QF Conference 2026, NUS Singapore

Timetable Overview

Program for Day 1 (10 June 2026 - Wednesday)

Time	Venue	Room	Program and Speakers
Session 1 **9:00AM - 11:00AM**	Level 1 Auditorium 1 (Room A) Level 2 Seminar Rooms (Room B-E)		In each session, **the last listed person** is designated as the session chair.
		Room A	Track A: Numerical Methods for Derivatives and PIDE/SDEs Speakers: Qin Sheng, Matthias Grützner, Ariel Neufeld, **Wonjae Lee**
		Room B	Track B: Model Uncertainty in Finance: From Trading to Hedging and Pricing Speakers: Yun Zhao, Mateo Rodriguez Polo, Marco Rodrigues, **Evgeny Kolosov**
		Room C	Track C: Model-Free, Fractional, and Martingale Pricing Speakers: Qi Yihan, Jayen Tan, Libo Li, **Giorgia Bifronte**
		Room D	Track D: Stochastic Machine Learning for High-Dimensional Finance Speakers: Dongyoung Lim, Sizhou Wu, Xiaoyu Wang, **Ying Zhang**
		Room E	Track E: Trading, Market Making, and Market Microstructure Speakers: Youssef Ouazzani Chahdi, Hao Xing, **Yumin Lu**
11:00AM - 11:30AM	Level 2 Town Plaza		Coffee Break
11:30AM - 11:45AM	Level 1 Auditorium		Opening Remarks
Keynote 1 **11:45AM - 12:45PM**	Level 1 Auditorium 1		Keynote Lecture 1: Beatrice Acciaio — The geometry of the adapted Bures-Wasserstein space
12:45PM - 1:00PM	Town Plaza		Group Photo
1:00PM - 2:00PM	Level 2 Town Plaza		Lunch
Session 2 **2:00PM – 4:00PM**	Level 1 Auditorium 1 (Room A) Level 2 Seminar Rooms (Room B-E)		In each session, **the last listed person** is designated as the session chair.
		Room A	Track A: Sustainable and Energy Finance Speakers: Thomas Kloster, Ludger Overbeck, Asma Khedher, **Wolfgang Härdle**
		Room B	Track B: Recent Advances in Stochastic Games Speakers: Chiara Rossato, Filippo Beretta, **Sturmius Tuschmann**
		Room C	Track C: Credit, Model Risk, and Volatility Speakers: Chang-Shu Chung, Luna Rigby, **Simon Feistle**
		Room D	Track D: Generative Models Speakers: Rama Cont, Anand Deo, **Philipp Schmocker**
		Room E	Track E: Exotic options and their numerical methods. Speakers: Yamazaki Kazutoshi, Kyunghyun Park, Daniel Chee, **Xiaolu Tan**
4:00PM - 4:30PM	Level 2 Town Plaza		Coffee Break
Keynote 2 **4:30PM - 5:30PM**	Level 1 Auditorium 1		Keynote Lecture 2: Steven Kou — Axiomatic Frameworks in Finance
5:45PM	Level 1 SRC pick-up pt		Board Bus to dinner venue
6:30PM	Jade @ Fullerton Hotel		Dinner

Program for Day 2 (11 June 2026 - Thursday)

Time	Venue	Room	Program and Speakers
Keynote 3 **9:30AM - 10:30AM**	Level 1 Auditorium 1		Keynote Lecture 3: Nizar Touzi — Some new developments on continuous time Principal-Agent
10:30AM - 11:00AM	Level 2 Town Plaza		Coffee Break
Session 3 **11:00AM - 1:00PM**	Level 1 Auditorium 1 (Room A) Level 2 Seminar Rooms (Room B-E)		In each session, **the last listed person** is designated as the session chair.
		Room A	Track A: Learning, Switching, Cyber-Risk, and Sentiment Signals Speakers: Umut Cetin, Junyan Ye, **Bingyan Han**
		Room B	Track B: Reinforcement Learning Speakers: Kailiang Liu, Hoa Nguyen, Richard Weiss, **Yongqi Zhou**
		Room C	Track C: Machine Learning meets Financial Networks Speakers: Yunming Hui, Mengzhong Ma, Daniel Liebau, **Simon Trimborn**
		Room D	Track D: Rough and High-Frequency Volatility Dynamics Speakers: Carsten Chong, Doerte Kreher, **Masaaki Fukasawa**
		Room E	Track E: Financial Networks, Systemic Risk, and Market Equilibrium Speakers: Rosalio G. Artes Jr., Li Ke, Quan Hoang Anh, **Paolo Guasoni**
1:00PM - 2:00PM	Level 2 Town Plaza		Lunch
Panel Discussion **2:00PM-3:00PM**	Level 1 Auditorium 1		Panel Discussion: AI in Finance
3:00PM			Group Photo

Program for Day 3 (12 June 2026 - Friday)

Time	Venue	Room	Program and Speakers
Session 4 **9:00AM - 11:00AM**	Level 1 Auditorium 1 (Room A) Level 2 Seminar Rooms (Room B-E)		In each session, **the last listed person** is designated as the session chair.
		Room A	Track A: Recent Advances in Portfolio Selection and Financial Technology Speakers: Chen Yang, Xiaofei Shi, Shuaijie Qian, **Qianyu Liu**
		Room B	Track B: Machine learning approach to computational finance Speakers: Fengyi Yuan, Xiaoli Wei, Yanwei Jia, **Gaozhan Wang**
		Room C	Track C: Safe Learning and Portfolio Decision-Making Speakers: Himanshu Choudhary, Paul Mangers Bastian, Jin Qian, **Luca De Gennaro Aquino**
		Room D	Track D: Stochastic Processes, Swaps, and Digital Markets Speakers: Cornelius Hanel, Yining Ding, Jonghwa Park, **Yutian Zhou**
		Room E	Track E: Portfolio Choice, Retirement, and Asset Allocation Speakers: Birungi Criscent, Manish, Ho Man Tai, **Kasper Larsen**
11:00AM - 11:30AM	Level 2 Town Plaza		Coffee Break
Keynote 4 **11:30AM - 12:30PM**	Level 1 Auditorium 1		Keynote Lecture 4: Huyen Pham — Learning Generative Dynamics with Soft Law Constraints: a McKean-Vlasov FBSDE approach
12:30PM - 1:30PM	Level 2 Town Plaza		Lunch
Session 5 **1:30PM - 3:30PM**	Level 1 Auditorium 1 (Room A) Level 2 Seminar Rooms (Room B-E)		In each session, **the last listed person** is designated as the session chair.
		Room A	Track A: Quantitative Finance and FinTech Speakers: Ling Qin, Yilie Huang, Heqing Huang, **Min Dai**
		Room B	Track B: Explainable AI for FinTech Speakers: Dennis Thumm, Hannah L. H. Lai, Wenhan Gao, **Paolo Recchia**
		Room C	Track C: Volatility Forecasting and Partial-Information Portfolio Models Speakers: Zhensi Tan, Li Hao, Christian Tezza, **Dannin Eccles**
		Room D	Track D: Digital Assets, FinTech, and Decentralised Finance Speakers: Huei-Wen Teng, Joerg Osterrieder, Daniel Traian Pele, **Stefan Lessmann**
		Room E	
3:30PM - 4:00PM	Level 2 Town Plaza		Coffee Break
Keynote 5 **4:00PM - 5:00PM**	Level 1 Auditorium 1		Keynote Lecture 5: Damir Filipovic — Joint Estimation of Conditional Mean and Covariance for Unbalanced Panels
5:00PM	Level 1 Auditorium		Closing Remarks

Talk Titles & Abstracts

Keynote Lectures

Keynote 1: Beatrice Acciaio

Title	The geometry of the adapted Bures-Wasserstein space
Affiliation	ETH Zürich
E-Mail	beatrice.acciaio@math.ethz.ch
Track/Session	Keynote 1
Abstract	The Adapted Wasserstein distance has been proven to be a suitable metric for comparing stochastic processes while accounting for the evolution of information in time, especially for the purpose of stochastic dynamic optimization in a context of model uncertainty. In this talk I will first recall the motivation and intuition behind such distance for general processes, and then restrict the attention to the space of Gaussian processes. I will introduce the adapted Bures-Wasserstein space, analogue of the classical Bures-Wasserstein space in optimal transport for the setting of stochastic processes, and provide a geometric analysis of such space. This includes full description of tangent spaces and geodesic curves, as well as explicit characterization of solutions to the adapted transport problem. I will conclude by presenting a sensitivity analysis w.r.t. the adapted Bures-Wasserstein distance. Based on joint works with D. Bartl, A. Grass, S. Hou and G. Pammer.

Keynote 2: Steven Kou

Title	Axiomatic Frameworks in Finance
Affiliation	Boston University
E-Mail	kou@bu.edu
Track/Session	Keynote 2
Abstract	This talk will survey recent research on axiomatic frameworks in finance. Although axiomatic frameworks may not lead to optimal outcomes, they can yield surprising representations based on reasonable axioms. Examples include (1) credit rating criteria, (2) risk-sharing rules, and (3) income inequality indices and tail risk measures.

Keynote 3: Nizar Touzi

Title	Some new developments on continuous time Principal-Agent
Affiliation	New York University
E-Mail	nizar.touzi@nyu.edu
Track/Session	Keynote 3
Abstract	Building on the reduction method of Sannikov and of Cvitanic–Possamaï–NT, we show that, for a broad class of problems, the Principal's problem can be further characterized through the solution of a forward-backward SDE. This approach makes it possible to treat the case of a Principal facing an optimally regime-switching Agent, which had remained open in the previous literature.

Keynote 4: Huyen Pham

Title	Learning Generative Dynamics with Soft Law Constraints: a McKean-Vlasov FBSDE approach
Affiliation	École Polytechnique
E-Mail	huyen.pham@polytechnique.edu
Track/Session	Keynote 4
Abstract	We propose a generative framework for learning stochastic dynamics from endpoint and intermediate distributional observations. The method formulates generation as a McKean-Vlasov control problem in which terminal and time-marginal laws are enforced through soft energy constraints. The associated optimality system is a forward-backward stochastic differential equation (FBSDE) whose backward component receives a

	continuous drift induced by the marginal law penalties. This provides a principled alternative to hard interpolation and optimal transport between observed distributions: the model learns a stochastic path law whose dynamics remain globally coupled through the mean-field objective. We derive the reduced FBSDE system for quadratic control cost and constant diffusion, connecting terminal and marginal law flat derivatives to score-like training signals. The resulting neural solver is evaluated on low-dimensional distributional benchmarks, where it recovers smooth stochastic paths matching prescribed marginal laws. In a higher dimensional ALAE latent space, endpoint supervision is sufficient to transport non-smiling faces toward smiling ones, reflecting the intrinsic regularity of the latent representation. We then study articulated human motion generation on a curated AMASS low-to-high transition dataset, using SMPL-H pose sequences and reduced pose representations. The experiments show that soft marginal law constraints can produce coherent stochastic trajectories whose intermediate distributions follow the observed evolution of human motion.

Keynote 5: Damir Filipovic

Title	Joint Estimation of Conditional Mean and Covariance for Unbalanced Panels
Affiliation	Ecole Polytechnique Fédérale de Lausanne (EPFL)
E-Mail	damir.filipovic@epfl.ch
Track/Session	Keynote 5
Abstract	We develop a nonparametric, kernel-based joint estimator for conditional mean and covariance matrices in large and unbalanced panels. The estimator is supported by rigorous consistency results and finite-sample guarantees, ensuring its reliability for empirical applications. We apply it to an extensive panel of monthly US stock excess returns from 1962 to 2021, using macroeconomic and firm-specific covariates as conditioning variables. The estimator effectively captures time-varying cross-sectional dependencies, demonstrating robust statistical and economic performance. We find that idiosyncratic risk explains, on average, more than 75% of the cross-sectional variance.

Session 1

Track A: Numerical Methods for Derivatives and PIDE/SDEs

Contributed Talk: Qin Sheng

Title	Exponential Splitting Operators and Error Bounds for Multiphysics and Option Pricing Applications
Affiliation	Baylor University
E-Mail	qin_sheng@baylor.edu
Track/Session	Session 1; Track A
Abstract	Exponential splitting methods, together with other advanced numerical techniques and adaptive grids, have been widely used to solve multiphysics and financial optimization problems. These approaches have demonstrated strong capabilities in handling major classes of mathematical models, including the Black-Scholes-Merton equations that govern the evolution of financial derivative prices. Current research challenges in this area include developing more accurate solution approximations for high-dimensional financial models and improving the computational efficiency of algorithms, particularly when deep neural network architectures are incorporated. There is a growing demand for advanced deep splitting frameworks and rigorous error estimates that can overcome the curse of dimensionality while ensuring numerical stability and preserving key financial invariants. This talk is based on a collaborative research effort with several scholars in the field, conducted under the 2026 NSF-JSPS Research Program.

Contributed Talk: Matthias Grützner

Title	Efficient Sobolev approximations for Partial Integro-Differential Equations and Stochastic Differential Equations with Jumps in high dimensions
Affiliation	ETH Zürich
E-Mail	matthias.gruetzner@math.ethz.ch
Track/Session	Session 1; Track A
Abstract	The curse of dimensionality is well known in numerical mathematics as an issue of a many classical grid-based models; the computational effort scales exponentially in the dimension of the considered space. A

	common approach is moving the problem into the stochastic setting through a suitable Feynman-Kac equation and tackling it with a Monte Carlo simulation, which famously does not suffer the curse of dimensionality. The key contribution of our research is the derivation of explicit error bounds in Sobolev norm, yielding bounds not only on the function itself, but also its derivative. Further, we extend the classical model to include jumps in the stochastic differential equation, that could e.g. represent market crashes, which are modelled by a Poisson point process.

Contributed Talk: Ariel Neufeld

Title	Deep Learning algorithm for solving high-dimensional nonlinear PIDEs in finance
Affiliation	NTU Singapore
E-Mail	ariel.neufeld@ntu.edu.sg
Track/Session	Session 1; Track A
Abstract	We present a (random) neural networks based algorithm which can efficiently solve high-dimensional nonlinear partial integro-differential equations (PIDEs) and apply this algorithm to price high-dimensional financial derivatives under default risk. We provide a full error analysis of our algorithm as well as empirically demonstrate that our algorithm can approximately solve nonlinear PIDEs in 10’000 dimensions within seconds. This talk is based on joint work with Philipp Schmocker and Sizhou Wu.

Contributed Talk: Wonjae Lee

Title	Fourier neural operators for non-markovian processes
Affiliation	Department of Mathematics, Seoul National University
E-Mail	phwjlee117@gmail.com
Track/Session	Session 1; Track A
Abstract	This talk introduces an operator-based neural network, the mirror-padded Fourier neural operator (MFNO), designed to learn the dynamics of stochastic systems. MFNO extends the standard Fourier neural operator (FNO) by incorporating mirror padding, enabling it to handle non-periodic inputs. We demonstrate that MFNOs can approximate solutions of path-dependent stochastic differential equations and Lipschitz transformations of fractional Brownian motions to an arbitrary degree of accuracy. Our theoretical analysis builds on Wong--Zakai type theorems and various approximation techniques. Empirically, the MFNO exhibits strong resolution generalization--a property rarely seen in standard architectures such as LSTMs, TCNs, and DeepONet. Furthermore, our model achieves performance that is comparable or superior to these baselines while offering significantly faster sample path generation than classical numerical schemes.

Track B: Model Uncertainty in Finance: From Trading to Hedging and Pricing

Invited Session: Yun Zhao

Title	Reinforcement Learning for Speculative Trading Under Exploratory Framework
Affiliation	Imperial College London
E-Mail	yun.zhao23@imperial.ac.uk
Track/Session	Session 1; Track B
Abstract	We study a speculative trading problem within the exploratory reinforcement learning (RL) framework of Wang et al. [2020]. The problem is formulated as a sequential optimal stopping problem over entry and exit times under general utility function and price process. We first consider a relaxed version of the problem in which the stopping times are modeled by the jump times of Cox processes driven by bounded, non-randomized intensity controls. Under the exploratory formulation, the agent's randomized control is characterized via the probability measure over the jump intensities, and their objective function is regularized by Shannon's differential entropy. This yields a system of the exploratory HJB equations and Gibbs distributions in closed-form as the optimal policy. Error estimates and convergence of the RL objective to the value function of the original problem are established. Finally, an RL algorithm is designed, and its implementation is showcased in a pairs-trading application.

Invited Session: Mateo Rodriguez Polo

Title	Equilibrium Prices Under Hidden Markov Fundamentals
Affiliation	ETH Zürich
E-Mail	mateo.rodriguezpolo@math.ethz.ch

Track/Session	Session 1; Track B
Abstract	We study equilibrium asset prices in a economy in which dividend growth is driven by a hidden finite-state Markov chain and investors have Epstein--Zin recursive preferences. We first characterise a large class of equilibria and show that the Markovian equilibrium is selected by a stability requirement. We show that the positive Markovian solution is unique, strictly increasing and analyze its convexity. Finally, we derive option prices under the risk-neutral filtering dynamics. We find that learning makes stock return volatility belief-dependent and generates a sign-changing implied-volatility skew: low beliefs make upside calls relatively expensive, whereas high beliefs make downside puts relatively expensive.

Invited Session: Marco Rodrigues

Title	Robust Hedging of American Options Under Model Uncertainty
Affiliation	TU Berlin and Weierstrass Institute
E-Mail	marco-rod@outlook.com
Track/Session	Session 1; Track B
Abstract	We study the problem of hedging American options when the dynamics of the underlying assets are uncertain. Our approach is based on constructing a Snell envelope under a sublinear expectation, which allows us to characterize the robust superhedging price. We establish a dual representation of this price and prove the existence of a minimal superhedging strategy. The framework accommodates both continuous processes and processes with jumps and nonvanishing diffusion. As an application, we discuss financial market models in which uncertainty is described through constraints on the drift, volatility, and jump behavior of the underlying assets.

Invited Session: Evgeny Kolosov

Title	On Arbitrage-Free Prices of American Options
Affiliation	ETH Zürich
E-Mail	evgeny.kolosov@math.ethz.ch
Track/Session	Session 1; Track B
Abstract	One of the key questions in robust finance is the study of arbitrage-free models consistent with observed derivative prices. Most of the literature focuses on the case where the observations are prices of European options. It is well known that the full curve of European option prices for given maturities determines the marginal distributions of the underlying asset price, and the existence of a compatible arbitrage-free model is equivalent to these distributions being in convex order. However, when observing prices of American options, the problem becomes significantly more complex, as they do not provide complete information regarding the marginal distributions. In this work, we investigate the conditions for the existence of arbitrage-free models calibrated to American option prices. This leads to a novel extension of the convex order concept, which we term “biased convex order.” We explore the properties of this new order and provide a Strassen-type result to characterize it.

Track C: Model-Free, Fractional, and Martingale Pricing

Contributed Talk: Qi Yihan

Title	Empirical Analysis of the Model-Free Valuation Approach: Hedging Gaps, Conservatism, and Trading Opportunities
Affiliation	NUS Singapore
E-Mail	e1124131@u.nus.edu
Track/Session	Session 1; Track C
Abstract	This paper studies the empirical performance of model-free valuation approaches for financial derivatives using historical option price data from several S&P 500 constituents over 2018–2022. We evaluate the realized gap between model-free super-hedging strategies and the ex post payoff of the derivative, with the aim of quantifying how conservative model-free methods are in practice. We compare the model-free approach with industry-standard benchmarks, including the Heston model, and find that model-free hedging is only moderately more conservative while retaining the advantage of model independence. Based on the empirical distribution of hedging gaps, we further construct a trading strategy that exploits the conservatism embedded in model-free bounds. The results suggest that model-free methods can deliver informative hedging and valuation benchmarks while also generating economically meaningful trading opportunities with explicit downside-risk control.

Contributed Talk: Jayen Tan

Title	Arbitrage-free valuation in fractional Brownian motion markets
Affiliation	Cornell University
E-Mail	jt734@cornell.edu
Track/Session	Session 1; Track C
Abstract	No-arbitrage pricing models rely exclusively on semimartingale processes to preclude arbitrage opportunities and ensure compatibility with Itô calculus. This restriction, however, excludes a wide array of empirically relevant dynamics, including persistency, anomalous diffusion, and local fractionality, thereby limiting applicability to some real-world complexities. We develop the risk-neutral valuation framework for derivatives whose underlyings are driven by fractional Brownian motion, a non-semimartingale generalization of classical Brownian motion that captures these phenomena. Motivated by empirical evidence of persistent temperature dynamics, we apply the framework to the weather derivatives market.

Contributed Talk: Libo Li

Title	Vulnerable European and American options in a hazard process model
Affiliation	University of New South Wales
E-Mail	libo.li@unsw.edu.au
Track/Session	Session 1; Track C
Abstract	Vulnerable options of a European and American style are studied in a hazard process setup. We show that the upper price of a vulnerable European option coincides with the price of an American option where the holder's exercise times are constrained to the support of the hazard process. For a vulnerable American option, we show that the upper price equals the price of an American option with unrestricted exercise times, whereas the lower price coincides with the price of a game option where the issuer's exercise times are constrained to the support of the hazard process.

Contributed Talk: Giorgia Bifronte

Title	On Schrödinger and Bass Martingales
Affiliation	University of Vienna
E-Mail	giorgia.bifronte@univie.ac.at
Track/Session	Session 1; Track C
Abstract	A Bass martingale attempts to be as Gaussian as possible, as measured by an adapted Wasserstein distance, while interpolating prescribed initial and terminal marginals. A martingale Schrödinger bridge attempts the same, but measured in the relative entropy sense. Beyond their theoretical interest, these martingales have been proposed for the purpose of model calibration in finance. In this talk we present new findings on martingale Schrödinger bridges, including a discrete- to continuous-time correspondence, a numerical algorithm for their computation, and a connection to the celebrated Föllmer processes. We also discuss some of the similarities / differences, between Bass and Schrödinger martingales.

Track D: Stochastic Machine Learning for High-Dimensional Finance

Invited Session: Dongyoung Lim

Title	AI-Accelerated Calibration of Financial Models
Affiliation	Ulsan National Institute of Science & Technology (UNIST)
E-Mail	dlim@unist.ac.kr
Track/Session	Session 1; Track D
Abstract	Financial model calibration is a central component of modern financial modeling. However, due to its formulation as an inverse problem, calibration is often computationally expensive and becomes increasingly challenging for advanced models. In this work, we propose a novel AI-driven approach that enables fast and accurate calibration of financial models. Our method significantly reduces computational cost while maintaining high fidelity. In particular, we address the inherent ill-posedness of the calibration problem by introducing a principled learning-based framework that stabilizes the inverse mapping and improves robustness.

Invited Session: Sizhou Wu

Title	Modern Numerical Methods for Semi-linear PIDEs using Monte Carlo and Random Neural Networks for High-dimensional Option Pricing
Affiliation	Shanghai University of Finance and Economics
E-Mail	wusizhou@sufe.edu.cn
Track/Session	Session 1; Track D
Abstract	In this talk, we introduce two modern numerical algorithms to approximate semi-linear parabolic partial integro-differential equations (PIDEs). The first algorithm is the so-called multilevel Picard (MLP) approximation based on Monte Carlo simulation, and the other one is the random deep-splitting algorithm using random neural networks. For both algorithms, we provide full error analysis and complexity analysis. Furthermore, we present a numerical example in which both algorithms are applied to price high-dimensional financial derivatives in a jump-diffusion model.

Invited Session: Xiaoyu Wang

Title	Generalized EXTRA stochastic gradient Langevin dynamics
Affiliation	Hong Kong University of Science and Technology (Guangzhou)
E-Mail	xiaoyuwang@hkust-gz.edu.cn
Track/Session	Session 1; Track D
Abstract	Langevin algorithms are popular Markov Chain Monte Carlo methods for Bayesian learning, particularly when the aim is to sample from the posterior distribution of a parametric model, given the input data and the prior distribution over the model parameters. Their stochastic versions such as stochastic gradient Langevin dynamics (SGLD) allow iterative learning based on randomly sampled mini-batches of large datasets and are scalable to large datasets. However, when data is decentralized across a network of agents subject to communication and privacy constraints, standard SGLD algorithms cannot be applied. Instead, we employ decentralized SGLD (DE-SGLD) algorithms, where Bayesian learning is performed collaboratively by a network of agents without sharing individual data. Nonetheless, existing DE-SGLD algorithms induce a bias at every agent that can negatively impact performance; this bias persists even when using full batches and is attributable to network effects. Motivated by the EXTRA algorithm and its generalizations for decentralized optimization, we propose the generalized EXTRA stochastic gradient Langevin dynamics, which eliminates this bias in the full-batch setting. Moreover, we show that, in the mini-batch setting, our algorithm provides performance bounds that significantly improve upon those of standard DE-SGLD algorithms in the literature. Our numerical results also demonstrate the efficiency of the proposed approach.

Invited Session: Ying Zhang

Title	Robust SGLD algorithm for solving non-convex distributionally robust optimisation problems
Affiliation	Hong Kong University of Science and Technology (Guangzhou)
E-Mail	yingzhang@hkust-gz.edu.cn
Track/Session	Session 1; Track D
Abstract	We develop a Stochastic Gradient Langevin Dynamics (SGLD) algorithm tailored for solving a certain class of non-convex distributionally robust optimisation problems. By deriving non-asymptotic convergence bounds, we build an algorithm which for any prescribed accuracy $\varepsilon>0$ outputs an estimator whose expected excess risk is at most ε. As a concrete application, we employ our robust SGLD algorithm to solve the (regularised) distributionally robust Mean-CVaR portfolio optimisation problem using real financial data. We empirically demonstrate that the trading strategy obtained by our robust SGLD algorithm outperforms the trading strategy obtained when solving the corresponding non-robust Mean-CVaR portfolio optimisation problem using, e.g., a classical SGLD algorithm. This highlights the practical relevance of incorporating model uncertainty when optimising portfolios in real financial markets.

Track E: Trading, Market Making, and Market Microstructure

Contributed Talk: Youssef Ouazzani Chahdi

Title	Exact Conditional Simulation of Point Processes for Counterfactual Market Impact Evaluation
Affiliation	CentraleSupélec
E-Mail	youssef.ouazzanichahdii@gmail.com
Track/Session	Session 1; Track E
Abstract	Market impact is defined as the difference between the observed price trajectory under a given execution strategy and the counterfactual trajectory that would have prevailed without it. Since this counterfactual is

	unobservable, estimating market impact requires simulating alternative paths under the same realized market randomness. We address this by studying the conditional simulation of point processes under perturbed intensities. Given an observed Poisson process, we characterize the conditional law of the underlying Poisson random measure and derive an explicit conditional structure. This yields an exact, event-driven algorithm that reconstructs counterfactual paths on a common randomness source, enabling rigorous pathwise market impact estimation.

Contributed Talk: Hao Xing

Title	Latent fragility and clustered withdrawals in bank runs
Affiliation	Boston University Questrom
E-Mail	haoxing@bu.edu
Track/Session	Session 1; Track E
Abstract	We develop a dynamic mean-field game of bank runs to explain how heterogeneous and gradual depositor shocks generate clustered withdrawals. Depositors differ in private states affecting withdrawal incentives and face strategic complementarities because the bank's failure intensity rises with the cumulative withdrawn share. The model's central mechanism is latent fragility: vulnerable depositors accumulate inside the bank and may wait collectively until aggregate fragility becomes sufficiently severe, at which point they withdraw together in a concentrated wave. We establish the existence and ordering of equilibria in a non-stationary environment and show that stronger strategic complementarities lead to earlier and larger withdrawal clusters. In a tractable two-type model, we characterize the timing and size of clustered withdrawals in closed form. Introducing a common aggregate state generates multiple withdrawal waves by coordinating the exit of accumulated vulnerable depositors. The paper contributes to the bank-run literature by endogenizing the timing and magnitude of clustered withdrawals in a heterogeneous-agent dynamic setting. This is a joint work with Jodi Dianetti, Giorgio Ferrari, and Yunzhi Hu

Contributed Talk: Yumin Lu

Title	Optimal market making for a polymarket topic
Affiliation	University College London
E-Mail	ucahyyl@ucl.ac.uk
Track/Session	Session 1; Track E
Abstract	We study a market making problem for a polymarket topic in prediction market, featuring inventory penalisation and state-dependent order-arrival intensities. We derive the dynamics of the implied probability process and formulate the associated HJB equation under a quadratic value function. An approximation method leads to Feynman-Kac representation of value function. Numerical sensitivity graphs are analyzed to illustrate how the optimal bid/ask spread respond to changes in time and asset prices.

Session 2

Track A: Sustainable and Energy Finance

Invited Session: Thomas Kloster

Title	The fine structure of electricity price volatility
Affiliation	Aarhus BBS
E-Mail	tkk@econ.au.dk
Track/Session	Session 2; Track A
Abstract	We conduct the first rigorous study of electricity price volatility for the full panel of electricity prices across three European generation zones. By interpreting the observed day-ahead prices as local averages of a latent price process governed by a stochastic partial differential equation, we develop estimators of the weekly integrated variance. The inherently infinite dimensional setting introduce several complications that are not relevant in the conventional finite dimensional semimartingale setting, and we spend considerable effort in dealing with these. In particular, we must account for both mean-reversion in prices and semigroup-smoothing in the estimated variance. We provide a detailed decomposition and interpretation of the empirical estimates across three vastly different European generation zones, namely Germany, Norway, and Spain. Our findings indicate that each zone has very different drivers of volatility, and that the impact of generation variables differs considerably. We document that leverage effects appear to be present at first sight, but disappear once we condition on suitable state variables, thereby showing that electricity price volatility does not generally exhibit asymmetric responses to price shocks.

Invited Session: Ludger Overbeck

Title	Stochastic Modelling and Forecasting with Applications to Wind Derivatives: The Australian Case
Affiliation	Justus-Liebig-University Gießen
E-Mail	ludger.overbeck@math.uni-giessen.de
Track/Session	Session 2; Track A
Abstract	Wind electricity generation in Australia is accelerating, and South Australia has one of the highest penetration rates globally, so pricing and hedging risks associated with wind power generation are becoming of critical importance. This paper proposes stochastic models for the dynamics and attributes of wind capacity utilisation, a typical underlying asset of wind derivatives. Using daily wind generation data from the five states of the Australian national electricity market (NEM), we will identify the nature of wind capacity utilisation (generation/capacity) in terms of seasonality, mean reversion, and roughness. We employ rough and Lévy stochastic models to capture these dynamics and gauge their efficiency for calibration applications. Finally, we will assess the forecasting performance of the proposed models. This research will offer practical insights into the role of wind derivatives in the energy transition and the integration of wind energy into the electricity markets.

Invited Session: Asma Khedher

Title	Lévy-Driven CARMA Processes on Cones in Banach Spaces and Applications to Space-Time Modelling
Affiliation	University of Amsterdam
E-Mail	A.Khedher@uva.nl
Track/Session	Session 2; Track A
Abstract	In this paper, we examine continuous-time autoregressive moving-average (CARMA) processes on Banach spaces driven by Lévy subordinators. We establish their existence and cone invariance, investigate their first- and second-order moment structure, and derive explicit conditions for stationarity. In particular, we introduce measure- and density-valued CARMA processes by formulating CARMA dynamics on spaces of finite signed measures and suitable function spaces. By appropriately selecting the input, transition, and output operators in the underlying state-space representation, we obtain cone-preserving CARMA dynamics that remain within the cone of positive measures or densities over a spatial domain. We further demonstrate how such cone-valued infinite-dimensional CARMA processes provide flexible models for the evolution of functionals of spatio-temporal random fields. Finally, we relate our framework to existing CARMA-type models in both finite and infinite dimensions, highlighting its generality and applicability.

Invited Session: Wolfgang Härdle

Title	Stable Coin Risk
Affiliation	Humboldt-Universität zu Berlin
E-Mail	haerdle@hu-berlin.de
Track/Session	Session 2; Track A
Abstract	An analysis one the stability and systemic risk of major stable coins from 2020–2025 is presented. Employing the Financial Risk Meter (FRM) https://www.theida.net/financial-risk-meter-2/ we study 11 leading fiat-, crypto-, and commodity-backed stable coins alongside key macro factors. We show that stable coins remain largely resilient in normal periods but exhibit sharp vulnerabilities during tail events such as COVID-19, UST/LUNA, FTX, the SVB-induced USDC de-peg, and regulatory shocks. FRM results indicate that risk flows primarily from macro markets into stable coins rather than between coins. Portfolio analysis demonstrates substantial volatility reduction through diversification. Overall, stable coins act as stability anchors—until peg credibility erodes.

	See this CL: https://quantinar.com/course/976/stable-coin-risks?q=stable coin

Track B: Recent Advances in Stochastic Games

Invited Session: Chiara Rossato

Title	Optimal incentives under mean-variance preferences
Affiliation	ETH Zürich
E-Mail	chiara.rossato@math.ethz.ch
Track/Session	Session 2; Track B
Abstract	We study a continuous-time principal-agent problem in which a principal hires a sophisticated agent and promises a lump-sum payment at a finite terminal time. To address the moral hazard problem faced by the principal, we build on recent work by Cvitanić, Possamaï, and Touzi, as well as Hernández and Possamaï, and show that, without loss of generality, the analysis can be restricted to a class of revealing contracts for which the agent's optimal effort can be characterised explicitly. This reduction allows us to use the backward stochastic differential equation associated with the agent's continuation utility as a controlled state variable for the principal, and to reformulate the moral hazard problem as a standard stochastic control problem with stochastic target constraints.

Invited Session: Filippo Beretta

Title	Closed-loop equilibria in leader–follower games with private and common infor...
Affiliation	ETH Zürich
E-Mail	filippo.beretta@math.ethz.ch
Track/Session	Session 2; Track B
Abstract	We study a linear-quadratic Stackelberg stochastic differential game between a more informed leader and a follower. Similar models are attracting increasing attention in Finance, with a wide range of recent applications ranging from price impact to climate finance. Despite this momentum, the mathematical tools to tackle these problems are still limited and lack a unified framework. It is well known that even in full-information settings, different notions of admissible strategies may generate different Stackelberg equilibria. In contrast with most of the existing partial-information literature, we focus on closed-loop equilibria, which are more natural for applications and better suited for filtering techniques. By restricting the leader to linear strategies, we establish the existence of an equilibrium. Moreover, we characterize the equilibrium through a system of forward–backward ordinary differential equations, obtained via the deterministic maximum principle.

Invited Session: Sturmius Tuschmann

Title	Stochastic games on large sparse graphs
Affiliation	Imperial College London
E-Mail	s.tuschmann23@imperial.ac.uk
Track/Session	Session 2; Track B
Abstract	We introduce a framework for stochastic games on large sparse graphs. Players are indexed by the vertices of locally finite graphs, allowing both finite and countably infinite populations, while asymptotics are described through local weak convergence of marked graphs. Under a contraction condition, we prove existence and uniqueness of Nash equilibria and establish exponential decay of correlations with graph distance. We further show that global equilibria can be approximated by truncated local games, and can even be reconstructed exactly on subgraphs given information on their boundary. Finally, we prove convergence of Nash equilibria along locally weakly convergent graph sequences. Based on joint work with Eyal Neuman.

Track C: Credit, Model Risk, and Volatility

Contributed Talk: Chang-Shu Chung

Title	Correcting Biases in Corporate Bond Pricing Models: Risk Perception, Mispricing, and Arbitrage
Affiliation	Baruch College, City University of New York
E-Mail	changshu.chung@baruch.cuny.edu
Track/Session	Session 2; Track C
Abstract	Identifying mispricing exposures across systematic and idiosyncratic dimensions raises a nontrivial question: how can these risks be effectively hedged out-of-sample? This paper develops effective risk-management strategies for corporate bond portfolios and emphasizes the importance of correcting biases in structural pricing models caused by financial constraints. I propose two debiasing methods---incorporating risk perception and rank-difference pricing errors---to address structural issues after delta hedging, reducing portfolio risk by 46%, over 70% relative to standalone long positions. The information ratio rises from 1.23 based on raw spread errors to 2.95 achieving a 90% risk reduction. Despite limited arbitrage opportunities, effective stock–bond hedging depends on precise share management, with pricing-error signals guiding convergence and neutralizing structural risks to time temporary idiosyncratic mispricing.

Contributed Talk: Luna Rigby

Title	If Not Now, Then When? Model Risk in the Optimal Exercise of American Options
Affiliation	WU Vienna University of Economics and Business
E-Mail	luna.rigby@wu.ac.at
Track/Session	Session 2; Track C
Abstract	Model risk arises from the misspecification of probabilistic models used for pricing and hedging derivatives. While model risk for European-style claims has been widely studied, much less attention has been given to American-style derivatives and the associated optimal stopping problems. We analyze model risk in the optimal exercise of an American put option using the benchmark methodology of Hull and Suo (2002). The true data-generating process is assumed to follow a Heston stochastic volatility model. We compare the optimal exercise strategy of an investor who correctly uses the Heston model with those of investors who instead use misspecified Black–Scholes or Dupire local volatility models. Optimal exercise boundaries are computed numerically via finite difference methods. Stochastic volatility dynamics and return–volatility correlation are found to have a substantial impact on optimal exercise behavior across models, creating a source of model risk. As this behavior is not transmitted to exercise strategies determined by misspecified models, even if such models are fully calibrated to European option prices, calibration fails to mitigate model risk in this context. This issue persists under frequent recalibration of a misspecified model.

Contributed Talk: Simon Feistle

Title	Filtering of Two-Factor Stochastic Volatility Models with Autoregressive Gamma Dynamics
Affiliation	University of St. Gallen
E-Mail	simon.feistle@unisg.ch
Track/Session	Session 2; Track C
Abstract	We propose a new filtering-based approximate maximum likelihood estimator for two-factor stochastic volatility (SV) models with latent autoregressive gamma (ARG) variance processes, based solely on return observations. Building on Bates (2006, 2012), our approach models the joint conditional transition law of the latent variance components to address the extreme volatility dynamics seen in estimations of two-factor ARGSV models. Empirically, our method implies a more plausible two-factor ARGSV structure and outperforms a recent two-factor GARCH benchmark in forecasting realized variance. These findings suggest that multi-factor ARGSV models estimated from returns alone can provide both econometric tractability and improved forecasting performance.

Track D: Generative Models

Contributed Talk: Rama Cont

Title	Data driven hedging in generative models
Affiliation	University of Oxford
E-Mail	rama.cont@maths.ox.ac.uk
Track/Session	Session 2; Track D
Abstract	We propose a nonparametric data-driven methodology for hedging using generative models. In contrast with model-based hedging approaches relying on sensitivity analysis of model pricing functions, our approach uses a conditional generative model trained on market data to simulate realistic market scenarios given current market conditions and computes hedge ratios which minimize risk across these scenarios. The approach incorporates transaction costs, leads to an optimal selection of hedging instruments, and adapts to market conditions. We illustrate the effectiveness of this methodology for hedging option portfolios using VolGAN, a generative model for implied volatility surfaces. The out-of-sample performance of the method matches and improves over delta and delta-vega hedging, without retraining the model for more than 4 years after the training period.

Contributed Talk: Anand Deo

Title	SSGEN: Learning Rare Financial Stress Distributions under Multivariate Heavy Tails
Affiliation	Indian Institute of Management Bangalore
E-Mail	deo.anand.a@gmail.com
Track/Session	Session 2; Track D
Abstract	We develop a theory for tail-preserving stress generation in multivariate heavy-tailed financial systems. The target is the conditional law of the full risk-factor vector under extreme multi-loss stress, rather than a single scalar risk measure. We first characterize when a sequence of approximations is asymptotically correct: in the heavy-tailed setting, a generator is tail-preserving precisely when it preserves the limiting tail measure, or equivalently the radial scaling and angular spectrum. This criterion implies asymptotic accuracy of rare-event probabilities for all admissible loss functions, and shows that misspecifying the limiting angular law necessarily distorts some tail event. As a constructive implementation, we propose Self-Similar Generative Estimation (SSGEN). The method combines an explicit Pareto radial component with a nonparametrically learned angular law from intermediate exceedances, reducing rare-scenario generation to a compact-domain learning problem followed by radial extrapolation. We prove total variation convergence of the generated stressed law, with error decomposing into statistical learning error and finite-threshold tail bias, and derive downstream guarantees for stress summaries and tail-weighted risk functionals. We further study plausibility through reverse stress testing, showing that SSGEN preserves asymptotically most-likely stress mechanisms even when multiple competing mechanisms coexist. Numerical experiments in multi-loss systems illustrate coherent rare-stress generation in regimes where direct empirical conditioning is statistically infeasible.

Contributed Talk: Philipp Schmocker

Title	The convenient setting for measurable and integrable functions
Affiliation	ETH Zürich
E-Mail	philipp.schmocker@math.ethz.ch
Track/Session	Session 2; Track D
Abstract	In a convenient setting, we introduce measurable spaces as Frölicher spaces, which yields the category generated by measurable maps on the real line that is equivalent to the category of quasi-Borel spaces. Using the properties of Frölicher spaces, we study functions and probability measures on these spaces, and show that the corresponding category is complete, cocomplete, and cartesian closed. Moreover, the L^p-spaces over these measurable spaces admit a projective limit representation induced by curve reparametrization. In addition, we introduce Gaussian measures on convenient vector spaces and give some applications in quantitative finance. This talk is based on joint work with Philipp Harms and Peter W. Michor.

Track E: Exotic options and their numerical methods

Invited Session: Yamazaki Kazutoshi

Title	Non-zero-sum optimal stopping game with continuous versus periodic exercise opportunities
Affiliation	University of Queensland
E-Mail	k.yamazaki@uq.edu.au
Track/Session	Session 2; Track E
Abstract	We introduce a new non-zero-sum game of optimal stopping with asymmetric exercise opportunities. Given a stochastic process modelling the value of an asset, one player observes and can act on the process continuously, while the other player can act on it only periodically at independent Poisson arrival times. The first one to stop receives a reward, different for each player, while the other one gets nothing. We study how each player balances the maximisation of gains against the maximisation of the likelihood of stopping before the opponent. In such a setup, driven by a Lévy process with positive jumps, we not only prove the existence, but also explicitly construct a Nash equilibrium with values of the game written in terms of the scale function. Numerical illustrations with put-option payoffs are also provided to study the behaviour of the players' strategies as well as the quantification of the value of available exercise opportunities. Joint with José Luis Pérez and Neofytos Rodosthenous.

Invited Session: Kyunghyun Park

Title	Robust Exploratory Stopping Under Ambiguity in Reinfocement Learning
Affiliation	Nanyang Technological University
E-Mail	kyunghyun.park@ntu.edu.sg
Track/Session	Session 2; Track E
Abstract	We introduce a continuous-time robust reinforcement learning framework for op timal stopping under ambiguity. In this framework, an agent chooses a robust exploratory stopping time motivated by two objectives: robust decision-making under ambiguity and learn ing about the unknown environment. Here, ambiguity refers to considering multiple probability measures dominated by a reference measure, reflecting the agent's awareness that the reference measure representing her learned belief about the environment would be erroneous. Using the g-expectation framework, we reformulate the optimal stopping problem under ambiguity as a robust exploratory control problem with Bernoulli distributed controls. We then characterize the optimal Bernoulli distributed control via backward stochastic differential equations and, based on this, construct the robust exploratory stopping time that approximates the optimal stopping time under ambiguity. Last, we establish a policy iteration theorem and implement it as a reinforcement learning algorithm. Numerical experiments demonstrate the convergence, robustness, and scalability of our reinforcement learning algorithm across different levels of ambiguity and exploration. This is based on joint work with Hoi Ying Wong and Junyan Ye (CUHK)

Invited Session: Daniel Chee

Title	Entropy regularization for American options and singular reflected BSDEs.
Affiliation	University of New South Wales
E-Mail	daniel@localedge.com.au
Track/Session	Session 2; Track E
Abstract	In this talk, we extend the entropy regularization framework proposed in Chee et al. 2025 for Bermudan options to American options. Our approach is based on the randomized stopping representation for American options and incorporates entropy regularization within a penalized BSDE formulation. We analyse the convergence properties of the resulting scheme and propose a policy improvement algorithm to demonstrate its numerical feasibility. As the penalization parameter tends to infinity, we show that the scheme converges to a reflected BSDE with a logarithmically singular.

Invited Session: Xiaolu Tan

Title	Unbiased simulation of Asian options
Affiliation	The Chinese University of Hong Kong
E-Mail	xltan@math.cuhk.edu.hk
Track/Session	Session 2; Track E
Abstract	We provide an extension of the unbiased simulation method for stochastic differential equations (SDEs) developed by Henry-Labordère et al. [Ann Appl Probab., 27 (2017), pp. 1–37] to a class of path-dependent dynamics, pertaining to Asian options. In our setting, both the payoff and the SDE coefficients depend on the (weighted) average of the process or, more precisely, on the integral of the solution to the SDE against a continuous function with finite variation. In particular, this applies to the numerical resolution of the class of path-dependent PDEs whose regularity, in the sense of Dupire, was studied previously.

Session 3

Track A: Learning, Switching, Cyber-Risk, and Sentiment Signals

Contributed Talk: Umut Cetin

Title	Optimal cyber-risk mitigation
Affiliation	LSE
E-Mail	U.Cetin@lse.ac.uk
Track/Session	Session 3; Track A
Abstract	We construct a class of piecewise deterministic Markov processes with resetting inspired by self-exciting cyber attacks. The construction provides a framework to study the optimal intervention mechanism in practical situations with sufficient generality. In particular we analyse the ergodic behaviour of this Markov process under some hypothesis reminiscent from ruin theory. Joint work with G. Callegaro and B. D'Auria

Contributed Talk: Junyan Ye

Title	Duality and DeepMartingale for High-Dimensional Optimal Switching: Computable Upper Bounds and Approximation-Expressivity Guarantees
Affiliation	Department of Statistics, The Chinese University of Hong Kong
E-Mail	junyanye@link.cuhk.edu.hk
Track/Session	Session 3; Track A
Abstract	We study finite-horizon optimal switching with discrete intervention dates on a general filtration, allowing continuous-time observations between decision dates, and develop a deep-learning-based dual framework with computable upper bounds. We first derive a dual representation for multiple switching by introducing a family of martingale penalties. The minimal penalty is characterized by the Doob martingales of the continuation values, which yields a fully computable upper bound. We then extend DeepMartingale from optimal stopping to optimal switching and establish convergence under both the upper-bound loss and an \$L^2\$-surrogate loss. We also provide an expressivity analysis: under the stated structural assumptions, for any target accuracy \$\varepsilon>0\$, there exist neural networks of size at most \$c d^{q}\varepsilon^{-r}\$ whose induced dual upper bound approximates the true value within \$\varepsilon\$, where \$c\$, \$q\$, and \$r\$ are independent of \$d\$ and \$\varepsilon\$. Hence, the dual solver avoids the curse of dimensionality under the stated structural assumptions. For numerical assessment, we additionally implement a deep policy-based approach to produce feasible lower bounds and empirical upper--lower gaps. Numerical experiments on Brownian and Brownian--Poisson models demonstrate small upper--lower gaps and favorable performance in high dimensions. The learned dual martingale also yields a practical delta-hedging strategy.

Contributed Talk: Bingyan Han

Title	Continuous-time Online Learning via Mean-Field Neural Networks
Affiliation	HKUST (Guangzhou)
E-Mail	bingyanhan@hkust-gz.edu.cn
Track/Session	Session 3; Track A
Abstract	We study continuous-time online learning where data are generated by a diffusion process with unknown coefficients. The learner employs a two-layer neural network, continuously updating its parameters in a non-anticipative manner. The mean-field limit of the learning dynamics corresponds to a stochastic Wasserstein gradient flow adapted to the data filtration. We establish regret bounds for both the mean-field limit and finite-particle system. Our analysis leverages the logarithmic Sobolev inequality, Polyak-Lojasiewicz condition, Malliavin calculus, and uniform-in-time propagation of chaos. Under displacement convexity, we obtain a constant static regret bound. In the general non-convex setting, we derive explicit linear regret bounds characterizing the effects of data variation, entropic exploration, and quadratic regularization. Finally, our simulations demonstrate the outperformance of the online approach and the impact of network width and regularization parameters.

Track B: Reinforcement Learning

Invited Session: Kailiang Liu

Title	Multi-Agent Reinforcement Learning for Large-scale Dynamic Operating Rooms Scheduling under Uncertainty
Affiliation	NUS Singapore
E-Mail	liukl@u.nus.edu
Track/Session	Session 3; Track B
Abstract	Dynamic surgical scheduling within a workday is a complicated decision-making process that critically influences hospital resource utilization, with direct consequences for revenue, reputation, patient satisfaction, and staff well-being. A critical challenge lies in the inherent uncertainty of surgery durations and the unpredictable arrival of emergency patients, which severely complicates real-time scheduling decisions. We develop a reinforcement learning (RL) framework to optimize operating rooms (ORs) schedules under uncertainty, explicitly accounting for both elective and non-elective patients arriving throughout the day. Modeling the problem as a Markov Decision Process, we introduce a multi-agent formulation where each OR acts as an autonomous decision-maker. Each agent learns to assign patients to operating rooms and determine their start times, enabling scalable learning in high-dimensional action spaces.

Invited Session: Hoa Nguyen

Title	Optimal Market Making under Model Uncertainty: A Reinforcement Learning Approach
Affiliation	NUS Singapore
E-Mail	hoa.ng@nus.edu.sg
Track/Session	Session 3; Track B
Abstract	We study optimal market making in order-driven electronic markets under model uncertainty. In particular, we address distributional risk in the market dynamics through a Sinkhorn-based ambiguity set. We formulate the problem as a tractable Markov decision process and develop a model-based robust reinforcement learning method for computing quoting strategies that remain effective under changing market conditions. The framework incorporates realistic market-microstructure variables while retaining a clear economic interpretation of states, actions, and rewards. We demonstrate the effectiveness of the proposed method through simulations and empirical studies, highlighting the value of accounting for distributional uncertainty in market-making applications. More broadly, our work provides a practical framework for combining robust optimization and reinforcement learning in quantitative finance.

Invited Session: Richard Weiss

Title	Offline Multi-Objective Policy Evaluation
Affiliation	NUS Singapore
E-Mail	richard.weiss@nus.edu.sg
Track/Session	Session 3; Track B
Abstract	Existing offline policy evaluation methods typically focus on scalar return functions and require separate estimators for each reward component. Not only does this lead to redundant computation. Complex environment

	dynamics need to be re-discovered each time, when evaluating with respect to a new reward component. This is why many stochastic gradient descent based policy evaluation methods can lead to inconsistent estimation across reward components. Stationary Distribution Correction Estimation (DICE) based methods can jointly approximate all policy value components simultaneously, by using a shared scalar marginalised importance sampling ratio, called the stationary distribution correction (SDC). This claim hinges on the observation that the SDC is independent of the reward function. Building on already established DICE methods DualDICE (Behavior-Agnostic Estamation of Discounted Stationary Distribution Corrections), GradientDICE (Rethinking Generalized Offline Estimation of Stationary Values), Off-Policy Evaluation via the Regularized Lagrangian, we propose a unifying regularised Lagrangian min-max saddle point objective. By choosing certain hyper-parameters, we can recover numerous existing methods. Furthermore, we conduct a theoretical analysis on these hyper-parameters, providing convergence and consistency results both for the discounted and un-discounted setting.

Invited Session: Yongqi Zhou

Title	Interaction-Limited Safe Continuous-Time RL for Dynamical Medical Treatment
Affiliation	NUS Singapore
E-Mail	yongqi.zhou@u.nus.edu
Track/Session	Session 3; Track B
Abstract	Dynamic medical treatment requires deciding treatment intensity and intervention timing, while patient states evolve continuously and adverse events may occur between clinical interactions. Most existing treatment learning methods assume fixed schedules or enforce safety only at discrete decision points. We propose Interaction-Limited Safe Continuous-Time Reinforcement Learning, a framework that jointly optimizes treatment administration and clinical interaction timing under trajectory-level safety constraints. Our key idea is to reformulate the continuous-time treatment problem as an option-based semi-Markov decision process, where each option specifies a continuous-time treatment policy and its duration. We develop a safety-tightening mechanism showing that suitably constructed constraints at interaction times guarantee safety over the full continuous-time trajectory with high probability. We further establish finite-sample guarantees for policy learning from logged treatment trajectories and introduce a practical data-driven conservative surrogate. Experiments show that the proposed adaptive interaction-timing mechanism improves both safety and treatment effectiveness over equidistant interaction schemes across different safe policy optimization methods.

Track C: Machine Learning meets Financial Networks

Invited Session: Yunming Hui

Title	How Users Gain Attention in Online Social Networks: Evidence from Reddit's Me...
Affiliation	University of Amsterdam
E-Mail	y.hui@uva.nl
Track/Session	Session 3; Track C
Abstract	Meme assets are financial assets whose prices are shaped by online communities and investor attention rather than traditional fundamentals. This study examines how user attention develops in Reddit meme coin communities, asking which factors drive attention, how their importance changes across market conditions, and what types of users receive attention. We analyze 103,882 posts, 609,904 comments, and 128,697 users from 10 Reddit communities linked to seven meme coins between September 2024 and March 2025. Using a heterogeneous graph with users, posts, comments, subreddits, and market conditions, we build a graph-embedding-based prediction model and assess factor importance through permutation importance with time-series cross-validation. Results show that network topology and posting time are consistently important predictors of attention, while content and market information vary across market states. User-level characteristics show little association with attention. Attention is highly concentrated among few users, who either sustain influence over time or experience brief attention spikes.

Invited Session: Mengzhong Ma

Title	Listing Premium in Crypto-Asset Markets
Affiliation	Nanyang Technological University
E-Mail	mengzhon001@e.ntu.edu.sg
Track/Session	Session 3; Track C
Abstract	This paper systematically documents the listing premium in the cryptocurrency markets. We find statistically significant and substantial premium for listings on top-tier exchanges: the average cumulative abnormal return (CAR) around the listing date is 7.2% for top-tier exchanges, compared with only 0.1% for bottom-tier exchanges. We pay special attention to two structural features of the crypto market: a crypto asset can be listed on multiple exchanges either sequentially over time ("sequential listings") or simultaneously in time ("simultaneous listings"). By focusing on the overlapping private information revealed through listings of the same asset on different exchanges, we address these two features associated with multi-venue listings by extending the two-announcement conditional event-study framework of Nayak and Prabhala (2001) to a multivariate

	setting. With the innovative framework, we find that a substantial share of abnormal returns reflects market reactions to implied uplistings, as if investors were partially observing the actual uplistings themselves. Our findings refine the interpretation of private information conveyed by IPOs and cross-listings on sub-top venues while demonstrating the practical value of our empirical framework for analyzing simultaneous multiple announcements.

Invited Session: Daniel Liebau

Title	On the Viability Bounds of Agentic Payments: Inference, Verification, and Se...
Affiliation	ESSEC Singapore
E-Mail	liebau@rsm.nl
Track/Session	Session 3; Track C
Abstract	This paper formalises three cost components of autonomous agent transactions on public blockchains (inference, verification, and settlement) and derives a non-monotone automation frontier. Verification cost scales superlinearly because interactive failure modes grow combinatorially with the highest relevant interaction order k*. Under autonomous agents, trustless verification, and blockchain settlement, this superlinearity creates a bounded automation interval: agentic execution is cheaper than human execution only for intermediate-complexity tasks. The upper bound follows a root-compression law: each unit increase in k* compresses the viable range by a root of the cost ratio. Modular decomposition reduces superlinear verification to linear scaling.

Invited Session: Simon Trimborn

Title	High-Frequency Realized Volatility Forecasting with Influential Assets
Affiliation	University of Amsterdam
E-Mail	s.trimborn@uva.nl
Track/Session	Session 3; Track C
Abstract	Accurate forecasting of realized volatility in high-frequency financial markets is essential for risk management, portfolio allocation, and derivative pricing, yet remains challenging due to complex cross-asset dependencies and limited effective sample sizes. This paper investigates the role of influential assets in shaping volatility dynamics and introduces a framework that captures dominant lead-lag structures within large cross-sections of assets. Using 1-minute and 5-minute data, we construct weekly models to forecast one-step-ahead realized volatility in the subsequent week. Our approach explicitly identifies and incorporates influential assets whose dynamics exert a disproportionate impact on the broader system. We find that these influencer structures are persistent over time and extend beyond sectoral boundaries. Empirical results show that accounting for influential assets leads to substantial improvements in forecasting accuracy compared to standard benchmark models. These findings underscore the importance of modeling dominant lead-lag relationships in high-frequency realized volatility and provide new evidence on how influential assets drive the propagation of risk across financial markets.

Track D: Rough and High-Frequency Volatility Dynamics

Invited Session: Doerte Kreher

Title	Not-so-rough volatility
Affiliation	Humboldt University
E-Mail	Kreher@math.hu-berlin.de
Track/Session	Session 3; Track D
Abstract	We propose a method for option pricing in rough Heston models by expanding the solution to the fractional Riccati equation around $H = 1/2$. Starting from the closed-form solution of the standard Heston Riccati equation, higher-order terms are obtained recursively by solving linear ODEs. We prove analyticity of the fractional Riccati solution, implying convergence of the approximation to the characteristic function under usual constraints. Numerically, we obtain fast and accurate implied volatilities down to $H = 0.2$ depending on the maturity. This is joint work with Paul Hager.

Invited Session: Carsten Chong

Title	Intraday Volatility Dynamics
Affiliation	The Hong Kong University of Science and Technology

E-Mail	carstenchong@ust.hk
Track/Session	Session 3; Track D
Abstract	Return-based spot volatility estimates are noisy and prone to biases. In this paper, we present a robust inference method for the autocorrelation of intraday volatility changes that overcomes the latency of volatility and potential biases due to jumps and microstructure noise in observed prices. Our procedure builds on two results: First, we show that the limiting autocorrelation function of increments of a continuous-time stochastic process over shrinking horizons is parametrically determined and must coincide with that of fractional Gaussian noise. Second, we derive limit theorems that characterize the infill asymptotic limit and distribution of realized autocovariances of spot variance estimators. The two results combine to yield a robust and feasible generalized method of moments estimator of the autocorrelation of high-frequency volatility increments. In an empirical application to SPY transaction data, we document widespread predictability of short-term intraday changes in both latent and realized volatility. This is joint work with Marc Hoffmann, Mathieu Rosenbaum and Grégoire Szymanski.

Invited Session: Masaaki Fukasawa

Title	On the Skew Stickiness Ratio
Affiliation	Osaka University
E-Mail	masaaki.fukasawa@gmail.com
Track/Session	Session 3; Track D
Abstract	The skew stickiness ratio is a statistic that captures the joint dynamics of an asset price and its volatility. We derive a representation formula for this quantity using the Ito–Wentzell and Clark–Ocone formulae, and we apply it to analyze its asymptotics under Bergomi-type stochastic volatility models.

Track E: Financial Networks, Systemic Risk, and Market Equilibrium

Contributed Talk: Rosalio G. Artes Jr.

Title	Multilayer Financial Networks: A Graph-Theoretic Framework for Systemic Risk and Resilience
Affiliation	Mindanao State University - Tawi-Tawi College of Technology and Oceanography
E-Mail	rosalioartes@msutawi-tawi.edu.ph
Track/Session	Session 3; Track E
Abstract	The growing interconnectedness of financial institutions has intensified the need for mathematical frameworks capable of representing heterogeneous channels of dependence and contagion. Traditional single-layer network models are often unable to capture the structural complexity of modern financial systems, in which institutions are simultaneously linked through interbank exposures, common asset holdings, derivative contracts, and liquidity dependencies. This paper develops a purely theoretical graph-theoretic framework for multilayer financial networks, with particular emphasis on systemic risk and resilience. A multilayer representation is introduced through a supra-adjacency structure that preserves both intra-layer and inter-layer interactions. Based on this construction, spectral characteristics of the associated supra-Laplacian are employed to examine connectivity, fragility, and structural robustness. A multilayer centrality measure and a resilience index are formulated to provide mathematically interpretable indicators of institutional importance and network stability. The framework establishes a rigorous basis for analyzing how coupling across financial layers may either strengthen resilience through redundancy or intensify systemic vulnerability through amplified contagion channels. The study contributes a formal foundation for future analytical and empirical work in network-based financial stability assessment.

Contributed Talk: Li Ke

Title	A Consensus-ADMM Framework with Local Price Copies for Solving Market Equilibrium Models
Affiliation	NUS Singapore
E-Mail	like06@u.nus.edu
Track/Session	Session 3; Track E
Abstract	We propose a deep consensus-ADMM framework with local price copies for solving continuous-time financial equilibrium models with heterogeneous agents. The main computational challenge in solving these equilibrium models is to enforce individual optimality and market clearing simultaneously along stochastic paths. Our method addresses this issue by decomposing the equilibrium problem into agent-level and market-level subproblems. In the agent step, each agent solves a local optimization problem that determines its consumption and investment strategies, as well as local copies of market prices. The market block then freezes the agent networks and updates the consensus prices to reduce market clearing conditions, while staying close to the agents' local copies through ADMM penalties.

	The proposed framework is implemented with neural networks and trained on fixed stochastic path pools. We demonstrate the method on two continuous-time equilibrium models. The first is the complete-market heterogeneous-agent exchange economy studied by Cvitanić, Jouini, Malamud, and Napp. The second is the equilibrium model with restricted stock-market participation studied by Basak and Cuoco. Across both examples, the method produces economically interpretable investment and consumption strategies and achieves small market-clearing residuals. The results indicate that the consensus-ADMM framework is a flexible deep learning framework for computing continuous-time financial equilibria.

Contributed Talk: Quan Hoang Anh

Title	A Perturbation Approach to Equilibrium Asset Pricing in Restricted Markets
Affiliation	NUS Singapore
E-Mail	e1353502@u.nus.edu
Track/Session	Session 3; Track E
Abstract	This talk studies equilibrium asset pricing in restricted markets where two logarithmic-utility agents have homogeneous time preferences but face different trading opportunities. We consider a market in which aggregate dividends follow a geometric Brownian motion and two stocks split the total dividend through a mean-reverting dividend-share process, with one agent able to trade both stocks and the bond while the other has access only to one stock and the bond. When the dividend-share process is deterministic, the model admits an explicit complete-market baseline, including the interest rate, stock prices, consumption shares, portfolio positions, and zero-bond holdings. To study the incomplete-market case, we introduce small dividend-share volatility and apply a perturbation expansion around this baseline. This method provides a tractable way to approximate equilibrium objects and isolate the effects of market restrictions on consumption, stock prices, interest rates, portfolios, money accounts, and Sharpe ratios. The expansion shows that the interest rate is unchanged at first order but decreases at second order, stock-price distortions arise only at second order, and the restricted agent becomes a lender in the bond market. The Sharpe ratio comparison between restricted and unrestricted assets remains nontrivial and depends on long-run dividend shares and the initial wealth distribution.

Contributed Talk: Paolo Guasoni

Title	Dynamic General Equilibrium with Transaction Costs
Affiliation	Dublin City University
E-Mail	paolo@guasoni.it
Track/Session	Session 3; Track E
Abstract	This paper develops a continuous-time general equilibrium model of proportional transaction costs with two risk-averse agents who face opposite, time-varying hedging needs. The model endogenizes the safe rate, stock price, return volatility, equity premium, Sharpe ratio, consumption, and portfolios. Small-cost asymptotics deliver closed-form equilibrium formulas and an explicit no-trade region driven by the shadow value of trading. Transaction costs weaken risk sharing, raise precautionary saving, and lower the safe rate. The lower rate transmits to asset prices: it raises stock volatility, equity premia, and Sharpe ratios, while linking all corrections through a single interest-rate effect.

Session 4

Track A: Recent Advances in Portfolio Selection and Financial Technology

Invited Session: Chen Yang

Title	Equilibrium Liquidity Provision in Automated Market Makers: A Present-Value Approach
Affiliation	The Chinese University of Hong Kong
E-Mail	cyang@se.cuhk.edu.hk
Track/Session	Session 4; Track A
Abstract	Decentralized exchanges facilitate trustless trading of digital assets via Automated Market Makers (AMMs), leveraging blockchain technology to eliminate traditional intermediaries. In an AMM, liquidity providers

	passively earn transaction fees paid by liquidity takers and arbitrageurs but are subject to losses from adverse selection as well as the forgone convenience yields generated by holding the assets. From an investment perspective, the pro-rata sharing rule utilized by Uniswap v2 implies that the marginal benefit of providing liquidity in a pool equals the average present value of the pool. Furthermore, no entry costs indicate that the marginal cost of providing liquidity equals the prevailing token pair value as measured on centralized exchanges. Consequently, the equilibrium pool size in the AMM is determined by the tradeoff between this marginal benefit and marginal cost. Calibrating the model to the ETH-USDT token pair using Uniswap and Binance data, we demonstrate that the model-generated equilibrium pool dynamics are consistent with empirical observations.

Invited Session: Xiaofei Shi

Title	Active Portfolio Management for Market Diversity and Dispersion
Affiliation	University of Toronto
E-Mail	xf.shi@utoronto.ca
Track/Session	Session 4; Track A
Abstract	Equal-weighted portfolios have historically delivered episodes of substantial outperformance relative to capitalization-weighted benchmarks. It is well known in stochastic portfolio theory that the relative return of any dynamic allocation between the equal-weight and capitalization-weighted portfolios can be expressed solely in terms of a market diversity process and a dispersion process. This observation leads us to introduce a class of stochastic diversity–dispersion models, in which market diversity evolves as an Itô diffusion whose volatility is driven by a latent dispersion factor. Within this framework, we formulate an active mean–variance problem for the allocation process between the equal-weight and capitalization-weighted portfolios, expressed in terms of their relative log return and penalizing both active risk and quadratic maintenance and trading costs. We derive the optimal strategy via a linear forward–backward SDE and obtain an explicit "aiming in front of a moving target" representation of the optimal trading rate, in the spirit of Gârleanu and Pedersen. We then calibrate the stochastic diversity and dispersion model to S&P 500 data on diversity and dispersion and illustrate the performance of the resulting strategy under realistic maintenance and trading costs.

Invited Session: Shuaijie Qian

Title	A Deep-Learning Approach for Solving HJB Equations from Stochastic Control
Affiliation	Hong Kong University of Science and Technology
E-Mail	sjqian@ust.hk
Track/Session	Session 4; Track A
Abstract	In this talk, we present a deep learning–based approach for numerically solving Hamilton–Jacobi–Bellman (HJB) equations arising from infinite-horizon regular stochastic control problems. Exploiting the control structure of the problem, we adopt an actor–critic framework to decouple policy evaluation from policy improvement. The proposed approach is robust and readily extendable to high-dimensional HJB equations. Our contributions are twofold. First, we establish the numerical convergence of the proposed scheme. Second, our method is more robust and exhibits faster convergence than the approach of Duarte, Duarte, and Silva (The Review of Financial Studies, 2024).

Invited Session: Qianyu Liu

Title	Optimal Investment with Quadratic Transaction Costs in a Multi-factor and Stochastic Interest Rate Environment
Affiliation	The Chinese University of Hong Kong
E-Mail	qyliu@link.cuhk.edu.hk
Track/Session	Session 4; Track A
Abstract	This paper studies an optimal investment problem in a comprehensive market environment featuring geometric Brownian motion type asset dynamics, stochastic interest rates, return and volatility factors, and time-varying, state-dependent quadratic transaction costs. Under a finite-horizon mean-quadratic variation objective, we derive a semi-analytical solution for the value function and optimal trading strategy and establish the well-posedness of the associated system of nonlinear PDEs along with the verification theorem. The optimal strategy retains a tractable structure that echoes the core economic principle of trading gradually toward an aim portfolio as proposed in Garleanu and Pedersen (2013). Specifically, the aim portfolio is adjusted for hedging demands against the interest rate risk, and the trading speed reflects an intertemporal trade-off between immediate rebalancing costs and future market liquidity. Additionally, we propose a deep learning scheme to enable efficient implementation in high-dimensional settings based on observable market data and estimated factor models.

Track B: Machine learning approach to computational finance

Invited Session: Fengyi Yuan

Title	Solving dynamic portfolio selection problems via score-based diffusion models
Affiliation	The Chinese University of Hong Kong (Shenzhen)
E-Mail	yuanfengyi@cuhk.edu.cn
Track/Session	Session 4; Track B
Abstract	In this paper, we tackle the dynamic mean-variance portfolio selection problem in a {\it model-free} manner, based on (generative) diffusion models. We propose using data sampled from the real model $\mP$ (which is unknown) with limited size to train a generative model $\mQ$ (from which we can easily and adequately sample). With adaptive training and sampling methods that are tailor-made for time-series data, we obtain quantification bounds between $\mP$ and $\mQ$ in terms of the adapted Wasserstein metric $\cA\cW_2$. Importantly, the proposed adapted sampling method also facilitates {\it conditional sampling}. In the second part of this paper, we provide the stability of the mean-variance portfolio optimization problems in $\cA\cW_2$. Then, combined with the error bounds and the stability result, we propose a policy gradient algorithm based on the generative environment, in which our innovative adapted sampling method provides approximate scenario generators. We illustrate the performance of our algorithm on both simulated and real data. For real data, the algorithm based on the generative environment produces portfolios that beat several important baselines, including the Markowitz portfolio, the equal weight (naive) portfolio, and S\&P 500.

Invited Session: Xiaoli Wei

Title	Continuous-time q-learning for mean-field control with common noise
Affiliation	Harbin Institute of Technology
E-Mail	xiaoli.wei@hit.edu.cn
Track/Session	Session 4; Track B
Abstract	This paper investigates the continuous-time counterpart of the Q-function for entropy-regularized mean-field control (MFC) with controlled common noise, coined as q-function by Jia and Zhou (2023) in the single agent's model. We first show that, under discretely sampled actions, the value function in the exploratory formulation converges to the one in the relaxed control formulation as the time grid refines. Leveraging the relaxed control formulation, we derive the exploratory Hamilton-Jacobi-Bellman (HJB) equation, in which the controlled common noise gives rise to an additional nonlinear functional of policy, rendering the policy iteration intricate. Under certain concavity condition, we establish the existence and uniqueness of the optimal one-step policy iteration via a first-order condition using the partial linear functional derivative with respect to policy. The policy improvement at each iteration is verified by relating to an entropy-regularized optimization problem over the space of policies. In the mean-field setting, we introduce the integrated q-function (Iq-function) defined on the state distribution and the policy, and it is shown that an optimal policy is identified as a two-layer fixed point to the argmax operator of the Iq-function. Finally, we provide the explicit characterization of an optimal policy as a Gaussian distribution in the general linear-quadratic (LQ) setting.

Invited Session: Yanwei Jia

Title	Data-Driven Merton's Strategies via Policy Randomization
Affiliation	The Chinese University of Hong Kong
E-Mail	yanweijia@cuhk.edu.hk
Track/Session	Session 4; Track B
Abstract	We study Merton's expected utility maximization problem in an incomplete market, characterized by a factor process in addition to the stock price process, where the functional forms of the model primitives are unknown. The agent under consideration is a price taker who has access only to the stock and factor value processes and the instantaneous volatility. We propose an auxiliary problem in which the agent can invoke policy randomization according to a specific class of Gaussian distributions, and prove that the mean of its optimal Gaussian policy solves the original Merton problem. With randomized policies, we are in the realm of continuous-time reinforcement learning (RL) recently developed in Wang et al. (2020) and Jia and Zhou (2022a,b, 2023), enabling us to solve the auxiliary problem in a data-driven way without having to estimate the model primitives. Specifically, we establish a policy improvement theorem based on which we design both online and offline actor-critic RL algorithms for learning Merton's strategies. A key insight from this study is that RL in general and policy randomization in particular are useful beyond the purpose for exploration—they can be employed as a technical tool to solve a problem that cannot be otherwise solved by mere deterministic policies. Finally, we present simulation and empirical studies in a stochastic volatility environment to demonstrate the decisive outperformance of the devised RL algorithms in comparison to the conventional model-based, plug-in method. This is a joint work with Min Dai, Yuchao Dong, and Xunyu Zhou.

Invited Session: Gaozhan Wang

Title	Reinforcement Learning for Entropy-Regularized Stochastic Control with Model Uncertainty
Affiliation	University of Southern California
E-Mail	gaozhanw@usc.edu
Track/Session	Session 4; Track B
Abstract	We propose a reinforcement-learning (RL) scheme for solving the entropy-regularized Hamilton–Jacobi–Bellman (HJB) equation arising in stochastic control with model uncertainty. Compared with the standard policy-iteration algorithm(PIA) discussed in our earlier work, our approach overcomes two key limitations: (i) the classical PIA depends on unknown model parameters, rendering it impractical; and (ii) it requires derivatives of the value function, and since an L2 (or uniform) approximation of the value function does not lead to a good approximation of its derivatives, the converence is not guaranteed. Our method assumes only that the diffusion (volatility) coefficient is known—an empirically reasonable requirement—while allowing the drift to be unknown. The algorithm operates entirely via coupled fixed-point maps that are estimable from simulation and trials without derivative access. We prove convergence and obtain rate of convergence, and we corroborate the theory with comprehensive numerical experiments. As a further illustration, we also extend the algorithm framework into a special case of volatility control in one dimension. This talk is based on joint work with Jin Ma, Jianfeng Zhang, and Xunyu Zhou.

Track C: Safe Learning and Portfolio Decision-Making

Contributed Talk: Himanshu Choudhary

Title	A CVaR-Constrained Safe Reinforcement Learning Framework with Action Repair for Practical Portfolio Optimization
Affiliation	Indian Institute of Technology Mandi, India
E-Mail	ch.himanshu1199@gmail.com
Track/Session	Session 4; Track C
Abstract	Portfolio optimization refers to allocating capital across multiple assets to attain an optimal balance between potential returns and associated risks under dynamic market conditions. While Deep Reinforcement Learning (DRL) has shown significant promise in learning adaptive allocation strategies through continuous interaction with market, most existing approaches struggle to incorporate practical investment constraints and robust risk measures. To address these limitations, this paper introduces a constraint-aware and risk-sensitive DRL framework for portfolio optimization that operates within a Constrained Markov Decision Process (CMDP) settings. The proposed model embeds Conditional Value-at-Risk (CVaR) as a hard safety constraint using an Augmented Lagrangian approach on the agent to prevent high-loss outcomes. Additionally, a repair mechanism is introduced to adjust the agent’s actions which ensures the resulting portfolio weights remain feasible and compliant with real-world investment constraints. Our approach is evaluated using real-world stock data from multiple global indices and consistently outperforms existing DRL-based and traditional portfolio optimization strategies, achieving superior returns while strictly adhering to defined risk constraints.

Contributed Talk: Paul Mangers Bastian

Title	Stochastic factors can matter: improving robust growth under ergodicity
Affiliation	London School of Economics
E-Mail	p.j.mangers-bastian@lse.ac.uk; paul.mangers@gmail.com
Track/Session	Session 4; Track C
Abstract	Drifts of asset returns are notoriously difficult to model accurately and, yet, trading strategies obtained from portfolio optimization are very sensitive to them. To mitigate this well-known phenomenon we study robust growth-optimization in a high-dimensional incomplete market under drift uncertainty of the asset price process X, under an additional ergodicity assumption, which constrains but does not fully specify the drift in general. The class of admissible models allows X to depend on a multivariate stochastic factor Y and fixes (a) their joint volatility structure, (b) their long-term joint ergodic density and (c) the dynamics of the stochastic factor process Y. A principal motivation of this framework comes from pairs trading, where X is the spread process and models with the above characteristics are commonplace. Our main results determine the robust optimal growth rate, construct a worst-case admissible model and characterize the robust growth-optimal strategy via a solution to a certain partial differential equation (PDE). We demonstrate that utilizing the stochastic factor leads to improvement in robust growth complementing the conclusions of the previous study by Itkin et. al. (arXiv:2211.15628 [q-fin.MF], forthcoming in Finance and Stochastics), which additionally robustified the dynamics of the stochastic factor leading to Y-independent optimal strategies. Our analysis leads to new financial insights, quantifying the improvement in growth the investor can achieve by optimally incorporating stochastic factors into their trading decisions. We illustrate our theoretical results on several numerical examples including an application to pairs trading.

Contributed Talk: Jin Qian

Title	Dynamic Signal Precision in Rationally Inattentive Decision-Making
Affiliation	Lehigh University
E-Mail	jiq220@lehigh.edu
Track/Session	Session 4; Track C
Abstract	This paper develops a continuous-time rational inattention portfolio model in which signal precision, rather than being fixed at the initial date, is endogenously adjusted over time in response to evolving uncertainty. This generates a dynamic learn–decision–value loop and introduces an additional dynamic effect in the value function through the intertemporal tradeoff between attention reward and attention cost. The paper first establishes long-run global convergence toward the baseline precision regime over a fixed horizon. As the terminal date approaches, the structural thresholds for intensified attention move upward, so increasingly high belief variance is required to justify costly monitoring, and optimal precision eventually returns to the baseline level. The analysis then characterizes transitional dynamics through a complete six-case classification across low-, moderate-, and high-risk environments: degenerate adjustment, latent-risk monitoring, immediate rescue, routine monitoring, mild over-attention, and short-run correction. These cases are distinguished by endogenous structural thresholds and generate three broad belief-variance trajectory patterns: monotone downward, monotone upward, and non-monotone. Beyond the belief-variance dynamics, the paper further identifies four canonical dynamic precision paths: monotone decline to the baseline level, upper-bound plateau followed by relaxation, interior peak followed by decline, and rise to the upper-bound plateau followed by relaxation. These paths describe how optimal attention intensity responds to evolving uncertainty through baseline monitoring, intensified monitoring, upper-bound rescue, and eventual normalization. Finally, the paper characterizes the duration of transitional regimes, measuring how long intensified monitoring, correction, or rescue dynamics persist before the system returns to baseline precision. Overall, the model provides a dynamic theory of endogenous attention allocation in which information precision evolves jointly with belief uncertainty, portfolio choice, and continuation value.

Contributed Talk: Luca De Gennaro Aquino

Title	Equilibrium investment under dynamic preference uncertainty
Affiliation	Reykjavik University
E-Mail	lucaa@ru.is
Track/Session	Session 4; Track C
Abstract	We study a continuous-time portfolio choice problem for an investor whose state-dependent preferences are determined by an exogenous factor that evolves as an Itô diffusion process. Since risk attitudes at the end of the investment horizon are uncertain, terminal wealth is evaluated under a set of utility functions corresponding to all possible future preference states. These utilities are first converted into certainty equivalents at their respective levels of terminal risk aversion and then (nonlinearly) aggregated over the conditional distribution of future states, yielding an inherently time-inconsistent optimization criterion. We approach this problem by developing a general equilibrium framework for such state-dependent preferences and characterizing subgame-perfect equilibrium investment policies through an extended Hamilton-Jacobi-Bellman system. This system gives rise to a coupled nonlinear partial integro-differential equation for the value functions associated with each state. We then specialize the model to a tractable constant relative risk aversion specification in which the preference factor follows an arithmetic Brownian motion. In this setting, the equilibrium policy admits a semi-explicit representation that decomposes into a standard myopic demand and a novel preference-hedging component that captures incentives to hedge against anticipated changes in risk aversion. Numerical experiments illustrate how features of the preference dynamics -most notably the drift of the preference process and the correlation between preference shocks and asset returns- jointly determine the sign and magnitude of the hedging demand and the evolution of the equilibrium risky investment over time.

Track D: Stochastic Processes, Swaps, and Digital Markets

Contributed Talk: Cornelius Hanel

Title	Generalized Langevin Dynamics
Affiliation	ETH Zürich
E-Mail	chanel@ethz.ch

Track/Session	Session 4; Track D
Abstract	A key result in generative modeling is that complex distributions can be sampled by reversing diffusion processes, with the reverse dynamics governed by the score function. We generalize Langevin dynamics intrinsically in the framework of symmetric Dirichlet forms and extend the construction to metric measure spaces. The resulting theory replaces classical gradients and stochastic differential equations by carré du champ operators and derivations.

Contributed Talk: Yining Ding

Title	Choice of Collateral Currency in Differential Swaps
Affiliation	The University of Sydney
E-Mail	ydin6541@uni.sydney.edu.au
Track/Session	Session 4; Track D
Abstract	The role of collateral in derivative pricing has evolved beyond credit risk mitigation, particularly following the global financial crisis, when funding costs and basis spreads became central to valuation practices. This development coincided with the transition from the London Interbank Offered Rate (LIBOR) to risk-free rates (RFRs) and the increasing standardization of collateralised trading. We study the valuation and hedging of a class of differential swaps referencing backward-looking averages of overnight rates, with SOFR swaps appearing as a particular instance. The focus is on the impact of the collateral currency. Extending earlier results Ding et al. [Math. Finance 36 (2026), pp.~180--202], we allow the collateral account to be denominated in a currency different from that of the contractual cash flows and derive explicit pricing and hedging strategies using a futures-based replication approach. We show that the choice of collateral currency can have a non-trivial effect on both valuation and risk management. In particular, foreign-currency collateral can introduce additional risk exposures even when contractual cash flows are entirely denominated in the domestic currency. Numerical study demonstrates that collateral effects can lead to significant valuation adjustments and therefore need to be properly incorporated in modern multi-currency modelling frameworks.

Contributed Talk: Jonghwa Park

Title	Convergence of diffusion processes in the adapted weak topology
Affiliation	Carnegie Mellon University
E-Mail	jonghwap@andrew.cmu.edu
Track/Session	Session 4; Track D
Abstract	The Wasserstein distance equips the space of probability measures with a rich geometric structure. While this structure has found numerous applications across pure and applied mathematics, it is not well suited for dynamic optimization problems arising in finance, as it fails to capture the temporal structure of stochastic processes. The adapted Wasserstein distance addresses this limitation by incorporating the causal strucuture inherent in the flow of time, making it a natural tool for comparing stochastic processes. In this talk, we discuss the compactness criterion for diffusion processes in the adapted weak topology. We show that a sequence of diffusion processes converge in the adapted weak topology if and only if it converges in the weak topology and its transition kernels are sufficiently regular. This talk is based on a joint work with Martin Larsson and Johannes Wiesel.

Contributed Talk: Yutian Zhou

Title	Arbitrage on decentralized exchanges
Affiliation	The Chinese University of Hong Kong
E-Mail	1155167851@link.cuhk.edu.hk
Track/Session	Session 4; Track D
Abstract	Decentralized exchanges using automated market makers create arbitrage opportunities with centralized exchanges, where gas fees and transaction ordering are critical. Existing models largely overlook competition among arbitrageurs, despite price discrepancies being public information. We develop the first equilibrium

	model of gas fee competition between two arbitrageurs under three transaction reversion settings: no-revert, auto-revert, and selectable-revert. We show that pure symmetric equilibria do not exist, but unique mixed equilibria can be characterized. Comparative analysis reveals that under low inventory risk, the no-revert setting favors arbitrageurs in terms of profit, while auto-revert and selectable-revert settings enhance market efficiency. Under high inventory risk, the no-revert and selectable-revert settings dominate the auto-revert setting in both profitability and efficiency. Using data from Binance and Uniswap V2, we empirically confirm that arbitrageurs face positive inventory risk and validate our model's implications: gas fees increase with price discrepancies and liquidity, while trading amounts rise with both price discrepancies and gas fees.

Track E: Portfolio Choice, Retirement, and Asset Allocation

Contributed Talk: Criscent, Birungi

Title	Optimal Retirement and Annuitization via Stochastic Vitality
Affiliation	Department of Mathematics and Statistics, Concordia University, Montreal QC, CANADA
E-Mail	criscent.birungi@concordia.ca
Track/Session	Session 4; Track E
Abstract	Traditional lifecycle models typically rely on deterministic, age-dependent mortality laws, which limit their ability to capture the heterogeneity and randomness inherent in individual health trajectories. This paper introduces a vitality-based mortality structure in which death occurs when a latent health reserve first hits zero, allowing mortality risk to evolve stochastically and endogenously over the life cycle. Modeling longevity as a first-passage event generates rich dispersion in survival outcomes and links economic decisions directly to the dynamics of health. Within this framework, we characterize optimal consumption, labor supply, portfolio choice, and annuitization using stochastic control. The results show that vitality volatility has clear behavioral implications: higher volatility strengthens precautionary motives at low vitality, raising consumption and labor effort and delaying annuitization. As vitality improves, the influence of volatility diminishes and policies converge. Despite the richer health dynamics, the optimal risky portfolio share remains the constant Merton fraction, implying that vitality risk primarily affects intertemporal allocation rather than financial risk-taking. Overall, the vitality-based approach provides a biologically grounded and behaviorally consistent explanation for heterogeneity in retirement timing and annuity demand. It highlights the importance of modeling health uncertainty explicitly in lifecycle finance and offers a flexible foundation for analyzing retirement behavior under personalized longevity risk.

Contributed Talk: Manish

Title	Sentiment-Informed Predictive Modeling for Optimal Asset Allocation and Strategic Commodity Exposure
Affiliation	Indian Institute of Technology Roorkee, India
E-Mail	manish@mfs.iitr.ac.in
Track/Session	Session 4; Track E
Abstract	The integration of artificial intelligence (AI) has revolutionized portfolio management by enabling adaptive, predictive, and risk-aware investment frameworks. This study conducts a comparative analysis of seven AI-enhanced asset allocation models, each derived from a specific risk measure, alongside a modified Black-Litterman (BL) model. A key innovation is the integration of forecasts from the best-performing deep learning (DL) model, selected from multiple architectures and trained on a hybrid dataset comprising historical market data, technical indicators, and sentiment measures that capture investor behavior. These forecasts provide forward-looking expected returns within risk-based portfolios and serve as the view vector in the BL framework. The proposed portfolios are constructed from ten U.S. sectoral ETFs and twelve commodity futures spanning energy, metals, and agriculture. By systematically integrating these commodities into ETF portfolios, the study evaluates their contribution to diversification and hedging effectiveness. The results show that forecast-driven portfolios consistently outperform traditional benchmarks, delivering superior risk-adjusted returns with effective downside protection. The findings further emphasize the importance of aligning portfolio construction with specific risk–return preferences and asset-class configurations. Notably, portfolios integrating energy exposures deliver the strongest risk-adjusted returns and hedging performance, whereas agricultural commodities offer comparatively weaker diversification benefits and, in some cases, dilute performance. Overall, these results underscore the value of sentiment-informed AI forecasts and strategic commodity exposures in enhancing portfolio resilience and allocation efficiency.

Contributed Talk: Ho Man Tai

Title	A Variational Approach to Portfolio Choice
Affiliation	University of Sydney
E-Mail	homan.tai@sydney.edu.au
Track/Session	Session 4; Track E
Abstract	In this talk, we discuss a variational approach to solve consumption--investment problems for infinite-horizon isoelastic investors with any number of assets and stochastic investment opportunities driven by a scalar diffusion. We identify a strictly convex functional whose unique minimizer coincides with the optimal wealth–consumption ratio and yields the optimal portfolio explicitly. We demonstrate the method's effectiveness by computing optimal consumption--investment policies through gradient descent in models with nonlinear predictability, stochastic volatility, and interest-rate risk. The method combines verification with numerical tractability and yields accurate solutions in analytically intractable models.

Contributed Talk: Kasper Larsen

Title	Existence of an equilibrium with limited stock market participation and power utilities.
Affiliation	Rutgers University
E-Mail	kl756@math.rutgers.edu
Track/Session	Session 4; Track E
Abstract	We prove existence and uniqueness of a solution to a singular and path-dependent Riccati-type ODE. As an application, we use the ODE solution to prove existence of a Radner equilibrium with homogenous power-utility investors in the limited participation model from Basak and Cuoco (1998). To ensure that all traders survive in the long run, we conclude by giving a model variation with heterogenous time-preferences. This is joint work with Paolo Guasoni (Università di Bologna and Dublin City University), Giovanni Leoni (CMU), and Heeyoung Kwon (Rutgers)

Session 5

Track A: Quantitative Finance and FinTech

Invited Session: Ling Qin

Title	Optimal Bonding Function in Decentralized Exchanges: A Convex Analysis Approach
Affiliation	ShanghaiTech University
E-Mail	qinling@shanghaitech.edu.cn
Track/Session	Session 5; Track A
Abstract	Decentralized exchanges rely on automated market makers to provide liquidity passively. This exposes liquidity providers to both stale-price arbitrage and reverse-trade arbitrage. In this paper, we use tools from the calculus of variations and convex optimization to study the optimal design of bonding functions, the core mechanism of AMMs, in the presence of these two forms of arbitrage. We also discuss the economic implications of our findings.

Invited Session: Yilie Huang

Title	ART for Diffusion Sampling: A Reinforcement Learning Approach to Timestep Scheduling
Affiliation	The Hong Kong Polytechnic University
E-Mail	yh2971@columbia.edu
Track/Session	Session 5; Track A
Abstract	We consider time discretization for score-based diffusion models to generate samples from a learned reverse-time dynamic on a finite grid. Uniform and hand-crafted grids can be suboptimal given a budget on the number of time steps. We introduce Adaptive Reparameterized Time (ART), which controls the clock speed of a reparameterized time variable to redistribute computation along the sampling trajectory while preserving

	the terminal time, with the objective of minimizing the aggregate Euler discretization error. We derive a randomized companion ART-RL that recasts ART as a continuous-time reinforcement learning problem with Gaussian policies, and prove a two-directional bridge between the two: the deterministic ART optimum lifts to an optimal Gaussian policy, and conversely any optimal Gaussian policy must recover the ART control through its mean. This bridge turns continuous-time actor--critic learning into a principled, rather than heuristic, route to the deterministic timestep optimum. Within the official EDM pipeline, ART-RL improves FID on CIFAR--10 across a wide range of budgets; after one-time offline training, the distilled deterministic schedule transfers without retraining to AFHQv2, FFHQ, and ImageNet at no extra inference cost.

Invited Session: Heqing Huang

Title	Dynamic Debt Policies and Cash Management
Affiliation	National University of Singapore
E-Mail	hheqing@u.nus.edu
Track/Session	Session 5; Track A
Abstract	We develop a dynamic corporate finance model in which firms issue and repurchase debt, raise external equity, and manage cash reserves, subject to financing costs, cash-carrying costs, and bankruptcy risk. These frictions discipline firms' incentives to ratchet up leverage while still permitting active leverage adjustment. Integrating cash management into the capital structure problem generates a precautionary savings motive: firms may delever and build liquidity buffers to preserve financial flexibility. A key feature of our framework is the presence of nonlinear debt issuance costs, under which equity value and debt price are jointly determined by a coupled system of nonlinear PDEs. We provide analytical characterizations and efficient numerical methods to solve the model.

Invited Session: Min Dai

Title	When Reinforcement Learning Aligns with Estimate-Then-Plug-In? Insights from Continuous-Time Portfolio Selection
Affiliation	The Hong Kong Polytechnic University
E-Mail	mindai@polyu.edu.hk
Track/Session	Session 5; Track A
Abstract	Traditional dynamic decision-making under uncertainty typically follows an estimate-then-plug-in approach: specifying a model, estimating its parameters from historical data, and solving the resulting stochastic control problem analytically or numerically. In contrast, modern reinforcement learning (RL) directly learns optimal policies from data without estimating model parameters. Continuous-time RL methods have recently gained traction in financial decision-making, but it remains unclear when RL outperforms the traditional approach and why. We investigate these questions through continuous-time portfolio optimization using a relaxed control framework. Without transaction costs, we derive an analytical solution and develop a model-free q-learning algorithm to learn optimal trading strategies. Our theoretical results show that given continuous stock price data, the RL-learned policy matches the estimate-then-plug-in solution. Incorporating transaction costs, we find this equivalence still holds. Numerical experiments with synthetic data confirm these findings across various models. However, empirical studies demonstrate that our RL method significantly outperforms the traditional approach in realistic settings, likely because real-world stock price dynamics exhibit complex features not fully captured by parametric models. This is a joint work with Yanwei Jia and Zhichao Lu.

Track B: Explainable AI for FinTech

Invited Session: Dennis Thumm

Title	Causal Regime-Switching Models for Decision-Making Under Uncertainty in European Electricity Markets
Affiliation	NUS Singapore
E-Mail	dennis.thumm@u.nus.edu
Track/Session	Session 5; Track B
Abstract	European electricity markets operate under increasing uncertainty driven by renewable integration, fuel price volatility, and cross-border interconnections. Accurate and interpretable price forecasting is critical not only for prediction, but for operational decision-making in trading, dispatch, and risk management. However, existing approaches either prioritise predictive accuracy without interpretability, or impose fixed structural assumptions that fail to capture regime-dependent market behaviour. We propose a Causal Regime-Switching deep state space model (CaRS) that jointly learns (i) latent market regimes, (ii) dynamic price evolution, and (iii) regime-specific causal relationships among key price drivers. The framework integrates Markov regime switching with data-driven causal structure learning, enabling the identification of state-dependent price formation mechanisms in electricity markets. The model is evaluated on 12 interconnected European markets over 11 years of hourly day-ahead price data using a curated Tier-1 feature set of 21-25 variables per market (price, load, weather, calendar, day-ahead generation and demand forecasts), selected by a four-method importance consensus from a wider 222-feature unified pipeline. CaRS matches the directional accuracy of a Lasso-per-regime baseline (0.738 vs. 0.737) and achieves the lowest mean absolute error across all methods (0.505 vs. 0.580 for XGBoost) on one-hour-ahead returns, while providing interpretable regime-specific causal

	graphs that align with economic principles such as the merit order effect. A spillover variant with GATv2 attention and cross-border lagged-price features recovers the documented European inter-market coupling without a physical-adjacency prior. In contrast to black-box models, the framework identifies how and when key drivers - gas prices, renewable generation, demand - dominate price dynamics under different market conditions. The model demonstrates stable directional accuracy across forecasting horizons (1--24 hours) on the Sweden case study, suggesting robustness for short-term operational decision-making. The results highlight a fundamental trade-off between accuracy and interpretability and demonstrate that causal, regime-aware models can provide actionable insights for energy market participants, including improved understanding of price drivers, enhanced risk management, and more informed decision-making under uncertainty.

Invited Session: Hannah L. H. Lai

Title	Adaptive Cross-Sectional Forecasting of Implied Volatility Surfaces via Oblique Local Linear Forests
Affiliation	NUS Singapore
E-Mail	hannahlhlai@nus.edu.sg
Track/Session	Session 5; Track B
Abstract	We study the predictability of single-stock implied volatility surfaces (IVS) using an econometric learning framework that integrates functional dimension reduction with adaptive cross- sectional pooling. Daily IVS are represented via an Ad-hoc Black–Scholes sieve, and the dynamics of its coefficients are modeled using oblique decision forests with local linear leaves. Oblique splits are constructed from singular-value projections of predictor–response covariance, allowing multivariate partitioning in the joint space of lagged surface coefficients and firm characteristics, while local linear estimation refines within-group dynamics. Using over 30 million option observations from 225 S&P 500 firms (2015–2021), we document three findings. First, adaptive nonlinear pooling improves multi-day IVS forecasts relative to random walk, linear VAR, and standard random forest benchmarks, particularly at longer horizons and during volatile periods. Second, oblique partitioning produces more parsimonious trees, indicating efficient representation of cross-sectional heterogeneity. Third, predictive drivers exhibit horizon dependence: option-based features dominate short-run dynamics, whereas firm characteristics drive medium- and long-horizon predictability. Our results highlight the importance of structured cross-sectional heterogeneity in shaping implied volatility dynamics.

Invited Session: Wenhan Gao

Title	Innovation Value Discrepancy and Its Role in Shaping Firms' Short-Term Gains and Sustainable Growth
Affiliation	NUS Singapore
E-Mail	wenhan_gao@u.nus.edu
Track/Session	Session 5; Track B
Abstract	This study examines how divergence between the scientific and economic dimensions of innovation is associated with firms' subsequent growth and valuation dynamics. Using a comprehensive dataset of 7,506 publicly listed U.S. firms and more than 2.5 million granted patents from 1976 to 2022, we develop a dual-value classification framework that distinguishes four types of innovation—Aurora, Mirage, Success, and Invisible—based on scientific significance and market-based economic recognition. We document systematic differences in firm outcomes across these innovation types. Firms associated with Aurora-type innovations, those with high scientific value but relatively weak contemporaneous market recognition, tend to exhibit stronger long-term growth in market capitalization and sustained improvements in industry ranking. In contrast, Mirage-type innovations, characterized by strong initial market reactions but lower scientific impact, are associated with short-term performance gains that are less persistent over time. Success-type innovations, where scientific and economic value are aligned, display comparatively stable performance patterns. These patterns are robust across alternative performance measures, classification schemes, and model specifications, including tests addressing industry heterogeneity and potential endogeneity concerns. Overall, the results indicate that discrepancies between the scientific significance of innovation and its contemporaneous market recognition are systematically related to firms' subsequent growth and valuation trajectories.

Invited Session: Paolo Recchia

Title	Quantum Algorithms for Derivative Pricing
Affiliation	NUS Singapore and Standard Chartered Bank Singapore
E-Mail	paolo_re@nus.edu.sg
Track/Session	Session 5; Track B
Abstract	Numerical integration is a central computational task in many scientific and industrial applications. In quantitative finance, in particular, it is an essential tool for the pricing of derivative securities and for risk-management calculations, where high-dimensional expectations are commonly estimated through Monte Carlo methods. Despite their flexibility and broad applicability, classical Monte Carlo algorithms remain computationally demanding: achieving an accuracy ε typically requires a number of function evaluations scaling as $O(1/\varepsilon^2)$. Quantum Monte Carlo methods based on Quantum Amplitude Estimation can quadratically improve this dependence. However, the standard quantum Monte Carlo framework provides a speedup over randomized sampling methods and does not directly exploit the deterministic structure underlying quasi-

	random Monte Carlo integration. In this work, we introduce a Quantum Quasi-Monte Carlo algorithm that combines low-discrepancy nets with Quantum Amplitude Estimation. The proposed method prepares a deterministic quasi-random point set coherently in superposition and estimates the corresponding finite average through amplitude estimation. Classical quasi-Monte Carlo methods can achieve significantly improved convergence over Monte Carlo for low-dimensional functions with finite Hardy-Krause variation, with an error controlled by the discrepancy of the generated point set. Our quantum construction does not yield an asymptotic improvement over classical quasi-Monte Carlo, since the total error separates into a discretization error, determined by the finite net, and a quantum estimation error. Instead, we prove the existence of a pre-asymptotic advantage window: for a target accuracy that would classically require 2^q low discrepancy points, one can prepare a higher-resolution net of size 2^Q, with Q>q, in superposition and reach the same accuracy using significantly fewer function queries. This window can be controlled by tuning the circuit resolution and amplitude-estimation parameters, making the approach relevant for practical regimes where the number of queries is finite rather than asymptotically large.

Track C: Volatility Forecasting and Partial-Information Portfolio Models

Contributed Talk: Zhensi Tan

Title	Dynamic Model Selection for Volatility Forecasting: A Reinforcement Learning Approach
Affiliation	NUS Singapore
E-Mail	tantan_zs@163.com
Track/Session	Session 5; Track C
Abstract	Accurately forecasting financial volatility remains a core challenge due to non-linear structures and regime shifts across different time scales. To address this challenge, we develop a two-stage reinforcement learning (RL) framework designed for adaptive volatility forecasting. The first stage establishes a heterogeneous candidate pool consisting of traditional econometric baselines and deep learning architectures. This pool serves as the action space for the RL agent, which dynamically selects the optimal model based on market states. Empirical evaluations using high-frequency financial data reveal a frequency-dependent performance divergence: while the RL framework remains highly competitive with the optimal individual baseline at the daily frequency, it yields significant performance gains over the best single model in the intraday hourly forecasting scenario. In conclusion, our RL-driven framework establishes a robust dynamic model selection paradigm that enhances forecasting accuracy.

Contributed Talk: Li Hao

Title	Portfolio Choice under Preference Switching
Affiliation	NUS Singapore
E-Mail	e1311647@u.nus.edu
Track/Session	Session 5; Track C
Abstract	This talk studies a consumption-investment problem in which an investor's CRRA risk aversion changes at a deterministic or exogenous stochastic time. Such a switch may represent intergenerational transfer, mandate revision, or another future event that alters the investor's risk attitude. In a Black-Scholes market, we ask how the possibility of a future preference change affects current consumption and portfolio choice. We compare three formulations. The benchmark ignores the future switch before it occurs and leads to a piecewise Merton rule. An anticipative formulation evaluates the successor's post-switch consumption plan using the current utility index, thereby preserving the original homogeneity structure. A second anticipative formulation instead uses the successor's value function as the continuation payoff, which creates a mismatch between the pre- and post-switch homogeneity classes. For an exogenous Cox switching time, a small-intensity expansion shows that homogeneity preservation affects consumption but leaves the Merton risky share unchanged at first order. By contrast, homogeneity mismatch produces a non-homothetic value correction and a first-order distortion of the risky share. The leading investment effect therefore arises from the interaction between switching risk and preference mismatch.

Contributed Talk: Christian Tezza

Title	A GARCH model with two volatility components and two driving factors
Affiliation	University of Bologna & ABN AMRO bank
E-Mail	christian.tezza@unibo.it; tezzachris@gmail.com
Track/Session	Session 5; Track C
Abstract	We introduce a novel GARCH model that integrates two sources of uncertainty to better capture the rich, multi-component dynamics often observed in the volatility of financial assets. This model provides a quasi closed-form representation of the characteristic function for future log-returns, from which semi-analytical formulas for option pricing can be derived. A theoretical analysis is conducted to establish sufficient conditions for strict stationarity and geometric ergodicity, while also obtaining the continuous-time diffusion limit of the model. Empirical evaluations, conducted both in-sample and out-of-sample using S&P500 time series data, show that our model outperforms widely used single-factor models in predicting returns and option prices.

Contributed Talk: Dannin Eccles

Title	Endogenous Emergence of MACD in Partial-Information Portfolio Optimization
Affiliation	University of Chicago
E-Mail	dannin@uchicago.edu
Track/Session	Session 5; Track C
Abstract	We consider a class of partial-information portfolio optimization problems in which the drift of a risky asset is driven by two latent stochastic factors evolving at distinct time scales. We show that the filtered estimate of the latent mean-reversion level can be represented in terms of fast and slow exponentially weighted moving averages of prices, yielding a Moving Average Convergence Divergence (MACD)-type signal. Under logarithmic, power, and exponential utility, we derive explicit candidate optimal strategies and establish admissibility and verification results under suitable horizon conditions. In particular, the results provide a mathematical foundation for the endogenous emergence of a MACD-type signal as an estimator of latent drift information contained in observed price paths.

Track D: Digital Assets, FinTech, and Decentralised Finance

Invited Session: Huei-Wen Teng

Title	Trading BTC with robust pricing kernels
Affiliation	NYCU, Taiwan
E-Mail	venteng@gmail.com
Track/Session	Session 5; Track D
Abstract	This study develops a robust approach for estimating Bitcoin pricing kernels under stochastic volatility models with jumps. Bitcoin's growing importance in modern financial markets highlights the need for reliable tools to capture its heavy tails and abrupt price movements. The pricing kernel represents the ratio between discounted risk-neutral probabilities and their physical counterparts, forming the core link between preferences, valuation, and derivative pricing. Because its construction depends on accurately estimating both densities, standard methods often fail for Bitcoin. Our robust estimator produces stable kernels and delivers improved trading performance when applied to practical BTC strategies.

Invited Session: Joerg Osterrieder

Title	Quantifying Economic Narratives from Six Million News Headlines
Affiliation	Twente University & Chur University
E-Mail	josterri@googlemail.com
Track/Session	Session 5; Track D
Abstract	This talk introduces a data-driven framework to quantify economic narratives using six million real-time US news headlines collected over more than a decade. Custom LLM-based embeddings and clustering yield approximately 350 distinct narratives, including themes such as market volatility, inflation, geopolitical risk, and COVID-19. Daily narrative intensity measures show clear and interpretable patterns, such as the sharp expansion of COVID-19 coverage in early 2020 and repeated surges in volatility-related narratives during market stress. Narrative time series and their cross-sectional exposure across firms illustrate how large-scale text analysis captures shifts in public economic discourse with high temporal resolution. The results demonstrate how narrative construction at scale provides a structured way to analyse headline flow, identify emerging themes, and map their relevance to financial markets.

Invited Session: Daniel Traian Pele

Title	One Scalar, Two Jobs: Conformal Diagnosis and Correction of Tail Risk Forecasts
Affiliation	Bucharest University of Economic Studies; Institute for Economic Forecasting, Romanian Academy
E-Mail	danpele@ase.ro
Track/Session	Session 5; Track D

Abstract	This paper develops a model-free framework for diagnosing and correcting miscalibrated tail risk forecasts under temporal dependence. I introduce a one-sided conformal recalibration of Value-at-Risk that delivers a single signed scalar, applied post-hoc to any forecaster without retraining or distributional assumptions. The scalar plays two roles: it restores Basel-compliant coverage and serves as a continuous diagnostic of tail miscalibration in return units, with sign indicating direction and magnitude indicating severity. A finite-sample coverage guarantee is established under beta-mixing, with stationary GARCH satisfying the assumptions. The framework is evaluated across foundation models and classical benchmarks on a diverse cross-section of assets. Rolling conformal correction substantially improves Basel Green Zone classification and dominates higher-complexity recalibration alternatives, documenting a tail-sparsity regime in which low-dimensional corrections are statistically favored.

Invited Session: Stefan Lessmann

Title	Credit Risk Modeling in the AI Era: An Empirical Analysis
Affiliation	Humboldt-Universität zu Berlin
E-Mail	stefan.lessmann@hu-berlin.de
Track/Session	Session 5; Track D
Abstract	Predictive models play a pivotal role in credit risk management, guiding critical financial decisions through accurate estimation of default probabilities and losses. Their performance influences the stability of the financial system and the profitability of lending operations. The continuous improvement of risk models and modeling practices is therefore a top priority. Rapid advances in AI, particularly large language models and other foundation models, have disrupted predictive modeling paradigms across many fields. This paper investigates the potential of tabular data foundation models for probability of default (PD) and loss given default (LGD) forecasting in credit risk management. It evaluates various datasets, benchmark algorithms, performance indicators, and experimental conditions to clarify the benefits of pretraining across relevant risk modeling challenges. The empirical results show the effectiveness of foundation models for PD modeling. For LGD, however, classical machine learning approaches trained from scratch on specific datasets remain a challenging benchmark that current foundation models cannot surpass. Ongoing experiments extend the empirical study to verify the generality and robustness of the results and uncover the boundary conditions of these models.