary pricing model is an effective way of treating model uncertainty and link this approach with Model Predictive Control (MPC) and reinforcement learning (RL). We show that any auxiliary pricing model capable of calibrating the cross section of market prices of hedging instruments and satisfying an identifiability condition may be used in this way to compute adequate robust dynamic hedging strategies without explicit knowledge of market dynamics. This procedure captures exposures to latent risk factors through parameter recalibration and leads to a self-financing strategy with predictable loss. We provide explicit formulas for hedge ratios in dynamically recalibrated models, and provide a detailed decomposition of the gain of the dynamically hedged position. We argue that this approach is an effective recipe for dynamic hedging under model uncertainty and pleads in favor of computationally exible models with good inversion properties. Our results provide insights for model design and clarify the role of parameter calibration in risk management.

Hedging: Holding Stocks, Trading Bonds

Paolo Guasoni

Dublin City University, Ireland

University of Bologna, Italy

Why do rational investors trade, and which assets do they choose? Classical no-trade theorems show that new information can move prices without inducing trade. This paper contrasts and qualifies that logic. In a dynamic economy with stochastic growth, several long-lived agents with heterogeneous risk aversion, time preferences, and income streams make consumption and portfolio decisions while trading stocks, a long-term bond, and short-term loans with one another. The closed-form equilibrium reveals that agents do not trade stocks, even though stock returns vary over time and remain predictable. Instead, they trade the long-term bond continuously to hedge growth shocks intertemporally, despite sharing common information and beliefs. Yet agents would cease trading if they gained access to a set of bond funds that add no spanning power, since dynamic strategies in the long bond replicate them. The model identifies long-term bonds as the natural asset class for hedging growth shocks.

Neural Chaos: Minimax-Optimal Approximation of Predictable Processes without Iterated Integrals or Fixed Bases

Anastasis Kratsios

McMaster University, Canada

Predictable square-integrable processes are central in continuous-time stochastic control, reinforcement learning, and financial applications. Wiener-chaos-type expansions provide powerful representations, but their computational use is limited by large chaos dictionaries and high-order iterated stochastic integrals. We address this bottleneck with NeuralChaos, a finitely parameterized class of H_T -valued processes that uses only finitely many evaluations of the driving Brownian motion while preserving predictability and square-integrability by construction.

We prove that NeuralChaos achieves minimax-optimal nonlinear approximation rates, equivalently best N -term sparse approximation rates. For smooth and compressible predictable processes, these rates are sublinear, and we show that compressibility is generic: a process drawn from H_T according to a natural sub-Gaussian sampling model is compressible with high probability.

In contrast, we show that finite-dimensional Markovian neural-SDE-type parametrizations are not representative of H_T : their time-discretized predictable surrogates form a meagre and Gaussian-null subset of H_T . Numerical experiments on hedging and calibration problems for stochastic-volatility models illustrate the practical effectiveness of the NeuralChaos architecture.

Strategic Trading when Agents Forecast the Forecasts of Others: Existence and Convergence

Kasper Larsen

Rutgers University, USA

Foster and Viswanathan (1996) extend the discrete-time setting of Kyle (1985) to multiple informed traders who have partial information about the stock's terminal dividend. We prove that an equilibrium exists in the setting of Foster and Viswanathan (1996). We conclude by discussing convergence to the continuous-time version given in Back, Cao, and Willard (2000).

This is joint work with Jin Choi (UNIST)

Dual Approaches to Stochastic Control via SPDEs and the Pathwise Hopf Formula

Mathieu Lauriere

New York University Shanghai, China

We develop dual approaches for continuous-time stochastic control problems, enabling the computation of robust dual bounds in high-dimensional state and control spaces. Building on the dual formulation proposed in [L. C. G. Rogers, *SIAM Journal on Control and Optimization*, 46 (2007), pp. 1116–1132], we first formulate the inner optimization problem as a stochastic partial differential equation (SPDE); the expectation of its solution yields the dual bound. Curse-of-dimensionality-free methods are proposed based on the Pontryagin maximum principle and the generalized Hopf formula. In the process, we prove the generalized Hopf formula, first introduced as a conjecture in [Y. T. Chow, J. Darbon, S. Osher, and W. Yin, *Journal of Computational Physics* 387 (2019), pp. 376–409], under mild conditions. Numerical experiments demonstrate that our dual approaches effectively complement primal methods, including the deep BSDE method for solving high-dimensional PDEs and the deep actor-critic method in reinforcement learning. Joint work with Jiefei Yang (NYU Shanghai).

Particle Method for Mean-field Control Problems with Partial Observation

Xiaolu Tan

The Chinese University of Hong Kong, Hong Kong SAR

We consider a class of mean-field control problems under partial observation. Utilizing the classical reference probability approach, we reformulate the problem as a mean-field control problem with common noise-adapted controls. Building on this formulation, we introduce a particle approximation method and analyze both its convergence and the rate of convergence with respect to the number of particles.

Flexible Information Acquisition in the Kyle Model

Hao Xing

Boston University, USA

We study an information acquisition problem in which an informed trader acquires costly information prior to trading in the Kyle equilibrium. The cost of information acquisition is represented by an entropy cost. Regardless of the prior distribution of the asset payoff, continuous signals are optimal. Moreover, any continuously distributed signal, together with an associated logit-type posterior distribution of the payoff, yields the same ex-ante value for the informed trader, the same distribution of posterior expected payoff, and the same unconditional distribution of the informed trader's trading strategy. Consequently, a normally distributed signal can be adopted without loss of generality. We further show that when the information acquisition cost increases or the volatility of noise trades decreases, the variance of the posterior expected payoff declines, the profit potential from trading diminishes, meanwhile the posterior expected payoff increasingly resembles a normal distribution, and the information leakage cost from trading decreases.