SIAM 2025

Searchable Abstracts Document

SIAM Conference on Financial Mathematics and Engineering (FM25)

July 15–18, 2025

Hyatt Regency Miami, Miami, FL U.S.

This document was current as of July 8, 2025. Abstracts appear as submitted.

SIAM
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: 800-447-7426 (U.S. & Canada) +1-215-382-9800 (Worldwide)
meetings@siam.org

IP1

Volatility and Order Flow: A Tale of Two Fractional Brownian Motions

Price volatility and order flow dynamics are two commonly used measures of the intensity of market fluctuations that are evidently closely related. However, on the one hand, the rough nature of the volatility is an undeniable stylized fact. On the other hand, smooth fractional Brownian motions have been proposed to capture the autocorrelation of the order flow, in contrast to the paradigm of rough volatility. We discuss how to resolve this apparent contradiction, providing new microstructural foundations for the joint dynamics of prices, volatility and volumes. This is joint work with Johannes Muhle-Karbe, Youssef Ouazzabi Chahdi and Grgoire Szymanski.

Mathieu Rosenbaum
Ecole Polytechnique
mathieu.rosenbaum@polytechnique.edu

IP2

Machine Learning for Stochastic Control and Games: From Foundations to Multi-Agent Intelligence

This talk presents recent advances at the intersection of machine learning, stochastic control, and multi-agent systems, with applications in finance, climate change, and decision-making under limited information or adversarial settings. I begin with a high-level overview based on our recent review of deep learning methods for high-dimensional stochastic control and differential games, highlighting core algorithmic paradigms and theoretical challenges. Building on this foundation, I describe recent work on strategic deception in multi-agent environments, where learning agents operate under adversarial conditions involving goal inference, deception, and detection. I then discuss a global climate-economy model formulated as a mean field control problem over a continuum of heterogeneous regions subject to aggregate shocks, enabling the study of coordinated policy design amid interregional diversity and model ambiguity. These examples illustrate the interplay between modern AI methods and the theory of strategic decision-making in complex dynamic systems.

Ruimeng Hu
University of California, Santa Barbara
Dept. of Mathematics, and Dept. of Stat. and Applied Prob.
rhu@ucsb.edu

IP3

Model Risk Hedging Through Distributionally Robust Sensitivity

Distributionally robust optimization studies the worst deviation of an evaluation functional on the Wasserstein ball centered at the model of interest. We derive explicit sensitivity analysis under marginal and martingale constraints which provide first order semi-static hedge against model risk.

Nizar Touzi
Ecole Polytechnique
nt2635@nyu.edu

IP4

Arbitraging on Decentralized Exchanges

Decentralized exchanges (DEXs) are alternative venues to centralized exchanges to trade cryptocurrencies (CEXs) and have become increasingly popular. An arbitrage opportunity arises when the exchange rate of two cryptocurrencies in a DEX differs from that in a CEX. Arbitrageurs can then trade on the DEX and CEX to make a profit. Trading on the DEX incurs a gas fee, which determines the execution priority of the trade. We study a gas-fee competition game between two arbitrageurs who maximize their expected profit from trading. We derive the unique symmetric mixed Nash equilibrium and find that (i) the arbitrageurs may choose not to trade when the arbitrage opportunity is small; (ii) the probability of the arbitrageurs choosing a higher gas fee is lower; (iii) the arbitrageurs pay a higher gas fee and trade more when the arbitrage opportunity becomes larger and when liquidity becomes higher. The above findings are consistent with our empirical study. This is a joint work with Chen Yang and Yutian Zhou.

Xuedong He
The Chinese University of Hong Kong
Department of Systems Engineering and Engineering Management
xdhe@se.cuhk.edu.hk

IP5

A Stochastic Gordon-Loeb Model for Optimal Cybersecurity Investment Under Clustered Attacks

We develop a continuous-time stochastic model for optimal cybersecurity investment under the threat of cyberattacks. The arrival of attacks is modeled using a Hawkes process, capturing the empirically relevant feature of clustering in cyberattacks. Extending the Gordon-Loeb model, each attack may result in a breach, with breach probability depending on the system's vulnerability. We aim at determining the optimal cybersecurity investment to reduce vulnerability. The problem is cast as a two-dimensional Markovian stochastic optimal control problem and solved using dynamic programming methods. Numerical results illustrate how accounting for attack clustering leads to more responsive and effective investment policies, offering significant improvements over static and Poisson-based benchmark strategies. Our findings underscore the value of incorporating realistic threat dynamics into cybersecurity risk management. Joint paper with: C. Fontana, C. Hillairet and B. Ongarato.

Giorgia Callegaro
University of Padova
gcallega@math.unipd.it

IP6

Data-Driven Dynamic Factor Modeling via Manifold Learning

Dynamic factor modeling operates under the premise that a few latent factors drive the co-movements of a multivariate time series, and are extensively used in finance, including asset pricing, risk management, and portfolio optimization. Most dynamic factor models are assumed to be linear, despite in many settings real-world data show that the relation between factors and the observed time series is nonlinear. In this work, we introduce a dynamic factor framework where a response variable $y(t)$ in R^m depends, possibly nonlinearly, on a high-dimensional set of covariates $x(t) \in R^d$ without imposing any parametric model on the joint covariate dynamics. To uncover the underlying structure in a purely data-driven manner, we leverage diffusion maps, a nonlinear manifold learning technique introduced by Coifman and Lafon [2006], to extract low-dimensional embeddings that preserve explanatory power for the response time series $y(t)$. Concentration inequalities ensure the convergence of time averages and, under mild assumptions, we show that the dynamics of the diffusion map embeddings are approximately linear. This approximate linearity enables the use of Kalman filtering to generate the latent dynamic covariate embeddings and to predict the response variable directly from the diffusion map embedding space. We demonstrate two applications of our framework. The first application is nowcasting, where our method can be used to improve the real-time estimation precision of lower frequency response variables using observed higher frequency covariates. The second application is stress testing equity of portfolios, using a rich set of financial and macroeconomic factors drawn from the Federal Reserve's supervisory scenarios. Unlike standard scenario analysis (SSA)—which assumes zero conditional expectation for unstressed factors given the scenario—we introduce a novel conditional sampling procedure that captures dynamic correlations between stressed and unstressed risk factors. Through historical backtests covering three major financial crises, we show that our data-driven stress testing approach achieves reductions in mean absolute error (MAE) of up to 52% and 57%, relative to SSA and PCA respectively, in scenario-based portfolio return prediction. (based on joint work with Graeme Baker and Jose Antonio Sidaoui Gali)

Agostino Capponi
Columbia University
ac3827@columbia.edu

IP7
Path-Dependent Volatility

Recent empirical studies, extending the ARCH literature, have shown that the volatility of financial markets is mostly path-dependent, i.e., is very well explained by the history of past asset returns. After reviewing various arguments that all point to path-dependent volatility (PDV), we will present recent advances in PDV modeling. In particular, we will present new, simple, intuitive, low-dimensional Markovian PDV models which, unlike classical stochastic volatility models, use observable factors: a weighted average of past returns, and a weighted average of past returns squared. We will show that despite their simplicity, these models have a large predictive power and produce rich, intricate joint spot/volatility dynamics which capture many important stylized facts about volatility: strong endogeneity, a mix of short and long memory of realized variance and price trend, leverage effect, volatility clustering, roughness at the daily scale, Zumbach effect, jump-like behavior in absence of actual jumps, time-asymmetry of large volatility spikes, realistic implied volatility and implied volatility-of-volatility smiles. We will also show that these PDV models, which use a small number of parameters that all have a clear financial interpretation, can jointly fit S&P 500 smiles, VIX smiles, and the historical asset price dynamics, thus reconciling historical and risk-neutral calibrations. Exogenous volatility shocks can be included to produce path-dependent stochastic volatility models. These new PDV models also raise interesting mathematical questions (wellposedness, nonnegativity, tractability) that will also be discussed.

Julien Guyon
ENPC, Institut Polytechnique de Paris, France and NYU Tandon
julien.guyon@enpc.fr

IP8
Mean Field Games and Control

Mean field games are an active area of research with a wide range of applications, including in quantitative finance and financial economics. In recent years, considerable effort has been devoted to identifying conditions that ensure uniqueness of mean field game equilibria beyond the classical LasryLions monotonicity regime, as well as to developing selection criteria in cases where uniqueness fails. For Markovian problems in finite-player continuous-time dynamic games, uniqueness of the Markov perfect equilibrium (also known as feedback or Markovian Nash equilibrium) can often be established via the uniqueness of the corresponding Nash system. Remarkably, in discrete time, uniqueness can be recovered when the time steps are sufficiently small. In this talk, I will outline these results and ideas, and discuss an illustrative example: a two-state Kuramoto synchronization game.

Mete Soner
Princeton University
soner@princeton.edu

CP1
Overcoming Misconceptions About Local Volatility: Exact Prices Lead to Sound Continuous Markovian Models

For several reasons, local volatility is often considered insufficient for handling the many nuances of modern derivative markets. We prove that it is possible to design a local volatility model without the common drawbacks identified in the literature that is fully viable for industrial uses in option markets. In this paper, we introduce what to our knowledge is the first ever local volatility model with identifiable marginal distributions, leading to analytical option pricing formulae, as well as explicit implied and local volatility asymptotics. We prove the existence and uniqueness of a local volatility SDE strong solution, whose peculiarity lies in the diffusion coefficient being both locally unbounded and non-Lipschitzian on the half-space. In light of these findings, we revisit three claims of the extant literature on local volatility: i) the implied volatility smile shifts in the opposite direction compared to the spot price moves; ii) local variances stemming from jump-type option pricing models lead to ill-posed local volatility surfaces; iii) positive, continuous no-arbitrage pricing models need to exhibit "path roughness" to be consistent with the observed small time-to-maturity at-the-money implied volatility skew divergence. We demonstrate that none of these apply to the model presented and thoroughly detail the reasons why.

Michele Azzone
Politecnico di Milano
Michele.azzone@polimi.it

Lorenzo Torricelli, Marco Vitelli
University of Bologna

lorenzo.torricelli2@unibo.it, marco.vitelli3@unibo.it

CP1

Fast Smile Calibration in Discrete and Continuous Time Via Sinkhorn Algorithms

We propose an arbitrage-free multimarginal discrete-time model, constructed using Sinkhorn-based algorithms and calibrated to market data. The model is then extended to continuous time via a purely forward Markov functional approach. Our methodology is computationally efficient, and we present extensive empirical results across multiple assets to demonstrate its robustness and applicability.

Florian Bourgey
Bloomberg L.P.
fbourgey@bloomberg.net

CP1

Co-Trading Networks for Modeling Dynamic Interdependency Structures and Estimating High-Dimensional Covariances in Us Equity Markets

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 matrices, 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.

Yutong Lu
University of Oxford
yutong.lu@stats.ox.ac.uk

CP1

Feedback Effects and Endogenous Volatility: a Fixed-Point Approach

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Lakshithe Wagalath, Yann Braouezec
IESEG School of Management
wagalath@gmail.com, y.braouezec@ieseg.fr

CP1

Small-Time Central Limit Theorems for Stochastic Volterra Integral Equations and Their Implications on Volatility Derivatives

We study small-time central limit theorems for stochastic Volterra integral equations with Hölder continuous coefficients and general locally square integrable Volterra kernels. We prove the convergence of the finite-dimensional distributions, a functional CLT, and limit theorems for smooth transformations of the process, which covers a large class of Volterra kernels that includes rough models based on Riemann-Liouville kernels with short- and long-range dependencies. To illustrate our results, we derive asymptotic pricing formulae for digital calls on the realized variance in three different regimes. The latter provides a robust and model-independent pricing method for small maturities in rough volatility models. Finally, for the case of completely monotone kernels, we introduce a flexible framework of Hilbert space-valued Markovian lifts and derive analogous limit theorems for such lifts.

Kristof Wiedermann
TU Wien (Vienna University of Technology)
kristof.wiedermann@tuwien.ac.at

Martin Friesen
Dublin City University
martin.friesen@dcu.ie

Stefan Gerhold
Financial and Actuarial Mathematics
Vienna University of Technology
sgerhold@fam.tuwien.ac.at

CP1

Why is the Volatility of Single Stocks so Much Rougher than that of the SP500 ?

The Nested Factor Model, introduced by *Bouchaud et al.*, explains asset return fluctuations through common factors representing economic sectors and residuals. These residuals share a dominant volatility mode with the factors, alongside idiosyncratic modes unique to each residual. This model implies correlated log-volatilities between factors and residuals. Focusing on a single-factor case, we use a S-fBM process (introduced by *Peng, Bacry, and Muzy*) with a Hurst exponent $H \simeq 0.11$ as the dominant common mode. The residuals incorporate both this common mode and idiosyncratic components with $H \simeq 0$. This setup preserves the Nested Factor Models properties and aligns with *Peng et al.*'s observation that stock indices exhibit larger Hurst exponents compared to individual stocks. Our analysis shows that in the Nested log S-fBM Factor Model, single-stock Hurst exponents correspond to idiosyncratic volatility modes, while the index Hurst exponent reflects the common mode. We propose a statistical method to estimate the Hurst exponent from stock return dynamics, providing theoretical guarantees and accurate results as the number of stocks N grows. Additionally, we demonstrate that the factor can be interpreted as an index, constructed by weighting single stocks with specific coefficients. This framework bridges the Nested Factor Model with empirical observations on stock return dynamics.

Othmane Zarhali
Université Paris Dauphine - CNRS
othmane.zarhali@dauphine.eu

Cécilia Aubrun
Capital Fund Management
cecilia.aubrun@cfm.com

Emmanuel Bacry
Université Paris Dauphine - CNRS
bacry@ceremade.dauphine.fr

Jean-Philippe Bouchaud
Capital Fund Management
jean-philippe.bouchaud@cfm.com

Jean-François Muzy
Université de Corse - CNRS
muzy_j@univ-corse.fr

CP2

Deep Learning Meets Queue-Reactive: A Framework for Realistic Limit Order Book Simulation

The Queue-Reactive model introduced by [huang2015simulating] has become a standard tool for limit order book modeling, widely adopted by both researchers and practitioners for its simplicity and effectiveness. We present the Multidimensional Deep Queue-Reactive (MDQR) model, which extends this framework in three ways: it relaxes the assumption of queue independence, enriches the state space with market features, and models the distribution of order sizes. Through a neural network architecture, the model learns complex dependencies between different price levels and adapts to varying market conditions, while preserving the interpretable point-process foundation of the original framework. Using data from the Bund futures market, we show that MDQR captures key market properties including the square-root law of market impact, cross-queue correlations, and realistic order size patterns. The model demonstrates particular strength in reproducing both conditional and stationary distributions of order sizes, as well as various stylized facts of market microstructure. The model achieves this while maintaining the computational efficiency needed for practical applications such as strategy development through reinforcement learning or realistic backtesting.

Hamza Bodor
Paris 1 Panthéon-Sorbonne
bodor.hamza@gmail.com

CP2

An Efficient Deep Learning Monte Carlo Method Based on Walk on Spheres with Application to Risk Pricing

Monte Carlo methods have been widely used in evaluation of financial risks in the past three decades. While the main use of Monte Carlo methods was for linear PDEs via Feynman-Kac formula, they have also been used to generate solutions for semi-linear and fully non-linear PDEs by using backward stochastic differential equations (BSDEs), which arise in stochastic optimal control and applications such as utility maximization. A few recent research papers combined Monte Carlo and BSDEs with deep learning to obtain better evaluation estimates. In this study, we improve the efficiency of the latter methods by employing Walk on Spheres (WoS) method for simulation of Brownian motion paths. We show that this new method, called WoS-NN, provides precise and rapid global solutions and gradient approximations for PDEs. A typical experimental result demonstrated that for linear elliptic equations, the proposed WoS-NN method provides accurate field estimations, reducing 76.32% errors while using only 8% of path samples compared to the conventional Monte Carlo, which saves abundant computational time and resource consumption.

Arash Fahim
Department of Mathematics
Florida State University
afahim@fsu.edu

Michael Mascagni, Silei Song
Florida State University Speaker:
mvmascagni@fsu.edu, ssong4@fsu.edu

CP2

Liquidity Provision and Rebate Design in Option Market

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Zihan Guo
University of Oxford
zihan.guo@maths.ox.ac.uk

CP2

Semi-Supervised Bayesian Log-Signature Gans for Irregular Time Series Classification in Credit Card Fraud Detection

In this paper, we present a novel deep generative semi-supervised model for time series classification, demonstrated through its application to detecting fraudulent credit card users from individual transaction data. As data streams grow in scale and complexity, traditional methods often require large labeled datasets, struggle with irregular sampled time series of varying lengths and rely solely on point estimates without capturing uncertainty. We address these challenges by extending conditional Generative Adversarial Networks (GANs) for targeted data augmentation, integrating Bayesian inference for probabilistic predictions and uncertainty quantification, and leveraging log-signatures for robust feature encoding. We propose a new Wasserstein distance-based loss that aligns generated and real unlabeled samples while maximizing classification accuracy on labeled data in a semi-supervised setting. Experiments on the BankSim dataset show superior accuracy in low-labeled scenarios and competitive performance with ample labels. Overall, these findings highlight the potential of GAN-driven semi-supervised learning with log-signatures for scalable, high-fidelity time series classification across diverse domains.

David Hirnschall, Kurt Hornik
Vienna University of Economics and Business
david.hirnschall@wu.ac.at, kurt.hornik@wu.ac.at

CP2

Hedging Derivatives Using Tree Search and Neural Networks

Interpreting finance as a stochastic game, we investigate the application of advanced game-playing systems for the hedging of financial derivative contracts. Specifically, we compare the performance of our AlphaZero-based system in hedging and pricing financial derivatives to the state-of-the-art methodology, deep hedging. Deep hedging relies on reinforcement learning and gradient descent optimiza-

tion techniques, which can falter in market environments characterized by non-convex Q-functions. These non-convexities, arising from transaction costs or constraints such as capital or regulatory restrictions, can lead deep hedging to converge to suboptimal local minima. In this study, we construct specific market environments to highlight the limitations of deep hedging and demonstrate that AlphaZero consistently identifies optimal hedging strategies in scenarios where deep hedging fails. Our results further show that AlphaZero exhibits higher sample efficiency, achieving lower test loss with fewer training samples. This feature is critical in real-world derivative markets, where data availability is inherently limited, and controlling overfitting is paramount. This work contributes to the development of AI-driven solutions for robust risk management and pricing in challenging market conditions.

Oleg Szehr
IDSIA
oleg.szehr@idsia.ch

Matteo Maggiolo
VIGILANT ANALYTICS
matteo.maggiolo@vigilant-analytics.ch

Miroslav Strupl
IDSIA
miroslav.strupl@idsia.ch

CP2

Deep Neural Network Methods for Parabolic Variational Inequalities in Finance: A Dual Approach to Optimal Stopping and Control

Motivated by recent progress in applying deep neural networks to high-dimensional PDEs, we propose a streamlined approach for solving a mixed optimal stopping and control problem in finance. We begin by developing a deep neural network framework that approximates the solution of parabolic variational inequalities, using loss functions that directly incorporate the variational inequality. Crucially, we prove that there exist neural networks whose losses converge to zero, ensuring they approximate the true solution in a suitable Sobolev norm. For practical implementations, we extend these results to bounded domains, ensuring that our neural network surrogates satisfy boundary conditions while maintaining convergence guarantees. We then apply our approach to a specific mixed optimal stopping and control problem. By leveraging duality, we convert the nonlinear generator of the primal problem into a linear parabolic operator in the dual formulation. A key step in this process is recovering the primal value function from the dual neural network solutionan outcome made possible by our Sobolev norm analysis. Finally, we illustrate the versatility and accuracy of our method with numerical experiments for both power and non-HARA utilities, and discuss practical considerations such as domain truncation and sampling strategies. Our results underscore the potential of deep neural networks as a reliable and efficient tool for variational inequalities in optimization and control problems.

Yun Zhao, Harry Zheng
Imperial College London
yun.zhao23@imperial.ac.uk, h.zheng@imperial.ac.uk

CP3

High Dimensional Robust Shrinkage Covariance Matrix Estimator

We propose a method for constructing a new high-dimensional covariance estimator that is robust against change points in the underlying latent covariance process. The method leverages the group-fused Lasso method by [BleakleyVert:11] that approximates common change points in multiple co-occurring signals by solving multidimensional total variation problems. We use two proxies for the underlying latent covariance matrix - first is the Sample Covariance Matrix and second is the Ledoit-Wolf Shrinkage matrix and extend them to robust high dimensional covariance estimators by detecting these change points in a series of underlying proxies. Our method consistently estimates changepoints and corresponding covariance proxies under general consistency conditions for shrinkage estimators and broad noise assumptions. Simulation studies show that our method significantly outperforms the benchmarks like the Linear Shrinkage (LS) and Quadratic Inverse Shrinkage (QIS) estimators. Empirical results using TAQ data and NYSE stock portfolios confirm improvements in metrics such as the net Sharpe ratio. The paper concludes with extensive simulations and empirical analyses, showcasing the superiority of the proposed method over existing benchmarks.

Greeshma Balabhadra
Stonybrook University
greeshma.balabhadra@stonybrook.edu

El Mehdi Ainasse
Bloomberg
elmehdi.ainasse@gmail.com

Pawel Polak
Department of Applied Mathematics and Statistics
Stony Brook University
pawel.polak@stonybrook.edu

CP3

Determining Window Size for Topological Data Analysis of Financial Time Series

This study focuses on applications of Topological Data Analysis (TDA) to detect early warning signals of financial crashes. The method involves delay-coordinate embedding and windowing to convert a one-dimensional time series into a sequence of point clouds, for which persistence landscapes and their norms are computed. Earlier papers show the norms of the persistence landscapes spike before a financial crash. In this works we use the wavelet transform to determine the optimal window size in order for TDA to signal financial crashes. Stock price time series are highly dynamic and non-stationary, often characterized by abrupt drops alongside complex patterns across multiple time scales. The wavelet transform decomposes these series into distinct frequency components, effectively capturing both short-term fluctuations, such as sudden price falls, and long-term trends. We analyze the impact of different window sizes on persistence landscapes generated by TDA, identifying scales that are most effective in detecting and predicting downward trends. Our findings highlight the ability of wavelet-based feature extraction to enhance the detection of critical behaviors, like financial crashes, while improving the interpretability and accuracy of stock forecasting models. This framework offers a novel approach to integrating wavelet transform and TDA, providing critical insights into time series dynamics for financial applications and beyond.

Ruslan Gokhman, Marian Gidea
Yeshiva University
rgokhman@mail.yu.edu, Marian.Gidea@yu.edu

CP3

How Minimum Performance Thresholds Bias Backtests: Bayesian Estimation for Sharpe Ratios Under Selection Bias

It is generally accepted that high Sharpe ratios are preferred when comparing investment strategies, and that strategies which fail to reach a minimum performance threshold will be rejected. But what is the impact of this filter? We study how performance thresholds bias in-sample Sharpe ratios upwards, and develop a Bayesian framework for estimating Sharpe ratios that corrects for the bias induced by this performance thresholding. Additionally, by leveraging a dataset of observed in-sample and out-of-sample Sharpe ratios we are able to use empirical Bayes to calibrate our model, yielding improved estimates of expected future performance for new candidate strategies. This framework can help researchers make better-informed decisions by providing realistic confidence intervals for out-of-sample Sharpe ratio expectations. This work complements existing approaches in the literature on multiple testing, e.g. de Prado, Harvey, Chen.

Joseph Mulligan, Johannes Muhle-Karbe, Antoine Jacquier
Imperial College London
j.mulligan21@imperial.ac.uk, j.muhle-karbe@imperial.ac.uk, a.jacquier@imperial.ac.uk

CP4

A New Actor-critic Machine Learning Method for Solving High-dimensional HamiltonJacobiBellman Partial Differential Equations

We introduce a new actor-critic machine learning algorithm for solving the Hamilton-Jacobi-Bellman equation from stochastic control theory in high-dimensions. Thearchitecture of the critic (i.e. estimated value function) is structured so as to have the boundary condition always perfectly satisfied (rather than being merely included in the loss) and utilizes a biased gradient. The actor (i.e. estimated optimal control) is trained by minimizing the integral of the Hamiltonian over the domain. When combined with training at points sampled from a fixed measure (rather than from controlled SDE paths), this setup allows us to show that the training dynamics of the actor-critic pair in the infinite-width limit converge in a Sobolev-type space ($L^2 \times H^2$) to a certain infinite-dimensional ordinary differential equation. In particular, we prove under a convexity assumption that any fixed point of this limiting ODE is a solution of the original HJB equation. We also provide numerical examples of the algorithm solving stochastic control problems accurately in very high dimensions.

Jackson S. Hebner
Mathematical Institute
University of Oxford
jackson.hebner@worc.ox.ac.uk

CP4

Long-Term Investment for Partially Observed Stock Prices and the Azema-Yor Process

In this talk we analyze the problem of determining the optimal growth rate of the certainty equivalent under partial information with portfolios under drawdown constraints. Asset prices are modelled through a Brownian diffusion with coefficients governed by stochastic exogenous factors. These factors are non observable directly and decisions are taken only with information about the stock prices. Using the Kalman filter we deal with the source of incomplete information, and through the theory of Riccati algebraic equations an explicit form of the optimal growth rate is given. The optimal investment strategy is obtained by adapting the theory of Azema-Yor processes.This is a joint work with Erick Trevio.

Daniel Hernandez-Hernandez
Centro de Investigacion en Matematicas, Guanajuato, Mexico
dher@cimat.mx

CP4

A Stochastic Control Approach for Market Making in a Hawkes Process Driven Limit Order Book Market

Solutions to the Market Making problem in a limit order book (LOB) driven financial market typically comprises of several strong assumptions such as Brownian motion of the mid-price process, fill probability as smooth functions of distance from mid-price etc. In this work, we construct a LOB using the Compound Hawkes Process following [Jain, Muzy, Bacry et al. 2025] and form a stochastic optimal control (SOC) problem with the objective of Market Making with quadratic penalty costs. We state the existence, uniqueness of the SOCs solution and solve the Hamilton-Jacobi-Bellman equations using a numerical method. The main contribution of our work is that the mid-price process is stated as a multidimensional point process driven differential equation, the LOB fill rates are implicitly stated in the Hawkes dynamics which is dependent on the queue-depletions of preceding orders of the agents orders. Finally we compare our method to baselines and depict our methods performance against a model a la [Avellenada and Stoikov 2008] which makes several assumptions of the LOB dynamics.

Konark Jain
University College London
konark.jain.23@ucl.ac.uk

CP4

On the Solvability of Second-Order Backward Stochastic Volterra Integral Equations and Equilibrium Hjb Equations

This paper addresses the solvability of a broad class of nonlocal second-order backward stochastic Volterra integral equations (2BSVIEs) featuring two temporal parameters. These equations arise in the characterization of equilibrium strategies and corresponding value functions for time-inconsistent (TIC) stochastic control problems, where agents' preferences violate Bellman's principle of optimality. In such contexts, our formulation extends the scope of existing work by allowing both the drift and volatility of the underlying state process to be controllable, and considering objective functionals that depend on both the initial time and state. The comprehensive nature of our 2BSVIE framework requires moving away from a purely probabilistic approach for demonstrating solvability, directing us instead towards an analytical method grounded in partial differential equations (PDEs). Specifically, we employ a continuity method and Banach's fixed-point arguments within custom-designed Banach spaces to establish the well-posedness and regularity of solutions for a class of PDEs with nonlocality in both time and space (nPDEs). Subsequently, we derive a Feynman–Kac-type formula to establish a relationship between the solutions of the 2BSVIEs and the nPDEs, thereby proving the solvability of the general 2BSVIEs. These results advance our understanding of open problems in equilibrium HJB equations and TIC controls. Finally, we present a globally solvable financial example.

Chi Seng Pun
Nanyang Technological University
cspun@ntu.edu.sg

Qian Lei
Nanyang Technological University, Singapore
qian.lei@ntu.edu.sg

CP4

Optimal Transport, Informed Trading, and Stochastic Liquidity

We consider a Kyle-Back model where the noise trading volatility is stochastic and there are multiple traded assets. This may be viewed as a multidimensional generalization of Collin-Dufresne and Fos (2016). We start with the causal optimal coupling between the fundamental price of the assets and the Wiener process which drives the noise trades. By considering dual formulation of this problem, we discover the equilibrium minimizes an average of the initial market depth and the noise traders slippage costs. The dual problem admits a solution in one dimension thanks to a result from Ekren, Mostowski, and Zitkovic (2022), but we look towards the primal problem to extend this solution to the multidimensional case.

Lu Vy
Florida State University
luvy@umich.edu

Ibrahim Ekren
University of Michigan
U.S.
tbd

CP5

Some Computations for Optimal Execution with Monotone Strategies

We study an optimal execution problem in the infinite horizon setup. Our financial market is given by the Black-Scholes model with linear price impact. The main novelty of the current note is that we study the constrained case where the number of shares and the selling rate are non-negative processes. For this case we give a complete characterization of the value and the optimal control via a solution of a non-linear ordinary differential equation (ODE). Furthermore, we provide an example where the non-linear ODE can be solved explicitly. Our approach is purely probabilistic.

Yan Dolinsky
Department of Statistics
Hebrew University of Jerusalem
yan.dolinsky@mail.huji.ac.il

CP5

Path Dependent Trading Strategies Using Signatures and Kernels

In this work, we explore the role of path-dependencies in optimal trading problems. We begin by quantifying how the interplay between information, predictions, PnL, trading costs, and variance is inherently path-dependent, significantly influencing downstream optimisation tasks. To address these challenges, we extend classical methods beyond traditional Markovian assumptions. We derive solutions to optimal trading problems in a data-driven manner by parameterizing the strategy as a general function on path space. Specifically, we provide solutions for both linear functions of signatures and functions in a reproducing kernel Hilbert space (RKHS). By utilising Signatures and Kernels, we solve path-dependent optimisations in closed form, providing solutions that are both simple and intuitive.

Owen Futter
Imperial College London
o.futter21@imperial.ac.uk

CP5

Optimal Execution under Incomplete Information

We study optimal liquidation strategies under partial information for a single asset within a finite time horizon. We propose a model tailored for high-frequency trading, capturing price formation driven solely by order flow through mutually stimulating marked Hawkes processes. The model assumes a limit order book framework, accounting for both permanent price impact and transient market impact. Importantly, we incorporate liquidity as a hidden Markov process, influencing the intensities of the point processes governing bid and ask prices. Within this setting, we formulate the optimal liquidation problem as an impulse control problem. We elucidate the dynamics of the hidden Markov chains filter and determine the related normalized filtering equations. We then express the value function as the limit of a sequence of auxiliary continuous functions, defined recursively. This characterization enables the use of a dynamic programming principle for optimal stopping problems and the determination of an optimal strategy. It also facilitates the development of an implementable algorithm to approximate the original liquidation problem. We enrich our analysis with numerical results and visualizations of candidate optimal strategies.

Yadh Hafsi
Université Paris-Sacaly
yadh.hafsi@universite-paris-saclay.fr

Etienne Chevalier, Vathana Ly Vath
Université Paris-Saclay
etienne.chevalier@univ-evry.fr, vathana.lyvath@ensiie.fr

CP5

A Bayesian Approach to Generate Distribution-based Signals in Pairs Trading

We introduce a novel approach to improving the precision and adaptability of trading signals in pairs trading. Our method derives the full conditional distribution of the hedge ratio and utilizes its quantiles as confirmation thresholds for trading signals generated within the standard cointegration framework. We apply this approach to selected asset pairs across the U.S. and Brazilian markets, demonstrating its effectiveness through empirical analysis. Our findings indicate that the proposed Bayesian hierarchical model significantly enhances trading performance and risk management compared to traditional cointegration-based strategies. By adopting a distribution-based framework, our approach not only enables more timely and adaptive trading signals but also improves pair selection by effectively filtering out false positives in cointegration tests, as demonstrated through simulations.

Allan Quadros
Kansas State University
University of North Florida
allanvcq@gmail.com

Michael Higgins
Kansas State University
mikehiggin@ksu.edu

Brian Silverstein
University of South Carolina
brian.silverstein@moore.sc.edu

CP5

Fluid-Limits of Fragmented Limit-Order Markets

Maglaras, Moallemi, and Zheng (2021) have introduced a flexible queueing model for fragmented limit-order markets, whose fluid limit remains remarkably tractable. In the present study we prove that, in the limit of small and frequent orders, the discrete system indeed converges to the fluid limit, which is characterized by a system of coupled nonlinear ODEs with singular coefficients at the origin. Moreover, we establish that the fluid system is asymptotically stable for an arbitrary number of limit order books in that, over time, it converges to the stationary equilibrium state studied by Maglaras et al. (2021).

Yonatan Shadmi
Technion
shadmi@campus.technion.ac.il

Eyal Neuman
Imperial College
e.neumann@imperial.ac.uk

Johannes Muhle-Karbe
Imperial College London
j.muhle-karbe@imperial.ac.uk

CP5

Signal-Adaptive Optimal Execution Quotes

This talk focuses on optimal execution strategies for sequentially placing limit orders in a limit order book at specific quote prices. Unlike previous studies that primarily derive optimal trading speed within the Almgren-Chriss framework, we address the problem by focusing on how limit orders should be placed over time. Our model considers that price quotes in the limit order book may be influenced by signals and incorporates both execution risks and price impact. We frame the optimal execution problem with four objectives: execution without risk aversion (Case I), execution with running inventory risk (Case II), execution with exponential utility (Case III), and execution with both running inventory risk and exponential utility (Case IV). By deriving the corresponding HJB equations for each case, we demonstrate that these problems converge to a specific type and can be solved explicitly, leading to fully explicit solutions for optimal execution across all cases.

Fenghui Yu
Department of Mathematics, The University of Hong Kong
fenghui.yu@tudelft.nl

CP6

Managing Climate Risks for Sustainable Development

Extreme weather from climate change threatens sustainable development by destroying a society's productive base, including manufactured, human and natural capital. We model how society should manage these threats given a sustainable-development constraint that social welfare be non-decreasing over time in expectation. The planner learns about the consequences of climate change for extreme weather and optimally revises consumption and investment plans. State-dependent consumption taxes and investment subsidies decentralize the first-best allocation. Increasing extreme-weather risk leads the sustainable-development constraint to bind, triggering consumption taxes to subsidize investments in the productive base. We apply our model to country-level tropical cyclones. Compared to a no-learning or constant-risk environment, higher sustainable-development taxes and subsidies are needed. Bad climate news surprisingly leads to higher asset valuations because society has to build up more capital to ensure sustainable development.

Heqing Huang
National University of Singapore
Department of Mathematics
hheqing@u.nus.edu

Neng Wang
Cheung Kong Graduate School of Business
nw2128@columbia.edu

Harrison Hong
Columbia University and NBER
hh2679@columbia.edu

CP6

Market Coupling Mechanisms for Intraday Electricity Markets with Limited Transmission Capacities

In this talk, we study n interacting limit order markets, assuming that transactions between any two markets are limited as it is for example the case in the European intraday electricity market SIDC. While the generalised processor sharing discipline serves as an efficient market coupling mechanism for markets without transmission constraints, we propose a new class of market coupling mechanisms for markets with transmission constraints. To this end, we first introduce a novel class of Skorokhod reflection problems with two non-linear moving boundaries depending on the constraining process itself. Having established existence and uniqueness of the (extended) Skorokhod problem, we can then describe the market dynamics by a novel class of multidimensional SDEs with oblique reflections and can compare different market coupling mechanisms via simulations.

Dörte Kreher, Rahama Sani Abdullahi
HU Berlin
kreher@math.hu-berlin.de, rahamasani1996@gmail.com

CP6

Pricing Energy Spread Options with Variance Gamma-Driven Ornstein-Uhlenbeck Dynamics

We consider the pricing of energy spread options for spot prices following an exponential Ornstein-Uhlenbeck process driven by a sum of independent multivariate variance gamma processes. This gives rise to naturally multivariate, mean-reverting, infinite activity spot price dynamics. As energy is a nontraded asset, the Esscher transform is used to obtain the risk-neutral measure. Within this class of driving processes, we focus on the weak variance alpha-gamma process and show it is not closed under the Esscher transform, but changes into a sum of variance gamma processes. By deriving an analytic expression for the characteristic function of the innovation term, we then obtain a pricing formula for forwards and apply the Hurd-Zhou method to price bivariate spread options. Lastly, we demonstrate how the model should be both estimated on spot prices under the real world measure and calibrated on forward or call prices under the risk-neutral measure, and provide numerical results for the pricing of spread options.

Kevin W. Lu
Australian National University
kevin.lu@anu.edu.au

Tim Leung
University of Washington
timleung@uw.edu

CP6

Modeling the Impact of Climate Transition on Real Estate Prices

In this work, we propose a model to quantify the impact of the climate transition on a property in housing market. We begin by noting that property is an asset in an economy. That economy is organized in sectors, driven by its productivity which is a multidimensional Ornstein-Uhlenbeck process, while the climate transition is declined thanks to the carbon price, a continuous deterministic process. We then extend the sales comparison approach and the income approach to valuate an energy inefficient real estate asset. We obtain its value as the difference between the price of an equivalent efficient building following an exponential Ornstein-Uhlenbeck as well as the actualized renovation costs and the actualized sum of the future additional energy costs. These costs are due to the inefficiency of the building, before an optimal renovation date which depends on the carbon price process. Finally, we carry out simulations based on the French economy and the house price index of France. Our results allow to conclude that the order of magnitude of the depreciation obtained by our model is the same as the empirical observations.

Lionel Sopgoui
LPSM Paris
Université Paris Cité
sopgoui@lpsm.paris

CP7

Eliminating Maturity Dependency of Parameters in the Mixed Bergomi Model

The mixed Bergomi model, introduced by Lorenzo Bergomi in Smile Dynamics III, is celebrated for its ability to reproduce the upward-sloping smiles of VIX options. However, its reliance on maturity-dependent parameters makes it less parsimonious and potentially prone to overfitting. In this work, we enhance the model's parsimony by eliminating the parameter maturity dependency while maintaining its descriptive power. Specifically, we employ exponential spline interpolation techniques to capture the term structure of parameters, leveraging observed consistencies in daily term structure patterns. To validate our approach, we use a fast quantization method to jointly calibrate the parsimonious model to VIX futures and options over extended timescales. Our findings demonstrate the development of a more computationally efficient and interpretable volatility modelling framework, offering improved dynamical performance compared to the conventional mixed Bergomi model.

Nelson Kyakutwika
Stellenbosch University
20349912@sun.ac.za

Mesias Alfeus
Department of Statistics and Actuarial Science
Stellenbosch University
mesias@sun..ac.za

Erik Schlogl
Quantitative Finance Research Centre
University of Technology, Sydney
Erik.Schlogl@uts.edu.au

CP7

Polynomial-Based Schemes for Signature Kernels and Applications in Finance

Signature kernels have become a powerful tool in kernel methods for sequential data. In "The Signature Kernel is the solution of a Goursat PDE", the authors introduce a kernel trick showing that, for continuously differentiable paths, the signature kernel satisfies a hyperbolic PDE of Goursat type in two independent time variables. While finite difference methods have been explored for this PDE, they suffer from accuracy and stability issues when handling highly oscillatory inputs. In this work, we propose two advanced numerical schemes that approximate the solution using polynomial representations of boundary conditions, employing either approximation or interpolation techniques. We prove the convergence of the polynomial approximation scheme and demonstrate experimentally that both methods achieve several orders of magnitude improvement in mean absolute percentage error (MAPE) over finite difference schemes, without increasing computational complexity. Finally, we demonstrate the practical significance of these methods in quantitative finance (alpha generation, pricing, and hedging), and discuss their potential integration into broader mathematical finance applications.

Francesco Piatti
Imperial College London
francesco.piatti19@imperial.ac.uk

Thomas Cass
Imperial College London
Alan Turing Institute
thomas.cass@imperial.ac.uk

CP7

Dynamical Low Rank Approximation Applied to Financial SDEs

The Dynamical Low Rank Approximation (DLRA) is a time-dependent reduced-order model (ROM) that has demonstrated significant benefits in uncertainty quantification, showing computational efficiency and accuracy. Its strength lies in constructing a surrogate model with time-evolving deterministic and stochastic bases, enabling it to closely track the dynamics of the system under study. For stochastic differential equations (SDEs), a rigorous mathematical framework was introduced in [Kazashi, Y., Nobile, F., & Zoccolan, F. (2024). Dynamical Low-Rank Approximation for Stochastic Differential Equations. Mathematics of Computation], whereas its numerical properties are elaborated in [Kazashi, Y., Nobile, F., & Zoccolan, F. (2025). Numerical Methods for DLRA Applied to SDEs, Part I & II, in preparation]. This talk highlights practical applications in asset pricing. High-dimensional SDE models are extensively used in finance, but their discretization often suffers from the curse of dimensionality. DLRA mitigates this issue by capturing the dominant features of financial systems. Unlike the majority of ROMs, DLRA is computed "on the fly"; hence, it can be updated with real-time market data without an offline procedure. Numerical results concerning the high-dimensional Heston model and rough volatility approximations will be presented [Kazashi, Y., Nobile, F., Van Dieren E. & Zoccolan, F. (2025). Dynamical Low-Rank Approximation applied to financial problems, in preparation].

Fabio Zoccolan
Institute of Mathematics, EPFL
fabio.zoccolan@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Yoshihito Kazashi
University of Strathclyde
y.kazashi@strath.ac.uk

Eliott Van Dieren
Institute of Mathematics, EPFL
eliott.vandieren@epfl.ch

CP7

Numerical Approximation of RBSDEs Via Regularization

In this paper, we study the convergence of numerical solutions for Reflected Backward Stochastic Differential Equations (RBSDEs) using a regularization approach. We establish the order 1 convergence between the continuous regularized solution and the reflected solution, in full generality, as a function of the regularization parameter. The convergence between the continuous regularized solution and the corresponding RBSDE is obtained in both the almost sure and the $\mathbb{L}^p(\mathcal{F})$-sense ($p \geq 2$). Additionally, we derive the convergence rate for the discretized version of the regularized BSDE under mild regularity conditions. To illustrate these results, we validate the convergence rate through numerical experiments in the context of American put options.

Ankush Agarwal
University of Western Ontario
aagarw93@uwo.ca

Emmanuel Gobet
Ecole Polytechnique
France
emmanuel.gobet@polytechnique.edu

Yihan Zou
University of Glasgow
yihan.zou@outlook.com

CP8

On a Multivariate Extension for Copula-Based Conditional Value at Risk

Copula-based Conditional Value at Risk (CCVaR) is defined as an alternative version of the classical Conditional Value at Risk (CVaR) for multivariate random vectors intended to be real-valued. We aim to generalize the (CCVaR) to several dimensions ($d \geq 2$) when the dependence structure is given by an Archimedean copula. While previous research was focused on the bivariate case leaving the multivariate version unexplored, we derive an almost closed-from formula for CCVaR when the copula is Archimedean. We verify under which conditions the CCVaR is a coherent risk measure. We present numerical experiments for estimation of this CCVaR with real data and compare it with previous classical notions of Value at Risk (VaR) and CVaR.

Andres Mauricio Molina Barreto
Keio University
Faculty of Economics
molina.mauricio@keio.jp

CP8

Bank Liquidity Management and Payout Policy under Peer Pressure

We present a theoretical model that examines the effects of peer pressure on the trade-off that bank managers face when deciding whether to accumulate reserves or to pay shareholders. We show that under high peer pressure, banks reduce payouts and increase cash reserves, which reduces the probability of default and improves financial stability. Using data from the Federal Reserve's Y-9C report from 1987-2020, we find that a one standard deviation increase in peer pressure corresponds to a 2%-12% rise in cash reserves and a 1%-17% decline in the dividend payout ratio, relative to their respective averages, thereby leading to improved bank risk profiles.

Yuri Saporito
Fundacao Getulio Vargas, Brazil
yuri.saporito@gmail.com or yuri.saporito@fgv.br

CP8

Applications of the Second-Order Esscher Pricing in Risk Management

In this talk, we consider a stock process modelled by the exponential of a quasi-left-continuous semimartingale. In a

multidimensional setting, we define the Esscher transform of the second order in a continuous-time framework. As a first result we obtain that the second-order Esscher transform is a martingale measure and we provide a martingale condition. We show that for the same financial model we have two characterizations of the Esscher measure, the exponential Esscher measure and the linear Esscher measure, respectively. To explore the application significance of the second-order Esscher pricing model in option pricing and risk management, we split the study into two main parts. First, we focus on the constant jump diffusion case, analyzing the behavior of option prices as a function of the second-order parameter and the resulting pricing intervals. Using real data, we perform a dynamic delta hedging strategy, illustrating how risk managers can determine an interval of value-at-risks, granting flexibility in pricing based on additional information. We compare our pricing interval to other jump-diffusion models, showing its comprehensive risk factor incorporation. The second part extends the second-order Esscher pricing to more complex models, including the Merton and Kous Double Exponential jump-diffusion. We derive option prices using the fast Fourier transform method and provide practical formulas for European call and put options under these models.

Michèle Vanmaele
Ghent University
Michele.Vanmaele@UGent.be

Tahir Choulli, Ella Elazkany
University of Alberta
tchoulli@ualberta.ca, elazkany@ualberta.ca

CP8

Decision-Making Frameworks for Cyber Resilience of Financial Networks

This talk introduces decision-making frameworks for evaluating systemic cyber risks and the effectiveness of risk mitigation strategies in hypothetical cyber incident scenarios. These frameworks are based on three fundamental components: (1) a set of acceptable network configurations, (2) a set of interventions aimed at risk mitigation, and (3) a cost function that quantifies the trade-offs associated with these interventions. The analysis utilizes a unique dataset on outsourcing activities by German funds to investigate two key cyber scenarios: a widespread cyber-attack and the disruption of a third-party provider. The results contribute to the ongoing discussion on effective micro- and macroprudential cyber risk mitigation strategies and the development of meaningful criteria for network acceptability.

Alexander Voß
Leibniz University Hannover
alexander.voss@insurance.uni-hannover.de

CP8

Total Value Adjustments for Counterparty Credit Risk in Multicurrency Setting

In a global economy, financial institutions operate in different currencies. In the context of total valuation adjustments (XVA), they can either fund or post collateral in different currencies. Some attention has been addressed to the extension of the different adjustments included in XVA from the single currency to the multicurrency setting. In this work, we propose appropriate models to compute the XVA in a multicurrency setting by means of dynamic hedging methodologies. Besides the stochastic evolution of the assets in different currencies, the presence of stochastic intensities of default and the consideration of constant or stochastic exchange rates are assumed when computing the XVA associated to some options contracts. These models can be formulated in terms of (non)linear parabolic partial differential equations (PDEs) or in terms of expectations. When the number of stochastic factors is not greater than two, we propose a Lagrange-Galerkin scheme for solving the PDEs, combined with fixed point techniques for the nonlinear problems. For problems that include more than two underlying stochastic factors (assets, intensities of default, and/or stochastic FX rates), we propose the use of Monte Carlo simulations applied to the formulations based on expectations, combined with a Picard method and the more efficient Multilevel Picard iteration (MPI) scheme for the nonlinear cases. We apply these techniques to different European style options.

Carlos Vázquez Cendón
University of Coruña
University of Coruña
carlosv@udc.es

Iñigo Arregui, Roberta Simonella
University of A Coruña
arregui@udc.es, r.simonella@udc.es

CP8

A Network Approach to Macroprudential Buffers

I use network modelling of systemic risk to set macroprudential buffers from an operational perspective. I focus on the countercyclical capital buffer, an instrument designed to protect the banking sector from periods of excessive growth associated with a build-up of system-wide risk. I construct an indicator of financial vulnerability with a model of fire sales, which captures the spillover losses in the system caused by deleveraging and joint liquidation of illiquid assets. Using data on the U.S. bank holding companies, I show that the indicator is informative about the build-up of vulnerability and can be useful for setting the countercyclical capital buffer.

Yuliang Zhang
London School of Economics and Political Science
y.zhang300@lse.ac.uk

CP9

Mean-Variance Equilibria in Continuous Time

We revisit the classical topic of mean-variance equilibria in the setting of continuous time, where asset prices are driven by continuous semimartingales. We show that under mild assumptions, a mean-variance equilibrium corresponds to a quadratic equilibrium for different preference parameters. We then use this connection to study a fixed-point problem that establishes existence of mean-variance equilibria. Our results rely on fine properties of mean-variance hedging as well as a novel stability result for quadratic BSDEs. The talk is based on joint work with Christoph Czichowsky, Martin Herdegen and David Martins.

Nikolaos Constantinou
University of Warwick
Department of Statistics
n.constantinou@warwick.ac.uk

Christoph Czichowsky

London School of Economics and Political Science
c.czichowsky@lse.ac.uk

Martin Herdegen
University of Warwick
m.herdegen@warwick.ac.uk

David Martins
ETH Zürich
Department of Mathematics
davidptmartins@gmail.com

CP9

Optimal Control of Stochastic Delay Differential Equations and Applications to Portfolio Optimization and Optimal Advertising

Optimal control problems of Markovian stochastic differential equations (SDEs) have been deeply studied; however, many real-world applications in finance and economics necessitate the consideration of path-dependent dynamics. In this talk, we consider optimal control problems of (path-dependent) SDEs with delays. Examples we consider are portfolio optimization and optimal advertising problems with delays. The dynamic programming approach poses difficult challenges since the Hamilton-Jacobi-Bellman equation is a partial differential equation on a Hilbert space with an unbounded operator. We show how to overcome these to solve the control problem. The talk is based on: F. de Feo, S. Federico, A. Swiech, "Optimal control of stochastic delay differential equations and applications to path-dependent financial and economic models", SIAM J. Control Optim. 62 (2024) no. 3, 14901520 F. de Feo, "Stochastic optimal control problems with delays in the state and in the control via viscosity solutions and applications to optimal advertising and optimal investment problems", Decis. Econ. Finance (2024) 31 pp. F. de Feo, A. Swiech, "Optimal control of stochastic delay differential equations: Optimal feedback controls, J. Differ. Equ. 420 (2025) 450-508 F. de Feo, A. Swiech, L. Wessels, "Stochastic optimal control in Hilbert spaces: $C^{1,1}$-regularity of the value function and optimal synthesis via viscosity solutions", arXiv:2310.03181 (2023)

Filippo De Feo
Technische Universität Berlin
filippo.defeo@yahoo.it

CP9

Cardinality-Constrained Portfolio Optimization with Clustering

Cardinality-constrained portfolio optimization aims to determine optimal investment strategies while limiting the total number of stocks. However, the NP-hard nature of cardinality constraints makes this optimization problem intractable in high-dimensional settings. We introduce new portfolio optimization frameworks that use clustering methods to reduce dimensionality while enforcing asset limits within each cluster and optimizing allocation proportions to balance risk and return. We apply hierarchical clustering to group stocks by the correlations of their residual returns, to achieve internally similar clusters. By dividing the complex problem into smaller subproblems, we can solve each one individually and approximate the standard optimization in a reduced space. Our empirical results on the SP 500 show that the proposed methods are effective in portfolio management and outperform various benchmarks. Additionally, by optimizing the number of selected assets and the total cluster weights in each group, we simplify asset selection while achieving performance comparable to traditional mean-variance portfolios. This highlights the potential of strategic clustering as a more effective framework for large-scale portfolio management under real-world constraints.

Amirmohammad Ebrahimi
University of Florida
a.ebrahimi@ufl.edu

Hamed Amini
University of Florida
Department of Industrial and Systems Engineering
aminil@ufl.edu

Hongcheng Liu
Associate Professor
University of Florida
liu.h@ufl.edu

CP9

Kendall Correlation Coefficients for Portfolio Optimization

Markowitz's optimal portfolio relies on the accurate estimation of correlations between asset returns, a difficult problem when the number of observations is not much larger than the number of assets. Using powerful results from random matrix theory, several schemes have been developed to 'clean' the eigenvalues of empirical correlation matrices [1,2]. By contrast, the (in practice equally important) problem of correctly estimating the eigenvectors of the correlation matrix has received comparatively little attention. Here we discuss a class of correlation estimators generalizing Kendall's rank correlation coefficient [3] which improve the estimation of both eigenvalues and eigenvectors in data-poor regimes. Using both synthetic and real financial data, we show that these generalized correlation coefficients yield Markowitz portfolios with lower out-of-sample risk than those obtained with rotationally invariant estimators. Central to these results is a property shared by all Kendall-like estimators but not with classical correlation coefficients: zero eigenvalues only appear when the number of assets becomes proportional to the square of the number of data points. [1] J. Bun, J.-P. Bouchaud, and M. Potters, My beautiful laundrette: Cleaning correlation matrices for portfolio optimization (2016) [2] D. Bartz, Cross-validation based nonlinear shrinkage (2016) [3] H. E. Daniels, The Relation Between Measures of Correlation in the Universe of Sample Permutations (1944).

Tomas Espana
CentraleSupélec
tomas.espana@student-cs.fr

Victor LeCoz, Matteo Smerlak
Capital Fund Management
victor.lecoz@cfm.com, matteo.smerlak@cfm.com

CP9

Portfolio Selection in Contests

In an investment contest with incomplete information, a finite number of agents dynamically trade assets with idiosyncratic risk and are rewarded based on the relative ranking of their terminal portfolio values. We explicitly characterize a symmetric Nash equilibrium of the contest and rigorously verify its uniqueness. The connection between the reward structure and the agents portfolio strategies is examined. A top-heavy payout rule results in an equilibrium portfolio return distribution with high positive skewness, which suffers from a large likelihood of poor performance. Risky asset holding increases when competition intensifies in a winner-takes-all contest.

Yumin Lu, Alex Tse
University College London
ucahyyl@ucl.ac.uk, alex.tse@ucl.ac.uk

CP10

Ambiguity Meets Merton: A Segmented Approach to Risk Premiums in Portfolio Theory

This study reimagines Mertons continuous-time portfolio optimization framework by incorporating segmented and discontinuous risk premiums that adapt to varying levels of market uncertainty and ambiguity. The model introduces thresholds z_1 and z_2, which segment uncertainty into low, medium, and high ranges, each governed by distinct risk premiums. These segments capture abrupt changes in risk tolerance and premium demands driven by ambiguity aversion. Optimal portfolio strategies are derived using the Hamilton-Jacobi-Bellman (HJB) equation, illustrating how ambiguity-averse investors adjust allocations across segmented uncertainty levels. By incorporating ambiguity into risk premium modeling, the study reveals significant deviations from Mertons original solution, demonstrating how ambiguity reshapes dynamic risk responses and portfolio strategies. Furthermore, the framework examines the sensitivity of expected utility to deviations in risk premiums, offering insights into how small shifts in market conditions or investor expectations influence portfolio decisions. It emphasizes the critical role of ambiguity in shaping investment behavior, delivering practical insights for portfolio optimization and financial policy design. By addressing the segmented and dynamic nature of risk premiums, this work advances both theoretical and applied perspectives on managing risk and ambiguity in financial markets.

Inhwa Kim
Harris-Stowe State University
ihyahootn@gmail.com

CP10

Equilibrium Asset Pricing with Epstein-Zin Stochastic Differential Utility

We revisit the classical problem of equilibrium asset pricing in a continuous-time complete-market setting, but in the case where investors' preferences are described by Epstein-Zin Stochastic Differential Utility. The market is comprised of a riskless bond and a risky asset, where the latter pays continuously a stochastic dividend stream. The equilibrium is characterised by a system of strongly coupled Forward-Backward Stochastic Differential Equations (FBSDEs). This is joint work in progress with Dr. Martin Herdegen.

Andreea Popescu, Martin Herdegen
University of Warwick
andreea.popescu@warwick.ac.uk,
m.herdegen@warwick.ac.uk

CP10

A New Approach For The Kyle-Back Strategic Insider Equilibrium Problem

Kyle-Back (KB) model studies the price impact of a strategic inside trader. It has been widely studied by math finance community. Most of the works assume certain Markovian structure for pricing rule of the market maker in order to use the PDE approach. We provide a new formulation for the KB model that does not assume Markovian structure for pricing rule. We find an example such that only non-Markovian pricing rule exists for equilibrium. Under our new formulation, when insiders signal follows a discrete distribution, the equilibrium of the KB model are characterized by the solutions of a Forward-Backward Stochastic Differential Equations. We further study the set value of the insider's payoff. By applying the so-called duality property of the set value, we reformulate KB model to an optimal control problem thus allow us to further apply the rich results in optimal control theory. The well known bridge strategy with linear pricing function is an element of our set value.

Bixing Qiao, Jianfeng Zhang
University of Southern California
bqiao@usc.edu, jianfenz@usc.edu

CP10

Bio-Value-at-Risk: A Concept to Assessing the Implications of Biodiversity Risks on Portfolio Management using Geospatial Analysis

Nature-related risks, such as biodiversity loss and carbon footprints, have finally been recognized as critical business risks, yet their integration into corporate risk management frameworks, valuations, and portfolio management remains challenging. While sustainable finance initiatives and regulatory frameworks increasingly mandate transparent reporting of nature-related risks and impacts, existing approaches often rely on top-down estimations of global exposures, limiting their precision and practical applicability. This study introduces a bottom-up, scenario-based methodology to quantify companies impacts on nature by leveraging geospatial analysis. By overlaying geospatial biodiversity variables derived from global remote sensing data with spatial footprints of asset-level operational data, we assess localized nature impacts attributable to companies physical assets. These impacts are then translated into financial loss estimates at the site level and aggregated at the company level. Unlike traditional approaches, this method enables a more granular and consistent evaluation of financial and environmental risks, facilitating portfolio risk management through targeted investment and credit decisions rather than wholesale industry exclusions. To demonstrate the combined environmental and financial materiality of this approach, we apply it to the mining sector as there is significant scientific literature on its environmental impacts and on remote sensing to detect these impacts.

Peter Schwendner, Jan-Alexander Posth, Patrick Laube, Tomasz Orpiszewski
Zurich University of Applied Sciences
scwp@zhaw.ch, posh@zhaw.ch, laup@zhaw.ch, orpi@zhaw.ch

CP11

Mean Field Analysis for Mitigating Overlearning in Stochastic Control

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Boris Baros
Oxford University
boris.baros@gmail.com

CP11

Spanning Multi-Asset Payoffs With Relus

We propose a distributional formulation of the spanning problem of a multiasset payoff by vanilla basket options. This problem is shown to have a unique solution if and only if the payoff function is even and absolutely homogeneous, and we establish a Fourier-based formula to calculate the solution. Financial payoffs are typically piecewise linear, resulting in a solution that may be derived explicitly, yet may also be hard to exploit numerically. One-hidden-layer feedforward neural networks instead provide a natural and efficient numerical alternative for discrete spanning. We test this approach for a selection of archetypal payoffs and obtain better hedging results with vanilla basket options compared to industry-favored approaches based on single-asset vanilla hedges. (Joint work to appear in *Mathematical Finance*)

Sebastien Bossu
New York University
sbossu@charlotte.edu

Stephane Crepey, Hoang-Dung Nguyen
Universite Paris-Cite
stephane.crepey@lpsm.paris, hdnguyen@lpsm.paris

CP11

Graph Analytics and Machine Learning for Anomaly Detection in Financial Transactions

The accuracy of classification algorithms in detecting fraudulent financial activity is critical in assisting human analysts in the task of preventing financial crime. Graph-based anomaly detection methods represent the state-of-the-art in analyzing connectivity patterns within monetary transaction networks and identifying suspicious behaviors. This talk will present algorithms designed to detect anomalies using graph partitioning methods for highly imbalanced datasets, which provide fully interpretable results, and Graph Neural Networks that incorporate the latest research conducted in message passing architectures. The effectiveness of these approaches is demonstrated in experiments with data that simulate real-world financial behaviour, and are infused with a variety of anomalous money laundering topologies.

Dimosthenis Pasadakis
Università della Svizzera italiana
dimosthenis.pasadakis@usi.ch

Madan Sathe
Partner, Forensics, Deloitte AG, Zurich
msathe@deloitte.ch

Olaf Schenk
Università della Svizzera italiana
Switzerland
olaf.schenk@usi.ch

CP11

Deep Dual Optimal Stopping: Addressing the Curse of Dimensionality

We propose a novel deep-learning approach to the duality of optimal stopping problems and provide a tight upper bound for the value function in a possibly high-dimensional setting. This deep dual optimal stopping (DDOS) framework is supported with theoretical grounds, including: (1) an approximate computational scheme in a continuous-time setting, (2) convergence and expressivity guarantees under structural conditions, and (3) a theoretical explanation for the variance reduction of the deep upper bound. Our framework features an independent primal-dual algorithm, which solves the dual optimal stopping problem without incurring the value function. Our numerical experiments demonstrate the effectiveness of our approach in terms of convergence, expressivity, and stability performance.

Junyan Ye
The Chinese University of Hong Kong
junyanye@link.cuhk.edu.hk

Hoi Ying Wong
Department of Statistics,
The Chinese University of Hong Kong
hywong@cuhk.edu.hk

CP12

Optimal Income Drawdown and Investment with Longevity Basis Risk

We investigate a utility maximisation problem for a pension scheme which offers an income drawdown policy. Apart from market risk, we specifically focus on longevity basis risk which arises when the forces of mortality of the schemes target population and the reference population of a longevity bond used for hedging, are not perfectly correlated. By modelling the forces of mortality of the reference and target populations as stochastic affine processes, we derive analytic solutions for the relevant investment strategy and benefit withdrawal rate. Our model also accounts for dependencies between mortality rate fluctuations and stock prices. Extensive numerical results demonstrate that the longevity bond acts as an effective hedging instrument, even in the presence of longevity basis risk. Our work adds a new dimension to the management of pension schemes by considering the interaction between financial markets and demographic risks. The results show that when stock price movements are positively correlated with mortality rate changes, the stock can provide some longevity risk hedging and offer a positive longevity risk premium. On the other hand, when stock price movements are negatively correlated with mortality rate changes, it is optimal to invest less weight of the fund in the stock, as investing in the stock under these circumstances can decrease welfare.

Ankush Agarwal
University of Western Ontario
aagarw93@uwo.ca

Christian Ewald
Umea University
christian.ewald@umu.se

Yongjie Wang

University of Glasgow
yongjie.lizzy.wang@gmail.com

CP12

Revisiting Mertons Model: A Dual Calibration Approach Using Black-Scholes-Merton and Binomial Trees

We explore credit risk pricing by modeling equity as a call option and debt as the difference between the bond price and a put option on the asset value, following the framework of the Merton model. Our approach involves two key steps: computing implied volatility and calibrating the probability and drift parameters. Using the classical Merton credit risk model, which employs the Black-Scholes-Merton (BSM) option pricing formula, we calibrate the volatility parameter within a continuous-time framework for both asset and equity. To calibrate the mean and probability surfaces, we introduce a recombining binomial tree in the real (natural) world, treating the initial asset price as a fixed value. The volatility used is the calibrated asset volatility for a specific level of moneyness and across all maturities, which then informs the implied mean return and implied probability. A unique relationship is developed between the risk-neutral and natural-world parameters, enabling us to construct implied probability and implied drift surfaces for practical applications.

Jagdish Gnawali
Texas Tech University
jgnawali@ttu.edu

Abootaleb Shirvani
Kean University
ashirvan@kean.edu

Svetlozar Rachev
Texas Tech University
zari.rachev@ttu.edu

Frank Fabozzi
Carey Business School, Johns Hopkins University
fabozzi321@aol.com

CP12

Multi-Task Dynamic Pricing in Credit Market with Contextual Information

We study the dynamic pricing problem faced by a broker seeking to learn prices for a large number of credit market securities, such as corporate bonds, government bonds, loans, and other credit-related securities. A major challenge in pricing these securities stems from their infrequent trading and the lack of transparency in over-the-counter (OTC) markets, which leads to insufficient data for individual pricing. Nevertheless, many securities share structural similarities that can be exploited. Leveraging these insights, we propose a multi-task dynamic pricing framework that leverages the shared structure across securities to enhance pricing accuracy. In the OTC market, a broker wins a quote by offering a more competitive price than rivals. The broker's goal is to learn winning prices while minimizing expected regret against a clairvoyant benchmark. We model each security using a d-dimensional feature vector and assume a linear contextual model for the competitor's pricing, with parameters unknown a priori. We propose the Two-Stage Multi-Task (TSMT) algorithm: first, an unregularized MLE over pooled data to obtain a coarse parameter estimate; second, a regularized MLE on individual securities to refine the parameters. We show that the TSMT achieves a regret bounded by $\tilde{O}(\delta_{\max}\sqrt{TMd}+Md)$, outperforming both fully individual and fully pooled baselines, where M is the number of securities and $\delta_{\max}$ quantifies their heterogeneity.

Adel Javanmard
Marshall School of Business
University of Southern California,
javanma@usc.edu

Jingwei Ji
University of Southern California
jingweij@usc.edu

Renyuan Xu
Department of Finance and Risk Engineering
New York University
x2364@nyu.edu

CP12

Isogeometric Analysis for the Pricing of Financial Derivatives with Nonlinear Models: Convertible Bonds and Options

Computational efficiency is essential for enhancing the accuracy and practicality of pricing complex financial derivatives. In this paper, we discuss Isogeometric Analysis (IGA) for valuing financial derivatives, modeled by two nonlinear Black-Scholes PDEs: the Leland model for European call with transaction costs and the AFV model for convertible bonds with default options. We compare the solutions of IGA with finite difference methods (FDM) and finite element methods (FEM). In particular, very accurate solutions can be numerically calculated on far less mesh (knots) than FDM or FEM, by using non-uniform knots and weighted cubic NURBS, which in turn reduces the computational time significantly.

Rakhymzhan Kazbek
Nazarbayev University
r.kazbek@astanait.edu.kz

CP12

Efficient Calibration of the Shifted Square-Root Diffusion Model to Credit Default Swap Spreads Using Asymptotic Approximations

We derive a closed-form approximation for the credit default swap (CDS) spread in the two-dimensional shifted square-root diffusion (SSRD) model for its efficient calibration, using asymptotic coefficient expansion technique to approximate solutions of nonlinear partial differential equations. Specifically, we identify the Cauchy problems associated with two terms in the CDS spread formula that lack analytical solutions and derive asymptotic approximations for these terms. Our approximation does not require the assumption of uncorrelated interest rate and default intensity processes as typically required for calibration in the SSRD model. Through several numerical experiments using market data on CDS spreads, we demonstrate the accuracy and efficiency of our proposed formula when calibrating the SSRD model.

Ying Liao
University of Glasgow
2636546L@student.gla.ac.uk

Ankush Agarwal
University of Western Ontario
aagarw93@uwo.ca

CP13

Synthetic Data For Portfolios: A Throw of the Dice Will Never Abolish Chance

Simulation methods have always been instrumental in finance, and data-driven methods with minimal model specification – commonly referred to as generative models – have attracted increasing attention, especially after the success of deep learning in a broad range of fields. However, the adoption of these models in financial applications has not kept pace with the growing interest, probably due to the unique complexities and challenges of financial markets. This paper aims to contribute to a deeper understanding of the limitations of generative models, particularly in portfolio and risk management. To this end, we begin by presenting theoretical results on the importance of initial sample size, and point out the potential pitfalls of generating far more data than originally available. We then highlight the inseparable nature of model development and the desired use case by touching on a paradox: generic generative models inherently care less about what is important for constructing portfolios (in particular the long-short ones). Based on these findings, we propose a generative pipeline for multivariate returns tested on a large universe of US equities. Moreover, we insist on the need for more delicate evaluation methods, and suggest, through an example of mean-reversion strategies, a method designed to identify poor models for a given application based on regurgitative training, i.e. retraining the model using the data it has itself generated.

Adil Rengim Cetingoz
Université Paris 1 Panthéon-Sorbonne
Pantheon-Sorbonne University
adil-rengim.cetingoz@etu.univ-paris1.fr

Charles-Albert Lehalle
Ecole Polytechnique
charles-albert.lehalle@polytechnique.edu

CP13

What Drives Corporate Savings?

Why do some firms save more than others? We revisit this important, widely studied question by developing a tractable continuous-time model for financially constrained firms facing costly external equity. In addition to including the standard building blocks for a q theory of investment (Hayashi 1982; Abel and Eberly 1994): capital accumulation, capital adjustment costs, and a persistent stochastic productivity process, we assume that instantaneous profits (cashflows generated by productive capital) are stochastic. This assumption is a key difference between our model and widely used quantitative corporate finance models, e.g., Hennessy and Whited (2007) and Riddick and Whited (2009). Firm value V depends on capital stock K, cash balance W, and productivity A. Consider the alternative that instantaneous profits are (locally) deterministic. Then the firms cash balance at the end of the period would also be deterministic. This substantially reduces its precautionary savings demand, as locally the risk-neutral firm (as long as it is not too cash strapped) faces little liquidation risk and hence holding cash has limited value. Mathematically, whether (conditional) instantaneous profits are stochastic implies whether the second-order derivative $\frac{\partial^2 V}{(\partial W)^2}$ term appears in the Hamilton-Jacobi-Bellman (HJB) equation, and economically this term has a first-order effect on precautionary savings demand.

Ling Qin
The Chinese University of Hong Kong
lingqin001@cuhk.edu.hk

Nan Chen
The Chinese University of Hong Kong
Hong Kong
nchen@se.cuhk.edu.hk

Xavier Giroud
Columbia University
xavier.giroud@gsb.columbia.edu

Neng Wang
Cheung Kong Graduate School of Business
nw2128@columbia.edu

CP13

A Unified Framework for Time-Series and Factor Momentum Strategies

Time-series momentum strategies adjust positions based on return trends but often degrade during rapid market shifts and suffer from crowding effects. This study addresses these issues by integrating time-series and factor momentum, leveraging persistence in factor to improve risk-adjusted returns. Our approach combines time-series and cross-sectional signals for a robust strategy. Time-series signals are processed using a masked multi-head attention mechanism to capture temporal dependencies, while cross-sectional signals are analyzed with Instrumented Principal Component Analysis (IPCA) to extract latent factors guiding asset-specific weights based on factor momentum. The strategy integrates time-series and factor momentum (TSFM) through a gating layer that combines outputs from the attention mechanism and IPCA. The unified return is:

$$r^{\text{TSFM}}_{t,t+1} = \frac{1}{N_t}\sum_{i=1}^{N_t} X^{(i)}_t \frac{\sigma_{tgt}}{\sigma^{(i)}_t} r^{(i)}_{t,t+1}, \quad X^{(i)}_t = f(u^{(i)}_t;\Theta),$$

where f is a neural network generating trading signals from input features $u^{(i)}_t$, with parameters Θ optimized to minimize loss. Backtesting on S&P 100 assets demonstrates that the proposed model effectively captures both time-series and factor momentum effects. Results show a 20% improvement in Sharpe ratios compared to traditional strategies, highlighting the models robust performance across different market states.

Jaemin Song, Eungpyo Kim
Department of Data Science, Chonnam National University
cantor@jnu.ac.kr, eungpyo0219@jnu.ac.kr

Jaegi Jeon
Chonnam National University, Department of Data Science
jaegijeon@jnu.ac.kr

CP13

Trading Perfect Risk Budgeting for Portfolio Re-

turns

The Markowitz portfolio problem aims to find the portfolio weights that maximize the expected return while minimizing the portfolio volatility. Risk budgeting is a portfolio strategy where each asset contributes a prespecified amount to the aggregate risk of the portfolio. In this work, we propose new optimization models that induce a compromise between these two strategies, and in particular allow to relax the strict risk budget allocation to achieve higher expected portfolio returns.

Rodrigo Targino
School of Applied Mathematics, Getulio Vargas Foundation
Brazil
rodrigo.targino@fgv.br

CP13

Chesformer: Quantitative Financial Factor Model with Chebyshev Positioning

A large number of quantitative factors are designed with mathematical formulas, which aim to catch market sentiment by using data on stocks. In recent years, some machine learning methods have tried to catch the connection between long and high-dimension input stock data sequences. In this paper, a new transformer-based model named Chesformer is introduced. Chesformer contains a new Chebyshev positioning module instead of the initial sin and cosine function positioning to better capture the relationship with different time series. Based on three and half year backtest, the factors based on various training scales show robustness performance in balance return and risk. Chesformer gives a novel idea on catch market sentiment by finding connections in stocks' complex data.

Zhaofeng Zhang
University of Michigan
Department of Mathematics
zhangzf@umich.edu

MS1

Designing Automated Market Makers for Derivative Contracts

Prediction markets allow traders to bet on potential future outcomes. These markets exist for weather, political, sports, and economic forecasting. Financial derivatives can be viewed as a kind of prediction on the outcome of some underlying assets. With this viewpoint, a derivatives market is a type of prediction market that permits hedging via the underlying. Within this talk, we consider a decentralized framework for prediction markets using automated market makers (AMMs) for derivatives contracts.

Zachary Feinstein
Stevens Institute of Technology
zfeinste@stevens.edu

MS2

Transfer Learning for Functional Linear Models

Transfer learning (TL) is a powerful paradigm in machine learning and data science, allowing models trained on a source domain to enhance estimation and prediction in a related but data-scarce target domain. This talk presents a transfer learning framework for functional linear models. First, we introduce an interpretable similarity measure between domains based on reproducing kernel Hilbert space (RKHS). It connects the type of transferred knowledge to the employed RKHS, allowing flexibility in customizing the specific knowledge to be transferred for the coefficient functions. Second, building on the offset transfer learning paradigm, two algorithms are proposed: one conducts the transfer when positive sources are known, while the other leverages aggregation techniques to achieve robust transfer without prior information about the sources. We establish asymptotic lower bounds for this learning problem and show the proposed algorithms enjoy a matching upper bound. These analyses provide statistical insights into factors that contribute to the dynamics of the transfer. The effectiveness of the proposed algorithms is demonstrated via a stock return prediction application.

Haotian Lin
Department of Statistics
Pennsylvania State University
hzl435@psu.edu

MS2

Anomaly Detection on Time Series with Signatures

We present SigMahaKNN, a novel pipeline designed for scoring anomalous streams based on clean corpus anomaly detection. The conformance scores of new samples are derived from a combination of path signatures with the Mahalanobis distance. We show that such an approach preserves desirable invariances, namely to affine transformations of the data and appending metadata. Our pipeline is versatile and can be used with a wide array of non-stationary, multi-modal tick data with complex missingness patterns. We showcase its effectiveness in detecting time series obtained from deep generative models for financial time series and voice cloning.

Lingyi Yang
University of Oxford
lingyi.yang@maths.ox.ac.uk

MS2

Modeling Probability Density Functions As Data Objects

Analyzing distributions as data objects leads to models that are adaptive and coherent to various functional forms of distributions. Recent developments in distributional modeling are reviewed with a particular focus on the probabilistic and statistical analysis of probability density functions. Density functions are treated as data objects for which suitable notions of the center of distribution and variability are discussed. Special attention is given to nonlinear methods that respect the constraints density functions must obey. Regression, time series and spatial models are discussed. The exposition is illustrated with a data example of forecasting corporate default probabilities.

Chao Zhang
Quant Risk Analytics
Bloomberg
czhang867@bloomberg.net

MS3

Sharp Propagation of Chaos for Mean Field Langevin Dynamics and Mean Field Games

We establish the sharp rate of propagation of chaos for

McKean–Vlasov equations with coefficients that are nonlinear in the measure argument; we then apply our results to Wasserstein gradient flows and mean field games. Our arguments combine a version of the BBGKY hierarchy with ideas from the literature on weak propagation of chaos and analysis on the space of measures.

Manuel Arnese
Columbia University
ma4339@columbia.edu

MS3

Interacting Particle Systems on Sparse W-Random Graphs

We consider a general interacting particle system with interactions on a random graph and study the large population limit of this system. When the sequence of underlying graphs converges to a graphon, we show convergence of the interacting particle system to a so called graphon stochastic differential equation. This is a system of uncountable many SDEs of McKean-Vlasov type driven by a continuum of Brownian motions. We make sense of this equation in a way that retains joint measurability and essentially pairwise independence of the driving Brownian motions of the system by using the framework of Fubini extensions. The convergence result is general enough to cover nonlinear interaction, as well as various examples of sparse graphs. Moreover, we extend the results to the unbounded graphon case.

Carla Crucianelli
Princeton University
cc1197@princeton.edu

MS3

Mean-field Games with Partial Information and Relative Performance Concerns

I will introduce a mean-field game with partial information for optimal choice under relative performance concerns. Both utilities and coupling functions are general. I will derive the master equation and provide suitable solutions for both the value of the game and the optimal processes. Representative examples will be presented. (Joint work with P.Souganidis.)

Thaleia Zariphopoulou
The University of Texas at Austin
zariphop@math.utexas.edu

MS4

Memory in Quantitative Finance: a Helicopter Tour

Serving as a kickoff for the minisymposium "Memory in Quantitative Finance," this talk provides an accessible overview of memory-aware tools and their applications in financial markets. We will introduce Volterra processes, continuous-time analogs of weighted moving averages, and path signatures, sequences of iterated integrals, and explore their roles in capturing memory effects. These approaches provide powerful mathematical and numerical tools for addressing challenges in volatility modeling, derivative pricing, portfolio optimization, and risk management in path-dependent settings. The talk will highlight key theoretical insights, practical applications, and open challenges, setting the stage for deeper discussions throughout the minisymposium.

Eduardo Abi Jaber
Ecole Polytechnique
eduardo.abi-jaber@polytechnique.edu

Christian Bayer
Weierstrass Institute
christian.bayer@wias-berlin.de

MS4

Stochastic Volterra Equations on Convex Domains

We will present in this talk sufficient conditions on the kernel and on the coefficients to get the existence of a solution that stays in a convex domain. The underlying tool is an approximation scheme that also stays in this domain. Applications include: a comparison result for scalar SVEs, existence of solutions possibly with a jump component, weak second-order approximation schemes for SVEs with multifactor kernels such as the multi-factor approximation of the rough Heston model.

Aurelien Alfonsi
CERMICS, Ecole Nationale des Ponts et Chaussees
aurelien.alfonsi@enpc.fr

MS4

Volterra Stein-Stein Models and $\sqrt{\text{Fredholm Determinants}}$

This talk focuses on the Volterra Stein-Stein model, for which a formula for the characteristic function of the joint log price and integrated variance was derived in [Abi Jaber, Finance and Stochastics, 2022]. Complex discontinuities, arising from the mismanagement of branch cuts in complex square roots and logarithms, can lead to numerical inaccuracies and instability in Fourier-based pricing methods. Building on this, we present an alternative formula expressed in terms of the complex square root of a Fredholm determinant, allowing for significantly more accuracy and improved computational efficiency in the pricing of derivatives by Fourier inversion techniques. In the complex domain, the determinant may cross the negative real axis, which is handled by our formula by counting the number of times it does. However, this number is a priori unknown. Therefore, we address two key problems: (1) determining whether such crossings occur and, if so, where; and (2) proposing efficient numerical methods to compute the number of crossings. The implications of these results are demonstrated through comparisons of accuracy and computational efficiency between competing approaches.

Maxime Guellil
Ecole Polytechnique, France
maximeguellil@gmail.com

Eduardo Abi Jaber
Ecole Polytechnique
eduardo.abi-jaber@polytechnique.edu

MS5

Mean Field Games of Controlled Jump Intensity, and Applications to Cryptocurrency Mining Competition Models

In this talk, we discuss Nash equilibrium existence results for Mean Field Games (MFG) with dynamics given by jump processes of controlled intensity, with mean field interactions occurring via the controls. As part of the analysis, we establish a convergence theorem showing that discrete-time MFG equilibria converge to their continuous-time counterparts, motivating time-discretization-based numerical schemes for such models. We then apply our results to a continuous-time model for cryptocurrency mining. A key mechanism that maintains the integrity of major cryptocurrencies is the presence of miners, who compete to solve sequential brute-force computational challenges and receive payouts when successful. Mining is computationallyand thus energyintensive, requiring miners to trade off the cost of high electricity consumption against the probability of receiving a payout. These competition dynamics are captured in a crypto mining MFG model introduced by Z. Li, A. M. Reppen, and R. Sircar. We provide equilibrium existence results for this model under general wealth utility functions, and demonstrate qualitatively similar numerical solutions for the corresponding discrete-time version. Under power utility, we observe preferential attachment behavior: over time, a small group of wealthy miners consume large amounts of electricity and come to dominate the game, while most miners exit once their wealth falls below a threshold.

Nicolas Garcia
Princeton University
ng6303@princeton.edu

MS5
Optimal Data-Driven Trading of Renewables in Intraday Energy Markets

The integration of weather-dependent renewable energy sources into the European power system amplifies imbalance risks and intraday price volatility, necessitating advanced tools to effectively mitigate these risks. In this work, we introduce a data-driven, continuous-time stochastic optimal control framework tailored for intraday trading on continuous exchanges. Our approach incorporates deterministic forecasts of wind generation and prices into an It jump-diffusion model. By formulating and solving the corresponding HamiltonJacobiBellman partial integro-differential equation, we determine optimal trading strategies aimed at maximizing expected revenues while limiting liquidity risks and imbalance penalties. Numerical experiments conducted under realistic market conditions demonstrate the robustness and profitability of our proposed framework.

Michael Samet
RWTH Aachen
samet@uq.rwth-aachen.de

Chiheb Ben Hammouda
Utrecht University
c.benhammouda@uu.nl

Raul Tempone
KAUST
RWTH Aachen University
raul.tempone@kaust.edu.sa

MS6
Entropic Adapted Wasserstein Distance on Gaussians

The adapted Wasserstein distance is a metric for quantifying distributional uncertainty and assessing the sensitivity of stochastic optimization problems on time series data. A computationally efficient alternative to it, is provided by the entropically regularized AW-distance. Suffering from similar shortcomings as classical optimal transport, there are only few explicitly known solutions to those distances. Recently, GunasingamWong provided a closed-form representation of the AW-distance between real-valued stochastic processes with Gaussian laws. We further extend their work in two directions, by considering multidimensional (Rd-valued) stochastic processes with Gaussian laws and including the entropic regularization. In both settings, we provide closed-form solutions and characterize Gaussian optimizers. The talk is based on joint work with Beatrice Acciaio and Gudmund Pammer.

Songyan Hou, Beatrice Acciaio
ETH Zurich
songyan.hou@math.ethz.ch,
beatrice.acciaio@math.ethz.ch

Gudmund Pammer
TU Graz
gudmund.pammer@tugraz.at

MS6
Causal Transport on Path Space

Causal optimal transport and its induced Wasserstein-type distance on the space of probability measures on path space has been the focus of much research recently, owing to its efficacy in handling dynamic optimization problems. We investigate these (bi-)causal transports between weak solutions of stochastic differential equations by providing a full characterization of these (bi-)causal Monge maps and couplings. For example, we show that bicausal Monge maps of d-dimensional Wiener measures are actually stochastic integrals of rotation-valued integrands. As an application, we give necessary and sufficient conditions for bicausal couplings to be induced by Monge maps and show that such bicausal Monge transports are dense in the set of bicausal couplings between laws of SDEs with regular coefficients.

Fang Rui Lim, Rama Cont
University of Oxford
fang.lim@maths.ox.ac.uk, rama.cont@maths.ox.ac.uk

MS6
Johannes Wiesel -Tbd

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Johannes Wiesel
Carnegie Mellon University
wiesel@cmu.edu

MS8
Diffusion Factor Models: Generating High-Dimensional Returns with Factor Structure

Financial scenario simulation is crucial for risk management, portfolio optimization, and stress testing, yet it remains challenging under small-data constraints. Effective simulation methods must integrate economic structure with advanced generative learning techniques to ensure robustness and interpretability. We propose a principled approach for financial scenario simulation based on diffusion model that incorporates factor structure in data. Grounded in asset pricing theory yet novel in the generative diffusion model literature, our method addresses the challenges of the curse of dimensionality and data scarcity in financial markets. More specifically, by leveraging the factor structure, we decompose the score function—a key component in diffusion models—using time-varying orthogonal projections. Incorporating this decomposition into the design of neural network architectures, we establish non-asymptotic error bounds for score estimation and generated distribution, which surpass the dimension-dependent limits in the classical diffusion model literature. Simulation studies show that our method outperforms classical methods in recovering the latent subspace in small-data regimes. We empirically demonstrate the economic significance of our framework in constructing mean-variance optimal portfolios and factor portfolios. This work presents the first theoretical integration of factor models with diffusion models.

Minshuo Chen
Princeton University
minshuochen@princeton.edu

Renyuan Xu
New York University
renyuanx@usc.edu

Yumin Xu, Ruixun Zhang
Peking University
2301110088@pku.edu.cn, zhangruixun@pku.edu.cn

MS8

Learning Mean-Field Optimal Control via Neural Networks

In this talk, we discuss our recent development of neural network-based learning algorithms designed to efficiently solve mean-field optimal control problems. Our approach leverages particle-based approximations and the inherent symmetry among agents to achieve computational tractability even in high-dimensional state spaces. Although our methodology does not directly employ reinforcement learning techniques, we compare and contrast our mean-field approach to traditional reinforcement learning methods, highlighting how the mean-field perspective naturally captures interactions among large populations of agents and facilitates understanding collective behaviors such as algorithmic collusion or manipulation. Through rigorous convergence analysis and illustrative numerical experiments, we demonstrate that our method effectively approximates optimal strategies. Additionally, we present a universal approximation result in Wasserstein spaces, reinforcing the robustness and general applicability of our neural network-based framework.

Qinxin Yan
Princeton University
qy3953@princeton.edu

MS9

Nash Allocations for Systemic Risk Measures

Systemic risk measures aggregate the risks from multiple financial institutions to find system-wide capital requirements. Though much attention has been given to assessing the level of systemic risk, less has been given to allocating that risk to the constituent institu- tions. We propose a Nash allocation rule that is inspired by game theory. Intuitively, to construct these capital allocations, the banks compete in a game to reduce their own capital requirements while, simultaneously, maintaining system-level acceptability. Existence and uniqueness for these Nash allocations are provided under different systemic risk settings. We demonstrate the efficacy of these Nash allocations with numerical case studies.

agin Ararat
Bilkent University
cararat@bilkent.edu.tr

Zachary Feinstein
Stevens Institute of Technology
zfeinste@stevens.edu

MS9

Set Valued PDEs and Games

Building upon the dynamic programming principle for set valued functions arising from many applications, we will present a new notion of set valued PDEs. The key component is a set valued Itô formula, characterizing the flows on the surface of the dynamic sets. In the contexts of multivariate control problems, we establish the wellposedness of the set valued HJB equations, which extends the standard HJB equations in the scalar case to the multivariate case. As an application, we discuss moving scalarization, constructed using the classical solution of the set valued HJB equation. Additionally, we introduce the concept of set values for games under Nash equilibrium, along with the corresponding PDE, and explore its geometric properties. The talk is based on joint work with Jianfeng Zhang and an ongoing work joint with Nizar Touzi and Jianfeng Zhang.

Melih Iseri
University of Michigan
iseri@umich.edu

MS9

Computing Systemic Risk Measures with Graph Neural Networks

This paper investigates systemic risk measures for stochastic financial networks of explicitly modelled bilateral liabilities. We extend the notion of systemic risk measures from Biagini, Fouque, Fritelli and Meyer-Brandis (2019) to graph structured data. In particular, we focus on an aggregation function that is derived from a market clearing algorithm proposed by Eisenberg and Noe (2001). In this setting, we show the existence of an optimal random allocation that distributes the overall minimal bailout capital and secures the network. We study numerical methods for the approximation of systemic risk and optimal random allocations. Our proposition is to use permutation equivariant architectures of neural networks such as graph neural networks (GNNs) and a class that we name (extended) permutation equivariant neural networks ((X)PENNs). The performance of these architectures is benchmarked against several alternative allocation methods. The main feature of GNNs and (X)PENNs is that they are permutation equivariant with respect to the underlying graph data. In numerical experiments, we find evidence that these permutation equivariant methods are superior to other approaches.

Lukas Gonon
Imperial College London

l.gonon@imperial.ac.uk

Thilo Meyer-Brandis
University of Munich
Germany
meyerbra@math.lmu.de

Niklas Weber
University of Munich
weber@math.lmu.de

MS9

Consumption–Investment Problems Under Time-Varying Incomplete Preferences

I will discuss the martingale solution to optimal consumption–investment problems in a multi-asset, multi-commodity market where investors have time-varying incomplete preferences represented by multi-utility functions. A set-valued stochastic process is introduced for the dynamics of multi-utility indices, and utility maximization becomes a multi-criteria problem with possibly function-valued criteria. Optimal consumption policies are solved with a modified linear scalarization technique allowing for infinite stochastic dimensions. Using Malliavin calculus with stochastic geometry, optimal investment policies are found to be generally set-valued, each of whose selectors admits a four-way decomposition involving an additional indecisiveness risk-hedging portfolio. This study touches on new directions for optimal consumption–investment choices in the presence of incomparability and time variation, signaling potentially testable assumptions on the variability of asset prices. Simulation methods for set-valued processes are studied for how solved optimal policies can be computed in practice.

Weixuan Xia
University of Southern California
weixuanx@usc.edu

MS11

From Online Vector Balancing to Mean Field Stochastic Control

Online vector balancing problems (OVBPs) are a central task in discrepancy theory. Given a finite stream of vectors $v_1, \ldots, v_n \in \mathbb{R}^d$, the goal is to adaptively choose signs $\varepsilon_k \in \pm 1$ so that the signed sum $\sum_{k=1}^{n} \varepsilon_k v_k$ is small in every coordinate. Such problems have been studied extensively from an algorithmic point of view. We examine the OVBP from the new perspective of stochastic control. Under i.i.d. Gaussian inputs $v_k \sim \mathcal{N}(0, I_d)$ and under the standard scaling $n/d \to T \in (0, \infty)$, we conjecture that the mean-field scaling limit is characterized by a continuous-time stochastic control problem of novel type. In it, one steers a Brownian motion with an L^2 constrained drift to minimize the L^∞ norm of the state at the terminal time T. We provide heuristics supporting this convergence and partial results that establish the value of the limiting control problem as a lower bound for the asymptotic value of the OVBP. A full proof of the conjecture would yield the exact asymptotic value of the OVBP and suggest new avenues for algorithm design based on stochastic control.

Christian Fiedler
Columbia University
christian.fiedler@columbia.edu

MS11

Synchronization Games

We propose a new mean-field game model with two states to study synchronization phenomena, and we provide a comprehensive characterization of stationary and dynamic equilibria along with their stability properties. The game undergoes a phase transition with increasing interaction strength. In the subcritical regime, the uniform distribution, representing incoherence, is the unique and stable stationary equilibrium. Above the critical interaction threshold, the uniform equilibrium becomes unstable and there is a multiplicity of stationary equilibria that are self-organizing. Under a discounted cost, dynamic equilibria spiral around the uniform distribution before converging to the self-organizing equilibria. With an ergodic cost, however, unexpected periodic equilibria around the uniform distribution emerge.

Felix Höfer
Princeton University
fh6011@princeton.edu

MS11

Mean Field Convergence Via Hamilton-Jacobi Equations on the Wasserstein Space

The study of differential equations on the Wasserstein space of probability measures has become an active area of research in recent years. One motivating factor is that many relevant observable quantities in interacting agent systems have mean field approximations that solve such equations, as has been borne seen in mean field games/control theory as well as applications in statistical physics. A well-posedness theory for these equations can thus lead to new methods for studying such mean field models. In this talk, we present some well-posedness results for Hamilton-Jacobi equaitons on Wasserstein space, and explain how these results lead to qualitative and quantitative mean field convergence results. This is joint work with Samuel Daudin (Universit Paris-Cit) and Joe Jackson (University of Chicago).

Benjamin Seeger
UNC Chapel Hill
bseeger@unc.edu

MS11

Mean Field Games with Common Noise Via Malliavin Calculus

We present a simpler proof of the existence of equilibria for a class of mean field games with common noise, where players interact through the conditional law given the current value of the common noise rather than its entire path. By extending a compactness criterion for Malliavin-di?erentiable random variables to processes, we establish existence of strong equilibria, where the conditional law and optimal control are adapted to the common noise filtration and defined on the original probability space. Notably, our approach only requires measurability of the drift and cost functionals with respect to the state variable. Based on a joint work with S. Wang.

Ludovic Tangpi
Princeton University

ludovic.tangpi@princeton.edu

MS12

Path-Dependent Ruin and Dividend Problems in a Spectrally Negative Lvy Model

We study a version of De Finetti's optimal dividend problem driven by a spectrally negative Lvy process, where the bankruptcy time is path dependent in the sense that we consider a function $\omega(x)$ such that $\omega(x)dt$ represents the intensity of bankruptcy. We approach this problem using the generalized scale functions introduced in Li & Palmowski (2018) to derive semi-explicit expressions for the quantities involved. Our goal is to prove that an optimal dividend strategy exists, and that it lies within the set of barrier strategies. This is joint work with Jean-Franois Renaud (UQAM).

Dante Mata
Department of Mathematics
Universite du Quebec a Montreal
mata_lopez.dante@uqam.ca

MS12

A Multidimensional Self-exciting Process with Dependencies (MSPD) for Insurance and FinancialApplications

The compound Poisson process is commonly used to model risk. This process assumes independence between the counting process (a Poisson process) and the claims (independent and identically distributed random variables), which makes the calculation of the first two cumulants straightforward. However, these assumptions restrict the applicability of this model to certain types of risks. To overcome these limitations, we propose a similar model where the counting process is a Hawkes process, with its intensity influenced by the occurrence of claims. To this end, we present a general process that not only addresses this problem but also opens the door to envisioning other applications.

Thomas Peyrat
Finance Group
Center for Research in Economics and Statistics
thomas.peyrat@exiompartners.com

Anthony Reveillac
INSA de Toulouse
anthony.reveillac@insa-toulouse.fr

Caroline Hillairet
ENSAE-Paris, CREST, France
Caroline.Hillairet@ensae.fr

MS12

Optimal Control of Discontinuous Processes in Random Environments

One of the primary methods for solving optimization problems involves determining a set of necessary conditions for any optimal solution. These conditions can become sufficient under additional convexity assumptions on the objective/constraint functions. Pontryagins Maximum Principle (PMP) expands this idea to optimal control problems, i.e., optimization problems in infinite dimensional spaces. Namely, PMP states that for an optimally controlled dynamical system, there exists an adjoint set of equations; together, they both solve a two-point boundary value problem plus a maximum condition for a function called the Hamiltonian. In this talk, we present a stochastic PMP for dynamical systems affected by a measure-valued random environment, such that the state process is assumed to be driven by a diffusion with self-exciting jumps. Then, we apply our results to the problem of Mean-Variance Portfolio Selection with Regime Switching. This is a joint work with Daniel Hernndez Hernndez (CIMAT).

Joshué Ricalde
ETH Zürich, Mathematics Department
joshue.ricalde@math.ethz.ch

Daniel Hernández
Department of Probability and Statistics
Centro de Investigación en Matemáticas, A.C.
dher@cimat.mx

MS13

A Crash-Course on Signatures

Given a (possibly random) continuous path $X = (X_t)_{t\in[0,T]}$ taking values in $\mathbb{R}^d$, its *signature* $\mathbb{X}^{<\infty}_{s,t}$ is defined as the collection of all iterated integrals

$$\mathbb{X}^{(i_1,\dots,i_n)}_{s,t} = \int_{s<t_1<\cdots,t_n} \mathrm{d}X^{i_1}_{t_1}\cdots \mathrm{d}X^{i_n}_{t_n}$$

for $0 \le s \le t \le T$ and $1 \le i_1,\dots,i_n \le d$. We tacitly assume that X is either a (possibly random) geometric rough path (and the integral defined as rough integral) or a semimartingale (with the integral defined in the Stratonovich sense). It is well known that the signature essentially characterizes the path. In addition, linear functionals on the signature form a sub-algebra of the continuous functions on path-space. As a consequence, a universal approximation result holds in terms of linear functionals of the signature. The signature is usually considered as an element of a large, non-commutative algebra $T((\mathbb{R}^d))$. This algebraic structure allows us to work efficiently and elegantly with linear and non-linear functionals of the signature. In this talk, we introduce the concepts of signatures and the algebraic framework. We explain the universal approximation property, and show how signatures are powerful tool for analysis and numerics in the context of processes with memory, I.e., non-Markov processes. Time permitting, we will show examples from stochastic optimal control.

Christian Bayer
Weierstrass Institute
christian.bayer@wias-berlin.de

MS13

Signatures for Volatility: Pricing and Hedging

In this talk we introduce Signature volatility models, where the stochastic volatility is modeled by a (possibly infinite) linear combination of the path signature of the time augmented driving Brownian motion. First, we will discuss the universality of this framework by providing explicit series expansions to certain stochastic path-dependent integral equations in terms of this path signature, which encompasses a large class of stochastic linear Volterra and delay equations and in particular the fractional Brownian motion with a Hurst index $H \in (0,1)$. Second, we will highlight the tractability of this class of models for Fourier pricing and hedging using the joint characteristic functional of the log-price and integrated variance known up to some infinite-dimensional Riccati equation.

Louis-Amand Gérard
Université Paris 1 Panthéon-Sorbonne, Paris, France
louisamandgerard@gmail.com

Eduardo Abi Jaber
Ecole Polytechnique
eduardo.abi-jaber@polytechnique.edu

Yuxing Huang
NYU
yh6004@nyu.edu

MS13

Signature Methods for Non-Markovian Optimal Stopping Problems

In this talk, we introduce novel methods for solving the optimal stopping problem in non-Markovian settings, relying on so-called *path-signatures*. These methods can be applied to price American options under rough volatility, where the underlying price process exhibits memory and is therefore non-Markovian. We propose two complementary procedures: First, a generalization of the well-known Longstaff-Schwartz algorithm based on signature regression, to obtain lower bounds on the option value. Second, a solution to Rogers' dual formulation by expressing Doob-martingales as stochastic integrals with signature integrands, which provides upper bounds. These approaches are supported by a global signature approximation result for L^p-functionals on rough-path spaces, from which we deduce convergence of the algorithms. We present several numerical examples in non-Markovian frameworks, including American options in the rough Heston and rough Bergomi models, and compare the performance of signature-based methods using linear, deep, and kernel learning techniques.

Luca Pelizzari
Weierstrass Institute and TU Berlin, Berlin, Germany
pelizzari@wias-berlin.de

MS14

Developing Financial Technologies from Decentralized Finance Ideas

In this talk we discuss how the Automated Market Makers (AMM's) can help produce a universal Liquidity Token and how this product is in fact equivalent to issuing shares in an Exchange Traded Fund (ETF). We show how this architecture can bypass the need for Authorized Participants (AP's) and instead function in a more efficient way. The resulting system can naturally produce target weighted ETF's which are notoriously hard to maintain the stated weights. As a byproduct we discuss how Liquidity providers may earn fees that could overcome the so called impermanent loss.

Ionut Florescu
Stevens Institute of Technology
Department of Mathematical Sciences
ifloresc@stevens.edu

Zachary Feinstein
Stevens Institute of Technology
zfeinste@stevens.edu

MS14

Scores as Actions: a Framework of Fine-tuning Diffusion Models by Continuous-time Reinforcement Learning

Reinforcement Learning from human feedback (RLHF) has been shown a promising direction for aligning generative models with human intent and has also been explored in recent works for alignment of diffusion generative models. In this work, we provide a rigorous treatment by formulating the task of fine-tuning diffusion models, with reward functions learned from human feedback, as an exploratory continuous-time stochastic control problem. Our key idea lies in treating the score-matching functions as controls/actions, and upon this, we develop a unified framework from a continuous-time perspective, to employ reinforcement learning (RL) algorithms in terms of improving the generation quality of diffusion models. We also develop the corresponding continuous-time RL theory for policy optimization and regularization under assumptions of stochastic differential equations driven environment.

Wenpin Tang
Columbia University
wt2319@columbia.edu

MS14

Some Rough Paths Techniques in Reinforcement Learning

In this talk I will start by reviewing some classical results relating machine learning problems with control theory. I will mainly discuss some very basic notions of supervised learning as well as reinforcement learning. Then I will show how noisy environments lead to very natural equations involving rough paths. This will include a couple of motivating examples. In a second part of the talk I will try to explain the techniques used to solve reinforcement learning problems with a minimal amount of technicality. In particular, I will focus on rough HJB type equations and their respective viscosity solutions. If time allows it, I will give an overview of our current research program in this direction. This talk is based on a joint work with Estepan Ashkarian (Purdue, Math), Prakash Chakraborty (Penn State) and Harsha Honnappa (Purdue, Industrial Engineering).

Samy Tindel
Purdue University
stindel@purdue.edu

MS14

Continuous Policy and Value Iteration for Stochastic Control Problem and Its Convergence

We propose a policy iteration algorithm, in which the approximations of value function and the control are updated simultaneously, for both the entropy regularized relaxed control problem and the classical control problem, with infinite horizon. We show the policy improvement and the convergence to the optimal control. Since both the value function and the control are updated according to differential equations in a continuous manner, we also confirm the convergence rate of the proposed algorithm.

Gu Wang
Worcester Polytechnic Institute
gwang2@wpi.edu

Qi Feng
Florida State University
qfeng2@fsu.edu

MS15

Trading Wind Power and Green Hydrogen

We develop a bidding strategy for a hybrid power plant combining co-located wind turbines and an electrolyzer, constructing a price-quantity bidding curve for the day-ahead electricity market while optimally scheduling hydrogen production. Without risk management, single imbalance pricing leads to an all-or-nothing trading strategy, which we term betting". To address this, we propose a data-driven, pragmatic approach that leverages contextual information to train linear decision policies for both power bidding and hydrogen scheduling. By introducing explicit risk constraints to limit imbalances, we move from the all-or-nothing approach to a trading" strategy, where the plant diversifies its power trading decisions reducing the risk of each individual decision. Comparing our data-driven strategy with an oracle model that has perfect foresight, we show that the risk-constrained, data-driven approach delivers satisfactory performance.

Yannick Heiser
Technical University of Denmark
yahei@dtu.dk

MS16

Distributionally Robust Merton's Portfolio Problem

Merton's problem is a classical problem in continuous-time stochastic control in finance, where the objective is to find the optimal investment and consumption strategy for an individual seeking to maximize its utility over a given time horizon. However, such an optimal strategy may fail to perform reliably when facing environment ambiguity. To address this limitation, we revisit Merton's problem by incorporating methodologies from distributionally robust optimization (DRO) [LFG22, KSW24], which combines the "worst-case" (min-max) approach of robust optimization with the probabilistic insights of stochastic optimization. In this work, we aim to solve the continuous-time distributionally robust control problem, where the goal is to maximize the worst-case utility over all possible distributions of terminal wealth within a given ambiguity set. First, we investigate the ambiguity set constructed using the Wasserstein metric, emphasizing how the structure of the Wasserstein ball can be explicitly characterized through the solution of the Fokker–Planck equation. Second, we explore the ambiguity set constructed using the Kullback–Leibler divergence, highlighting the connection between robust control under this type of ambiguity set and risk-averse control with the entropic value-at-risk. By drawing on these two perspectives, we aim to establish both theoretical advancements and data-driven algorithms for this class of problems.

Zhiyang Deng
Stevens Institute of Technology
zdeng10@stevens.edu

MS16

Distribution Builder Method for Optimal Selling

The distribution builder method, first introduced by Sharpe et al. in 2000, allows investors to select their preferred distributions rather than relying on a classical utility function when searching for investment strategies, such as determining the optimal time to sell an asset. While expected utility theory has many desirable properties, its practical implementation faces a major shortcoming: the lack of consistent estimation of economic agents utility functions. In this talk, we present a distribution-based framework to guide optimal selling decisions under a diffusive model, taking into account ruin and default times. We develop a criterion to determine whether a given distribution is attainable. We then consider two scenarios: for unattainable distributions, we propose reasonable modifications that are both attainable and as close as possible to the original desired distribution; for attainable ones, we explore whether they can be improved upon using the concept of stochastic dominance.

Sixian Jin
California State University
sjin@csusm.edu

MS16

Finding the Nonnegative Minimal Solutions of Cauchy PDEs in a Volatility-Stabilized Market

The strong relative arbitrage problem in Stochastic Portfolio Theory seeks to generate an investment strategy that almost surely outperforms a benchmark portfolio at the end of a given time horizon. The highest relative return in relative arbitrage opportunities is characterized by the smallest nonnegative continuous solution of a Cauchy problem for a partial differential equation (PDE). However, solving this type of PDE poses analytical and numerical challenges, due to the high dimensionality and its multiple solutions. In this paper, we discuss numerical methods to address the relative arbitrage problem and the associated PDE in a volatility-stabilized market, using time-changed Bessel bridges. We present a practical algorithm and demonstrate numerical results through an example in volatility-stabilized markets.

Nicole Yang
Emory University
tianjiao.yang@emory.edu

MS16

Numerical Methods with Branching Brownian Motion

McKean's seminal work on branching Brownian motion (BBM) and the Fisher-KPP equation introduced a novel probabilistic framework for solving PDEs. This insight inspired further discoveries linking BBM to semi-linear PDEs with various nonlinearities, including a voting scheme that leverages genealogical tree data to handle more complex problems. There are high-dimensional PDE solvers built on random coding trees and neural networks. However, the numerical performance of these voting models remains under-explored. In this work, we combine voting mechanisms with neural networks. We benchmark our approach against established Deep BSDE methods and the deep branching framework, demonstrating comparable or superior results across diverse scenarios. We further examine computational accuracy, limitations, and potential optimization strategies. By integrating machine learning with BBM, we aim to offer a robust alternative for high-dimensional PDEs, with broad potential applications in financial mathematics and beyond.

Xiaoqian Zhu
University of California Santa Barbara
xiaoqianzhu@umail.ucsb.edu

MS18

Algorithmic Collusion and a Folk Theorem from Learning with Bounded Rationality

We prove a Folk theorem when players with bounded rationality learn as they play a repeated potential game. We use a dynamic generalization of smooth fictitious play with bounded m-memory strategies to model learning with bounded rationality that is consistent with learning by algorithms. In a repeated potential game with perfect monitoring, we use this learning model to show that for any feasible and individually rational payoff profile, if players have sufficient memory, are sufficiently patient, and best respond with sufficiently few mistakes, then the players have a non-zero probability of learning an m-memory strategy profile that achieves an average payoff close to the specified payoff profile for an appropriate continuation game. Moreover, the strategy profile learned is an m-memory ϵ-subgame perfect equilibrium of the repeated game. This finding demonstrates that competition authorities are correct in their concern about algorithmic collusion.

lvaro Cartea, Patrick Chang
University of Oxford
alvaro.cartea@maths.ox.ac.uk,
patrick.chang@maths.ox.ac.uk

José Penalva
Universidad Carlos III de Madrid
jpenalva@emp.uc3m.es

Harrison Waldon
Oxford-Man Institute, University of Oxford
harrison.waldon@eng.ox.ac.uk

MS18

Do Ai-Algorithmic Traders Lead to Market Instability? A Multi-Agent Reinforcement Learning Approach

We show that AI-powered algorithmic trading can lead to inefficient and unstable markets. We investigate the impact of algorithmic traders on market equilibrium using a multi-agent deep reinforcement learning framework. We develop a conceptual simulation framework to study the market equilibrium among informed speculators with asymmetric information under different market conditions. Our findings reveal that more informed algorithmic traders engage in manipulative behavior, leading to local market price bubbles and deviations from the fundamental price. The informed AI-traders exploit and actively distort the learning process of less informed AI-traders. These results are robust across different number of agents and buying power scenarios. Unlike existing studies on algorithmic collusion, our research highlights the novel finding of the emergence of local price bubbles. Given the rapid growth in AI-powered algorithmic trading, our results identify a new form of systematic risk.

Markus Pelger
Stanford University
Department of Management Science and Engineering
mpelger@stanford.edu

MS18

A Two Time-Scale Evolutionary Game Approach to Multi-Agent Reinforcement Learning and Its Application in Algorithmic Collusion Studies

In this paper, we propose a novel two-timescale evolutionary game approach for solving general-sum multi-agent reinforcement learning (MARL) problems. Unlike existing literature that requires solving Nash equilibria exactly or approximately in each learning episode, our new approach synthesizes three key design components. First, we introduce a simple perturbed best response-based protocol for policy updates, avoiding the computationally expensive task of finding exact equilibria at each step. Second, agents use fictitious play to update their beliefs about other agents' policies, relaxing the requirement for observable Q-values of all agents as in classical Nash Q-learning. Third, our algorithm updates policies, beliefs, and Q-values at two different timescales to address non-stationarity during learning. The new approach provably converges to ϵ-Nash equilibria of MARL problems without imposing the global optima or saddle point conditions, two restrictive assumptions typically needed in the existing literature. AI-powered algorithms are increasingly used in marketplaces for pricing goods and services. The numerical experiments in this paper show that the realization of collusive Nash equilibrium is highly path-dependent and hinges on many factors, such as agents' impatience and the algorithm's exploration degree. Moreover, we also find that intervention timing is crucial for the effectiveness of regulatory policies.

Nan Chen
The Chinese University of Hong Kong
Hong Kong
nchen@se.cuhk.edu.hk

Mingyue Zhong
Department of Systems Engineering and Engineering Management
The Chinese University of Hong Kong
myzhong@se.cuhk.edu.hk

Yumin Xu, Ruixun Zhang
Peking University
2301110088@pku.edu.cn, zhangruixun@pku.edu.cn

MS19

Resource Sharing with Local Time Interactions and Systemic Failure

We model resource sharing in financial networks using a system of reflecting Brownian motions on the positive half-line in which each particle has a drift toward the origin determined by the local times at the origin of all the particles. We show that if this local time drift is too strong, such systems can exhibit a break down in their solutions in that there is a time beyond which the system cannot be extended. In the finite particle case we give a complete characterization of the finite time breakdown. We consider the mean field limit of the system in the symmetric case and show that there is a McKean-Vlasov representation. If the drift is too strong, the solution to the corresponding Fokker-Planck equation has a blow up in its solution. We also establish the existence of stationary solutions to the McKean-Vlasov equation in the case where there is no breakdown of the system. This work is motivated by the study of liquidity in financial markets, and we exhibit a connection to the super-cooled Stefan problem.

Graeme Baker
Columbia University
gb2847@columbia.edu

Ben Hambly
University of Oxford
hambly@maths.ox.ac.uk

Philipp Jettkant
Imperial College London
p.jettkant@imperial.ac.uk

MS19

Reinforcement Learning on Graphs and Dynamic Capital Injection in Financial Networks

We consider a discrete-time, multi-period financial network model subject to a potential shock that triggers a contagion cascade. The contagion spread follows the clearing equilibrium introduced by Weber and Weske (2017), which extends the framework of Eisenberg and Noe by incorporating bankruptcy costs, cross-holdings, and fire sales. A central regulator can inject capital into the system at each time step to mitigate the contagion. The decision-making process is modeled as a partially observable Markov decision process (POMDP). This setup reflects the realistic assumption that a) financial networks can only be partially observed and b) capital injections for a particular bank have to be explainable to them without revealing sensitive information from other banks. We solve this complex optimal decision making problem by reinforcement learning techniques that facilitate Graph Neural Networks to parameterize the optimal control function.

Nils Detering, Florian Kneilmann
Heinrich Heine University Düsseldorf
nils.detering@hhu.de, florian.kneilmann@hhu.de

Thilo Meyer-Brandis
University of Munich
meyer-brandis@math.lmu.de

MS19

Robust Optimal Strategies for Two-Period Liquidation in Financial Systems

We study the problem of asset liquidation in financial systems. During financial crises, asset liquidation is often inevitable but can lead to substantial losses if a significant amount of illiquid assets are sold simultaneously at depressed prices. To tackle this challenge, we consider a two-period liquidation model that allows for preemptive liquidation before maturity and propose a worst-case approach to the associated liquidation problem. Specifically, we propose a robust optimal strategya tractable liquidation approach that maximizes the worst-case terminal value of liquid assets, taking into account the uncertainty of other banks liquidation decisions. In addition, we find that the unique Nash equilibrium is attained when all banks adopt our proposed strategy. We further demonstrate that the robust optimal strategy retains a similar form even when interbank liabilities are involved, where we consider both scenarios of full and partial network information. While our analysis builds upon a stylized model, our findings offer valuable guidelines for developing robust liquidation strategies that mitigate losses resulting from asset liquidation.

Dohyun Ahn, Hongyi Jiang
The Chinese University of Hong Kong
dohyun.ahn@cuhk.edu.hk, hyjiang@link.cuhk.edu.hk

MS19

Modelling Contagious Bank Runs

We develop a modelling framework for contagion in financial networks arising from bank runs. We discuss different strategies that institutions which experience a bank run might use to satisfy their liquidity needs and analyse and compare implications of these different strategies for systemic risk. We provide an application of our framework to financial stress testing.

Luitgard Veraart
Department of Mathematics
London School of Economics and Political Science
l.veraart@lse.ac.uk

MS20

Optimal Trading Across Coexisting Exchanges: Limit-Order Books and Automated Market Makers

We study optimal trading across a limit-order book and an automated market maker. We determine the effective liquidity available to a large trader on both venues, when arbitrageurs immediately exploit intra-exchange price differences but only gradually offset the price impact of aggregate trade imbalances. The resulting price impact dynamics are nonlinear, but risk-neutral optimization problems can still be reduced to simple pointwise maximizations by reformulating the problem in impact space. These results are illustrated with an empirical case study for ETH/USDC price and trade data from Binance and Uniswap. This is joint work with Agostino Capponi and Johannes Muhle-Karbe.

Anthony Coache
University of Toronto
a.coache@imperial.ac.uk

Agostino Capponi
Columbia University
ac3827@columbia.edu

Johannes Muhle-Karbe
Imperial College London
j.muhle-karbe@imperial.ac.uk

MS21

Neural Sdes for Levy-Ito Processes

This paper introduces a novel framework for modeling Lvy-It processes using neural stochastic differential equations (SDEs). In this work, we focus on the jump-only setting, where the process is a pure jump model and where the Lvy measure is progressively measurable, resulting in history dependent jump arrival rates and distribution. This enhances the models expressiveness as it allows us to capture dynamic dependencies on the current and past states of the process. We parametrize this dynamic dependence using deep neural networks. Moreover, we develop a specific compensator structure that facilitates the construction of a tractable likelihood function, which is crucial for efficient inference and model estimation. The proposed framework is particularly well-suited for modeling, e.g., limit order books (LOBs), where updates are inherently jump-driven and occur at high frequencies. We demonstrate the practical feasibility of our method by providing a generative model for the entire LOB, capturing its high-dimensional, dynamic structure and enabling realistic simulations of market behavior.

Vedant Choudhary
University of Toronto
vedant.choudhary@mail.utoronto.ca

MS21

Data-Driven Heath-Jarrow-Morton Models

We develop a data-driven version of Heath-Jarrow-Morton models in the context of interest rate modeling. We consider models driven by linear functionals of the yield curve, such as a vector of representative forward rates, possibly augmented by a set of economic factors, whose characteristics can be easily estimated from market data. We then parametrize the volatility function via neural networks, thus considering the framework of neural SPDEs. Their parameters are learned by calibrating the model to past market yield curves. This results in a data-driven arbitrage-free generation/prediction of yield curves. Our setup also allows for the possibility of expected jumps, a key feature in interest rate markets, which can arise due to monetary policy decisions. We illustrate our deep learning procedure by reconstructing and forecasting the Euro area yield curves.

Christa Cuchiero
University of Vienna, Austria
christa.cuchiero@univie.ac.at

MS21

Universal Randomised Signatures for Generative Time Series Modelling

Randomised signature has been proposed as a flexible and easily implementable alternative to the well-established path signature. In this talk, we employ randomised signature to introduce a generative model for financial time series data in the spirit of reservoir computing. Specifically, we propose a novel Wasserstein-type distance based on discrete-time randomised signatures. This metric on the space of probability measures captures the distance between (conditional) distributions. Its use is justified by our novel universal approximation results for randomised signatures on the space of continuous functions taking the underlying path as an input. We then use our metric as the loss function in a non-adversarial generator model for synthetic time series data based on a reservoir neural stochastic differential equation. We compare the results of our model to benchmarks from the existing literature. Based on joint work with Francesca Biagini and Niklas Walter.

Lukas Gonon
University of St. Gallen
lukas.gonon@unisg.ch

MS21

Tail-GAN: Learning to Simulate Tail Risk Scenarios

The estimation of loss distributions for dynamic portfolios requires the simulation of scenarios representing realistic joint dynamics of their components, with particular importance devoted to the simulation of *tail risk* scenarios. We propose a novel data-driven approach that utilizes Generative Adversarial Network (GAN) architecture and exploits the joint elicitability property of Value-at-Risk (VaR) and Expected Shortfall (ES). Our proposed approach is capable of learning to simulate price scenarios that preserve tail risk features for benchmark trading strategies, including consistent statistics such as VaR and ES. We prove a universal approximation theorem for our generator for a broad class of risk measures. In addition, we show that the training of the GAN may be formulated as a max-min game, leading to a more effective approach for training. Our numerical experiments show that, in contrast to other data-driven scenario generators, our proposed scenario simulation method correctly captures tail risk for both static and dynamic portfolios.

Chao Zhang
Hong Kong University of Science and Technology (Guangzhou)
chao.zhang94@outlook.com

MS22

Stationarizing Signatures

In this talk, we introduce the notion of stationary signature, which plays the role of the polynomials of Ornstein-Uhlenbeck processes in path space and opens the door to modeling non-linear stationary processes depending on the infinite past trajectory. Remarkably, stationary signatures enjoy many important properties of standard signatures, such as Chen's identity, the linearization property, and uniqueness of signature. An explicit formula la Fawcett is derived for expected stationary signatures. We illustrate the usefulness of this notion with several applications, including regression of stationary processes against the signature of extended Brownian motion, stationary process representation, signal learning, and stochastic volatility modeling.

Dimitri Sotnikov
Ecole Polytecnhique, Palaiseau, France
dimitri.sotnikov@gmail.com

Eduardo Abi Jaber
Ecole Polytechnique
eduardo.abi-jaber@polytechnique.edu

MS23

Model-independent Pricing of Bermudan Options

Suppose we are given the market prices of European put/call options. Naturally, any pricing model, that can be used to price other exotic derivatives, should be consistent with the market quotesof these vanilla claims. Among such models, can we identify those that minimize/maximize the no-arbitrage price of a Bermudan option? We study this problem with the help of optimal transport.

Dominykas Norgilas
North Carolina State University
dnorgil@ncsu.edu

MS23

Outperforming a Benchmark with the Alpha-

Bregman-Wasserstein Distance

We consider the problem of active portfolio management, where an agent aims at finding the portfolio with maximal expected utility of terminal wealth subject to deviation constraints from a benchmark portfolio. As the agent values gains and losses differently, they utilise an asymmetric divergence on the space of distribution. Moreover, the agent aims at outperforming the benchmark, thus penalises outcomes where the portfolio wealth is below that of the benchmarks. This is achieved by the recently introduced α-Bregman Wasserstein divergence, generalising the Bregman Wasserstein and the popular Wasserstein divergence. We prove existence and uniqueness of the optimal portfolio strategy and discuss when the strategy coincides with the Merton strategy. We further give explicit criteria when the divergence constraints and the budget constraints are binding.

Silvana Pesenti
Department of Statistical Sciences
University of Toronto
silvana.pesenti@utoronto.ca

MS23

Firm Policies and Varying Recession Risk

Recession risk fluctuates “before the storm’ of a recession. Incorporating this into a model of liquidity and investment allows firms to endogenously delay actions until recession risk increases. When small, a firm chooses to act early when risk is low to protect attractive investments as investment reduces cash and raises liquidation risk. A larger firm delays actions as it invests less and accumulates cash. But, an imminent recession shortens its saving time, necessitating quick precautionary measures. Thus, as recession risk rises, the larger firm responds more by issuing preemptively and cutting investments and payouts. We estimate and validate our model.

A. Max Reppen
Boston University
amreppen@bu.edu

MS23

Macroscopic Properties of Equity Markets and Portfolio Selection

Macroscopic properties of equity markets affect the performance of active equity strategies but many are not adequately captured by conventional models of financial mathematics and econometrics. After reviewing some macroscopic properties using the CRSP Database of the US equity market, we study proportional transaction costs in the context of stochastic portfolio theory (SPT). We establish conditions under which relative arbitrages - portfolios that beat the market - can be constructed under transaction costs. We illustrate the performance of these portfolios and investigate the empirical validity of the conditions.

Leonard Wong
University of Toronto
tkl.wong@utoronto.ca

Steven A. Campbell
Columbia University
sc5314@columbia.edu

MS24

Continuous-Time Persuasion by Filtering

We frame dynamic persuasion in a partial observation stochastic control Leader-Follower game with an ergodic criterion. The Receiver controls the dynamics of a multidimensional unobserved state process. Information is provided to the Receiver through a device designed by the Sender that generates the observation process. The commitment of the Sender is enforced. We develop this approach in the case where all dynamics are linear and the preferences of the Receiver are linear-quadratic. We prove a verification theorem for the existence and uniqueness of the solution of the HJB equation satisfied by the Receivers value function. An extension to the case of persuasion of a mean field of interacting Receivers is also provided. We illustrate this approach in two applications: the provision of information to electricity consumers with a smart meter designed by an electricity producer; the information provided by carbon footprint accounting rules to companies engaged in a best-in-class emissions reduction effort. In the first application, we link the benefits of information provision to the mispricing of electricity production. In the latter, we show that information provision is a substitute for best-in-class objective.

Ofelia Bonesini
London School of Economics
o.bonesini@lse.ac.uk

René Aïd
Université Paris-Dauphine, PSL Research University
rene.aid@dauphine.psl.eu

Giorgia Callegaro
University of Padova
gcallega@math.unipd.it

Luciano Campi
University of Milan
Italy
luciano.campi@unimi.it

MS24

Optimal Execution Among N Traders with Transient Price Impact

We study N-player optimal execution games in an Obizhaeva–Wang model of transient price impact. When the game is regularized by an instantaneous cost on the trading rate, a unique equilibrium exists and we derive its closed form. Whereas without regularization, there is no equilibrium. We prove that existence is restored if (and only if) a very particular, time-dependent cost on block trades is added to the model. In that case, the equilibrium is particularly tractable. We show that this equilibrium is the limit of the regularized equilibria as the instantaneous cost parameter ε tends to zero. Moreover, we explain the seemingly ad-hoc block cost as the limit of the equilibrium instantaneous costs. Notably, in contrast to the single-player problem, the optimal instantaneous costs do not vanish in the limit $\varepsilon \to 0$. We use this tractable equilibrium to study the cost of liquidating in the presence of predators and the cost of anarchy. Our results also give a new interpretation to the erratic behaviors previously observed in discrete-time trading games with transient price impact.

Steven A. Campbell
Columbia University
sc5314@columbia.edu

Marcel Nutz
Columbia University, U.S.
mnutz@columbia.edu

MS24

Markov-Nash Equilibria in Mean-Field Games under Model Uncertainty

We propose and analyze a framework for mean-field Markov games under model uncertainty. In this framework, a state-measure flow describing the collective behavior of a population affects the given reward function as well as the unknown transition kernel of the representative agent. The agent's objective is to choose an optimal Markov policy in order to maximize her worst-case expected reward, where worst-case refers to the most adverse scenario among all transition kernels considered to be feasible to describe the unknown true law of the environment. We prove the existence of a mean-field equilibrium under model uncertainty, where the agent chooses the optimal policy that maximizes the worst-case expected reward, and the state-measure flow aligns with the agent's state distribution under the optimal policy and the worst-case transition kernel. Moreover, we prove that for suitable multi-agent Markov games under model uncertainty the optimal policy from the mean-field equilibrium forms an approximate Markov-Nash equilibrium whenever the number of agents is large enough.

Johannes Langner
Leibniz Universität Hannover
johannes.langner@insurance.uni-hannover.de

Ariel Neufeld
Nanyang Technological University
Singapore
ariel.neufeld@ntu.edu.sg

Kyunghyun Park
NTU Singapore
kyunghyun.park@ntu.edu.sg

MS24

Propagation of Chaos for Graphon Games

We study stochastic differential games where player interactions occur on sparse random graphs, providing a framework that generalizes finite-player games to infinite-population graphon limits. The model incorporates dissipative drifts and unbounded graphons, overcoming traditional monotonicity conditions. Our main contributions include the existence and uniqueness of graphon equilibria under strict dissipativity, convergence results linking finite-player games to their graphon limits, and inverse convergence theorems. We demonstrate that as the graph size grows, the N-player equilibria converge to the limiting graphon equilibrium under sparse connectivity regimes, characterized by sharp asymptotic sparsity conditions. Applications to disjoint communities, star graphs, and power-law networks showcase the versatility of our approach. This work extends mean-field game theory to structured sparse interactions, with implications for large-scale systems in economics, social dynamics, and statistical physics.

Dylan Possamaï
ETH
dylan.possamai@math.ethz.ch

Ludovic Tangpi
Princeton University
ludovic.tangpi@princeton.edu

MS25

Effect of Nonsmooth Initial Conditions on High-Order Timestepping Methods for PDEs: Analysis and Smoothings

Nonsmooth payoff functions are common in financial contracts and pose difficulties in obtaining high-order solutions of the contract prices. In this work, we consider parabolic PDEs with initial conditions involving various types of nonsmoothness. We apply a fourth-order finite difference (FD) discretization on a spatial grid, and BDF4 time-stepping initialized with two steps of an explicit third-order Runge-Kutta (RK3) method and one step of BDF3. Using Fourier analysis on the discrete system, we prove that the low-order errors generated by RK3 for nonsmooth data in the high-frequency domain get damped away by BDF steps, which prevents spurious oscillations; while low-order errors in the low-frequency domain come from the low-order initial condition discretization. To achieve globally fourth-order convergence, we apply fourth-order smoothing to the initial conditions, and provide explicit formulas of the discretization. Novel smoothing techniques are introduced and compared to existing ones. We mathematically prove and numerically verify that fourth-order convergence is obtained. The analysis is easily generalizable to higher order methods. Numerical examples on the model PDE and various option pricing problems are also given to demonstrate the fourth-order convergence of our method. The techniques can be applied to general nonsmooth (but piecewise smooth) initial conditions, general placement of the discontinuity points on the grid and on nonuniform grids.

Christina Christara
University of Toronto
Department of Computer Science
ccc@cs.toronto.edu

Dawei Wang, Kirill Serkh
University of Toronto
dwang@cs.toronto.edu, kserkh@math.toronto.edu

MS25

An Efficient Numerical Method for American Options and Their Greeks under the Two-Asset Kou Jump-Diffusion Model

In this talk we consider the numerical solution of the two-dimensional time-dependent partial integro-differential complementarity problem (PIDCP) that holds for the value of American-style options under the two-asset Kou jump-diffusion model. Following the method of lines (MOL), we derive an efficient numerical method for the pertinent PIDCP. Here, for the discretization of the nonlocal double integral term, an extension is employed of the fast algorithm by [J. Toivanen, Numerical valuation of European and American options under Kou's jump-diffusion model, SIAM J. Sci. Comp., 30:1949-1970, 2008] in the case of the one-asset Kou jump-diffusion model. For the temporal discretization, we study a useful family of second-

order diagonally implicit Runge-Kutta (DIRK) methods. Their adaptation to the semidiscrete two-dimensional Kou PIDCP is obtained by means of an effective iteration introduced by [Y. d'Halluin, P. A. Forsyth, and G. Labahn, A penalty method for American options with jump diffusion processes, Numer. Math., 97:321-352, 2004] and [Y. d'Halluin, P. A. Forsyth, and K. R. Vetzal, Robust numerical methods for contingent claims under jump diffusion processes, IMA J. Numer. Anal., 25:87-112, 2005]. Ample numerical experiments are presented showing that the proposed numerical method achieves a favourable, second-order convergence behaviour to the American two-asset option value as well as to its Greeks Delta and Gamma.

Karel In 't Hout
Department of Mathematics
University of Antwerp
karel.inthout@uantwerpen.be

MS25

Numerical Valuation of European Options under Two-Asset Infinite Activity Exponential Lvy Models

A numerical method for approximating the values of European-style financial derivatives is proposed, based on a two-asset exponential Lvy model for the underlying dynamic. The scheme accounts for both finite and infinite activity in the jump component. According to the Feynman-Kac theorem, the fair derivative value satisfies a two-dimensional partial-integro-differential equation (PIDE). Our numerical scheme is structured in two main parts. First, it replaces the integro-differential operator with a finite-difference approximation, reducing the PIDE to a system of ordinary differential equations (ODEs). Second, it solves this system numerically by a suitable time-stepping method. Among others, two notable challenges are addressed in developing this method. First, the integral term in the PIDE leads to a large dense matrix in the system of ODEs. To efficiently handle this, a second-order implicit-explicit time-stepping method is proposed, enabling application the fast Fourier transform. Second, for jump processes with infinite activity, the jumps below a small threshold are approximated by a diffusive term, which conveniently removes the singularity in the integral term. Numerical experiments are presented illustrating the effectiveness of our numerical approach. This work is inspired by [I.R. Wang, J.W.L. Wan, and P.A. Forsyth, Robust numerical valuation of European and American options under the CGMY process, J. Comp. Finan., 10:31-69, 2007].

Massimiliano Moda
University of Antwerp
massimiliano.moda@uantwerpen.be

Fred Espen Benth
University of Oslo
fredb@math.uio.no

Karel In 't Hout
Department of Mathematics
University of Antwerp
karel.inthout@uantwerpen.be

Michèle Vanmaele
Ghent University
Michele.Vanmaele@UGent.be

MS25

Numerical Methods for Solving Pides Arising in Swing Option Pricing under a Two-Factor Mean-Reverting Model with Jumps

Swing options are widely traded derivative contracts in energy markets, especially the electricity market. This type of contracts gives the holder the right to, dynamically, buy electricity at a predetermined, fixed price and, hence, reducing exposure to strong price fluctuations. There are constraints on the amount that can be bought by the holder at each exercise date and over the whole lifetime of the contract. We focus on swing options with discrete and fixed-time exercise rights. The electricity price is modelled by a two-factor model: an Ornstein-Uhlenbeck process and a zero mean-reverting pure jump process. This model takes into account the mean-reverting property of the price, the occurrence of price spikes as well as the possibility of negative prices. We shall consider in this talk finite-activity jumps, notably Merton and Kou jumps. The price of the swing options under consideration are given by the solution of systems of parabolic partial integro-differential equations (PIDEs) that are convection-dominated and have non-smooth initial conditions. In this talk, we shall develop an effective finite difference approach to solve these challenging equations numerically and next investigate its convergence behaviour.

Mustapha Regragui
Ghent University
mustapha.regragui@ugent.be

Fred Espen Benth
University of Oslo
fredb@math.uio.no

Karel In 't Hout
Department of Mathematics
University of Antwerp
karel.inthout@uantwerpen.be

Michèle Vanmaele
Ghent University
Michele.Vanmaele@UGent.be

MS26

Oligopolistic Market Equilibrium and the Role of Noise Observability

We develop a continuous time equilibrium model of an oligopolistic market in which the insider's optimal strategy may include a nonzero martingale component derived from observed noise. This challenges the standard assumption of absolute continuity. To clarify this, we study a sequence of models: a standard discrete time framework, a noise observable variant, and their continuous time limits under different strategy classes. We show that insider access to noise information determines whether the continuous time limit permits martingale components. Our results highlight a structural link between discrete and continuous time under imperfect competition and the role of noise observability in shaping strategic trading behavior.

JinHyuk Choi
Ulsan National Institute of Science and Technology
(UNIST)

jchoi@unist.ac.kr

MS26

An Axiomatic Theory for Anonymized Risk Sharing

We study an axiomatic framework for anonymized risk sharing. In contrast to traditional risk sharing settings, our framework requires no information on preferences, identities, private operations and realized losses from the individual agents, and thereby it is useful for modeling risk sharing in decentralized systems. Four axioms natural in such a framework actuarial fairness, risk fairness, risk anonymity, and operational anonymity are put forward and discussed. We establish the remarkable fact that the four axioms characterizes the conditional mean risk sharing rule, revealing the unique and prominent role of this popular risk sharing rule among all others in relevant applications of anonymized risk sharing. Several other properties and their relations to the four axioms are studied, as well as their implications in rationalizing the design of Bitcoin mining pools.

Yang Liu
Chinese University of Hong Kong, Shenzhen
yangliu16@cuhk.edu.cn

MS26

Relaxed Equilibria for Time-Inconsistent Markov Decision Processes

We consider an infinite-horizon Markov decision process (MDP) that allows for general non-exponential discount functions, in both discrete and continuous time. Due to the inherent time inconsistency, we look for a randomized equilibrium policy (i.e., relaxed equilibrium) in an intrapersonal game between an agents current and future selves. When we modify the MDP by entropy regularization, a relaxed equilibrium is shown to exist by a nontrivial entropy estimate. As the degree of regularization diminishes, the entropy-regularized MDPs approximate the original MDP, which gives the general existence of a relaxed equilibrium in the limit by weak convergence arguments. As opposed to prior studies that consider only deterministic policies, our existence of an equilibrium does not require any convexity (or concavity) of the controlled transition probabilities and reward function. Interestingly, this benefit of considering randomized policies is unique to the time-inconsistent case.

Zhenhua Wang
Shandong University
zhenhuaw78@gmail.com

MS26

Closed-Form Portfolio Optimization under Generalized Garch Models

We introduce a new framework of models for asset prices, the Generalized Heston Nandi GARCH (GHN-GARCH) models. Our objective with this family of models is portfolio optimization under expected utility theory in discrete time for an investor with constant relative risk aversion (CRRA) utility. We demonstrate that the GHN-GARCH class of models permits closed-form solutions for optimal allocation and value function. We define and study a new model as an example of this class, the 4/2-HN-GARCH model. The model is inspired by the continuous-time 4/2 stochastic volatility model of Grasselli [8]. The model is shown to be capable of capturing many market scenarios. An empirical data analysis on S&P500 price index returns confirms the superior fit of the 4/2-HN GARCH over the embedded HN-GARCH and the homoscedastic model and indicates that the optimal allocation under the 4/2-HN-GARCH model differs significantly from the optimal allocations under the HN-GARCH respectively the Merton model. As a second example we briefly discuss the new Crisis-HN-GARCH model and an extension to hyperbolic absolute risk aversion (HARA) utilities.

Rudi Zagst
Technical University of Munich
zagst@tum.de

MS27

The Role of Agent's Preferences in Sannikov's Model

We study the contracting model originally introduced by [Y. Sannikov, A continuous-time version of the principalagent problem, 2008], building on the rigorous analyses provided in recent studies by [D. Possamaï and N. Touzi, Is there a golden parachute in Sannikov's principal-agent problem?, 2024], as well as [D. Possamaï and C. Rossato, Golden parachutes under the threat of accidents, 2024]. The principal hires an agent to perform a task and compensate him with running payments throughout the contract's duration, which concludes at a random time, potentially with a lump-sum payment upon termination. In the setting of a risk-averse agent, this work aims to go beyond the power utility grow discussed in the aforementioned studies, presenting results general enough to tackle the most classical utility functions as exponential, negative power and logarithmic utility. Special attention is then given to the case of a risk-neutral agent. This scenario, characterized by the explosive feature of the associated Hamiltonian, leads to a series of singular stochastic control problems. We provide a complete, explicit solution for this case, resulting in optimal contracts that differ substantially from those in the risk-averse framework. Notably, these contracts are characterized by a singular nature, with lump-sum payments at initiation (a 'welcome bonus') and the involvement of local times (a 'recurring bonus' once some objectives are reached) throughout.

Filippo Beretta
ETH Zurich
filippo.beretta@math.ethz.ch

Dylan Possamaï
ETH
dylan.possamai@math.ethz.ch

MS27

Existence and Regularity for Systems of Parabolic PDE's: Economic Applications

We establish the existence and $W^{1,2,p}$-regularity of solutions to Hamilton-Jacobi-Bellman systems of parabolic nonlinear partial differential equations. We then apply this result to (i) prove that the value function of optimal control, switching, and stopping problems is smooth under mild regularity conditions on the primitives, (ii) establish the existence of and characterize the Markov Perfect equilibria of continuous-time stochastic games when the state follows a diffusion process, (iii) identify the surplus maximizing public perfect equilibrium in symmetric continuous-time stochastic games, and (iv) revisit past results on control problems for general stochastic differential recursive utility.

Théo Durandard
University of Illinois at Urbana-Champaign
theod@illinois.edu

MS27

Moral Hazard with Limited Liability

We study a continuous-time principal-agent problem with moral hazard, where the contract consists of a lump-sum payment at the end of a fixed time interval. We consider a situation in which both parties have limited liability, meaning the contract satisfies both lower and upper bounds. We characterize the set of all admissible contracts and study the principal's problem by means of an HJB equation on a specific domain. This talk is based on joint work with Dylan Possama and Stéphane Villeneuve.

Daniel Kršek
ETH Zurich
daniel.krsek@math.ethz.ch

Dylan Possamaï
ETH
dylan.possamai@math.ethz.ch

Stephane Villeneuve
Toulouse School of Economics
stephane.villeneuve@tse-fr.eu

MS27

Optimal Contract and Consumption Habit

The paper examines a continuous-time principal-agent model in which agents preference exhibits habit formation over consumption. As agents concern over the standard of living strengthens, his continuation utility is less sensitive to current wealth but more sensitive to the standard of living, leading to lower demand for risk-sharing compensation. The optimal contract has lower pay-for-performance but incentivizes agents higher effort. In the Leland (1994) capital structure model, agents habit formation preference combined with the optimal contract lowers firms leverage and mitigates the debt-overhang problem.

Hao Xing
Boston University
haoxing@bu.edu

MS28

Optimal Power Plant Investments with Multiple Competing Project

In this talk, we explore a finite-horizon mixed optimal singular control-stopping problem within a stochastic framework. We frame the problem as a system of Hamilton-Jacobi-Bellman (HJB) equations with gradient constraints. Our results are then applied to the optimal power investment strategies, driven by a controlled Ornstein-Uhlenbeck process, incorporating volatility control and capital injections over time. We extend the analysis from a single-project scenario to a system of multiple correlated Ornstein-Uhlenbeck processes, where singular control is applied to each process. The results are illustrated with numerical simulations and provide the optimal investment strategy and stopping rules for each project.

Isabel Agostino
UC Berkeley
IEOR
iagostino@berkeley.edu

Thibaut Mastrolia
UC Berkeley
Mastrolia@berkeley.edu

MS28

Optimization and Learning for Mean-Field Games Via Occupation Measure

Mean-field games (MFGs) and multi-agent reinforcement learning (MARL) have become essential frameworks for analyzing interactions in large-scale systems. This talk presents recent advancements at the intersection of MFGs and MARL. We begin with a new framework MF-OMO (Mean-Field Occupation Measure Optimization), which reformulates Nash equilibria for discrete-time MFGs as a single optimization problem over occupation measures, offering a fresh characterization that enables the use of standard optimization algorithms to identify multiple equilibria without relying on restrictive assumptions. We also extend these results to continuous-time finite state MFGs. Building on the concept of occupation measures, we then introduce MF-OML (Mean-Field Occupation Measure Learning), the first fully polynomial online RL algorithm capable of finding approximate Nash equilibria in large population games beyond zero-sum and potential games. We establish regret bounds for the N-player games that can be approximated by MFGs under monotonicity conditions. Together, these advancements provide a comprehensive approach to characterizing and solving Nash equilibria in complex multi-agent environments.

Anran Hu
COLUMBIA UNIVERSITY
ah4277@columbia.edu

MS28

Dimension-Free Estimates on the N-player Nash System and Applications

This talk will be about stochastic differential games with many players, similar to those considered in mean field game (MFG) theory, but without any symmetry assumptions on the costs. We explain how to obtain estimates on the corresponding PDEs (the Nash system) which are independent of the number N of players. In order to obtain such estimates, we require (i) monotonicity conditions inspired by MFG theory and (ii) a smallness condition on the strength of interaction between any two distinct players (which is satisfied in the MFG case for large N). We apply our results on the Nash system to obtain useful estimates on the gap between three formulations of the game: closed-loop, open-loop, and distributed. These bounds confirm the intuition that if the interaction between two distinct players is small, then the choice of formulation shouldn't matter very much. When specialized to the MFG setting, they also provide a new tool to address the convergence problem and certain questions about universality. This is based on ongoing joint work with Marco Cirant and Davide Redaelli.

Joe Jackson
University of Chicago

jsjackson@uchicago.edu

MS28

Time-Inconsistent Stochastic Games with Mean-Variance Preferences

We investigate a time-inconsistent N-player game in continuous time, where each player's objective functional depends non-linearly on the expected value of the state process, including classic meanvariance models as a special case. We identify subgame-perfect Nash equilibria and characterize each player's value function using a system of coupled backward stochastic differential equations. Building on this, we further analyze the mean-field counterpart and its associated mean-field equilibria.

Chiara Rossato
ETH Zurich
chiara.rossato@math.ethz.ch

Dylan Possamaï
ETH
dylan.possamai@math.ethz.ch

MS29

Identifying Optimal Capacity Expansion and Differentiated Capacity Payments under Risk Aversion and Market Power: A Financial Stackelberg Game Approach

We investigate how capacity payments in combination with scarcity pricing of energy can ensure resource adequacy in electricity markets, defined as the ability of supply and other resources to provide enough energy and capacity to meet demand under steady-state operating conditions. This work generalizes models for determining capacity payments by deriving second-best discriminatory payments by resource type that account not only for the "missing money' market failure that arises from energy price caps, but also for market power in the capacity market and differences in risk tolerance among resource types that can arise from failures in risk and capital markets. A bi-level equilibrium-constrained optimization model is proposed to define second-best capacity payments in a static long-run setting, considering the impacts of those payments on the mix and cost of generation investment and energy outputs. The lower-level suppliers play a Nash game to determine the generation mix under a capacity payment scheme, while the upper-level regulator considers consumer welfare and resource adequacy. We introduce an equivalent formulation via a variational inequality approach, and find conditions for the solution to exist. Discriminatory payments are found to be second-best when there is market power in the investment game, price caps in energy markets and imperfections in risk markets that lead to diverse risk attitudes.

Maxim Bichuch
Department of Mathematics
SUNY Buffalo
mbichuch@buffalo.edu

MS29

Improving Abatement and Payoffs by (Coarse) Correlation

We consider a simple emission abatement game among greenhouse gas emitters, where agents decide how much of a pollutant to abate. Each agent incurs a private cost for abatement but benefits from the total reduction achieved by the system. Reformulating the problem as a mean field game (MFG), we find that Nash equilibria incentivize free-riding behavior, leading both to inefficient payoffs and low overall abatement levels. To address this, we explore alternative equilibrium concepts and introduce the notion of mean field coarse correlated equilibrium (CCE). CCEs feature a moderator who randomly selects a strategy profile for the players, correlating their strategies without requiring them to cooperate. We demonstrate that CCEs can achieve both higher abatement levels and improved payoffs compared to the Nash equilibrium of the MFG, highlighting their potential benefits in addressing collective action problems. This talk is based on a joint work with Luciano Campi (University of Milan) and Fanny Cartellier (University of Zurich).

Ferderico Cannerozzi
University of Bielefeld
¡federico.cannerozzi@uni-bielefeld.de¿

MS29

Hilbert-Space Valued LQ Mean Field Games

Mean field games (MFGs) were originally developed in finite-dimensional spaces. However, there are scenarios where Euclidean spaces do not adequately capture the essence of the problem, such as systems involving time delays. We present a comprehensive study of linear-quadratic (LQ) MFGs in Hilbert spaces, involving N agents with dynamics governed by infinite-dimensional stochastic equations. In this framework, both state and control processes of each agent take values in separable Hilbert spaces. All agents are coupled through the average state of the population which appears in their linear dynamics and quadratic cost functional. Specifically, the dynamics of each agent incorporates an infinite-dimensional noise, namely a Q-Wiener process, and an unbounded operator. The diffusion coefficient of each agent is stochastic involving the state, control, and average state processes. We first study the well-posedness of a system of N coupled semilinear stochastic evolution equations establishing the foundation of MFGs in Hilbert spaces. We then specialize to N-player LQ games described above and study the asymptotic behavior as the number of agents, N, approaches infinity. We develop an infinite-dimensional variant of the Nash Certainty Equivalence principle and characterize a unique Nash equilibrium for the limiting MFG. Finally, we demonstrate the resulting limiting best-response strategies form an ?-Nash equilibrium for the N-player game in Hilbert spaces.

Dena Firoozi
HEC Montreal
University of Montreal
dena.firoozi@hec.ca

MS30

Liquidity Provision to Informed and Uninformed Traders

We characterise the solutions to a continuous-time optimal liquidity provision problem in a market populated by informed and uninformed traders. In our model, the asset price exhibits fads —these are short-term deviations from the fundamental value of the asset. Conditional on the value of the fad, we model how informed traders and uninformed traders arrive in the market. The market maker knows of the two groups of traders but only observes the anonymous order arrivals. We study both, the complete information and the partial information versions of the control problem faced by the market maker. In such frameworks, we characterise the value of information, and we find the price of liquidity as a function of the proportion of informed traders in the market. Lastly, for the partial information setup, we explore how to go beyond the Kalman-Bucy filter to extract information about the fad from the market arrivals.

Adrien Mathieu
Mathematical Institute, University of Oxford
Oxford-Man Institute of Quantitative Finance
adrien.mathieu@maths.ox.ac.uk

MS30

Information Leakage and Opportunistic Trading Around the Fx Fix

We study opportunistic traders that try to detect and exploit the order flow of dealers hedging their net exposure to the FX fix. We also discuss how dealers can take this into account to balance not only risk and trading costs but also information leakage in an appropriate manner.

Johannes Muhle-Karbe
Imperial College London
j.muhle-karbe@imperial.ac.uk

Roel Oomen
Deutsche Bank
roel.oomen@db.com

Mateo Rodriguez Polo
ETH Zurich
mateo.rodriguezpolo@math.ethz.ch

MS30

Brokers, Informed Clients, and Equilibrium Trading

We study the interactions between a broker and her clients (an informed trader and an uninformed trader). In the first part of the talk we study the imperfect information case where the strategic agents acknowledge misspecification in their models. Then, we formulate a two-stage optimisation where both parties filter the information that they do not have. We conclude with an overview of the open problems within this framework. This is based on two works: one written jointly with A. Cartea and one written with A. Aqsha and F. Drissi.

Leandro Sanchez-Betancourt
Mathematical Institute, University of Oxford
Oxford-Man Institute of Quantitative Finance
sanchezbetan@maths.ox.ac.uk

lvaro Cartea, Alif Aqsha
University of Oxford
alvaro.cartea@maths.ox.ac.uk, alif.aqsha@maths.ox.ac.uk

Fayçal Drissi
Université Paris 1 Panthéon-Sorbonne
faycal.drissi@gmail.com

MS30

Broker-Trader Partial Information Nash-Equilibrium

We study partial information Nash equilibrium between a broker and an informed trader. In this model, the informed trader, who possesses knowledge of a trading signal, trades multiple assets with the broker in a dealer market. Simultaneously, the broker trades these assets in a lit exchange where their actions impact the asset prices. The broker, however, only observes aggregate prices and cannot distinguish between underlying trends and volatility. Both the broker and the informed trader aim to maximize their penalized expected wealth. Using convex analysis, we characterize the Nash equilibrium and demonstrate its existence and uniqueness. Furthermore, we establish that this equilibrium corresponds to the solution of a nonstandard system of forward-backward stochastic differential equations.

Xuchen Wu
University of Toronto
xuchen.wu@mail.utoronto.ca

Sebastian Jaimungal
University of Toronto, Canada
sebastian.jaimungal@utoronto.ca

MS31

Robust Financial Calibration: a Bayesian Approach for Neural SDEs

We consider a Bayesian framework for the calibration of financial models using neural stochastic differential equations (neural SDEs). The method is based on the specification of a prior distribution on the neural network weights and an adequately chosen likelihood function. The resulting posterior distribution can be seen as a mixture of different classical neural SDE models yielding robust bounds on the implied volatility surface. Both, historical financial time series data and option price data are taken into consideration, which necessitates a methodology to learn the change of measure between the risk-neutral and the historical measure. The key ingredient for a robust numerical optimization of the neural networks is to apply a Langevin-type algorithm, commonly used in the Bayesian approaches to draw posterior samples.

Eva Flonner
Vienna University of Economics and Business
eva.flonner@s.wu.ac.at

MS31

Rough Transformers: Lightweight Continuous-Time Sequence Modelling with Path Signatures

Time-series data in real-world settings typically exhibit long-range dependencies and are observed at non-uniform intervals. In these settings, traditional sequence-based recurrent models struggle. To overcome this, researchers often replace recurrent architectures with Neural ODE-based models to account for irregularly sampled data and use Transformer-based architectures to account for long-range dependencies. Despite the success of these two approaches, both incur very high computational costs for input sequences of even moderate length. To address this challenge, we introduce the Rough Transformer, a variation of the Transformer model that operates on continuous-time representations of input sequences and incurs significantly lower computational costs. In particular, we propose multi-view signature attention, which uses path signatures to augment vanilla attention and to capture both local and global (multi-scale) dependencies in the input data, while remaining robust to changes in the sequence length and sampling frequency and yielding improved spatial processing. We find that, on a variety of time-series-related tasks, Rough Transformers consistently outperform their vanilla attention counterparts while obtaining the representational benefits of Neural ODE-based models, all at a fraction of the computational time and memory resources.

Fernando Moreno-Pino
University of Oxford
fernando.moreno-pino@eng.ox.ac.uk

MS31

HighRankPath Development and Its Application in Financial Time Series Generation

Since the weak convergence for stochastic processes does not account for the growth of information over time which is represented by the underlying filtration, a slightly erroneous stochastic model in weak topology may cause huge loss in multi-periods decision making problems. To address such discontinuities, Aldous introduced the extended weak convergence, which can fully characterise all essential properties, including the filtration, of stochastic processes; however, it was considered to be hard to find efficient numerical implementations. In this talk, we introduce a novel metric called High Rank PCF Distance (HRPCFD) for extended weak convergence based on the high rank path development method from rough path theory, which also defines the characteristic function for measure-valued processes. We then show that such HRPCFD admits many favourable analytic properties which allows us to design an efficient algorithm for training HRPCFD from data and construct the HRPCF-GAN by using HRPCFD as the discriminator for conditional time series generation. Our numerical experiments on both hypothesis testing and generative modelling validate the out-performance of our approach compared with several state-of-the-art methods, highlighting its potential in broad applications of synthetic time series generation and in addressing classic financial and economic challenges, such as optimal stopping or utility maximisation problems.

Hao Ni
University College London
h.ni@ucl.ac.uk

MS31

Limit Order Book Simulation and Trade Evaluation with K-nearest-neighbor Resampling

We propose a novel k-nearest neighbor resampling procedure for estimating the performance of a policy from historical data containing realized episodes of a decision process generated under a different policy. We provide statistical consistency results under weak conditions: in particular, we avoid the common assumption of identically and independently distributed transitions and rewards and thus allow sampling of entire episodes. By focusing on feedback policies that depend deterministically on the current state in environments with continuous state-action spaces and system-inherent stochasticity affected by chosen actions, and relying on trajectory simulation similar to Monte Carlo methods, the proposed method is particularly well suited for stochastic control environments as those encountered in financial markets. We demonstrate, using historical limit order book (LOB) data, that our simulation method is capable of recreating realistic LOB dynamics and that synthetic trading within the simulation leads to a market impact in line with the corresponding literature. Furthermore, we show that in a benchmark comparison our method outperforms a deep learning-based algorithm for several key statistics.

Christoph Reisinger
University of Oxford
christoph.reisinger@maths.ox.ac.uk

MS32

Risk Aversion of Insider and Asymmetric Information

We analyse how the risk aversion of insider affects the equilibrium in insider trading model with dynamic information. In particular, we will consider a dynamic information Kyle-Back model under new assumptions: a) exponential utility preferences of the insider, b) non- Gaussianity of the signal, and c) price set by the market maker being a function of weighted signal which is not necessarily Gaussian either. We will discuss conditions on the weighting and pricing functions which ensure the existence of equilibrium and derive, under afore mentioned conditions, the equilibrium pricing and weighting functions, as well as insiders optimal trading strategy.

Albina Danilova
London School of Economics
a.danilova@lse.ac.uk

MS32

The Mean Field Schrdinger Problem, Characterization and Dual Formulation

The mean field Schrdinger problem (MFSP) is the problem of finding the most likely path of a McKean-Vlasov type particle with constrained initial and final configurations. It was first introduced by Backhoff et al. (2020), who studied its existence and long-time behavior. This talk aims to show how ideas from mean field control theory allow us to derive new interesting results on the MFSP. In particular, we study its existence, characterization, and dual reformulation via control techniques. This talk is based on a joint work with Ludovic Tangpi.

Camilo Hernández
University of Southern California
camilohe@usc.edu

MS32

A Forward-Backward Approach to Endogenous Distress Contagion

In this talk, I will introduce a dynamic model of a banking network in which the value of interbank obligations is continuously adjusted to reflect counterparty default risk. An interesting feature of the model is that the credit value adjustments increase volatility in times of distress, leading to endogenous default contagion between the banks. The counterparty default risk can be computed backwards in time from the maturity date, leading to a specification of the model in terms of a forward-backward stochastic differential equation (FBSDE), coupled through the banks' default times. While one can prove the existence of solutions with minimal and maximal default probabilities, the question of uniqueness is currently open. I will conclude the talk by discussing a characterisation of the maximal default probabilities through a cascade of partial differential equations (PDEs), each representing a configuration with a different number of defaulted banks. The domain of each PDE has a free boundary that coincides with the banks' default thresholds.

Philipp Jettkant
Imperial College London
p.jettkant@imperial.ac.uk

Andreas Sojmark
London School of Economics and Political Science
a.sojmark@lse.ac.uk

MS32

Asymptotics in Prediction Problems

Prediction with expert advice is one of the fundamental problems in online learning and sequential decision making with an exploration-exploitation trade-off. The problem is often analyzed in adversarial settings and has wide-ranging applications, including risk management, security training, and betting markets. In this talk, we will explore the long time behavior of prediction problems, formulated as zero-sum games between an adversary and a forecaster. In the first part, we consider a scenario where both the adversary and the forecaster have full observation of what happens. In this case, the asymptotic behavior is described by a nonlinear degenerate parabolic equation. Subsequently, we address a situation where the forecaster only has access to partial information, leading to a PDE defined on Wasserstein space. We will also discuss a comparison principle relevant to this latter equation.

Xin Xin Zhang
New York University
xz1662@nyu.edu

Erhan Bayraktar
University of Michigan
Department of Mathematics
erhan@umich.edu

Ibrahim Ekren, Ibrahim Ekren
University of Michigan
iekren@umich.edu, iekren@umich.edu

MS33

The Indifference Value of the Weak Information

We propose indi?erence pricing to estimate the value of the weak information. Our framework allows for tractability, quantifying the amount of additional information, and permits the description of the smallness and the stability with respect to small perturbations of the weak information. We provide sharp conditions for the stability with counterexamples. The results rely on a theorem of independent interest on the stability of the optimal investment problem with respect to small changes in the physical probability measure. We also investigate contingent claims that are indi?erence price invariant with respect to changes in weak information. We show that, in incomplete models, the class of information-invariant claims includes the replicable claims, and it can be strictly bigger. In particular, in complete models, all contingent claims are information invariant. We augment the results with examples and counterexamples.

Oleksii Mostovyi
University of Connecticut
Department of Mathematics
oleksii.mostovyi@uconn.edu

Fabrice Baudoin
Aarhus University
fbaudoin@math.au.dk

MS33

Dynamic Noisy Rational Expectations Equilibrium with Endogenous Information Acquisition Time

In this talk, we establish equilibrium in the presence of heterogeneous information. In particular, there is an insider who receives a private signal, an uninformed agent with no private signal, and a noise trader with semi price-inelastic demand. The novelty is that we allow the insider to decide (optimally) when to acquire the private signal. This endogenizes the entry time and stands in contrast to the existing literature which assumes the signal is received at the beginning of the period. Allowing for optimal entry also enables us to study what happens before the insider enters with private information, and how the possibility for future information acquisition both affects current asset prices and creates demand for information related derivatives. Results are valid in continuous time, when the private signal is a noisy version of the assets terminal payoff, and when the quality of the signal depends on the entry time.

Scott Robertson
Questrom School of Business
Boston University
scottrob@bu.edu

Jerome Detemple, Nikolaos Vingos
Questrom School of Business, Boston University
detemple@bu.edu, nvingos@bu.edu

MS33

Backward Martingale Transport Maps and Equilibrium with Insider

We consider an optimal transport problem with backward martingale constraint. The objective function is given by the scalar product of a pseudo-Euclidean space S. We show that the supremums over maps and plans coincide, provided that the law ν of the input random variable Y is atomless. An optimal map X exists if ν does not charge any c-c surface (the graph of a difference of convex functions) with strictly positive normal vectors in the sense of the S-space. The optimal map X is unique if ν does not charge c-c surfaces with nonnegative normal vectors in the S-space. As an application, we derive sharp conditions for the existence and uniqueness of equilibrium in a multi-asset version of the model with insider from Rochet and Vila [10]. In the linear-Gaussian case, we characterize Kyle's lambda, the sensitivity of price to trading volume, as the unique positive solution of a non-symmetric algebraic Riccati equation.

Dmitry Kramkov
Carnegie Mellon University
kramkov@cmu.edu

Mihai Sirbu
University of Texas at Austin

sirbu@math.utexas.edu

MS33

Log-Optimal Porfolios with Small Liability Stream

In the context of a continuous semimartingale financial model with a log-utility framework, we study the long-term behavior of the first and second order approximations of utility-based hedging strategies for a relatively small exogenous endowment. These approximations are derived from consecutive quadratic optimization problems and are intrinsically linked to the second- and fourth-order asymptotics of the investor's value function as the units of the endowment approaches zero. Concrete examples are also analyzed in the context of a factor model. This is a joint work of M. Anthropelos, C. Kardaras and C. Stefanakis.

Michail Anthropelos
University of Piraeus
Department of Banking and Financial Management
anthropel@unipi.gr

Constantinos Kardaras
London School of Economics
k.kardaras@lse.ac.uk

Constantinos Stefanakis
University of Piraeus
Department of Banking and Financial Management
kstefanakis@unipi.gr

MS34

A Dynamic Model for Open Banking

Open banking allows customers to share their financial data with third-party providers, such as fintech companies, thereby increasing competition and expanding access to financial services. While this shift promises benefits such as improved product offerings and expanded financial inclusion, it also introduces significant technical, social, and economic risks. We develop a continuous-time model using a search-and-match framework to analyze borrower heterogeneity and competition between banks and fintechs. Our analysis shows that open banking reshapes financial services by leveling the playing field between banks and fintechs. However, while it enhances competition, it could also over-empower fintechs, potentially reversing the benefits to borrowers.

Ranu Castaneda, Christoph Frei
University of Alberta
rcastane@ualberta.ca, cfrei@ualberta.ca

MS34

Quantifying the Value of Revert Protection

Revert protection is a feature provided by some blockchain platforms that prevents users from incurring fees for failed transactions. We study the economic implications and benefits of revert protection in the context of decentralized finance, specifically the setting of priority gas auctions and maximal extractable value. We develop a model in which searchers bid for a top-of-block arbitrage opportunity under varying degrees of revert protection. This model applies to a broad range of settings, including bundle auctions on L1s and priority ordering sequencing rules on L2s. We quantify, in closed form, how revert protection improves equilibrium auction revenue, market efficiency, and blockspace efficiency.

Brian Zhu
Columbia University
bzz2101@columbia.edu

Xin Wan
Uniswap Labs
xin@uniswap.org

Ciamac C. Moallemi
Graduate School of Business
Columbia University
ciamac@gsb.columbia.edu

Dan Robinson
Paradigm
dan@paradigm.xyz

Brad Bachu
Uniswap Labs
brad.bachu@uniswap.org

MS34

An Economic Model of the L1-L2 Interaction

We provide an economic model of the interaction between a Layer-1 (L1) blockchain and an associated Layer-2 (L2). Our main finding is that, even when the L1 blockchain features value-creating decentralized applications (dApps), there nevertheless exist realistic conditions such that both L1 blockchain investment and L1 cryptoasset market value converge to zero over time. Importantly, in the absence of the L2, the same conditions regarding L1 technology imply that both the L1 blockchain investment and the market value of the L1 cryptoasset grow without bound asymptotically. Overall, our results highlight the need for developers to ensure that the L1 cryptoasset remains relevant for L2 dApps.

Campbell Harvey
Duke University
cam.harvey@duke.edu

Fahad Saleh
University of Florida
fahad.saleh@ufl.edu

Ruslan Sverchkov
University of Warwick
ruslan.sverchkov@wbs.ac.uk

MS34

EconAgentic in DePin Markets: A Large Language Model Approach to the Sharing Economy of Decentralized Physical Infrastructure

The Decentralized Physical Infrastructure (DePin) market is revolutionizing the sharing economy by decentralizing physical assets and incentivizing node providers through token-based economics. As part of the distributed web economy, DePin represents a new model for resource allocation and infrastructure management. This paper introduces EconAgent, an AI-powered simulation framework utilizing Large Language Models (LLMs) to mimic human-like decision-making within DePin ecosystems. By modeling how decentralized agents respond to token incentives, invest in infrastructure, and adapt to changing market conditions, EconAgentic provides valuable insights into the efficiency, fairness, and sustainability of DePin. Our findings demonstrate the potential of AI-driven simulations to improve the design and operation of decentralized, tokenized sharing economies, offering a roadmap for the future of distributed web-based markets. This research contributes to the broader understanding of how AI and token economies can shape the next generation of online markets and decentralized infrastructure.

Luyao Zhang
Duke Kunshan University
luyao.zhang@dukekunshan.edu.cn

MS35

Functional It Calculus and Some Financial Applications

Path dependence is ubiquitous in finance, sometimes explicitly as the payoff of an exotic option may depend on the whole path of the asset price, not only at maturity, other times through the dynamics of the underlying (volatility, dividends). The framework to model path dependence is the Functional It Calculus and we review its basic concepts before offering a partial panorama of its applications: computation of the Greeks of path dependent options, perturbation analysis, volatility risk decomposition, Taylor expansion with signatures for fast computation of VaR and characterization of attainable claims, amongst other ones.

Bruno Dupire
Bloomberg LP, U.S.
bdupire@bloomberg.net

MS35

Functional Its Formula and Taylor Expansions for Non-Anticipative Maps of Cdlg Rough Paths

We rely on the approximation properties of the signature of cdlg rough paths to derive a functional Ito-formula for non-anticipative maps of rough paths. This leads to a functional extension of the Ito-formula for cdlg rough paths derived in [Friz and Zhang, 2018] which coincides with the change of variable formula formulated in [Dupire, 2009] whenever the functionals representations, the notions of the regularity of the functionals and the integration concepts can be matched. In contrast to these works, by using the concept of vertical Lie derivatives, we can also incorporate path functionals where the second order vertical derivatives do not commute as it is the case for typical signature functionals. As a byproduct, we show that sufficiently regular non-anticipative maps admit a functional Taylor expansion, leading to a far reaching generalization of the recently established results in [Dupire and Tissot-Daguette, 2022].

Francesca Primavera
University of Vienna
francesca.primavera@berkeley.edu

Christa Cuchiero
University of Vienna, Austria
christa.cuchiero@univie.ac.at

Xin Guo
University of California, Berkeley
xinguo@berkeley.edu

MS35

Universality of Signature Based Models

In recent years, signature based methods have been very successfully applied in mathematical finance. At the very heart of these methods are universal approximation theorems, establishing that continuous functionals can be approximated arbitrary well on compact sets by linear maps acting on signatures. However, in the context of mathematical finance, the restriction to compact sets causes a lack of theoretical justification for the use of signature based methods. In this talk, we provide various global universal approximation theorems in the L^p-sense with respect to the Wiener measure. In particular, we demonstrate that functionals on rough path space can be approximated globally in the L^p-sense w.r.t. the Wiener measure. This allows, for instance, to approximate solutions to stochastic differential equations driven by Brownian motions by signature based models, leading to a certain universality of signature based models for financial markets.

David Prömel, Mihriban Ceylan
University of Mannheim
proemel@uni-mannheim.de, mihriban.ceylan@uni-mannheim.de

MS35

Invariance View on Signature and a Dual Proof of Hambly-Lyons' Theorem

Signature transforms of continuous bounded variation curves appear in the theory of controlled or stochastic differential equations, in rough path theory and recently in machine learning of path dependent functionals. By Hambly-Lyons, Uniqueness for the signature of a path of bounded variation and the reduced path group, Annals of Mathematics, 2010 signature transforms characterize bounded variation paths up to tree like equivalences. We provide an invariant theoretic perspective on this deep result by identifying signature components as basic invariant polynomials with respect to a group of reparametrizations. We also relate this result to signature kernels and real analytic functions on continuous bounded variation paths. We furthermore provide a different perspective on tree like equivalence.

Josef Teichmann
Department of Mathematics
ETH Zurich
jteichma@math.ethz.ch

MS36

Total Positivity and Option Pricing

Call option prices in the Black-Scholes model, viewed as functions of strike and maturity, are totally positive of order two (TP2), meaning that the price ratio of a higher-strike call to a lower-strike call increases with maturity, with suitable adjustments for dividends and interest. We develop conditions for this property in other models, derive analogous conditions for puts, and discuss connections with the shape of the implied volatility surface. In market data for SP 500 options, we find that violations of TP2 are infrequent and usually correct within a couple of days. Trading against TP2 violations can be highly profitable when properly implemented.

Paul Glasserman
Columbia University
pg20@columbia.edu

Dan Pirjol
Stevens Institute of Technology
dpirjol@gmail.com

MS36

Semi-Analytical Pricing of American Options in Time-dependent Models

Semi-analytical pricing of American options in a time-dependent Ornstein-Uhlenbeck model was presented in [Carr-Itkin, 2020. It was shown that to obtain these prices

one needs to solve (numerically) a nonlinear Volterra integral equation of the second kind to find the exercise boundary (which is a function of the time only). Once this is done, the option prices follow. It was also shown that computationally this method is as efficient as the forward finite difference solver while providing better accuracy and stability. Later this approach called the "Generalized Integral transform" method has been significantly extended to various time-dependent one factor, [Itkin-Lipton-Muravey, 2021], and stochastic volatility [Carr-Itkin-Muravey, 2020], [Itkin-Muravey, 2020] models as applied to pricing barrier options. However, for American options, despite possible, this was not explicitly reported anywhere. In this paper our goal is to fill this gap and also discuss which numerical methods could be efficient to solve the corresponding Volterra equations, also including machine learning.

Andrey Itkin
New York University
aitkin@nyu.edu

MS36

Heuristics for Options Trading

The talk will examine the volatility and skew dynamics observed in a range of US stocks and indices over the past few years. We will interpret these numbers in the context of the broader market environment.

Amal Moussa
Department of Statistics, Columbia University
am2810@caa.columbia.edu

MS36

Volatility Parametrizations with Random Coefficients: Analytic Flexibility for Implied Volatility Surfaces

It is a market practice to express market-implied volatilities in some parametric form. The most popular parametrizations are based on or inspired by an underlying stochastic model, like the Heston model (SVI method) or the SABR model (SABR parametrization). Their popularity is often driven by a closed-form representation enabling efficient calibration. However, these representations indirectly impose a model-specific volatility structure on observable market quotes. When the market's volatility does not follow the parametric model regime, the calibration procedure will fail or lead to extreme parameters, indicating inconsistency. This article addresses this critical limitation - we propose an arbitrage-free framework for letting the parameters from the parametric implied volatility formula be random. The method enhances the existing parametrizations and enables a significant widening of the spectrum of permissible shapes of implied volatilities while preserving analyticity and, therefore, computation efficiency. We demonstrate the effectiveness of the novel method on real data from short-term index and equity options, where the standard parametrizations fail to capture market dynamics. Our results show that the proposed method is particularly powerful in modeling the implied volatility curves of short expiry options preceding an earnings announcement, when the risk-neutral probability density function exhibits a bimodal form.

Lech A. Grzelak
Assistant Professor, Utrecht University
L.A.Grzelak@uu.nl

Leonardo Perotti
Utrecht University, The Netherlands
l.perotti@uu.nl

MS37

Correct Implied Volatility Shapes and Reliable Pricing in the Rough Heston Model

We use modifications of the Adams method and very fast and accurate sinh-acceleration method of the Fourier inversion (iFT) (S.Boyarchenko and Levendorskiĭ, IJTAF 2019, v.22) to evaluate prices of vanilla options; for options of moderate and long maturities and strikes not very far from the spot, thousands of prices can be calculated in several msec. with relative errors of the order of 0.5% and smaller running Matlab on a Mac with moderate characteristics. We demonstrate that for the calibrated set of parameters in Euch and Rosenbaum, Math. Finance 2019, v. 29, the correct implied volatility surface is significantly flatter and fits the data very poorly, hence, the calibration results in op.cit. is an example of the *ghost calibration* (M.Boyarchenko and Levendorkiĭ, Quantitative Finance 2015, v. 15): the errors of the model and numerical method almost cancel one another. We explain how calibration errors of this sort are generated by each of popular versions of numerical realizations of iFT (Carr-Madan, Lipton-Lewis and COS methods) with prefixed parameters of a numerical method, resulting in spurious volatility smiles and skews. We suggest a general *Conformal Bootstrap principle* which allows one to avoid ghost calibration errors. We outline schemes of application of Conformal Bootstrap principle and the method of the paper to the design of accurate and fast calibration procedures.

Sergei Levendorskii
University of Leicester
levendorskii@gmail.com

MS38

An Optimal Stopping Problem Arising in Zero-Sum Stochastic Games of Stopper Vs. Controller

Type

We construct a saddle point in a class of zero-sum games between a stopper and a singular-controller. The underlying dynamics is a one-dimensional, time-homogeneous, singularly controlled diffusion taking values either on $\mathbb{R}$ or on $[0, \infty)$. The games are set on a finite-time horizon, thus leading to analytical problems in the form of parabolic variational inequalities with gradient and obstacle constraints. The saddle point is characterised in terms of two moving boundaries: an optimal stopping boundary and an optimal control boundary. These boundaries allow us to construct an optimal stopping time for the stopper and an optimal control for the singular-controller. Our method relies on a new link between the value function of the game and the value function of an auxiliary optimal stopping problem with absorption. We show that the smooth-fit condition at the stopper's optimal boundary (in the game), translates into an absorption condition in the auxiliary problem. This is somewhat in contrast with results obtained in problems of singular control with absorption and it highlights the key role of smooth-fit in this context.

Andrea Bovo
University of Torino
andrea.bovo@unito.it

Tiziano De Angelis
University of Turin
tiziano.deangelis@unito.it

MS38

Mean Field Optimal Stopping: Dynamic Programming and Neural Networks Approximation

The aim of this talk is to present recent advances in the study of mean field optimal stopping problems. More precisely, we focus on the analysis of these problems in presence of common noise, which is natural in many applications (e.g. the optimal liquidation of a portfolio of correlated assets). After extending the dynamic programming principle to this framework and deriving the corresponding dynamic programming equation, we will also present a numerical method to approximate the solution. Due to the high dimensional nature of this problem, our approach relies on a neural network approximation of the value function and the optimal stopping policy, combined with dynamic programming.

Mehdi Talbi
ETHZ
mtalbi@lpsm.paris

MS38

Stopping Problems with An Unknown State

We extend the classical setting of an optimal stopping problem under full information to include problems with an unknown state. The framework allows the unknown state to influence (i) the drift of the underlying process, (ii) the payoff functions, and (iii) the distribution of the time horizon. Since the stopper is assumed to observe the underlying process and the random horizon, this is a two-source learning problem. Assigning a prior distribution for the unknown state, standard filtering theory can be employed to embed the problem in a Markovian framework with one additional state variable representing the posterior of the unknown state. We provide a convenient formulation of this Markovian problem, based on a measure change technique that decouples the underlying process from the new state variable. Moreover, we show by means of several novel examples that this reduced formulation can be used to solve problems explicitly.

Yuqiong Wang
PhD student
Uppsala University
yuqw@umich.edu

MS39

Sequential Optimal Contracting in Continuous Time

In this talk, I will present a Principal-Agent problem in continuous time with multiple lump-sum payments (contracts) paid at different deterministic times. We use BSDEs to reduce the Stackelberg game between the Principal and the Agent to a standard stochastic optimal control problem. We apply our result to a benchmark model to investigate how different inputs affect the Principals value.

Guillermo Alonso Alvarez
University of Michigan
guialv@umich.edu

Erhan Bayraktar
University of Michigan
Department of Mathematics
erhan@umich.edu

Ibrahim Ekren, Ibrahim Ekren, Liwei Huang
University of Michigan
iekren@umich.edu, iekren@umich.edu, huanglw@umich.edu

MS39

A New Approach for Principal-Agent Problems

The recent work by [Cvitanic, Possama, and Touzi: Dynamic programming approach to principalagent problems, 2018] develops a general approach for continuous-time principal-agent problems, through dynamic programming and second-order BSDEs. Here, we provide an alternative formulation of the problem, which can be solved simply relying on the theory of BSDEs. This reformulation is strongly inspired by an important remark in [CPT18], namely that if the principal observes the output process X in continuous-time, she can compute its quadratic variation pathwise. While in previous works, this information is used in the contract, our reformulation consists in assuming that the principal could directly control this process, in a first-best fashion. The resolution approach for this alternative problem actually follows the line of the so-called Sannikovs trick in the literature on continuous-time principal-agent problems. We then show that the solution to this first-best formulation is identical to the solution of the original problem. More precisely, using the contracts form introduced in [CPT18] as forcing contracts, we highlight that this first-best scenario can be achieved even if the principal cannot directly control the quadratic variation. Nevertheless, we do not have to rely on the theory of 2BSDEs to prove that such contracts are optimal, as their optimality is ensured by showing that the first-best scenario is achieved.

Emma Hubert, Alessandro Chiusolo
Princeton University
eh3988@princeton.edu, ac8541@princeton.edu

MS39

Dynamic Contracts with Aspirational Preferences

This paper characterizes the optimal dynamical contract when the agent features aspirational utility as described by Diecidue and Van der Ven (2008). We extend the seminal dynamic principal-agent framework of Sannikov (2008) to characterize the optimal contract when the agent features an aspiration level. There are three important results. The compensation contract discontinuously jumps after good performance to reach the aspiration level. It is constant around the aspiration level for many histories, and only continuous to increase after a substantially good history is realized. Second, even though consumption exhibits very low PPS around the aspiration level, the net present value of consumption displays high PPS around that level. Third, the implemented effort is very high around the aspiration level. Because the agent is risk-loving in this region, the usual trade-off between risk-sharing and incentives is muted and the optimal contract features very high incentives. Right below the aspiration the agent is incentivized to reach the aspirational level by being exposed to a large share of the projects risk. Finally, we discuss the implications for optimal contracts with multiple aspirational thresholds and for instances in which the agent can engange in risk-taking.

Alejandro Rivera
University of Texas at Dallas
alejandro.riveramesias@utdallas.edu

Fernando Zapatero
Boston University
fzapa@bu.edu

Hongda Zhong
University of Texas at Dallas
zhonghongda@gmail.com

MS39

Optimal Fees for Optimal Liquidation on Lit and Dark Pools

We address the problem of a trader who aims to liquidate her inventory simultaneously in lit markets and dark pools, while the exchange seeks to establish an appropriate trading fee structure to maximize profits while minimizing the permanent price impact caused by the traders activities. To solve the traders optimal routing problem, we formulate the associated Hamilton-Jacobi-Bellman (HJB) equation. First, we propose a Feynman-Kac type representation for the value function and prove that the value function $h(t, q)$ preserves concavity with respect to inventory Q_t at any time. By solving this partial differential equation, we derive the optimal trading strategy. Finally, we approximate the optimal fees for both lit markets and dark pools and demonstrate that introducing fees can enhance market efficiency.

Hao Wang
UC Berkeley
IEOR
haowang013@berkeley.edu

Thibaut Mastrolia
UC Berkeley
Mastrolia@berkeley.edu

MS40

Convergence Analysis in Mean-Field Games with Jumps

We investigate N-player games and mean-field games with jumps by characterizing Nash and mean-field equilibria as solutions to backward stochastic differential equations (BSDEs). Using the BSDE formulation, we analyze the convergence of N-player value functions to their mean-field counterparts and provide explicit convergence rates, offering a quantitative understanding of this asymptotic behavior. Building on a work of Possama and Tangpi, this analysis extends their framework to incorporate mean-field games with jumps, offering new insights into the relationship between finite-player dynamics and mean-field limits in systems with discontinuities. This is based on joint work with Joshu Hel Ricalde Guerrero and Chiara Rossato.

Marco Rodrigues, Chiara Rossato
ETH Zurich
marco.rodrigues@math.ethz.ch,
chiara.rossato@math.ethz.ch

Joshué Ricalde
ETH Zürich, Mathematics Department
joshue.ricalde@math.ethz.ch

MS40

Non-Markovian Graphon Games

We study finite-player dynamic stochastic games with heterogeneous interactions and non-Markovian linear-quadratic objective functionals. We derive the Nash equilibrium explicitly by converting the first-order conditions into a coupled system of stochastic Fredholm equations, which we solve in terms of operator resolvents. When the agents' interactions are modeled by a weighted graph, we formulate the corresponding non-Markovian continuum-agent game, where interactions are modeled by a graphon. We also derive the Nash equilibrium of the graphon game explicitly by first reducing the first-order conditions to an infinite-dimensional coupled system of stochastic Fredholm equations, then decoupling it using the spectral decomposition of the graphon operator, and finally solving it in terms of operator resolvents. Moreover, we show that the

Nash equilibria of finite-player games on graphs converge to those of the graphon game as the number of agents increases. Finally, we apply our results to various stochastic games with heterogeneous interactions. This is joint work with Eyal Neuman.

Sturmius Tuschmann, Eyal Neuman
Imperial College London
s.tuschmann23@imperial.ac.uk, eyaln13@gmail.com

MS41

Uniform-in-Time Rates of Convergence to a Nonlinear Markov Chain for Mean-Field Interacting Jump Processes

We consider a system of N particles interacting through their empirical distribution on a finite space in continuous time. In the formal limit as $N \to \infty$, the system takes the form of a nonlinear (McKean-Vlasov) Markov process. This paper rigorously establishes this limit. Specifically, under the assumption that the mean field system has a unique, exponentially stable stationary distribution, we show that the weak error between the empirical measures of the N-particle system and the law of the mean field system is of order $1/N$ uniformly in time. Our analysis makes use of a master equation for test functions evaluated along the measure flow of the mean field system, and we demonstrate that the solutions of this master equation are sufficiently regular. We then show that exponential stability of the mean field system is implied by exponential stability for solutions of the linearized Kolmogorov equation with a source term. Finally, we show that our results can be applied to the study of mean field games and give a new condition for the existence of a unique stationary distribution for a nonlinear Markov chain.

Ethan Huffman
University of Michigan,
ehuf@umich.edu

Asaf Cohen
University of Michigan
asafc@umich.edu

MS41

Decentralization and Learning in Mean-Field Control

The solution and complexity of multi-agent control problems rely heavily on the information structure of the problem. Decentralization, where each agent utilizes only their local information variables, is a highly desirable property to reduce the necessity for coordination and complexity. However, achieving such decentralization is generally not possible in most multi-agent control problems. The mean-field control formulation, where agents are only "weakly" correlated through a mean-field term, offers a promising modeling framework in which decentralization can be achievable. In particular, the solution of an auxiliary control problem, where one considers the infinite population dynamics, can provide near-optimal decentralized strategies for large population control problems. The main focus of the talk will be on the limitations and varying degrees of decentralization in these scenarios. In the final part of the talk, I will focus on the learning aspect of the problem. I will present model-based learning methods that use linear function approximations for the dynamics and cost structures. We will see that during exploration, coordination may become inevitable, as fully decentralized learning approaches can fail to sufficiently explore and excite the state space. Joint work with Erhan Bayraktar and Nicole Bauerle

Ali Kara
Florida State University
akara@fsu.edu

Erhan Bayraktar
University of Michigan
Department of Mathematics
erhan@umich.edu

Nicole Bäuerle
Karlsruhe Institute of Technology
nicole.baeuerle@kit.edu

MS41

Cascade Equation for Stephan Problem As a Mean Field Game

The solutions to Stefan problem with Gibbs-Thomson law (i.e., with surface tension effect) are well known to exhibit singularities which, in particular, lead to jumps of the associated free boundary along the time variable. The correct times, directions and sizes of such jumps are only well understood under the assumption of radial symmetry, under which the free boundary is a sphere with varying radius. The characterization of such jumps in a general multidimensional setting has remained an open question until recently. In our ongoing work with M. Shkolnikov and Y. Guo, we have derived a separate (hyperbolic) partial differential equation referred to as the cascade equation whose solutions describe the jumps of the solutions to the Stefan problem without any symmetry assumptions. It turns out that a solution of the cascade equation corresponds to a maximal element of the set of all equilibria in a family of (first-order local) mean field games. In this talk, I will present and justify the cascade equation, will show its connection to the mean field games, and will prove the existence of a solution to the cascade equation. If time permits, I will also show how these results can be used to construct a solution to the Stefan problem itself.

Sergey Nadtochiy
Illinois Institute of Technology
sserco38@gmail.com

Mykhalyo Shkolnikov
Carnegie Mellon University
mshkolni@andrew.cmu.edu

Yucheng Guo
Princeton University
yg7348@princeton.edu

MS41

Uniform-in-time Weak Propagation of Chaos for

Consensus-based Optimization

We study the uniform-in-time weak propagation of chaos for the consensus-based optimization (CBO) method on a bounded searching domain. We apply the methodology for studying long-time behaviors of interacting particle systems developed in the work of Delarue and Tse (arXiv:2104.14973). Our work shows that the weak error has order O(1/N) uniformly in time, where N denotes the number of particles. The main strategies behind the proofs are the decomposition of the weak errors using the linearized Fokker-Planck equations and the exponential decay of their Sobolev norms. Consequently, our result leads to the joint convergence of the empirical distribution of the CBO particle system to the Dirac-delta distribution at the global minimizer in population size and running time in Wasserstein-type metrics.

Hongyi Zhou
University of Michigan, Ann Arbor
hongyizh@umich.edu

Erhan Bayraktar
University of Michigan
Department of Mathematics
erhan@umich.edu

Ibrahim Ekren, Ibrahim Ekren
University of Michigan
iekren@umich.edu, iekren@umich.edu

MS42

Delegated Portfolio Management with Random Default

We are considering the problem of optimal portfolio delegation between an investor and a portfolio manager under a random default time. We focus on a novel variation of the Principal-Agent problem adapted to this framework. We address the challenge of an uncertain investment horizon caused by an exogenous random default time, after which neither the agent nor the principal can access the market. This uncertainty introduces significant complexities in analyzing the problem, requiring distinct mathematical approaches for two cases: when the random default time falls within the initial time frame $[0, T]$ and when it extends beyond this period. We develop a theoretical framework to model the stochastic dynamics of the investment process, incorporating the random default time. We then analyze the portfolio manager's investment decisions and compensation mechanisms for both scenarios. For both cases, we demonstrate that the contracting problem can be resolved by examining the existence of solutions to integro-partial Hamilton-Jacobi-Bellman (HJB) equations. We develop a deep-learning algorithm to solve the problem with no access to the optimizer of the Hamiltonian function.

Alberto Gennaro
UC Berkeley
IEOR
alberto.gennaro@berkeley.edu

Thibaut Mastrolia
UC Berkeley
Mastrolia@berkeley.edu

MS42

Mean Field Optimal Stopping Problem with a Singular Coefficient

The work considers a mean field optimal stopping problem with a singular coefficient in objective function. The corresponding large population problem is deeply related with quickest detection problem, which aims to study efficient detection of abrupt changes in the statistical behavior of streaming data, and it is a fundamental problem arising in many fields of engineering, in finance, in the natural and social sciences, and even in the humanities. The goal is to establish the regularity of the value function, dynamic programming principle, limit theory, and derive the corresponding obstacle PDE problem (a variant of HJB equation).

Xihao He, Xihao He
University of Michigan, Ann Arbor
hexihao@umich.edu, hexihao@umich.edu

MS42

Controlled Forward-Backward Dynamics in Interbank Systems

In this talk, I will introduce a dynamic model of interbank borrowing and lending that takes place due to targeted liquidity levels, akin to the work of Capponi, Sun, and Yao. Departing from the existing literature, we will consider a given finite horizon with a target for the terminal time. Banks are then borrowing or lending according to the expected deviations from their targets conditionally on the current information, leading to a system of forward-backward dynamics. On top of this, we finally devise a linear-quadratic control problem, whereby banks can adjust their drift at a cost. We first study the existence of Nash equilibria in the case of a finite network and then we proceed to discuss existence and uniqueness in the simplified mean field regime. Several interesting insights are obtained when comparing with control problems for corresponding forward-only systems.

Andreas Søjmark
London School of Economics
Department of Statistics
a.sojmark@lse.ac.uk

Zeng Zhang
London School of Economics
z.zhang103@lse.ac.uk

MS42

Optimal Contract, Delegated Investment, and Information Acquisition

This paper examines a model of delegated investment within the framework of a noisy rational expectations equilibrium. Portfolio managers can acquire costly signals about asset payoffs but incur portfolio management costs. They receive compensation from delegated investors and make investment decisions on their behalf. The optimal contract includes a benchmark component that mitigates agency frictions arising from portfolio management costs. The precision of private signals chosen by portfolio managers is determined by equilibrium market conditions rather than the specifics of their individual contracts. When portfolio management costs decrease, both the performance and benchmark components of the optimal contract are reduced, portfolio managers acquire less precise private signals, but market price efficiency improves.

Yuyang Zhang
Boston University
Questrom School of Business
yyz@bu.edu

Hao Xing
Boston University
haoxing@bu.edu

MS43

Deep Signature Algorithm for Multidimensional Path-Dependent Options

In this work, we study the deep signature algorithms for path-dependent options. We extend the backward scheme in [Hur e-Pham-Warin. Mathematics of Computation 89, no. 324 (2020)] for state-dependent FBSDEs with reflections to path-dependent FBSDEs with reflections, by adding the signature layer to the backward scheme. Our algorithm applies to both European and American type option pricing problems while the payoff function depends on the whole paths of the underlying forward stock process. We prove the convergence analysis of our numerical algorithm with explicit dependence on the truncation order of the signature and the neural network approximation errors. Numerical examples for the algorithm are provided including: Amerasian option under the Black- Scholes model, American option with a path-dependent geometric mean payoff function, and the Shiryaevs optimal stopping problem.

Erhan Bayraktar
University of Michigan
Department of Mathematics
erhan@umich.edu

Qi Feng
Florida State University
qfeng2@fsu.edu

Zhaoyu Zhang
UCLA
zhaoyu@math.ucla.edu

MS43

Understanding the Commodity Futures Term Structure Through Signatures

Signature methods have successfully been used as a tool for feature extraction in statistical learning methods, notably in mathematical finance. The specific reason for their success is often much less clear, besides a general handwaving to path-dependence. This presentation aims to explain their success in a particular task, namely classifying commodity futures markets according to storability. We provide a regular perturbation of the signature of the futures term structure in terms of the convenience yield and identify the volatility of the convenience yield as major discriminant. This is joint work with H. P. Krishnan (SCT Capital Management).

Stephan Sturm
WPI
ssturm@wpi.edu

MS43

Nonlinear Occupied Pricing PDEs

We consider the optimal control of occupied processes which record all positions of the state process. Dynamic programming yields nonlinear equations on the space of measures. We develop the viscosity theory for this parabolic occupied PDE by proving a comparison result between sub and supersolutions, and thus provide a characterization of the value function as the unique viscosity solution. Toward this proof, an extension of the Crandall-Ishii-Lions Lemma, as well as finite-dimensional approximations, is established. We then discuss the pricing PDEs of financial derivatives where the payoff or model is contingent on the occupation measure. Examples include path-dependent claims in occupied volatility models, variance/timer options, and nonlinear PDEs arising from the uncertain volatility model. This is joint work with Mete Soner and Jianfeng Zhang.

Valentin Tissot-Daguette
Princeton University / Bloomberg LP
vtissotdague@bloomberg.net

Mete Soner
Princeton University
soner@princeton.edu

Jianfeng Zhang
University of Southern California
jianfenz@usc.edu

MS44

Does the Market Have Private Information on Crypto Risk? Evidence From a Blockchain Insurance Platform

As Bitcoin price reaches its new high in 2024, crypto risk management has become an increasingly concern for both investors and policymakers. In this study, we examine the effectiveness of decentralized insurance in enhancing crypto security, which allows individual investors to actively participate in the insurance business. Using data from the largest decentralized insurer specializing in crypto coverage, we find evidence that individual insurers possess additional information about the underlying riskiness of requested insurance covers and incorporate this information into their pricing decisions. Our analysis identifies multiple channels for these information sources, including historical loss events, third-party audit reports, and community discussions. This informational advantage is particularly pronounced during periods of crypto market booms. These findings shed light on the potential of decentralized insurance as a tool for risk management in the rapidly evolving crypto landscape.

Xiaochen Jing
University of Illinois Urbana-Champaign
xjing4@illinois.edu

MS44

On the Optimality of Linear Residual Risk Sharing

In this talk, we shall explore the optimal risk-sharing problem in the context of peer-to-peer insurance. Using the criterion of minimizing total variance, we find that the optimal risk-sharing strategy should take a linear form. Although linear risk-sharing strategies have been examined in the literature, our study uncovers a significant finding: to minimize total variance, the linear strategy should be applied to the residual risks rather than the original risks, as commonly adopted in existing studies. By comparing with the existing models, we demonstrate the advantage of the linear residual risk sharing model in variance reduction and robustness. Furthermore, we develop and study a number of new models by incorporating some constraints, to reflect desirable properties required by the market. With those constraints, the optimal strategies turn out to favor market development, such as incentivize participation and guarantee fairness. A relevant model is considered in the end, which establishes the connection among multiple optimization problems and provides insights on how to extend the models into a more general setup.

Jiajie Yang, Wei Wei
University of Illinois Urbana-Champaign
jiajiey3@illinois.edu, weiw@illinois.edu

MS44

On the Design of Optimal Multiple-peril Cyber Insurance

In todays insurance market, numerous cyber insurance products provide bundled coverage for losses resulting from different cyber events, including data breaches and ransomware attacks. Every category of incident has its own specific coverage limit and deductible. Although this gives prospective cyber insurance buyers more flexibility in customizing the coverage and better manages the risk exposures of sellers, it complicates the decision-making process in determining the optimal amount of risks to retain and transfer for both parties. This article aims to build an economic foundation for these incident-specific cyber insurance products with a focus on how incident-specific indemnities should be designed for achieving Pareto optimality for both the insurance seller and the buyer. Real data on cyber incidents are used to illustrate the feasibility of this approach. Several implementation improvement methods for practicality are also discussed.

Linfeng Zhang
Ohio State University
zhang.14673@osu.edu

Wing Fung Chong
Heriot-Watt University
alfred.chong@hw.ac.uk

Daniël Linders
University of Amsterdam
d.h.linders@uva.nl

Zhiyu Quan
University of Illinois at Urbana-Champaign
zquan@illinois.edu

MS44

Co-opetition in Reinsurance Markets: When Pareto Meets Stackelberg and Nash

We develop and solve a two-layer game to model coopetition, a strategic combination of competition and cooperation, in a reinsurance market consisting of one primary insurer and two reinsurers, in which all players are equipped with mean-variance preferences and the reinsurance contracts are priced under the variance premium principle. The insurer negotiates reinsurance contracts with the two reinsurers simultaneously, modeled by two Stackelberg games, and the two reinsurers compete for business from the same insurer by setting their own pricing rules, modeled by a non-cooperative Nash game. The combined Stackelberg-Nash game constitutes the first layer of the game model and endogenously determines the risk assumed by each reinsurer. The two reinsurers, then, participate in a cooperative risk-sharing game, forming the second layer of the game model, and seek Pareto-optimal risk-sharing rules. We obtain equilibrium strategies in closed form for both layers. The equilibrium of the Stackelberg-Nash game consists of two proportional reinsurance contracts, with the more risk-averse reinsurer assuming a smaller portion of the insurer's total risk. The Pareto-optimal risk-sharing rules further dictate that the more risk-averse reinsurer transfers a portion of its assumed risk to the less risk-averse reinsurer, at the cost of a positive side payment.

Jingyi Cao, Dongchen Li
York University
jingyic@yorku.ca, dcli@yorku.ca

Virginia Young
University of Michigan
Department of Mathematics
vryoung@umich.edu

Bin Zou
University of Connecticut
bin.zou@uconn.edu

MS45

Predictable Forward Performance Processes for Loss-Averse Agents

Predictable forward performance processes (PFPPs) provide a stochastic optimal control framework for agents who manage randomly evolving systems, where system dynamics can only be prescribed over short time horizons. In this work, we explore PFPP problems in terms of a loss-averse agent who periodically recalibrates her model, with preferences over wealth that may lack concavity. Our analysis demonstrates that these problems can be reformulated into a concavified version involving Fredholm integral equations of the first kind, thereby extending the Volterra integral equations derived in our earlier studies. We further establish the existence and uniqueness of solutions within a characteristic system under markets characterized by atomless pricing kernels. To illustrate our theoretical findings, we consider the special case of a conditionally complete Black-Scholes model, providing numerical analyses of the solutions.

Shida Duan
University of Miami
Department of Mathematics
shidaduan@miami.edu

MS45

The Price of Information

When an investor is faced with the option to purchase additional information regarding an asset price, how much should she pay? To address this question, we solve for the indifference price of information in a setting where a trader maximizes her expected utility of terminal wealth over a finite time horizon. If she does not purchase the information, then she solves a partial information stochastic control problem, while, if she does purchase the information, then she pays a cost and receives partial information about the asset's trajectory. We further demonstrate that when the investor can purchase the information at any stopping time prior to the end of the trading horizon, she chooses to do so at deterministic time(s).

Xiaofei Shi
University of Toronto
xf.shi@utoronto.ca

MS45

Optimal Sequential Hypothesis Testing under Costly Information Acquisition

We study the sequential testing problem of two alternative hypotheses when observations are costly. Mathematically, we frame this as an optimal stopping problem in which the decision-maker aims to minimize classification errors plus a running cost proportional to the duration of information acquisition. By formulating a value function that captures both the expected decision risk and the running cost, we derive a set of variational inequalities that govern the evolution of the optimal strategy. Our analysis encompasses two different informational regimes: Gaussian and Poisson. We characterize the continuation regions within which further information collection is optimal and derive monotonicity properties with respect to different parameters, including the noise level, running cost, difference between the two hypotheses, and the relative importance of false positive and false negative errors. For Poisson signals, we utilize a diffusion approximation to study the asymptotic regime of frequent small jumps. Overall, our model advances understanding of how observation costs reshape classical sequential hypothesis testing. It also offers practical guidelines for designing adaptive and non-adaptive testing schemes in medicine, engineering, and related fields under real-world resource constraints.

Luhao Zhang
Johns Hopkins University
luhao.zhang@jhu.edu

MS45

Martingale Embedding: Exploration and Stopping

We fully characterize the possible outcomes of exploration and stopping: all state-time joint distributions achieved by stopping some martingale process with bounded variation. Utilizing this characterization, we provide a general methodology for solving an optimal exploration-stopping problem where the stopping payoff depends on state and time arbitrarily. We reveal the close relation between the pattern of exploration and time preference and apply it to study competitive exploration contests.

Weijie Zhong
Stanford University
Graduate School of Business
weijie.zhong@stanford.edu

MS46

Rough Differential Equations for Volatility

We introduce a canonical way of performing the joint lift of a Brownian motion W and a low-regularity adapted stochastic process X, extending [Diehl, Oberhauser, and Riedel (2015). A Levy area between Brown- ian motion and rough paths with applications to robust nonlinear filtering and rough partial differential equa- tions]. Applying this construction to the case where X is a one-dimensional fractional Brownian motion (possi- bly correlated with W) completes the partial rough path of [Fukasawa and Takano (2024). A partial rough path space for rough volatility]. We use this to model rough volatility with the versatile toolkit of rough differential equations (RDEs), namely by taking the price and volatility processes to be the solution to a single RDE. The lead-lag scheme of [Flint, Hambly, and Lyons (2016). Discretely sampled signals and the rough Hoff process] is extended to our fractional setting as an approximation theory for the rough path in the correlated case. Conti- nuity of the solution map transforms this into a numerical scheme for RDEs. We provide rigorous convergence results for different types of lead-lag approximations. Finally, we numerically test this framework and use it to calibrate a simple new rough volatility model to market data. This is joint work with Ofelia Bonesini (LSE), Emilio Ferrucci (Oxford) and Antoine Jacquier (Imperial College London).

Ioannis Gasteratos
Imperial College
i.gasteratos@imperial.ac.uk

MS47

High-performance Option Pricing with Discrete Dividends

If implemented right, integral equation methods for American option pricing offer significant speed and convergence advantages over alternatives. So far, however, these methods have predominantly been applied to smooth diffusive processes, and less is known about their usefulness when the underlying asset pays discrete dividends and, as a result, has discontinuous dynamics. In this talk, we examine this question from both a theoretical and numerical perspective, for European, Bermudan, and American options. We aim for a consistent framework that handles both put and call payouts; allows for flexible specification of the dependence between dividends and stock price levels; and can systematically incorporate non-zero borrow costs.

Leif Andersen
Bank of America
leif.andersen@bofa.com

MS47

Random Deep Splitting Method for Nonlinear Parabolic PDEs and PIDEs for High-Dimensional Option Pricing

In this talk we present a randomized extension of the deep splitting algorithm introduced in [Beck,Becker,Cheridito,Jentzen,Neufeld (2021)] using random neural networks suitable to approximately solve both high-dimensional nonlinear parabolic PDEs and PIDEs with jumps having (possibly) infinite activity. We provide a full error analysis of our so-called random deep splitting method. Moreover, we present several numerical examples relevant in the context of pricing of financial derivatives under default risk. We empirically demonstrate in all examples that our random deep splitting method can approximately solve nonlinear PDEs and PIDEs in 10000 dimensions within seconds. This talk is based on joint work with Philipp Schmocker and Sizhou Wu.

Ariel Neufeld
Nanyang Technological University
Singapore
ariel.neufeld@ntu.edu.sg

MS47

W-Shaped Implied Volatility Smiles and the Gaussian Mixture Model

We discuss relationships between the shape of the risk-neutral distribution (RND) and the implied volatility. The number of crossings of the implied volatility with a given level is bounded from below by the number of crossings of the RND with a log-normal distribution, and from above by the number of payoffs which are priced the same way under the two distributions. These results are applied to the Gaussian mixture model and are shown to restrict the shapes of the smile generated by the model. These include W-Shaped event-driven smiles.

Dan Pirjol
Stevens Institute of Technology
dpirjol@gmail.com

Paul Glasserman
Columbia University
pg20@columbia.edu

MS47

Short-Maturity Expansion for the Spx and Vix Options in Correlated Local-Stochastic Volatility Models

We derive the short-maturity asymptotics for European and VIX option prices in local-stochastic volatility models where the volatility follows a general diffusion process. Both out-of-the-money (OTM) and at-the-money (ATM) asymptotics are considered. Using large deviations theory methods, the asymptotics for the OTM options are expressed as a two-dimensional variational problem, which is reduced to an extremal problem for a function of two real variables. This extremal problem is solved explicitly in an expansion in log-moneyness. We derive series expansions for the implied volatility for European and VIX options which should be useful for model calibration. We give explicit results for two classes of local-stochastic volatility models relevant in practice, with Heston-type and SABR-type stochastic volatility. The leading-order asymptotics for at-the-money options are computed in closed-form. The asymptotic results reproduce known results in the literature for the Heston and SABR models and for the uncorrelated local-stochastic volatility model. The asymptotic results are tested against numerical simulations for a local-stochastic volatility model with bounded local volatility.

Xiaoyu Wang
Hong Kong University of Science and Technology
xiaoyuwang421@gmail.com

MS48

Towards Professional Readiness of LLMs in Financial Regulations

In this talk, Xiao-Yang Liu will showcase their FinGPT an open-source counterpart of BloombergGPT, on financial regulations. In particular, the team's two-year efforts on benchmarking financial large language models, with a zooming in Financial Regulations. He will also share ongoing projects in GenAI Research on Open Finance at Columbia University.

Xiao-Yang Liu
Columbia University
xl2427@columbia.edu

MS48

Can Large Language Models Trade? When AI Agents Compete in Stock Markets

We develop and validate a first-of-its-kind framework that enables Large Language Models (LLMs) to function as trading agents in financial markets. Our methodology demonstrates how to implement diverse trading strategiesincluding value investors, momentum traders, market makers, and contrariansthrough carefully engineered prompts in controlled market environments. Results reveal three key methodological findings: First, we establish that LLM agents can maintain semantic consistency in their trading decisions while processing complex market information. Second, we validate their capability to perform core market functions through systematic evaluation of price discovery and liquidity provision mechanisms. Third, we identify and characterize potential vulnerabilities in LLM-based trading systems, particularly regarding prompt engineering and strategy adherence. The framework's open-source architecture provides a comprehensive toolkit for agent validation, strategy development, and systematic testing of LLM capabilities in market environments. This contribution enables practitioners to develop and validate trading strategies, while providing researchers with a novel platform for studying LLM behavior in complex multi-agent systems.

Alejandro Lopez-Lira
University of Florida
alejandro.lopez-lira@warrington.ufl.edu

MS48

Open-Finllms: Open Multimodal Large Language Models for Financial Applications

Financial LLMs have shown promise in advancing financial tasks and enabling domain-specific applications. However, they are constrained by insufficient financial corpora, weak multimodal capabilities for complex financial data, and being assessed in narrow evaluation scenarios, leaving them ill-suited for real-world financial applications. To address these limitations, we introduce Open-FinLLMs, the first open-source financial LLMs capable of tackling diverse financial tasks across text, tabular, time-series, and chart data, excelling in zero-shot, few-shot, and supervised fine-tuning settings. The suite includes FinLLaMA, pre-trained on the largest and most comprehensive 52-billion-token corpus integrating text, tabular, and time-series data to capture domain-specific knowledge; FinLLaMA-Instruct, fine-tuned with 573K financial instructions to enhance task adaptability; and FinLLaVA, enhanced with 1.43M multimodal tuning pairs datasets, enables robust cross-modal reasoning. We comprehensively evaluate Open-FinLLMs across 14 financial tasks, 30 datasets, and 4 multimodal tasks in zero-shot, few-shot, and supervised fine-tuning settings, introducing two new multimodal evaluation datasets. Our results show that Open-FinLLMs outperforms other financial LLMs, including GPT-4, across financial NLP, decision-making, and multi-modal tasks, highlighting their potential to tackle real-world challenges.

Jimin Huang, Lingfei Qian
The Fin AI
jimin.huang@thefin.ai, lfqian94@gmail.com

MS49

A Mean-Field Game of Market Entry: Portfolio Liquidation with Trading Constraints

We consider both N-player and mean-field games of optimal portfolio liquidation in which the players are not allowed to change the direction of trading. Players with an initially short position of stocks are only allowed to buy while players with an initially long position are only allowed to sell the stock. Under suitable conditions on the model parameters we show that the games are equivalent to games of timing where the players need to determine the optimal times of market entry and exit. We identify the equilibrium entry and exit times and prove that equilibrium mean-trading rates can be characterized in terms of the solutions to a highly non-linear higher-order integral equation with endogenous terminal condition. We prove the existence of a unique solution to the integral equation from which we obtain the existence of a unique equilibrium both in the mean-field and the N-player game.

Paul P. Hager
University of Vienna
paul.peter.hager@univie.ac.at

MS49

Stochastic Portfolio Theory with Price Impact

Stochastic portfolio theory (SPT) is a powerful framework for portfolio selection, which is well suited for tackling questions related to outperformance of a buy-and-hold benchmark, such as the market portfolio. Many theoretical and some empirical studies have obtained performance guarantees for a class of functionally generated portfolios in the frictionless setting, but must less is understood in markets with frictions. Here we introduce modern price impact models to SPT, which are broad enough to allow for non-linear impact and impact decay and can handle general semimartingale trading strategies. In this framework we extend certain celebrated results in frictionless SPT, such as the master formula, and discuss relative arbitrage with respect to the market portfolio.

David Itkin
London School of Economics and Political Science
d.itkin@lse.ac.uk

MS49

Stochastic Graphon Games with Memory

We study a class of finite-player stochastic games with heterogeneous interactions and non-Markovian objective functionals. We derive the Nash equilibrium explicitly by converting the first-order conditions into a coupled system of stochastic Fredholm equations, which we solve in terms of operator resolvents. When the agents' interactions are modeled by a weighted graph, we formulate the corresponding non-Markovian continuum-agent game, where interactions are modeled by a graphon. We also derive the Nash equilibrium of the graphon game explicitly by first reducing the first-order conditions to an infinite-dimensional coupled system of stochastic Fredholm equations, then decoupling it using the spectral decomposition of the graphon operator, and finally solving it in terms of operator resolvents. We apply our results to various classes of stochastic games with heterogeneous interactions, including systemic risk models with delays, price impact games and stochastic network games.

Eyal Neuman, Sturmius Tuschmann
Imperial College London

eyaln13@gmail.com, s.tuschmann23@imperial.ac.uk

MS49

In-Context Operator Learning for Linear Propagator Models

We study operator learning in the context of linear propagator models for optimal order execution problems with transient price impact la Bouchaud et al. (2004) and Gatheral (2010). Transient price impact persists and decays over time according to some propagator kernel. Specifically, we propose to use In-Context Operator Networks (ICON), a novel transformer-based neural network architecture introduced by Yang et al. (2023), which facilitates data-driven learning of operators by merging offline pre-training with an online few-shot prompting inference. First, we train ICON to learn the operator from various propagator models that maps the trading rate to the induced transient price impact. The inference step is then based on in-context prediction, where ICON is presented only with a few examples. We illustrate that ICON is capable of accurately inferring the underlying price impact model from the data prompts, even with propagator kernels not seen in the training data. In a second step, we employ the pre-trained ICON model provided with context as a surrogate operator in solving an optimal order execution problem via a neural network control policy, and demonstrate that the exact optimal execution strategies from Abi Jaber and Neuman (2022) for the models generating the context are correctly retrieved.

Tingwei Meng
University of California, Los Angeles, U.S.
tingwei@math.ucla.edu

Moritz Voss
UC Los Angeles
Department of Mathematics
voss@math.ucla.edu

Nils Detering
Heinrich Heine University Düsseldorf
nils.detering@hhu.de

Giulio Farolfi
University of California, Los Angeles
gfarolfi@math.ucla.edu

Stanley J. Osher
UCLA
sjo@math.ucla.edu

Georg Menz
University of California, Los Angeles
georg.menz@gmail.com

MS50

Optimal Wealth Annuitization under Consumption Loss Aversion Relative to Past Maximum Consumption

We consider an optimal consumption problem for a retiree who invests in a Black-Scholes market and is allowed to annuitize her wealth by purchasing single-premium immediate life annuity income at any time and for any amount. We assume that the retiree is loss averse, and that she measures her utility of consumption not in absolute terms, but relative to her past maximum consumption. In particular, the retiree's consumption preference is represented by a general S-shaped utility function of her consumption rate divided by its running maximum. We provide a detailed analysis for the retiree's stochastic optimal control problem and obtain the feedback form of the optimal consumption, investment, and annuitization policies. Our analysis relies on the concavification principle, and is conducted by reducing the HJB variational inequality (of the concavified problem) into a free-boundary PDE. Some new arguments are developed to cope with the added dimension due to annuitization policy, as well as the generality of the utility function. We discuss properties of our optimal policies and compare them with existing policies through numerical examples.

Bahman Angoshtari
University of Miami
bangoshtari@miami.edu

MS50

On the Distribution Constrained Multiple Optimal Stopping

We consider the distribution constrained multiple optimal stopping. We first establish a duality in this context, and then combine the duality with optional cross-section theorem to prove the monotonicity principle.

Shuoqing Deng
Hong Kong University of Science and Technology
masdeng@ust.hk

MS50

Continuous-time Risk-sensitive Reinforcement Learning via Quadratic Variation Penalty

This paper studies continuous-time risk-sensitive reinforcement learning (RL) under the entropy-regularized, exploratory diffusion process formulation with the exponential-form objective. The risk-sensitive objective arises either as the agents risk attitude or as a distributionally robust approach against the model uncertainty. Owing to the martingale perspective, the risk-sensitive RL problem is shown to be equivalent to ensuring the martingale property of a process involving both the value function and the q-function, augmented by an additional penalty term: the quadratic variation of the value process, capturing the variability of the value-to-go along the trajectory. This characterization allows for the straightforward adaptation of existing RL algorithms developed for non-risk-sensitive scenarios to incorporate risk sensitivity by adding the realized variance of the value process. Additionally, I highlight that the conventional policy gradient representation is inadequate for risk-sensitive problems due to the nonlinear nature of quadratic variation; however, q-learning offers a solution and extends to infinite horizon settings. Finally, I prove the convergence of the proposed algorithm for Mertons investment problem and quantify the impact of temperature parameter on the behavior of the learning procedure. I also conduct simulation experiments to demonstrate how risk-sensitive RL improves the finite-sample performance in the linear-quadratic control problem.

Yanwei Jia
The Chinese University of Hong Kong
yanweijia@cuhk.edu.hk

MS50

Primal-dual Linear Programming Framework for Continuous-time Finite-state Mean Field gGames

TBA

Jiacheng Zhang
Chinese University of Hong Kong
jiachengzhang@cuhk.edu.hk

MS51

Machine Learning For Master Equations in Mean Field Games

Mean field games have been introduced to study games with many players. Since their introduction, they have found numerous potential applications and the theory has been extensively developed. While forward-backward systems of partial or stochastic differential equations can be used to characterize Nash equilibria with a fixed initial distribution, the Master equation introduced by P.-L. Lions provides a tool to solve the problem globally, for any initial condition. However this equation is a partial differential equation posed on the space of measures, which raises significant challenges to solve it numerically. In this talk, we will present several computational methods that have been proposed to tackle Master equations. Theoretical convergence results and numerical experiments will be presented. Mostly based on joint work with Asaf Cohen and Ethan Zell.

Mathieu Lauriere
NYU Shanghai
ml5197@nyu.edu

MS51

An Alpha-Potential Game Framework for N-Player Dynamic Games

This paper proposes and studies a general form of dynamic N-player non-cooperative games called α-potential games, where the change of a player's value function upon her unilateral deviation from her strategy is equal to the change of an α-potential function up to an error α. This framework is shown to reduce the challenge of finding α-Nash equilibria (NE) for a dynamic game to minimizing the α-potential function. An analytical characterization of α-potential functions is established, with α depending on the magnitude of the asymmetry of value functions' 2nd-order derivatives. For stochastic differential games with a controlled-diffusion-type state dynamic, α is characterized in terms of the number of players, the choice of admissible strategies, and the intensity of interactions and the level of heterogeneity among players. Distributed games and games with mean field (MF) interactions are analyzed to highlight the dependence of α on general game characteristics beyond the MF paradigm, which focuses on the limit of N with homogeneous players. The optimization problem associated with the α-NE is embedded into a conditional McKeanVlasov control problem. A verification theorem is established to construct α-NE via solutions to an infinite-dimensional Hamilton-Jacobi-Bellman equation, which is reduced to a system of ordinary differential equations for linear-quadratic games.

Xinyu Li
University of California, Berkeley
xinyu_li@berkeley.edu

MS51

Online Optimization of Stochastic Differential Equations

We develop a new continuous-time stochastic gradient descent method for optimizing over the stationary distribution of stochastic differential equation (SDE) models. The algorithm continuously updates the SDE model's parameters using an estimate for the gradient of the stationary distribution. The gradient estimate is simultaneously updated, asymptotically converging to the direction of steepest descent. We rigorously prove convergence of our online optimization algorithm for linear SDE models as well as a class of nonlinear SDEs. Numerical results demonstrate that the online optimization method performs well for a wide range of nonlinear examples. The proof requires analysis of the fluctuations of the parameter evolution around the direction of steepest descent. We prove bounds for the solutions of a new class of Poisson partial differential equations, which are then used to analyze the parameter fluctuations in the algorithm. Our algorithm is applicable to a range of mathematical finance applications involving statistical calibration of SDE models and stochastic optimal control for long time horizons where ergodicity of the data and stochastic process is a suitable modeling framework. Numerical examples explore these potential applications, including learning a neural network control for high-dimensional optimal control of SDEs.

Justin Sirignano
Mathematics, University of Oxford
justin.sirignano@maths.ox.ac.uk

MS51

Stochastic Differential Games on Graphs

In this talk, we present a new model for stochastic differential games on graphs, aiming to bridge game theory with network structures to capture the influence of graph structures on strategic interactions. Our framework supports heterogeneous player interactions across general graph structures, extending current models to encompass more complex, network-driven dynamics. We establish two main results: firstly, we demonstrate the convergence of fictitious play, along with numerical estimates of convergence rates that reflect key aspects of the graph structure. Secondly, we provide a semi-explicit construction of the Nash equilibrium, validated through numerical simulations and offering a reliable computational baseline for future applications in deep learning. Building on our theoretical findings and leveraging recent developments in graph neural networks, we propose a graph-dependent, non-trainable modification of the neural network architecture. By integrating this interpretable novel architecture with state-of-the-art game-solving algorithms, we demonstrate through numerical experiments that this architecture achieves comparable performance to standard architectures while using fewer parameters. This is joint work with Ruimeng Hu and Jihao Long.

Haosheng Zhou
University of California, Santa Barbara
hzhou593@ucsb.edu

MS52

Topological K-Means Clustering in Reproducing Kernel Hilbert Spaces

We propose a new topological clustering methodology, based on generalizing an empirical risk minimization framework, using a reproducing kernel Hilbert space (RKHS) for vectorized persistent homology representations of point clouds. In contrast to conventional Euclidean-based clustering methods which address only pairwise similarity among data points, our new approach of topological K-means clusters data based on similarity of shapes which are exhibited by the local vicinity of each data point at multiple scales. Thereby, topological clustering systematically captures the inherent local and global higher order data characteristics that are otherwise inaccessible with Euclidean-based clustering. We summarize the extracted shape characteristics of each local vicinity in the form of a persistence diagram (PD) and embed the PDs into a RKHS, which induces a distance among shapes of local vicinities in Hilbert space. Our derived theoretical guarantees on stability and consistency of the topological partitions are the first theoretical results of this kind at the intersection of topological data analysis and statistical inference. Additionally, we establish a number of new theoretical results on bounds of covering numbers in Hilbert spaces which are of independent interest in statistical learning theory. We discuss applications of topological clustering to the cryptocurrency risk analytics.

Matthew F. Dixon
Illinois Institute of Technology
matthew.dixon@iit.edu

Yuzhou Chen
University of California at Riverside
yuzhouc@ucr.edu

Yulia Gel
Virginia Tech
ygl@vt.edu

MS52

Why Topological Data Analysis Detects Financial Bubbles

Topological Data Analysis (TDA) has emerged as a powerful methodology in time-series analysis and signal processing. TDA takes as input complex, multi-dimensional, noisy data, and provides as output information on the 'shape' of the data. Such information is complementary to the one provided by statistical methods. Recent applications of TDA include detection of critical transitions in financial time series, particularly of financial bubbles. The methodology relies on time-delay coordinate embedding, which is used to construct, from the time-series, a sequence of point-clouds in some Euclidean space. The dynamics of the point-clouds unveils subtle patterns in the time-series. The effectiveness of TDA in detecting financial bubbles has been supported by empirical evidence. Here, we present an argument for why TDA can detect financial bubbles. Our argument relies on the Log-Periodic Power Law Singularity (LPPLS) model for financial bubbles. This model assert that the time series exhibit certain oscillatory patterns when approaching a tipping point. These oscillations determine holes in the point-clouds, which can be quantified by TDA. When approaching the tipping point of a bubble, there are significant changes in the nature of the oscillations, and consequently in the TDA output. These changes can be captured via persistence homology and yield early warning signals.

Marian Gidea
Yeshiva University
Marian.Gidea@yu.edu

MS52

Sparse Portfolio Selection Via Topological Data Analysis Based Clustering

This paper uses topological data analysis (TDA) tools and introduces a data-driven clustering-based stock selection strategy tailored for sparse portfolio construction. Our asset selection strategy exploits the topological features of stock price movements to select a subset of topologically similar (different) assets for a sparse index tracking (Markowitz) portfolio. We introduce new distance measures, which serve as an input to the clustering algorithm, on the space of persistence diagrams and landscapes that consider the time component of a time series. We conduct an empirical analysis on the S&P index from 2009 to 2022, including a study on the COVID-19 data to validate the robustness of our methodology. Our strategy to integrate TDA with the clustering algorithm significantly enhanced the performance of sparse portfolios across various performance measures in diverse market scenarios.

Anubha Goel
Tampere university, Finland
anubha.goel@tuni.fi

Damir Filipovic
EPFL
Swiss Finance Institute
damir.filipovic@epfl.ch

Puneet Pasricha
Indian Institute of Technology Ropar

Rupnagar,Punjab - 140001,India
puneet.pasricha@iitrpr.ac.in

MS53

Convergence of the Markovian Iteration for Coupled FBSDEs Via a Differentiation Approach

In this presentation, we study the Markovian iteration for coupled forward-backward stochastic differential equations (FBSDEs) with fully coupled forward drift. The convergence of this scheme for FBSDEs through the Y process was previously studied by Bender and Zhang (2008) "Time discretization and Markovian iteration for coupled FBSDEs." Ann. Appl. Probab. 18 (1) 143 - 177, and it was noticed that the main challenge of having Z coupling lies in controlling the Lipschitz constant of the decoupling field uniformly in time steps and iterations when applying the fixed-point argument. In this presentation, we address this difficulty through a differentiation approach for the Z process. As a result, we establish the well-posedness of the discretization of the FBSDE with fully coupled drift and obtain the convergence of the Markovian iteration for this type of equation. Finally, we provide several numerical examples to illustrate our theoretical results.

Zhipeng Huang
Mathematical Institute, Utrecht University
The Netherlands
z.huang1@uu.nl

MS53

Numerical Approaches to Backward Propagation of Chaos

Based on a general result for stability of backward propagation of chaos, specific numerical schemes for the Lvy case will be discussed. More precisely, initially it will be discussed the proper framework for discrete-time approximations of Lvy driven McKean–Vlasov BSDEs, both in terms of (discrete) Mean-Field BSDEs and of (discrete) McKean–Vlasov BSDEs. This will be followed by an exploration of the analogous framework for continuous-time approximations. If time permits, further properties will be discussed.

Alexandros Saplaouras
Dept. of Mathematics, ETH Zurich, Switzerland
alexandros.saplaouras@math.ethz.ch

MS53

Data-Driven Neural Network Methods for Portfolio Diversification

This talk explores the application of a data-driven neural network (NN) approach for international diversification of investment portfolios, particularly for Australian investors and pensioners. Specifically, we consider a portfolio construction that includes stock and bond indices from both Australian and American markets. Utilizing neural network methods, we optimize diversified rebalancing strategies across these markets to identify the most effective asset allocation. This work is expected to contribute to improved financial security and potentially better retirement outcomes for Australian investors and pensioners.

Hao Zhou
School of Mathematics and Physics
University of Queensland
h.zhou3@uq.net.au

Duy-Minh Dang
School of Mathematics and Physics, University of Queensland
Australia
duyminh.dang@uq.edu.au

MS54

Arbitrages in Perpetual Contracts

Perpetual contracts, extremely popular in cryptocurrency markets, are financial instruments that do not have an expiration date. The price of a perpetual contract is linked to the underlying price through regular funding swaps between long and short positions. This funding rate incentivizes traders to narrow the gap between perpetual contract prices and the underlying prices. We establish no-arbitrage bounds applicable to both linear and inverse perpetual contracts. Furthermore, we demonstrate that under these funding specifications, if perpetual contract prices fall outside the no-arbitrage bounds, arbitrage opportunities can be realized through dynamic trading strategies. Notably, even with trading costs on leading centralized exchanges like Binance, a straightforward trading strategy on BTC and ETH outperforms the benchmark carry strategy, highlighting the profitability potential in exploiting deviations from no-arbitrage intervals.

Min Dai
The Hong Kong Polytechnic University
mindai@polyu.edu.hk

Linfeng Li
National University of Singapore
e0431380@u.nus.edu

Chen Yang
The Chinese University of Hong Kong
Hong Kong
cyang@se.cuhk.edu.hk

MS54

Existence of Optimal Contract for Principal-agent Problem with Quadratic Cost Function

With recent advances in the mathematics community, the continuous-time principal-agent problem, a special case of the Stackelberg game, can be reformulated as a classical stochastic control problem. However, the existence of an optimal contract remains an open question. In the Markovian setting, this issue reduces to the existence of a classical solution to the associated Hamilton-Jacobi-Bellman (HJB) equation. The main technical difficulty arises from the degeneracy of the HJB equation. In this work, we consider the case where the agents effort cost function is quadratic. By exploiting the specific structure of the problem, we construct a classical solution. Furthermore, based on this result, we illustrate the principals contract and the agents optimal effort through numerical analysis.

Shuaijie Qian
The Hong Kong University of Science and Technology
sjqian@ust.hk

MS54

Optimal Design of Automated Market Makers on Decentralized Exchanges

Automated market makers are a popular mechanism used on decentralized exchange, through which users trade assets with each other directly and automatically through a liquidity pool and a fixed pricing function. The liquidity provider contributes to the liquidity pool by supplying assets to the pool, and in return, they earn trading fees from investors who trade in the pool. We propose a model of optimal liquidity provision in which a risk-averse liquidity provider decides the amount of wealth she would invest in the decentralized market to provide liquidity in a two-asset pool, trade in a centralized market, and consume in multiple periods. We derive the liquidity provider's optimal strategy and the optimal design of the automated market maker that maximizes the liquidity provider's utility. We find that the optimal unit trading fee increases in the volatility of the fundamental exchange rate of the two assets. We also find that the optimal pricing function is chosen to make the asset allocation in the liquidity pool efficient for the liquidity provider.

Chen Yang
The Chinese University of Hong Kong
Hong Kong
cyang@se.cuhk.edu.hk

MS54

Stochastic Volatility Models with Sticky Drawdown and Drawup Processes: A Deep Learning Approach

We propose a new financial model, the stochastic volatility model with sticky drawdown and drawup processes (SVSDU model), which captures the features of stochastic volatility, winning and losing streaks commonly observed across financial markets but not simultaneously accounted for by existing financial models. However, due to the complex dynamics involving sticky boundaries and the high dimensionality of the SVSDU model, option pricing under the SVSDU model is challenging: closed-form pricing formulae are unavailable and numerical methods such as Monte Carlo simulation and numerical partial differential equation (PDE) methods are time-consuming. In this paper, we develop a deep learning solution to the high-dimensional parametric PDE for the European option prices, yielding an analytical map from time, state, and parameters to option prices, which leads to an efficient calibration algorithm. In the numerical experiments and empirical studies, we confirm the efficiency and accuracy of the deep learning approach, and the importance of incorporating both winning and losing streaks. Our study provides a new perspective on the stochastic modeling of financial time series and option pricing.

Gongqiu Zhang
The Chinese University of Hong Kong (Shenzhen)
zhanggongqiu@cuhk.edu.cn

Pingping Jiang
Soochow University
ppjiang@suda.edu.cn

Yuhao Liu
Chinese University of Hong Kong
oxhowardliu@outlook.com

MS55

FinAudio: A Benchmark for Audio Large Language Models in Financial Applications

Financial Large Language Models (FinLLMs), such as OpenFinGPT and BloombergGPT, have demonstrated strong potential in various financial services. Moving beyond purely language-centric models, Multimodal Financial Foundation Models (MFFMs) aim to handle diverse financial data sources, including text, audio, images, and video. In this talk, we present an overview of the emerging landscape of MFFMs, with a particular focus on financial audio data. We will share our experience building FinAudio, a benchmark designed to evaluate Audio LLMs in financial scenarios, and discuss its implications for the future of multimodal financial intelligence.

Yupeng Cao
Stevens Institute of Technology
ycao33@stevens.edu

MS55

FIN-T1: Fusion LLM Agent with Gradient-based Reinforcement Learning for Financial Trading

Large language models (LLMs) have shown remarkable performance not only in language generation but also in tasks requiring reasoning, planning, and real-time decision-making. Their capabilities in financial natural language processing (NLP)including sentiment analysis, financial knowledge extraction, and market forecastinghave been extensively validated. Recently, a growing research trend aims to extend LLMs beyond token prediction, reframing them as general-purpose decision engines capable of translating multi-modal observations into actions. Motivated by advances in large reasoning models, we propose FIN-T1, a novel Financial Large Decision Model (FLDM) that repurposes pre-trained LLMs as policy networks within an online, on-policy deep reinforcement learning (DRL) framework for real-time financial trading. Departing from the conventional token-to-token modeling paradigm, FIN-T1 introduces a token-to-action framework, leveraging LLMs language understanding to generate trading decisions. The architecture incorporates a lightweight policy head trained via Proximal Policy Optimization (PPO), with most transformer layers frozen to preserve financial domain knowledge. It processes multi-modal inputsstructured market data and unstructured financial textencoded as prompts, enabling reasoning under uncertainty. Empirical results demonstrate that FIN-T1 outperforms traditional DRL baselines in Sharpe ratio, adaptation speed, and generalization, positioning it as an effective language-based agent for dynamic and high-volatility financial environments.

Guojun Xiong
Harvard University
gixiong@g.harvard.edu

MS55

A Multi-Purpose LLM-Based AI Agents Framework for Financial Decision Making

Large language models (LLMs) have shown potential in complex financial tasks, but sequential financial decision-making remains challenging due to the volatile environment and the need for intelligent risk management. While LLM-based agent systems have achieved impressive returns, optimizing multi-source information synthesis and decision-making through timely experience refinement is underexplored. We introduce FinCon, an LLM-based multi-agent framework with conceptual verbal reinforcement for diverse financial tasks. Inspired by real-world investment firm structures, FinCon employs a manager-analyst hierarchy, enabling synchronized cross-functional agent collaboration towards unified goals via natural language interactions. Its dual-level risk-control component enhances decision-making by monitoring daily market risk and updating systematic investment beliefs through self-critique. These conceptualized beliefs provide verbal reinforcement for future decisions, selectively propagated to relevant agents, improving performance while reducing unnecessary peer-to-peer communication costs. FinCon generalizes well across tasks, including single-stock trading and portfolio management.

Yangyang Yu
Stevens Institute of Technology
yyu44@stevens.edu

MS56

Continuous-time Q-learning for Jump-diffusion Models Under Tsallis Entropy

This paper studies the continuous-time reinforcement learning in jump-diffusion models by featuring the q-learning (the continuous-time counterpart of Q-learning) under Tsallis entropy regularization. Contrary to the Shannon entropy, the general form of Tsallis entropy renders the optimal policy not necessary a Gibbs measure, where the Lagrange and KKT multipliers naturally arise from some constraints to ensure the learnt policy to be a probability density function. As a consequence, the characterization of the optimal policy using the q-function also involves a Lagrange multiplier. In response, we establish the martingale characterization of the q-function under Tsallis entropy and devise two q-learning algorithms depending on whether the Lagrange multiplier can be derived explicitly or not. In the latter case, we need to consider different parameterizations of the optimal q-function and the optimal policy and update them alternatively in an Actor-Critic manner. We also study two financial applications, namely, an optimal portfolio liquidation problem and a non-LQ control problem. It is interesting to see therein that the optimal policies under the Tsallis entropy regularization can be characterized explicitly, which are distributions concentrated on some compact support. The satisfactory performance of our q-learning algorithms is illustrated in each example.

Xiang Yu
Department of Applied Mathematics
The Hong Kong Polytechnic University
xiang.yu@polyu.edu.hk

MS56

Adaptive Sampling Via Diffusion Models, with Applications to Dynamic Portfolio Choice Problems

Aimed at studying optimal portfolio problems in a model-free way, we propose to use generative models. Our idea is to use data sampled from the real model $\mathbb{P}$ with limited size to train a generative model $\mathbb{Q}$, from which we can sample easily and sufficiently. Importantly, we require $\mathbb{P}$ and $\mathbb{Q}$ to be close in adapted Wasserstein metric $\mathcal{AW}_2$. We choose this metric because dynamic problems with adapted structures are stable in the adapted Wasserstein metric. In the first part of the work, we propose to train and sample the generative model in an adapted way, which gives the $\mathcal{AW}_2$ bounds between the output distribution and data distribution. Importantly, the proposed adapted sampling method also facilitates the *conditional sampling*. In the second part of this work, we provide the sensitivity of the mean-variance portfolio optimization problems in $\mathcal{AW}_2$. In the third part, we propose a data-driven algorithm that relies on the innovative sampling methods in the first part, and sensitivity in the second part provides a convergence guarantee. This is an ongoing work with Erhan Bayraktar.

Fengyi Yuan
Department of Mathematics
The University of Michigan
fengyiy@umich.edu

MS56

Optimal Investment Stopping Problem with Random Horizon

In this paper we discuss a utility maximization problem that includes optimal investment, optimal stopping, and exogenous stopping in a finite horizon setting, which results in a fully nonlinear variational inequality with additional nonhomogeneous term for the value function. We use the dual method to formulate a minimax control stopping problem, which results in a semilinear variational inequality for the dual value function. We prove the existence, uniqueness, regularity, monotonicity and growth conditions of the solution to the dual variational inequality by the penalization method, and establish the verification theorems for general utility functions. We also show the existence and uniqueness of optimal solution and characterize the dual free boundary with a nonlinear integral equation involving a backward stochastic differential equation (BSDE) for power utility.

jie xing

Guizhou University of Finance and Economics
xj3553@mail.gufe.edu.cn

Jingtang Ma
Southwestern University of Finance and Economics
mjt@swufe.edu.cn

Harry Zheng
Imperial College London
h.zheng@imperial.ac.uk

MS57

Algorithmic Market Making in Spot Precious Metals

The primary challenge of market making in spot precious metals is navigating the liquidity that is mainly provided by futures contracts. The Exchange for Physical (EFP) spread, which is the price difference between futures and spot, plays a pivotal role and exhibits multiple modes of relaxation corresponding to the diverse trading horizons of market participants. In this paper, we model the EFP spread using a nested Ornstein-Uhlenbeck process, in the spirit of the two-factor Hull-White model for interest rates. We demonstrate the suitability of the framework for maximizing the expected PL of a market maker while minimizing inventory risk across both spot and futures. Using a computationally efficient technique to approximate the solution of the Hamilton-Jacobi-Bellman equation associated with the corresponding stochastic optimal control problem, our methodology facilitates strategy optimization on demand in near real- time, paving the way for advanced algorithmic market making that capitalizes on the co-integration properties intrinsic to the precious metals sector.

Alexander Barzykin
HSBC
alexander.barzykin@hsbc.com

Philippe Bergault
Université Paris Dauphine
bergault.philippe@protonmail.com

Olivier Gueant
Université Paris 1 - Panthéon Sorbonne
CERMSEM
olivier.gueant@gmail.com

MS57

Fredholm Approach to Nonlinear Propagator Models

We formulate and solve an optimal trading problem with alpha signals, where transactions induce a nonlinear transient price impact described by a general propagator model, including a power-law decay. Using a variational approach, we demonstrate that the optimal trading strategy satisfies a nonlinear stochastic Fredholm equation with both forward and backward coefficients. We prove the existence and uniqueness of the solution under a monotonicity condition reflecting the nonlinearity of the price impact. Moreover, we derive an existence result for the optimal strategy beyond this condition when the underlying probability space is countable. We introduce a novel iterative scheme and establish its convergence to the optimal trading strategy. We also provide a numerical implementation of the scheme that illustrates its convergence, stability, and the effects of concavity on optimal execution strategies under exponential and power-law decays. Based on a joint work with E. Abi Jaber, N. De Carvalho, E. Neuman and S. Tuschmann.

Alessandro Bondi
LUISS Guido Carli
abondi@luiss.it

MS57

A Prisoner's Dilemma for Dealers

We consider a central trading desk which aggregates the inflow of clients' orders with unobserved toxicity, i.e. persistent adverse directionality. The desk chooses either to internalise the inflow or externalise it to the market in a cost effective manner. In this model, externalising the order flow creates both price impact costs and an additional market feedback reaction for the inflow of trades. The desk's objective is to maximise the daily trading P&L subject to end of the day inventory penalization. We formulate this setting as a partially observable stochastic control problem and solve it in two steps. First, we derive the filtered dynamics of the inventory and toxicity, projected to the observed filtration, which turns the stochastic control problem into a fully observed problem. Then we use a variational approach in order to derive the unique optimal trading strategy. We illustrate our results for various scenarios in which the desk is facing momentum and mean-reverting toxicity. Our implementation shows that the P&L performance gap between the partially observable problem and the full information case are of order 0.01% in all tested scenarios.

Robert Boyce
Imperial College London
r.boyce23@imperial.ac.uk

MS57

Concave Cross-Impact

The price impact of large meta orders is well known to be a concave function of their size. We discuss how to extend models consistent with this "square-root law' to multivariate settings with cross impact, where trading each asset also impacts the prices of the others. In this context, we derive consistency conditions that rule out price manipulation. These basic requirements make risk-neutral trading problems tractable and also naturally lead to parsimonious model specifications that can be calibrated to historical data. We illustrate this with a case study using proprietary CFM meta order data.

Natascha Hey
Ecole Polytechnique

hey@ladhyx.polytechnique.fr

MS58

Fourier Neural Network Approximation for Transition Densities in Finance and Stochastic Control

This talk introduces FourNet, a novel single-layer feedforward neural network (FFNN) designed to approximate transition densities with known Fourier transforms (characteristic functions). FourNet leverages a Gaussian activation function to enable exact Fourier and inverse Fourier transformations, drawing parallels with Gaussian mixture models. We establish its capacity to approximate transition densities in the L_2 sense with arbitrary accuracy using a finite number of neurons. Parameters are learned by minimizing a loss function derived from the characteristic function, complemented by a strategic sampling approach to enhance training efficiency. Applications of FourNet in stochastic control problems, such as Bermudan option pricing and portfolio optimization, demonstrate its versatility and robustness. Examples include common dynamics in quantitative finance, such as the Heston stochastic volatility model and the self-exciting Queue-Hawkes jump process.

Duy-Minh Dang
School of Mathematics and Physics, University of Queensland
Australia
duyminh.dang@uq.edu.au

MS58

Lower Your Risk with Leverage: a Data Driven Neural Network Approach

We formulate a multi-period portfolio optimization problem which relaxes the long-only restriction and instead imposes a bound constraint on leverage. We propose a novel relaxed-constraint neural network (RCNN) model to approximate the optimal control. Our proposed RCNN model transforms the original leverage-constrained problem into an unconstrained optimization. We prove mathematically that the proposed RCNN control model can approximate the optimal relaxed-constraint strategy with arbitrary precision. We further propose to compute the optimal outperforming strategy over a benchmark based on cumulative quadratic shortfall (CS). Using U.S. historical market data from Jan 1926 to Jan 2023, we computationally compare and assess the proposed neural network approach to the optimal leverage-constrained strategy and long-only strategy respectively. We demonstrate that the leverage-constrained optimal strategy can achieve enhanced performance over the long-only strategy in outperforming a benchmark portfolio. Our main practical result is that using leverage optimally can actually reduce risk. * The talk in based joint work with Chendi Ni and Peter Forsyth.

Yuying Li
School of Computer Science
University of Waterloo
yuying@uwaterloo.ca

MS58

Time-Inconsistent Dynamic Optimal Asset Allocation Strategy Using Neural Networks

In our optimal portfolio asset allocation research, we take the perspective of a trader, who is allowed to trade in a risk-free bond, in stocks, and also in options, at a set of trading dates. A high quality objective function is able to numerically represent the investor's preferences of how much risk is acceptable for a certain level of potential profit. Moreover, we introduce market frictions aspects of incomplete markets and trading constraints. Regarding the market, we add transaction costs as well as a non-bankruptcy constraint and for the trading strategies, we introduce leverage constraints. For the resulting time-inconsistent optimization problem , we rely on numerical approximations. The discrete counterpart (in time and/or in probability space) is approximated by letting deep neural networks represent the trading strategies and optimizing with a gradient decent type algorithm. A general framework results, where we can invest in multiple assets, and deal with quite general objective functions.

Cornelis Oosterlee
Utrecht University, the Netherlands
c.w.oosterlee@uu.nl

MS58

Pricing and Hedging the Prepayment Option of Mortgages Under Stochastic Housing Market Activity

Prepayment risk embedded in fixed-rate mortgages forms a significant fraction of a financial institution's exposure. The embedded prepayment option (EPO) bears the same interest rate risk as an exotic interest rate swap with a suitable stochastic notional. With a focus on penalty-free prepayment because of the contract owner's relocation to a new house, we model the prepayment option value as an EU-type of interest rate receiver swaption with stochastic maturity matching the stochastic time of relocation. This is a convenient representation since it allows us to compute the EPO value in terms of well-known pricing formulae for EU-type swaptions. We investigate the effect of stochastic housing market (HM) activity as the explanatory variable for the distribution of the relocation time, as opposed to the assumption of a deterministic HM activity. We prove the HM (co)variance drives the EPO price difference between the stochastic HM setting and its deterministic counterpart. The EPO exposure is hedged using market instruments based on Delta-Gamma replication. Furthermore, since the HM activity is a non-tradable risk factor, we perform non-standard actuarial hedging focusing on controlling the EPO exposure yield by risky HM scenarios. This is achieved by imposing a path-wise exposure replication, possibly focusing on specific regions of interest in the exposure distribution, such as its tails.

Leonardo Perotti
Utrecht University, The Netherlands

l.perotti@uu.nl

MS59

CLVR Ordering of Transactions on AMMs

Trading on decentralized exchanges via an Automated Market Maker (AMM) mechanism has been massively adopted, with a daily trading volume reaching $1B. This trading method has also received close attention from researchers, central banks, and financial firms, who have the potential to adopt it to traditional financial markets such as foreign exchanges and stock markets. A critical challenge of AMM-powered trading is that transaction order has high financial value, so a policy or method to order transactions in a "good" (optimal) manner is vital. We offer economic measures of both price stability (low volatility) and inequality that inform how a "social planner" should pick an optimal ordering. We show that there is a trade-off between achieving price stability and reducing inequality, and that policymakers must choose which to prioritize. In addition, picking the optimal order can often be costly, especially when performing an exhaustive search over trade orderings (permutations). As an alternative we provide a simple algorithm, Clever Look-ahead Volatility Reduction (CLVR). This algorithm constructs an ordering which approximately minimizes price volatility with a small computation cost. We also provide insight into the strategy changes that may occur if traders are subject to this sequencing algorithm.

Nir Chemaya, Robert McLaughl, Dahlia Malkhi, Dingyue Liu
UC Santa Barbara
nir@umail.ucsb.edu, robert349@ucsb.edu, dahliamalkhi@ucsb.edu, dingyueliu@ucsb.edu

MS59

Automated Market Makers As a Mechanism for Self-Balancing Etfs

This paper presents a novel framework for constructing and maintaining target-weighted Exchange Traded Funds (ETFs) using Automated Market Makers (AMMs). By embedding AMM logic into ETF design, we propose a decentralized, self-rebalancing mechanism that addresses key inefficiencies of traditional target-weighted ETFs, such as high operational costs and tracking error. Through historical replication of two benchmark fundsVanguards Global Equity Fund (VHGEX) and Balanced Index Fund (VBIAX)we demonstrate the effectiveness of the AMM-ETF model in maintaining portfolio weights and delivering superior risk-adjusted returns. The framework leverages arbitrage incentives and fee redistribution to create a scalable, cost-efficient alternative to conventional ETF structures, offering a promising avenue for both institutional and retail investors.

Sean O'Leary
Stevens Institute
soleary@stevens.edu

MS59

Robust Measurement of Adversarial Behavior

We introduce the study of measures of manipulative behavior in multi-agent systems when economic incentives are involved. We define manipulation-free games as games in which whenever a coalition benefits from a victim's participation, the victim weakly benefits from the coalition's participation. We define surveillance metrics as computable proxies for the harm manipulation causes users as a way to measure "how far' a game is from being manipulation-free. We demonstrate the effectiveness of our framework in a case study of DeFi and blockchain systems, which are salient as real-world, rapidly emerging multi-agent systems with financial incentives for malicious behavior, for participation in algorithmic and AI systems, and in highlighting the need for new methods to measure levels of manipulative behavior. We introduce a surveillance metric for decentralized exchanges that measures manipulation in this context and demonstrates its effectiveness both theoretically and in a natural experiment to the Uniswap DeFi ecosystem. We show that whenever agents minimize our surveillance metric, the underlying decentralized exchange transforms from a highly manipulative game into a manipulation-free game. This underlines that data-driven auditing tools can have concrete contributions to improving security in multi-agent systems.

Matheus V. Xavier Ferreira
Harvard University, U.S.
matheusventuryne@gmail.com

MS59

Optimal Exiting for Liquidity Provision in Constant Function Market Makers

Providing liquidity to constant function market makers is often less profitable or favorable than simply holding assets, primarily because of impermanent loss. Using an optimal stopping approach, we show the existence of liquidity provision strategies that are profitable (excluding infrastructural fees) relative to holding, in which a liquidity provider (LP) exits the pool when the price ratio first hits certain thresholds. For the constant product pricing function, we derive closed-form solutions for the optimal thresholds. We demonstrate that pricing functions can be designed to maximize the expected time an optimally-acting LP exits the pool and backtest our strategy on Uniswap v2 data.

Agostino Capponi, Brian Zhu
Columbia University
ac3827@columbia.edu, bzz2101@columbia.edu

MS60

Optimal Control When the Drivers Future Is Partially Known and Applications to Insider Trading

Insider trading requires the investigation of anticipative controls. We consider a controlled stochastic differential equation driven by a Brownian motion B and assume that its future is partially known in the sense that the control ϕ may depend at time $t \in [0,T]$ on the path $(B(s), s \in [0,\tau(t)])$ for $\tau(t) > t$. Assuming there exists a pathwise solution Y^ϕ, the objective is to maximize an expected terminal cost functional given the initial information on the drivers future:

$$\operatorname*{esssup}_{\phi} \mathbb{E}(g(Y^\phi(T)) \mid \mathcal{F}^B_{\tau(0)}).$$

The framework of constantly updated partial information does not fit into either the standard stochastic optimal control theory or a purely deterministic setting. We develop a dynamic programming approach using Dupire's path-dependent derivatives and a suitable rough functional It's formula. In particular we provide an HJB equation characterizing the inherently path-dependent value function. As an application we maximize an insider's expected utility of terminal wealth in a market with price impact. Here the insider knows the stock prices for the near future and the path-dependent value function is given explicitly.

Franziska Bielert
Technische Universität Berlin
bielert@math.tu-berlin.de

MS60

Universal Approximation with It-type Signatures

Recently, a plethora of data-driven methods based on the signature of a path are being developed for applications to mathematical finance. These methods are fundamentally based on universal approximation theorems,which state that continuous functionals on the path space are approximated by linear functionals on the signature. For financial applications, this requires to compute the signature by adopting Stratonovich integration. However from a modeling perspective, It integration is typically the preferred choice of stochastic integration. In this talk, we introduce a notion of the signature of the path using a unifying framework for pathwise Stratonovich-type and It-type integration. Extending the path by suitable quadratic variation terms, we are able to deduce a pathwise universal approximation theorem for linear functionals on the signature. This can be translated into the probabilistic setting for the It signature of continuous semimartingales, making it particularly suitable for financial modeling. This talk is based on joint work with M. Ceylan and D. J. Prmel.

Mihriban Ceylan
University of Mannheim
mihriban.ceylan@uni-mannheim.de

Anna Kwossek
University of Vienna
anna.paula.kwossek@univie.ac.at

David Prömel
University of Mannheim
proemel@uni-mannheim.de

MS60

Path Characteristic Functions and Convergence of Stochastic Processes

In this talk I will briefly introduce how to use the so called unitary development technique to build Path Characteristic Functions (PCF) for stochastic processes, and discuss the relation between PCF and the convergence behavior of the laws of stochastic processes. Furthermore, I will also illustrate some applications of such tool in the market model generation task.

Chong Liu
ShanghaiTech University
cliustudent@gmail.com

MS60

Global Universal Approximation of Functional Input Maps on Weighted Spaces

We introduce so-called functional input neural networks defined on a possibly infinite dimensional weighted space with values also in a possibly infinite dimensional output space. To this end, we use an additive family to map the input weighted space to the hidden layer, on which a non-linear scalar activation function is applied to each neuron, and finally return the output via some linear readouts. Relying on Stone-Weierstrass theorems on weighted spaces, we can prove a global universal approximation result on weighted spaces for continuous functions going beyond the usual approximation on compact sets. This then applies in particular to approximation of (non-anticipative) path space functionals via functional input neural networks. As a further application of the weighted Stone-Weierstrass theorem we prove a global universal approximation result for linear functions of the signature. We also introduce the viewpoint of Gaussian process regression in this setting and emphasize that the reproducing kernel Hilbert space of the signature kernels are Cameron-Martin spaces of certain Gaussian processes. This paves a way towards uncertainty quantification for signature kernel regression. Joint work with Christa Cuchiero and Josef Teichmann.

Philipp Schmocker
Nanyang Technological University
philippt001@e.ntu.edu.sg

Christa Cuchiero
University of Vienna, Austria
christa.cuchiero@univie.ac.at

Josef Teichmann
Department of Mathematics
ETH Zurich
jteichma@math.ethz.ch

MS62

MeanVariance Portfolio Selection by Continuous-Time Reinforcement Learning: Algorithms, Regret Analysis, and Empirical Study

We study continuous-time mean–variance portfolio selection in markets where stock prices are diffusion processes driven by observable factors that are also diffusion processes yet the coefficients of these processes are unknown. Based on the recently developed reinforcement learning (RL) theory for diffusion processes, we present a general data-driven RL algorithm that learns the pre-committed investment strategy directly without attempting to learn or estimate the market coefficients. For multi-stock Black–Scholes markets without factors, we further devise a baseline algorithm and prove its performance guarantee by deriving a sublinear regret bound in terms of Sharpe ratio. For performance enhancement and practical implementation, we modify the baseline algorithm into four variants, and carry out an extensive empirical study to compare their performance, in terms of a host of common metrics, with a large number of widely used portfolio allocation strategies on S&P 500 constituents. The results demonstrate that the continuous-time RL strategies are consistently among the best especially in a volatile bear market, and decisively outperform the model-based continuous-time counterparts by significant margins.

Yilie Huang
Columbia University
yh2971@columbia.edu

MS62

Asymptotic Behavior of Heterogeneous Join-the-Shortest-Queue-D System

We consider a large-scale supermarket model with heterogeneous dispatchers and servers. Tasks arriving at a dispatcher are routed to one of its neighborhood servers based on the Join-the-Shortest-Queue-d policy. Arrival rates, service rates, and the bipartite graph representing connections between dispatchers and servers are all heterogeneous. The large population limit of the server queue length processes is given by independent but heterogeneous McKean-Vlasov type nonlinear stochastic processes whose probability distributions are fully coupled. Law of large number results are established as the system size grows and the underlying bipartite graphs converge to a graphon in the cut metric. This is joint work with Yan-Han Chen and Arka Ghosh.

Ruoyu Wu
Iowa State University
Department of Mathematics
ruoyu@iastate.edu

MS63

Topological Identification of Agent Status in Information Contagions: Application to Financial Markets

Cascade models serve as effective tools for understanding the propagation of information and diseases within social networks. Nevertheless, their applicability becomes constrained when the states of the agents (nodes) are hidden and can only be inferred through indirect observations or symptoms. This study proposes a Mapper-based strategy to infer the status of agents within a hidden information cascades using expert knowledge. To verify and demonstrate the method we identify agents who are likely to take advantage of information obtained from an inside information network. We do this using data on insider networks and stock market transactions. Recognizing the sensitive nature of allegations of insider trading, we design a conservative approach to minimize false positives, ensuring that innocent agents are not wrongfully implicated. The Mapper-based results systematically outperform other methods, such as clustering and unsupervised anomaly detection, on synthetic data. We also apply the method to empirical data and verify the results using a statistical validation method based on persistence homology. Our findings highlight that the proposed Mapper-based technique successfully identifies a subpopulation of opportunistic agents within the information cascades. The adaptability of this method to diverse data types and sizes is demonstrated, with potential for tailoring for specific applications.

Anubha Goel
Tampere university, Finland
anubha.goel@tuni.fi

Henri Hansan, Juho Kanniainen
Tampere University
henri.hansan@tuni.fi, juho.kanniainen@tuni.fi

MS63

The Changing Shape of Financial Risk

Topological data analysis (TDA) has emerged as a powerful tool for analyzing complex financial data, offering unique insights into market behavior and potential predictive capabilities. The method of persistence homology allows computing topological features at all scales and ranking these features by their persistence, which is particularly relevant for analyzing complex systems when there is no information available on the underlying stochastic process. Key applications include risk assessment by identifying hidden patterns, early detection and prediction of financial crises; identification of market regimes and transitions; risk analysis and portfolio management; network analysis of financial markets; enhancement of algorithmic trading strategies. Challenges and future directions are discussed, including applications of Crocker plots for practical use in financial decision-making as well as on the integration of TDA with traditional ML techniques and AI Reasoning. This novel econometric approach opens new horizons in financial forecasting. While TDA shows great promise in finance, its application is still evolving. As TDA methodology continues to mature, it has the potential to become an invaluable tool in the financial analyst's toolkit, complementing traditional methods of econometric analysis.

Yuri Katz
Yeshiva University
yuri.katz@cumulantx.com

MS63

Topological Risk

Financial markets are noisy dynamic systems ideally suited to the strengths of Topological Data Analysis (TDA). Investors in financial markets seek to make pecuniary returns from the assets in their portfolios. Higher returns are viewed as a reward for taking a higher risk of making losses. The tools of TDA help evaluate two key sources of risk. Firstly, the risk associated with higher volatility. High volatility is a direct indicator of high risk. Secondly, the risk associated with correlation. Correlation between assets reduces diversification options. Risk is divided into that which can be modelled, and that which is treated as truly random. Leveraging the potential of TDA offers a means to understand more of the observed volatility in financial markets and hence to model more of the observed risk. This presentation explores risk in asset pricing, formalises how TDA can contribute to risk modelling, and offers a direction for further research on TDA within asset pricing.

Simon Rudkin
Data Science, University of Manchester
simon.rudkin@manchester.ac.uk

Wanling Rudkin
University of Exeter Business School
University of Exeter
w.rudkin@exeter.ac.uk

MT1

Optimal Transport in Finance

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Johannes Wiesel
Carnegie Mellon University
jwiesel@andrew.cmu.edu

MT2

Applied Reinforcement Learning in Finance

Reinforcement learning (RL) is a powerful framework for decision-making. In this introductory lecture, we discuss popular RL methodologies and their financial applications, starting with simple multi-armed and contextual bandit problems and progressing to full RL problems. Applications include (i) optimal execution, where RL strategies minimize execution costs, (ii) optimal trading, where RL strategies learn to exploit signals, and (iii) market making, where RL agents maximize cash collected from setting the spread and manage inventory. The lecture concludes with an introduction to `mbt_gym`, a Python module that provides gym environments for training RL agents in model-based optimal trading within limit order books. We explore how to compute RL policies for well-known algorithmic trading papers from the mathematical finance community and how `mbt_gym` can serve as a benchmark.

Leandro Sanchez-Betancourt
Mathematical Institute, University of Oxford
Oxford-Man Institute of Quantitative Finance
sanchezbetan@maths.ox.ac.uk

PP1

Lppls-Net: Advanced Neural Networks for Predicting Temporal Critical Points in Financial Systems

Forecasting temporal critical points in financial markets is a complex, yet vital task. We consider the Log-Periodic Power Law Singularity (LPPLS) model for the detection of critical points in financial time series. We study the problem of fitting the model parameters to datasets. We introduce a supervised neural network trained on synthetically generated datasets designed to mimic realistic financial time series, incorporating white noise, autoregressive noise, and their combinations with controlled amplitude levels. The trained model achieves high accuracy, with a margin of error below 10% in predicting critical times on previously unseen and real-world market crash data. To enhance its predictive performance, we further integrate Long Short-Term Memory (LSTM) networks, leveraging their ability to capture sequential memory and dynamic patterns in time series data. By combining synthetic noise-infused datasets with advanced neural architectures, this research advances machine learning applications in financial modeling, offering improved forecast accuracy and robustness. These contributions highlight the potential for early detection of financial crises, aiding decision making and risk mitigation in volatile markets.

Varun Biyyala, Marian Gidea
Yeshiva University
vbiyyala@mail.yu.edu, Marian.Gidea@yu.edu

PP1

Kendall and Generalized Correlation Coefficients

Markowitz's optimal portfolio relies on the accurate estimation of correlations between asset returns, a difficult problem when the number of observations is not much larger than the number of assets. Using powerful results from random matrix theory, several schemes have been developed to 'clean' the eigenvalues of empirical correlation matrices [1,2]. By contrast, the (in practice equally important) problem of correctly estimating the eigenvectors of the correlation matrix has received comparatively little attention. Here we discuss a class of correlation estimators generalizing Kendall's rank correlation coefficient which improve the estimation of both eigenvalues and eigenvectors in data-poor regimes. Using both synthetic and real financial data, we show that these generalized correlation coefficients yield Markowitz portfolios with lower out-of-sample risk than those obtained with rotationally invariant estimators [3]. Central to these results is a property shared by all Kendall-like estimators but not with classical correlation coefficients: zero eigenvalues only appear when the number of assets becomes proportional to the square of the number of data points. [1] J. Bun, J.-P. Bouchaud, and M. Potters, My beautiful laundrette: Cleaning correlation matrices for portfolio optimization (2016) [2] D. Bartz, Cross-validation based nonlinear shrinkage (2016) [3] T. Espana, V. LeCoz, M. Smerlak, Kendall Correlation Coefficients for Portfolio Optimization (2024).

Tomas Espana
CentraleSupélec
tomas.espana@student-cs.fr

Victor LeCoz, Matteo Smerlak
Capital Fund Management

victor.lecoz@cfm.com, matteo.smerlak@cfm.com

PP1

Reconstruction of Financial Networks under Netting Constraints

We consider the problem of reconstructing financial networks from partial information in settings where specific constraints arise from netting activities. We propose a sampling methodology in this context and discuss its properties. To illustrate the method, we apply it to stylized examples.

Markus Karl
London School of Economics and Political Science
m.karl@lse.ac.uk

Luitgard Veraart
Department of Mathematics
London School of Economics and Political Science
l.veraart@lse.ac.uk

PP1

A Simple RL-Inspired Method for Optimal Stopping in Perpetual American Put Options

We propose a streamlined reinforcement learning (RL)-inspired method for evaluating the optimal early-exercise boundary π^* in perpetual American put options via discrete-time simulation. Let Δt be the time step and S_t the asset price at time t. At each step, the value at $t+1$ is given by

$$\text{Value at } (t+1) = \begin{cases} K - S_{t+1}, & \text{if } S_{t+1} \le \pi_{t+1} \text{ (exercise)}, \\ V(S_{t+1}), & \text{if } S_{t+1} > \pi_{t+1} \text{ (hold)}. \end{cases}$$

Because the option is perpetual, the option value $V(S)$ depends only on S. Once a stationary π^* is found, the exercise decision remains the same for all future times. Thus, at equilibrium, $V(S)$ remains unchanged. To locate such an equilibrium, we define the following loss function:

$$\text{loss} = \Big(V(S) - e^{-r\,\Delta t}\,\mathbb{E}\big[\text{Value at } (t+1)\big]\Big)^2,$$

and seek a point where both the value function and the threshold remain constant, rather than directly maximizing a payoff at every step. Once converged, $\pi_{t+1} = \pi_t = \pi^*$ and $V(S)$ remains the same for all t. Preliminary numerical tests show that this learned threshold aligns closely with known analytic boundaries under the Black–Scholes model. This balance of simplicity and flexibility suggests that fundamental RL concepts can serve as a practical alternative for addressing optimal stopping problems in finance, even without relying on full-scale PDE-based methods.

Eungpyo Kim, Jaemin Song
Department of Data Science, Chonnam National University
eungpyo0219@jnu.ac.kr, cantor@jnu.ac.kr

Jaegi Jeon
Chonnam National University, Department of Data Science
jaegijeon@jnu.ac.kr

PP1

Estimation of Tail Risk Metrics Using Synthetic Bar Chart Images from a Diffusion Model

We introduce a novel framework for estimating tail risk metrics, including Value-at-Risk (VaR) and Expected Shortfall (ES), using synthetic bar chart images generated by a generative diffusion model. Unlike traditional approaches that rely on predefined parametric assumptions about asset price dynamics, our method employs a data-driven approach, leveraging diffusion models trained on historical data to generate realistic financial scenarios. By simulating extensive synthetic paths, the method constructs an empirical return distribution that enables accurate estimation of tail risk metrics. The framework also provides a foundation for incorporating intraday metrics, such as those derived from opening, high, and low prices, to capture a more comprehensive view of risk dynamics. We validate our approach through extensive experiments on synthetic data and real-world financial data from the S&P 500 index. The results demonstrate that the proposed method captures tail risk features with greater precision than Monte Carlo simulations, establishing it as a robust solution for risk management challenges.

Seungju Lee, Jinseong Park, Jaewook Lee
Department of Industrial Engineering
Seoul National University, South Korea
tonysj@snu.ac.kr, jinseong@snu.ac.kr, jaewook@snu.ac.kr

PP1

Revisiting Ensemble Methods for Financial Reinforcement Learning

Reinforcement learning has demonstrated great potential for performing financial tasks. However, it faces two major challenges: policy instability and sampling bottlenecks. In this paper, we revisit ensemble methods with massively parallel simulations on graphics processing units (GPUs), significantly enhancing the computational efficiency and robustness of trained models in volatile financial markets. Our approach leverages the parallel processing capability of GPUs to significantly improve the sampling speed for training ensemble models. The ensemble models combine the strengths of component agents to improve the robustness of financial decision-making strategies. We conduct experiments in both stock and cryptocurrency trading tasks to evaluate the effectiveness of our approach. Massively parallel simulation on a single GPU improves the sampling speed by up to $1,746\times$ using $2,048$ parallel environments compared to a single environment. The ensemble models have high cumulative returns and outperform some individual agents, reducing maximum drawdown by up to 4.17% and improving the Sharpe ratio by up to 0.21.

Nikolaus LaHolzerst Name, Keyi Wang, Xiao-Yang Liu
Columbia University
nh2677@columbia.edu, kw2914@columbia.edu,

xl2427@columbia.edu

PP1

Optimal Portfolio Choice with Cross-Impact

We consider a class of optimal portfolio choice problems in continuous time where the agent's transactions create both transient cross-impact driven by a matrix-valued Volterra propagator, as well as temporary impact. We formulate this problem as the maximization of a revenue-risk functional, where the agent also exploits available information on a progressively measurable price predicting signal. We solve the maximization problem explicitly in terms of operator resolvents, by reducing the corresponding first order condition to a coupled system of stochastic Fredholm equations of the second kind and deriving its solution. We then give sufficient conditions on the matrix-valued propagator so that the model does not permit price manipulation. We also provide an implementation of the solutions to the optimal portfolio choice problem and to the associated optimal execution problem. Our solutions yield financial insights on the influence of cross-impact on the optimal strategies and its interplay with alpha decays. This is joint work with Eduardo Abi Jaber and Eyal Neuman.

Sturmius Tuschmann
Imperial College London
s.tuschmann23@imperial.ac.uk

Eduardo Abi Jaber
Ecole Polytechnique
eduardo.abi-jaber@polytechnique.edu

Eyal Neuman
Imperial College London
eyaln13@gmail.com

PP1

Quantformer: from Attention to Profit with a Quantitative Transformer Trading Strategy

In traditional quantitative trading practice, navigating the complicated and dynamic market trend presents a persistent challenge. Fully capturing various market variables, including long-term information and essential signals that may lead to profit remains a difficult task for learning algorithms. In order to tackle this challenge, this paper introduces quantformer, an enhanced transformer architecture, to build investment factors. By transfer learning from sentiment analysis, quantformer not only exploits its original inherent advantages in capturing long-range dependencies and modeling complex data relationships, but is also able to solve tasks with numerical inputs and accurately forecast future returns over a given period. This work collects more than 5,000,000 rolling data of 4,601 stocks in the Chinese capital market from 2010 to 2019. The results of this study demonstrated the model's superior performance in predicting stock trends compared with other 100 factor-based quantitative strategies, and stable profitability under different trading scales. Notably, the model's innovative use of transformer-like model to establish factors, in conjunction with market sentiment information, has been shown to enhance the accuracy of trading signals significantly, thereby offering promising implications for the future of quantitative trading strategies.

Zhaofeng Zhang
University of Michigan
Department of Mathematics
zhangzf@umich.edu

Banghao Chen
National University of Singapore
chenbanghao@u.nus.edu

Shengxin Zhu
Beijing Normal University, BNU-HKBU UIC
shengxin.zhu@bnu.edu.cn

Nicolas Langrené
BNU-HKBU United International College
nicolaslangrene@uic.edu.cn