The 9th Asian Quantitative Finance Conference Schedule

Day 1 (Saturday, April 26, 2025)

08:00 - 8:45	Registration (First floor)			
08:45 - 9:00	Opening Address (Room 519)			
	Plenary Talks (Room 519)			
09:00 – 9:25	Martin Schweizer, ETH Dynamic Monotone Mean-variance Portfolio Optimisation with Independent Returns			
09:25 – 9:50	Johannes Muhle-Karbe, Imperial Liquidity and Asset Prices			
9:50 – 10:20	Coffee Break[1]			
	Parallel Sessions			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Xuedong He	Nan Chen	Zhongfei Li	Lilian Hu
Topics	Financial Technology	Machine Learning	Stochastic Control	Asset Pricing
10:20 – 10:45	Xuedong He, CUHK Arbitraging on Decentralized Exchanges	Nan Chen, CUHK A Two Time-Scale Evolutionary Game Approach to Multi-Agent Reinforcement Learning and Its Application in Algorithmic Collusion Studies	Zhongfei Li, SUSTech Optimal housing, consumption and portfolio choices	Lilian Hu, HKUST (GZ) New perspectives on analytic solvability of stochastic volatility models
10:45 – 11:10	Ning Cai, HKUST (GZ) Bootstrapping with Privacy Preservation	Ruixun Zhang, PKU Diffusion Factor Models: Generating High-Dimensional Returns with Factor Structure	Tianyang Nie, SDU Indefinite linear-quadratic large population problem with partial observation	Cheolmin Shin, SNU Two-sided sticky short rates with jumps
11:10 – 11:35	Xuefeng Gao, CUHK Reward-Directed Score-Based Diffusion Models via q-Learning	Chao Zhang, HKUST (GZ) Trading Volume Alpha	Yufei Zhang, Imperial α-Potential Games: A New Paradigm for N-player Dynamic Games	Wencan Xia, HKUST SOFR Futures Pricing under Affine Term Structure Model with Scheduled Jumps
11:35 – 12:00	Junhuan Zhang, BUAA Profitability of collusive sandwich attack in automated market maker exchanges	Jean-Loup Dupret, ETH Deep Learning for Continuous-time Stochastic Control with Jumps	Sang Hu, CUHK (SZ) Menuless and Preference-Free Screening Contracts for Fund Managers	Jianong Li, SUFE An Option-Based Framework for Pricing Bond Repos
12:00 – 13:00	Lunch[1] (Room 415 and 417)			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Chao Zhou	Jiawen Gu	Zuoquan Xu	Wei Xiong
Topics	Asset Pricing	Stochastic Control	Portfolio Selection	Quantitative Investment and Corporate Finance
13:00 – 13:25	Chao Zhou, Qube Research & Tech. Convergence analysis on the particle systems with centralized control	Jiawen Gu, SUSTech Optimal Pairs Trading with Transaction Costs in Infinite Horizon	Zuoquan Xu, PolyU Optimal portfolio selection with VaR and portfolio insurance constraints under the rank-dependent expected utility theory	Wei Xiong, Oxford A Study of Algorithmic Collusion in Multi-dealer-to-client Platforms
13:25 – 13:50	Xiaolu Tan, CUHK Unbiased simulation of Asian options	Shuaijie Qian, HKUST Existence of optimal contract for principal-agent problem with quadratic cost function	Chen Yang, CUHK Portfolio Selection with Time-Varying Taxation	Krsek Daniel, ETH Moral Hazard with Limited Liability Fillippo Beretta, ETH The Role of Agent's Preferences in Sannikov's Principal-agent Model
13:50 – 14:15	Lingxiao Zhao, PKU No Sparsity in Asset Pricing: Evidence from a Generic Statistical Test	Engel John C. Dela Vega, HKU Optimal Ratcheting of Dividends with Irreversible Reinsurance	Xiang Yu, PolyU An Extended Merton Problem with Relaxed Benchmark Tracking	Jianrui Li, TJU Optimal investment and consumption problem associated with stochastic volatility models
14:15 – 14:45	Coffee Break[1]			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Chenxu Li	Jin Liang	Dylan Possamai	Hongyi Jiang
Topics	Financial Derivatives	Asset Pricing	Game Theory	Risk Management and Financial Decision Making
14:45 – 15:10	Chenxu Li, PKU Inference for Volatility Roughness from Options	Jin Liang, Tongji Structural Credit Rating Model with Stochastic Volatility	Dylan Possamai, ETH A target approach to Stackelberg games	Hongyi Jiang, CUHK Robust Optimal Strategies for Two-Period Liquidation in Financial Systems
15:10 – 15:35	Xue Cheng, PKU Joint Pricing in SPX and VIX Derivative Markets with Composite Change of Time Models	Aoxiang Yang, PKU Volatility-Managed Volatility Trading	Alex Tse, UCL Portfolio Selection in Contests	Joseph Mulligan, Imperial How Minimum Performance Thresholds Bias Backtests: Bayesian Estimation for Sharpe Ratios under Selection Bias
15:35 – 16:00	Pingping Zeng, SUSTech Analytical solvability and exact simulation in models with affine stochastic volatility and Lévy jumps	Urban Ulrych, ETH Kernel Conditional Factor Models	Yuchao Dong, Tongji Quantitative Propagation of chaos for mean field randomized stopping in Continuous time	Fabian Sandmeier, UZH The Impact of Risk Transfers on Systemic Risk in Banking Networks Chenxi Xia, PKU Choquet Rating Criteria and Risk Consistency
16:00 – 16:25	Zhuo Jin, Macquarie University A Hybrid Deep Reinforcement Learning Method for Insurance Portfolio Management	Yang Liu, CUHK (SZ) Anonymized risk sharing	Chi Seng Pun, NTU Quantile-based Distributional Reinforcement Learning under Cumulative Prospect Theory and its Dynamic Optimality Characterization	Linfeng Li, NUS Arbitrages in Perpetual Contracts

1. Coffee Break, Lunch, & Reception Dinner are provided to all registered attendees, including plenary speakers, invited speakers, registered attendees, local organizers, and student helpers.
2. Each report lasts for 25 minutes, but the time for each student report in Session 4 is 18-25 minutes, depending on the number of speakers in the session.

Day 2 (Sunday, April 27, 2025)

	Plenary Talks (Room 519)			
09:00 – 9:25	Jussi Keppo, NUS Misaligned Clocks: Incentives for Differently Patient Boards and CEOs			
09:25 – 9:50	Hyeng Keun Koo, AU Optimal Consumption and Portfolio Rules with Dynamic Adjustment of Consumption Bounds			
9:50 – 10:20	Coffee Break[1]			
	Parallel Sessions			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Josef Teichmann	Harry Zheng	Mathieu Lauriere	Bo Wu
Topics	Machine Learning	Stochastic Control	Heterogeneous and Mean-field Models	Machine Learning
10:20 – 10:45	Josef Teichmann, Imperial How Minimum Performance Thresholds Bias Backtests: Bayesian Estimation for Sharpe Ratios under Selection Bias	Harry Zheng, Imperial Convergence of PPGM for Linear Quasi-Convex Problems	Mathieu Lauriere, NYU (Shanghai) Global Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models	Bo Wu, CUHK Reinforcement Learning for Continuous-Time Mean-Variance Portfolio Selection in a Regime-Switching Market
10:45 – 11:10	Min Dai, PolyU Data-Driven Merton's Strategies via Policy Randomization	Juan Li, SDU Controllability Concepts for Mean-Field Dynamics With Reduced-Rank Coefficients	Xun Li, PolyU Discrete-Time Mean-Variance Strategy Based on Reinforcement Learning	Florian Krach, ETH Neural Jump ODEs for Input-Output Systems
11:10 – 11:35	Hoi Ying Wong, CUHK Reinforcement learning without a market simulator: Learning the stopping intensity	Shuoqing Deng, HKUST Distribution-constrained optimal multiple stopping: A Root-type solution	Alexandros Saplaouras, ETH Stability of backward propagation of chaos	Tengyingzi Ma, ETH Markov Decision Process with Observation Cost in Food Production Chain Yuqiao Fang, SJTU Deep-Learning Earnings
11:35 – 12:00	Lingfei Li, CUHK Value-aware estimation of diffusion systems for continuous-time reinforcement learning: Surrogate loss and financial application	Jiacheng Fan, PolyU Realization utility, market regimes, and the disposition effect	Na Li, DUT Policy Iteration Reinforcement Learning Method for Continuous-Time Linear-Quadratic Mean-Field Control Problems	Patrick Lucescu, UZH Regularized Multi-Currency Expected Shortfall Portfolios
12:00 – 13:00	Lunch[1] (Room 415 and 417)			

1. Coffee Break, Lunch, & Reception Dinner are provided to all registered attendees, including plenary speakers, invited speakers, registered attendees, local organizers, and student helpers.
2. Each report lasts for 25 minutes, but the time for each student report in Session 4 is 20 or 25 minutes, depending on the number of speakers in the session.

Day 3 (Monday, April 28, 2025)

	Plenary Talks (Room 519)			
09:00 – 9:25	Zengjing Chen, SDU Nonlinear robust limit theorems			
09:25 – 9:50	Yanwei Jia, CUHK Mean--Variance Portfolio Selection by Continuous-Time Reinforcement Learning: Algorithms, Regret Analysis, and Empirical Study			
9:50 – 10:20	Coffee Break1			
	Parallel Sessions			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Yue Kuen Kwok	Wanting He	Yong Hyun Shin	Konrad J Mueller
Topics	Financial Derivatives	Actuarial Science	Portfolio Selection	Stochastic Control
10:20 – 10:45	Yue Kuen Kwok, HKUST (GZ) Efficient simulation of the SARB model	Wanting He, HKU Periodic Evaluation of DC Pension Fund: A Dynamic Risk Measure Approach	Cong Qin, SUFE Periodic Evaluation with Non-Concave Utility	Konrad J Mueller, Imperial What matters in deep stochastic control?
10:45 – 11:10	Gongqiu Zhang, CUHK (SZ) Pricing American Parisian Options under General Time-Inhomogeneous Markov Models	Wenyuan Li, HKU Robust control of optimal investment and reinsurance under the path-dependent effect	Yuwei Wang, CUHK PreFER: Interactive Robo-Advisor with Scoring Mechanism	Chiara Rossato, ETH Time-Inconsistent Stochastic Games with Mean-Variance Preferences
11:10 – 11:35	Yating Wan, TUFE A Tail-Driven Nonparametric Estimator for State Price Densities	Xin Zang, BJTU Random risk measures on the order statistics of correlated risks	Zhesheng Liu, PolyU Portfolio optimisation with proportional transaction costs and stochastic investment opportunities	Mateo Rodriguez Polo, ETH Time-Inconsistent Stochastic Control: State Dependency Marco Rodrigues, ETH When BSDEs Meet Semimartingales
11:35 – 12:00	Jintao Li, NUSRI-CQ Calibration for Stochastic Local Volatility Models	Aniq Rohmawati, ITB Dependence-Based Distortion Risk Measures for Insurance Portfolio Analysis	Yong Hyun Shin, SWU (Sookmyung) The Effects of Income and HARA Utility on Optimal Consumption, Investment, and Retirement Decisions	Sturmius Tuschmann, Imperial Stochastic Graphon Games with Memory
12:00 – 13:00	Lunch1 (Room 763)			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Nian Yang	Xiaoyang Zhuo	Kexin Chen	Liulei Sun
Topics	Asset Pricing	Financial Derivatives	Stochastic Control	Stochastic Control
13:00 – 13:25	Nian Yang, NJU A fast and universal framework for option pricing within 3D models	Hyungbin Park, SNU A BSDE approach to designing funding rates for perpetual swaps in cryptocurrency markets	Kexin Chen, PolyU Robust Dividend Policy: Equivalence of Epstein-Zin and Maenhout Preferences	Liulei Sun, PKU Representation theorem of temporal copula families and their corresponding Markov processes
13:25 – 13:50	Weilong Liu, SYSU Covariance Matrix Estimation for Positively Correlated Assets	Xiaoyang Zhuo, BIT Multiverse Equivalent Expectation Measures for Computing Moments of Contingent Claim Returns	Yihan Zou, UOG Numerical approximation of RBSDEs via regularization	Ji Uk Jang, SNU A discretization scheme for path-dependent FBSDEs
13:50 – 14:15	Guanglian Hu, USYD The Term Structure of Index Option Returns	Huansang Xu, NUS Pricing and Hedging Strategies for Cross-Currency Equity Protection Swaps		Jongjin Park, SNU Criticality of HJB equations in stochastic control problems
14:15 – 14:45	Coffee Break1			
	Session 1 (Room 519)	Session 2 (Room 421)	Session 3 (Room 423)	Session 4 (Room 425)
Chair	Xianhua Peng	Jaehyun Kim	Zhenhua Wang	Guangyi He
Topics	Financial Technology	Risk Management	Stochastic Control	Financial Technology
14:45 – 15:10	Xiaoli Wei, HIT Continuous-time q-learning algorithms for mean-field control problems	Jaehyun Kim, SNU Long-term decomposition of robust pricing kernels under G-expectation	Zhenhua Wang, SDU Convergence of Policy Iteration for Entropy-Regularized Stochastic Control Problems	Guangyi He, Imperial Distributional Adversarial Attacks and Training in Finance
15:10 – 15:35	Sixian Zhuang, CUHK Liquidity, exploration and inference-based learning to reduce transaction cost	Joshue Ricalde, ETH Mean-Field Games with Common Poissonian Noise	Julian Sester, NUS Non-concave distributionally robust stochastic control in a discrete time finite horizon setting	Songyan Hou, ETH Entropic Adapted Wasserstein Distance on Gaussians
15:35 – 16:00	Xianhua Peng, PKU A Risk Sensitive Contract-unified Reinforcement Learning Approach for Option Hedging		Hanwu Li, SDU Optimal Consumption for Recursive Preferences with Local Substitution under Risk	Owen Futter, Imperial Path Dependent Trading Strategies Using Signatures and Kernels

1. Coffee Break, Lunch, & Reception Dinner are provided to all registered attendees, including plenary speakers, invited speakers, registered attendees, local organizers, and student helpers.
2. Each report lasts for 25 minutes, but the time for each student report in Session 4 is 20 or 25 minutes, depending on the number of speakers in the session.

The Role of Agent's Preferences in Sannikov's Principal-agent Model

Fillippo Beretta

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

Bootstrapping with Privacy Preservation

Ning Cai

Bootstrapping has a wide range of applications in statistics, machine learning, and econometrics and typically assumes availability of the data. However, many data contain sensitive information in fields such as finance, economics, and life science. We propose a privacy-preserving bootstrap method that allows bootstrapping while protecting all parties' privacy and justify the privacy preservation theoretically.

Robust Dividend Policy: Equivalence of Epstein-Zin and Maenhout Preferences

Kexin Chen

The classic optimal dividend problem aims to maximize the expected discounted dividend stream over the lifetime of a company. Since dividend payments are irreversible, this problem corresponds to a singular control problem with a risk-neutral utility function applied to the discounted dividend stream. In cases where the company's surplus process encounters model ambiguity under the Brownian filtration, we explore robust dividend payment strategies in worst-case scenarios. We establish a connection between ambiguity aversion in a robust singular control problem and risk aversion in Epstein-Zin preferences. To do so, we first formulate the dividend problem as a recursive utility function with the EZ aggregator within a singular control framework. We investigate the existence and uniqueness of the EZ dividend problem. By employing Backward Stochastic Differential Equation (BSDE) representations

where singular controls are involved in the generators of BSDEs, we demonstrate that the EZ formulation is equivalent to the maximin problem involving risk-neutral utility on the discounted dividend stream, incorporating Meanhout's regularity that reflects investors' ambiguity aversion. Considering the equivalent Meanhout's preferences, we solve the robust dividend problem using a Hamilton-Jacobi-Bellman (HJB) approach combined with a variational inequality (VI). Our solution is obtained through a novel shooting method that simultaneously satisfies the VI and boundary conditions. This is a joint work with Kyunghyun Park and Hoi Ying Wong.

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

Nan Chen

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

Joint Pricing in SPX and VIX Derivative Markets with Composite Change of Time Models

Xue Cheng

The Chicago Board Options Exchange Volatility Index (VIX) is calculated from SPX options and derivatives of VIX are also traded in market, which leads to the so-called "consistent modeling" problem. This paper proposes a time-changed Lévy model for log price with a composite change of time structure to capture both features of the implied SPX volatility and the implied volatility of volatility. Consistent modeling is achieved naturally via flexible choices of jumps and leverage effects, as well as the composition of time changes. Many celebrated models are covered as special cases. From this model, we derive an explicit form of the characteristic function for the asset price (SPX) and the pricing formula for European options as well as VIX options. The empirical results indicate great competence of the proposed model in the problem of joint calibration of the SPX/VIX Markets.

Nonlinear robust limit theorems

Zengjing Chen

Motivated by multi-armed bandit problem and reinforcement learning, in this paper, we introduce a similar binary model in the context of nonlinear probabilities. This can be viewed as a nonlinear Bernoulli-like model and is motivated in modelling distribution uncertainties. It provides a new probabilistic understanding of the nonlinear probability theory. In one main result we obtain a generalized robust limit theorem for this model with mean-variance uncertainty, and give an explicit formula for the robust limit distribution. The limit is shown to depend heavily on the structure of the events or the integrating functions, which demonstrate the key signature of nonlinear structure. As applications, these limit theorems provide the theoretical foundation for statistical inferences and hypothesis testing.

Data-Driven Merton’s Strategies via Policy Randomization

Min Dai

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 all 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's 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. At last, we carry out both 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.

Moral Hazard with Limited Liability

Krsek Daniel

We study a continuous-time principal-agent problem with moral hazard, where the contract consists of a lump-sum payment at the end of a fixed time interval. We consider a situation in which both parties have limited liability, meaning the contract satisfies both lower and upper bounds. We characterize the set of all admissible contracts and study the principal’s problem

by means of an HJB equation on a specific domain. This talk is based on joint work with Dylan Possamaï and Stéphane Villeneuve.

Optimal Ratcheting of Dividends with Irreversible Reinsurance

Engel John C. Dela Vega

This work considers an insurance company that faces two key constraints: a ratcheting-dividend constraint and an irreversible reinsurance constraint. The company allocates part of its reserve to pay dividends to its shareholders while strategically purchasing reinsurance for its claims. The ratcheting dividend constraint ensures that dividend cuts are prohibited at any time. The irreversible reinsurance constraint ensures that reinsurance contracts cannot be prematurely terminated or sold to external entities. The dividend rate level and the reinsurance level are modelled as nondecreasing processes, thereby satisfying the constraints. The incurred claims are modelled via a Brownian risk model. The main objective is to maximize the cumulative expected discounted dividend payouts until the time of ruin. A characterization of the optimal value function as the unique viscosity solution of the associated Hamilton-Jacobi-Bellman equation is established. The existence of an optimal threshold strategy is shown. Finally, numerical examples are presented to illustrate the optimality conditions and optimal strategies.

Distribution-constrained optimal multiple stopping: A Root-type solution

Shuoqing Deng

We consider the problem of optimal multiple stopping where the stopping times should satisfy some distribution constraints. For a large class of cost functions, we reformulate the problem into a sequence of optimal stopping problem of a time-reversed process, explicitly construct the solution and verify the optimality using a martingale inequality. The methodology has links with the Root's solution to Skorokhod embedding problem and the inverse boundary hitting problem.

Quantitative Propagation of chaos for mean field randomized stopping in Continuous time

Yuchao Dong

The classical decision problem of optimal stopping has found an amazing range of applications from finance, statistics, marketing, and phase transitions to engineering. When extend the standard framework to multi-agent setting, one obtains mean field stopping problem when the number of agents goes to infinity. In this paper, we consider a mean field stopping system with a randomized stopping strategy. The wellposedness of such system is given together with some results on quantitative Propagation of chaos.

Deep Learning for Continuous-time Stochastic Control with Jumps

Jean-Loup Dupret

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

Realization utility, market regimes, and the disposition effect

Jiacheng Fan

We study a model of realization utility in a financial market where the drift of the stock switches between a bull and a bear market driven by a hidden Markov chain. In this model, an investor trading between a risky, regime switching stock and a risk-free bond seeks to determine an optimal sequence of purchasing and selling times for the stock to balance utility bursts experienced when realizing a gain or loss of the stock and utility derived from terminal wealth. We characterize the optimal value functions when holding the bond and stock respectively as the unique solutions to a system of coupled variational equalities and accordingly obtain optimal holding and selling regions for the bond and stock. The presence of market regimes and preferences incorporating realization utility have opposite effects on trading behavior: Models of realization utility without market regimes predict extreme versions of the disposition effect while a model incorporating market regimes but not realization utility predicts trend following strategies, i.e., the opposite of the disposition effect. Including both features generates a rich set of trading behavior including voluntarily realizing gains and losses with optimal liquidation points depending on the assessment of the market state, selling of the stock with or without immediate repurchase, and reasonable levels of the disposition effect and other statistics conforming with recorded behavior of individual investors.

Deep-Learning Earnings

Yuqiao Fang

We utilize a deep learning model, the Kolmogorov-Arnold Neural Network, to produce “permanent earnings” estimates and examine their incremental usefulness compared to Generally Accepted Accounting Principles (GAAP) earnings for equity valuation. Our model takes 58 detailed items from trailing-twelve-month (TTM) income and cash flow statements as input features and uses the average of the operating cash flows of the past 8 quarters and next 12 quarters as the outcome variable. The intuition is that a complex nonlinear combination of the 58 items that maps into past and future cash flows captures the “permanent earnings”. We demonstrate that deep-learning earnings are more persistent, less volatile, and less affected by the conservatism and earnings management biases inherent in GAAP earnings. The price-to-

earnings ratios derived from deep-learning earnings are more consistent across firms within the same industry. Furthermore, the deep-learning earnings provide additional explanatory power beyond GAAP earnings in predicting future cash flows and can predict future returns, even after accounting for the five factors identified by Fama and French (2015). Overall, our results suggest that deep-learning earnings more effectively capture recurring or "permanent earnings" than GAAP earnings.

Reward-Directed Score-Based Diffusion Models via q-Learning

Xuefeng Gao

We propose a new reinforcement learning (RL) formulation for training continuous-time score-based diffusion models for generative AI to generate samples that maximize reward functions while keeping the generated distributions close to the unknown target data distributions. Unlike most existing studies, our formulation does not involve any pretrained model for the unknown score functions of the noise-perturbed data distributions. We present an entropy-regularized continuous-time RL problem and show that the optimal stochastic policy has a Gaussian distribution with a known covariance matrix. Based on this result, we parameterize the mean of Gaussian policies and develop an actor–critic type (little) q-learning algorithm to solve the RL problem. A key ingredient in our algorithm design is to obtain noisy observations from the unknown score function via a ratio estimator. Numerically, we show the effectiveness of our approach by comparing its performance with two state-of-the-art RL methods that finetune pretrained models. Finally, we discuss extensions of our RL formulation to probability flow ODE implementation of diffusion models and to conditional diffusion models. This is a joint work with Jiale Zha and Xunyu Zhou.

Optimal Pairs Trading with Transaction Costs in Infinite Horizon

Jiawen Gu

We consider an optimal pairs trading problem with transaction costs in infinite horizon. The spread of two stocks is assumed to follow an OU process. We allow investors to hold any quantity of spreads during the investment period. The transaction costs are proportional to the quantity of spread traded. Our objective is to maximize the wealth increment while imposing a penalty on our holding quantity. The corresponding value function is governed by a Hamilton-Jacobi-Bellman variational inequality with gradient constraints. It is shown that after transformation the original problem is equivalent to a double obstacle problem. Based on the smooth fit principle, we derive a solution and the free boundaries can be obtained by solving the system of implicit equations. Finally we consider a specific example of penalty function. The corresponding numerical example is provided to illustrate the differences between scenarios with and without transaction costs.

Distributional Adversarial Attacks and Training in Finance

Guangyi He

Adversarial robustness is well-studied in deep learning, particularly in image classification, where small perturbations can mislead neural networks. Financial network strategies with time series input are also vulnerable to distributional adversarial attacks. By applying adversarial attacks and training, we solve a distributional robust optimization (DRO) problem, resulting in a more robust network toward both adversarial examples and model misspecification.

Periodic Evaluation of DC Pension Fund: A Dynamic Risk Measure Approach

Wanting HE

This paper introduces a novel framework for the periodic evaluation of Defined Contribution (DC) pension funds using dynamic risk measures and reinforcement learning. We propose an actor-critic algorithm with M-bootstrap to optimize investment-insurance strategies over the working horizon, incorporating dynamic conditional value-at-risk (CVaR) to manage tail risks and balance wealth accumulation, mortality risk, and future contributions. By integrating Lee-Carter mortality projections with exponential transitory jumps and renewal effects, we enhance longevity risk forecasting and demonstrate its impact on optimal strategies. Numerical results reveal that mortality projections significantly influence risk appetite and terminal wealth. Periodic evaluations promote stable and prudent strategies, whereas non-periodic evaluations encourage higher risk-taking under the same market scenario.

Arbitraging on Decentralized Exchanges

Xuedong He

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

Entropic Adapted Wasserstein Distance on Gaussians

Songyan Hou

The adapted Wasserstein distance is a metric for quantifying distributional uncertainty and

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

The Term Structure of Index Option Returns

Guanglian Hu

We study the term structure of index option returns by comparing realized option returns with their expected returns implied from option pricing models. We find that while the diffusive volatility risk premium can match 1-month index option returns well, it cannot fit the returns of long maturity options. The volatility jump risk premium shares the similar drawback in fitting the term structure of index option returns, though to a lesser extent. The jump risk premium is sufficient to explain the patterns in realized index option returns, but its success critically depends on a specification in which the variance of price jumps is priced. Finally, we present novel results on higher moments of option returns and also characterize the dynamics in the conditional expected index option returns.

New perspectives on analytic solvability of stochastic volatility models

Lilian Hu

Moment generating functions (mgf) of the terminal log-asset price and integrated variance of stochastic volatility models conditional on the terminal variance value are required in exact simulation algorithms and Fourier based algorithms for pricing European path dependent options. Broadie and Kaya (2006) initiated the analytic derivation of conditional mgf of integrated variance of the Heston model based on related analytic results for the Bessel bridge. Kang et al. (2017) and Zeng et al. (2023) employed different measure changes to obtain conditional mgfs of the Heston model, Wishart multidimensional stochastic volatility model, 4/2-model and Ornstein Uhlenbeck model. In this paper, we develop systematic and comprehensive measure change techniques that provide effective derivation procedures for the associated conditional mgfs. We establish an interesting linkage between conditional joint mgfs and their unconditional counterparts. We reveal the relations between our framework and other analytic derivation procedures; and examine analytic tractability of conditional mgfs under various types of stochastic volatility models. Interestingly, the conditional mgfs under the 4/2-model can be deduced from those under the Heston model via a measure change.

Menuless and Preference-Free Screening Contracts for Fund Managers

Sang Hu

We propose a family of incentive contracts that can attract some fund managers favored by all investors and deter any manager unfavored by some investors. This contracting problem has asymmetric information of hidden types and hidden knowledge of preferences as well as opportunity costs. In contrast to standard screening contracts, our contracts neither depend explicitly on the utilities of the managers and investors nor have a menu of choices. The contracts have two crucial components: (i) a first-loss deposit to offset some of the investors' losses and (ii) a liquidation boundary. A case study is also given.

A discretization scheme for path-dependent FBSDEs

JI UK JANG

This work studies a discretization scheme for solutions to forward-backward stochastic differential equations (FBSDEs) with path-dependent coefficients. We show the convergence of the Picard-type iteration to the FBSDEs solution and provide its convergence rate. Using this result, we establish a numerical method for solutions of second-order parabolic path-dependent partial differential equations. To achieve this, weak approximation of martingale representation theorem (Cont, Rama, and Yi Lu. “Weak approximation of martingale representations.” Stochastic Processes and their Applications 2016) is employed. Our results generalize the scheme for Markovian cases in (Bender, Christian, and Robert Denk. “A forward scheme for backward SDEs.” Stochastic processes and their applications, 2007). Furthermore, we present some numerical simulations for our scheme and error analysis. This research was conducted in collaboration with Hyungbin Park.

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

Yanwei Jia

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

in a volatile bear market, and decisively outperform the model-based continuous-time counterparts by significant margins.

Robust Optimal Strategies for Two-Period Liquidation in Financial Systems

Hongyi JIANG

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

A Hybrid Deep Reinforcement Learning Method for Insurance Portfolio Management

Zhuo Jin

This paper develops a hybrid deep reinforcement learning approach to manage an insurance portfolio for diffusion models. To address the model uncertainty, we adopt the recently developed modelling of exploration and exploitation strategies in a continuous-time decision-making process with reinforcement learning. We consider an insurance portfolio management problem in which an entropy-regularized reward function and corresponding relaxed stochastic controls are formulated. To obtain the optimal relaxed stochastic controls, we develop a Markov chain approximation and stochastic approximation-based iterative deep reinforcement learning algorithm where the probability distribution of the optimal stochastic controls is approximated by neural networks. In our hybrid algorithm, both Markov chain approximation and stochastic approximation are adopted in the learning processes. The idea of using the Markov chain approximation method to find initial guesses is proposed. A stochastic approximation is adopted to estimate the parameters of neural networks. Convergence analysis of the algorithm is presented. Numerical examples are provided to illustrate the performance of the algorithm.

Misaligned Clocks: Incentives for Differently Patient Boards and CEOs

Jussi Keppo

We study incentive contracts between a firm's shareholders (the board) and potential CEOs who

differ in their level of patience, which is private information. Through the contract, the board incentivizes the hired CEO to behave either patiently or impatiently, depending on her investment decisions—specifically, her choice among projects with varying completion times. First, we show that the optimal contract is linear. Second, we find that a sufficiently impatient board may design a contract that induces even a patient-type CEO to behave impatiently. Finally, when the contract can depend on the CEO's project choice, the board offers a pooling contract that directs both CEO types toward the same project. Altogether, this indicates that market interest rates not only reflect the patience of a representative investor and influence loan production, but also shape individual behavior through incentive design.

Long-term decomposition of robust pricing kernels under G-expectation

Jaehyun Kim

This study develops a BSDE method for the long-term decomposition of pricing kernels under the G-expectation framework. We establish the existence, uniqueness, and regularity of solutions to three types of quadratic G-BSDEs: finite-horizon G-BSDEs, infinite-horizon G-BSDEs, and ergodic G-BSDEs. Moreover, we explore the Feynman– Kac formula associated with these three types of quadratic G-BSDEs. Using these results, a pricing kernel is uniquely decomposed into four components: an exponential discounting component, a transitory component, a symmetric G-martingale, and a decreasing component that captures the volatility uncertainty of the G-Brownian motion. Furthermore, these components are represented through a solution to a PDE. This study extends previous findings obtained under a single fixed probability framework in [Hansen and Scheinkman, 2009], [Hansen, 2012] and [Qin and Linetsky, 2018] to the G-expectation context.

What matters in deep stochastic control?

Mueller Konrad J

TBA

Optimal Consumption and Portfolio Rules with Dynamic Adjustment of Consumption Bounds

Hyeng Keun Koo

We develop a comprehensive model to derive optimal consumption-investment choices for agents aiming to sustain spending power over an infinite horizon. Our model accounts for agents constrained by costly adjustable spending bounds and evaluates utility based on both current and future minimum allowable consumption levels, as well as actual consumption. Higher minimum levels assure investors that their consumption will not fall below a certain threshold, thereby generating utility from this anticipating security. However, adjusting these minimum levels incurs utility costs, leading to a non-trivial trade-off. The problem is

transformed into an optimal switching problem through its conversion to a dual problem. The optimal switching problem is analytically characterized as a two-dimensional double obstacle problem. Our explicit solutions reveal an empirically consistent consumption pattern and U-shaped portfolio choices, implying that investors exhibit loss aversion toward changes in consumption.

Neural Jump ODEs for Input-Output Systems

Florian Krach

In this talk we study the problem of (online) forecasting general stochastic processes using a path-dependent (PD) extension of the Neural Jump ODE (NJ-ODE) framework. The PD-NJ-ODE was the first framework to establish theoretical guarantees for the optimal prediction of irregularly observed time series data of generic, possibly non-Markovian or discontinuous, stochastic processes with incomplete observations, by utilising the reconstruction properties of the signature transform. In this work, we generalise the results, which were limited to predicting the exact same process that is observed, to input-output systems, where the observed and target processes can, but do not need to, coincide. Theoretically, this allows for weaker assumptions and practically this leads to training a model more specialised for the target task. We apply this generalised framework to stochastic filtering and parameter filtering problems. This is joint work with Jakob Heiss, Thorsten Schmidt and F´elix B. Tambe-Ndonfack.

Efficient simulation of the SARB model

Yue Kuen Kwok

We propose an efficient and reliable simulation scheme for the stochastic-alpha-beta-rho (SABR) model. The two challenges of the SABR simulation lie in sampling (i) the integrated variance conditional on terminal volatility and (ii) the terminal price conditional on terminal volatility and integrated variance. For the first sampling procedure, we analytically derive the first four moments of the conditional average variance, and sample it from the moment-matched shifted lognormal approximation. For the second sampling procedure, we approximate the conditional terminal price as a constant-elasticity-of-variance (CEV) distribution. Our CEV approximation preserves the martingale condition and precludes arbitrage, which is a key advantage over Islah's approximation used in most SABR simulation schemes in the literature. We then adopt the exact sampling method of the CEV distribution based on the shifted Poisson-mixture Gamma random variable. Our enhanced procedures avoid the tedious Laplace inversion algorithm for sampling integrated variance and non-e-cient inverse transform sampling of the forward price in some of the earlier simulation schemes. Numerical results demonstrate our simulation scheme to be highly efficient, accurate and reliable.

Global Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models

Mathieu Lauriere

We propose and compare new global solution algorithms for continuous time heterogeneous agent economies with aggregate shocks. First, we approximate the agent distribution so that equilibrium in the economy can be characterized by a high, but finite, dimensional non-linear partial differential equation. We consider different approximations: discretizing the number of agents, discretizing the agent state variables, and projecting the distribution onto a finite set of basis functions. Second, we represent the value function using a neural network and train it to solve the differential equation using deep learning tools. We refer to the solution as an Economic Model Informed Neural Network (EMINN). The main advantage of this technique is that it allows us to find global solutions to high dimensional, non-linear problems. We demonstrate our algorithm by solving important models in the macroeconomics and spatial literatures (e.g. Krusell and Smith (1998), Khan and Thomas (2007), Bilal (2023)). Joint work with Zhouzhou Gu, Sebastian Merkel, and Jonathan Payne.

Inference for Volatility Roughness from Options

Chenxu Li

Rough volatility models have recently gained significant attention in both academia and industry due to their potential advantages over traditional stochastic volatility models in option pricing and volatility forecasting. This paper focuses on statistical inference for volatility roughness using option data. As a theoretical foundation, we analyze the non-Markovian nature of volatility under a general rough volatility model by employing functional stochastic processes that capture the dependency of volatility on historical information. This approach allows us to derive a functional expression for option prices. We then obtain closed-form short-maturity asymptotics of the at-the-money shape characteristics of implied volatility surfaces. These results enable us to develop statistical methods for testing volatility roughness. Through Monte Carlo simulations, we validate the robustness and accuracy of these inference methods. Empirical analysis further confirms the pronounced volatility roughness, driven primarily by idiosyncratic shocks affecting volatility rather than common shocks influencing both the underlying asset's return and volatility.

Optimal Consumption for Recursive Preferences with Local Substitution under Risk

Hanwu Li

We explore intertemporal preferences that are recursive and account for local intertemporal substitution. First, we establish a rigorous foundation for these preferences and analyze their properties. Next, we examine the associated optimal consumption problem, proving the existence and uniqueness of the optimal consumption plan. We present an infinite-dimensional version of the Kuhn-Tucker theorem, which provides the necessary and sufficient conditions for optimality. Additionally, we investigate quantitative properties and the construction of the optimal consumption plan. Finally, we offer a detailed description of the structure of optimal consumption within a geometric Poisson framework.

An Option-Based Framework for Pricing Bond Repos

Jianong Li

From a unique proprietary data of bond-pledged repurchase agreements (repos) in China's interbank market, we find that interest rate risk and credit risk in bonds play key roles in pricing repos backed by these bonds, in contrast to the widespread belief that repos are secured and nearly risk-free. From a bond-option-theoretic perspective in continuous time, we build compatible models by studying in theory how the yield curve of repos backed by bonds (either default-free or defaultable, such as Treasury securities and corporate bonds) anchors at the yield curve of pledged bonds in a stochastic interest-rate environment under the arbitrage free principle. Additionally, we propose a novel analytical framework for pricing bond repos. Essentially, repo haircut is interpreted as the per-unit price of a call option on the pledged bond. Empirically, we implement our models to consistently price bond-pledged repos from the data. To better explain the data, we further incorporate the repo's potential bail-in (by cash borrower) or bail-out (by government or "national team"). Overall, repo-yield spread under our full model can be decomposed into three components: the compensations for interest-rate risk and default risk of pledged bonds, and the premium for the repo's bail-in/out potential. Repos are typically short-term in markets, and our models would also be useful for projecting longer-term arbitrage-free repo rates whose market data are not available.

Optimal investment and consumption problem associated with stochastic volatility models

Jianrui Li

This paper studies optimal consumption-investment problem for an agent who consumes and invests in a riskless asset and a risky one. The risky asset is characterized by a stochastic volatility model. The consumption strategy is only assumed to be a nondecreasing processes guaranteeing no bankruptcy. The investment proportion is constrained in a bounded set. Main challenges of this problem are to overcome the difficulties caused by Hölder continuous and super-linear growth coefficients induced from the stochastic volatility model. Then, the value function is proved to be a unique viscosity solution to a variational inequality.

Calibration for Stochastic Local Volatility Models

Jintao Li

As one of the most advanced asset pricing models currently available, the stochastic local volatility (SLV) model can simultaneously capture the stochastic nature of market volatility, leverage effects, and volatility clustering, providing a powerful tool for accurately describing market dynamics. Prior to applying this model in practice, the primary task lies in model calibration. This research focuses on the calibration of leverage functions within stochastic local volatility models. From the perspective of PDE inverse problem, we transform the calibration task into an optimal control problem incorporating Tikhonov regularization techniques. Unlike

the existing study (Saporito et al., 2019) whose objective function focuses on the L 2 distance between the theoretical and market local volatility, we adopt an optimization criterion from Chen et al. (2024)– to minimize the L 2 distance between the model's implied volatility and its market-observed counterpart. This objective function design substantially reduces the relative error in calibrated theoretical option prices, particularly for out-of-the-money options On the theoretical front, we establish the existence of a solution to the optimal control problem. Furthermore, we approximate the Gâteaux derivative of the objective function, which enables the development of a gradient-descent-like algorithm to solve the problem. Simulated and real market tests demonstrate superior performance of our method compared to the benchmark (Ren et al., 2007).

Controllability Concepts for Mean-Field Dynamics with Reduced-Rank Coefficients

Juan Li

In my talk, we consider different novel notions of exact controllability of mean-field linear controlled stochastic differential equations (SDEs). The key feature is that the noise coefficient is not required to be of full rank. We begin with showing that classical exact controllability with L^2-controls necessarily requires both rank conditions on noise. When this condition fails, we turn to similar rank requirements on drift allowing exact controllability by lowering the regularity of controls. When both the aforementioned rank conditions fail, we introduce and characterize a new notion of exact terminal controllability to normal laws (ETCNL). We also investigate a new class of Wasserstein-set-valued backward SDEs naturally associated to ETCNL.

Arbitrages in Perpetual Contracts

Linfeng Li

Perpetual contracts have gained widespread adoption in cryptocurrency markets, offering roughly delta-one exposure to underlying coins without a fixed expiration date. The perpetual contract price is anchored to the underlying index price through recurrent funding swaps between long and short positions. This funding rate incentivizes trades to narrow the gap between perpetual contracts and underlying index prices. Existing studies (Angeris et al., 2022; He et al., 2023; Ackerer et al., 2023) have studied how to price perpetual contracts. However, these works oversimplify the funding swap mechanism, limiting the practical applicability of their results. This paper establishes rigorous no-arbitrage bounds for both linear and inverse perpetual contracts by incorporating funding specifications and entry transaction costs. We demonstrate that arbitrage opportunities can be systematically exploited through dynamic trading strategies when perpetual contract prices deviate beyond these bounds. Additionally, we conduct an empirical study using data from Binance to validate the efficiency of our theoretical bounds.

Value-aware estimation of diffusion systems for continuous-time reinforcement learning: Surrogate loss and financial application

Lingfei Li

We estimate diffusion systems modeled by stochastic differential equations to solve optimal control problems in a continuous-time model-based reinforcement learning framework. Instead of minimizing a statistical loss, we perform estimation by looking for a model that minimizes the mismatch between the model-based value function and empirical rewards collected by a chosen policy. Compared with using classical statistical methods for model estimation, our value-aware approach can potentially reduce decision bias for the control problem. Nevertheless, the loss function in our approach is generally intractable to optimize due to the challenge in computing the model-based value function. In this paper, based on the backward stochastic differential equation satisfied by the model-based value function process, we develop a surrogate loss that can be efficiently optimized even in high dimensions. We prove convergence of the minimizer of the surrogate loss to that of the original loss and also analyze the error caused by time discretization. As an application, we consider the problem of liquidating stock portfolios and demonstrate the advantage of our estimation approach over regression in identifying weak signals from noisy environments.

Policy Iteration Reinforcement Learning Method for Continuous-Time Linear-Quadratic Mean-Field Control Problems

Na Li

This paper employs a policy iteration reinforcement learning (RL) method to study continuous-time linear-quadratic mean-field control problems in infinite horizon. The drift and diffusion terms in the dynamics involve the states, the controls, and their conditional expectations. We investigate the stabilizability and convergence of the RL algorithm using a Lyapunov Recursion. Instead of solving a pair of coupled Riccati equations, the RL technique focuses on strengthening an auxiliary function and the cost functional as the objective functions and updating the new policy to compute the optimal control via state trajectories. A numerical example sheds light on the established theoretical results.

Robust control of optimal investment and reinsurance under the path-dependent effect

Wenyuan Li

This paper investigates robust equilibrium investment and reinsurance strategies under a mean–variance criterion in a rough environment. To address decision-makers' concerns about model ambiguity, we introduce multiple alternative dominated models through changes of measures. In addition, we use a rough volatility model and a Hawkes process to capture the path dependence of the market. By adding auxiliary state variables, we degenerate a non-Markovian problem to a Markovian problem. Next, an explicit solution is derived for a path-dependent extended Hamilton–Jacobi–Bellman (HJB) equation. Our results show that individuals become

more risk-seeking in trading when considering path dependence, while more risk-averse when considering model ambiguity. Moreover, an individual's insurance demand increases with model ambiguity but remains steady with path dependence.

Discrete-Time Mean-Variance Strategy Based on Reinforcement Learning

Xun Li

This talk studies a discrete-time mean-variance model based on reinforcement learning. Compared with its continuous-time counterpart in (Zhou et.al 2020), the discrete-time model makes more general assumptions about the asset's return distribution. Using entropy to measure the cost of exploration, we derive the optimal investment strategy, whose density function is also Gaussian type. Additionally, we design the corresponding reinforcement learning algorithm. Both simulation experiments and empirical analysis indicate that our discrete-time model exhibits better applicability when analyzing real-world data than the continuous-time model.

Optimal housing, consumption and portfolio choices

Zhongfei Li

This paper investigates a housing, investment and consumption problem for a household who has stochastic income, consumes non-housing goods, invests in financial assets, and plan to buy a house for dwelling. When buying a house, the household is required to pay a fraction of the housing value as the down payment and finances the rest through a mortgage loan. The household obtains utility from the consumption (including rental housing) before buying a house, and from the consumption and owner-occupied housing after buying a house. The household aims to maximize the total expected discounted utility and to find the optimal consumption and investment strategies before and after buying a house, the optimal time to buy a house, and the optimal housing size. The optimization problem is solved by transforming it into a combined stochastic control and impulse control problem. The derived results show that the time preference, consumption preference, down payment ratio, and relative risk aversion coefficient play important roles. The optimal time to buy a house is the first time the house price hits the price threshold from above. The earlier the household buys a house, the smaller the house she tend to buy. If the household is extremely patient or impatient, she will postpone buying a house but will buy a larger one. The wealth in the investment account and the human capital raise consumption. The loan repayment reduces consumption after buying a house.

Structural Credit Rating Model with Stochastic Volatility

Jin Liang

In this paper, use corporate bonds as a means to assess the risk of credit rating migration with random volatility. According to the size of the company's assets, the company is divided into

high or low credit rating. The volatility of company assets regresses around different mean volatility under different credit ratings. By introducing a special zero-coupon coupon to hedge the risks caused by the randomness of volatility, a model by partial differential equation of the corporate bond value is derived, with continuous first-order partial derivatives about assets on the boundary of the credit rating migration. The theoretical analysis has been carried and through the ADI difference method, the numerical solution of the corporate bond value is obtained and then the influence of the parameters and the financial significance are discussed.

Commodity Market Liquidity in Two Days: The Impact of FOMC Announcements

Jingfeng Liu

I examine how scheduled Federal Open Market Committee (FOMC) announcements impact short-term liquidity, as documented in Kang et al. (2020). Commercials can earn an 8% annualized premium in the commodity futures market by providing liquidity to noncommercials ahead of scheduled FOMC announcements—twice the premium earned on non-FOMC days. Pre-FOMC heightened uncertainty risk deteriorates liquidity in the commodity futures market. Meanwhile, increased hedging motives prompt noncommercials to offer a higher uncertainty risk premium and a larger liquidity premium to incentivize commercials to make the market.

Covariance Matrix Estimation for Positively Correlated Assets

Weilong Liu

The comovement phenomenon in financial markets gives rise to decision scenarios where asset returns are positively correlated. This paper focuses on covariance matrix estimation for such scenarios, motivated by the observation of strong positive correlations in the monthly returns of factor-sorted portfolios. We demonstrate that fine-tuning the eigenvectors associated with weak factors within the rotation-invariant estimation framework can produce well conditioned covariance matrix estimates. We propose an Eigenvector Rotation Shrinkage Estimator (ERSE) that pairwise rotates eigenvectors while preserving orthogonality, equivalent to performing multiple linear shrinkage on two distinct eigenvalues. Empirical results from factor-sorted portfolios show that ERSE outperforms existing rotation-invariant estimators in reducing out-of-sample risk.

Anonymized risk sharing

Yang Liu

We study an axiomatic framework for anonymized risk sharing. In contrast to traditional risk sharing settings, our framework requires no information on preferences, identities, private operations and realized losses from the individual agents, and thereby it is useful for modeling risk sharing in decentralized systems. Four axioms natural in such a framework – actuarial

fairness, risk fairness, risk anonymity, and operational anonymity – are put forward and discussed. We establish the remarkable fact that the four axioms characterizes the conditional mean risk sharing rule, revealing the unique and prominent role of this popular risk sharing rule among all others in relevant applications of anonymized risk sharing. Several other properties and their relations to the four axioms are studied, as well as their implications in rationalizing the design of Bitcoin mining pools.

Portfolio optimisation with proportional transaction costs and stochastic investment opportunities

Zhesheng Liu

We determine the growth optimal portfolio under proportional transaction costs for an investor trading a risk-free asset and a risky asset with stochastic investment opportunities given by a linear diffusion. This problem has been introduced by Taksar, Klass and Assaf (1988) for asset prices with constant return rate and volatility. Despite extensive research, our results are the first that construct optimal trading strategies in continuous time beyond the restrictive setting of constant parameters. This allows us to investigate the tradeoff between active trading due to the random parameters and the proportional transaction costs. We solve this problem by explicitly constructing a shadow price process and provide the asymptotic expansion of the non-trade region, the stock-cash ratio and the proportion of wealth invested in the risky asset.

Regularized Multi-Currency Expected Shortfall Portfolios

Patrick Lucescu

This paper introduces a sparse and stable optimization approach for multi-currency asset allocation, where the investor minimizes the expected shortfall of their portfolio. We study the industry-standard currency overlay strategies, where the desired asset allocation is determined first, followed by the optimal currency allocation. Additionally, we propose a joint expected shortfall optimization approach that integrates assets and currencies, and analyze its performance compared to currency overlay strategies. Preliminary empirical results show that the joint optimization outperforms when the number of assets in the investment universe is relatively small, whereas the overlay approach is more effective for larger universes due to its dimensionality reduction. Both approaches achieve lower risk and higher risk-adjusted returns compared to non-regularized portfolios and naive benchmarks.

Markov Decision Process with Observation Cost in Food Production Chain

Tengyingzi Ma

Applying financial mathematics to microbial risk management offers innovative opportunities for advancing sustainable food systems. In this study, we enhance traditional Markov decision processes (MDPs) by incorporating observation costs, a critical yet often overlooked factor in

decision-making under uncertainty. By integrating these costs, we develop a robust framework to optimize food production chains, balancing microbial risk, product quality, and profitability. Using simulations of the milk powder production process, we model stochastic microbial risk, quality indices, and profit dynamics within a unified reward function. Our results demonstrate that this approach not only improves system stability but also achieves a threefold increase in profit compared to conventional methods. These outcomes underscore the transformative potential of decision-theoretic frameworks, particularly when augmented with observation costs, in enhancing food safety, operational efficiency, and economic sustainability.

Liquidity and Asset Prices

Johannes Muhle-Karbe

Do investors pay more for assets that are easy to buy and sell? Does the introduction of a transaction tax make markets less volatile? And what happens to interest rates when assets become easier to trade? Questions like this evidently cannot be answered by starting from a convenient stochastic process model for the asset prices at hand. Instead, to understand how prices depend on the liquidity of the market, they need to be derived from first principles by matching the supply and demand of the investors trading them. However, the corresponding fixed-point problems are intractable even in the simplest toy models. In this talk, we discuss how to overcome this impasse by using probabilistic methods to derive closed-form results in the practically relevant regime of large (but finite) liquidity. (Joint work in progress with Paolo Guasoni and Xiaofei Shi).

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

Joseph Mulligan

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

Indefinite linear-quadratic large population problem with partial observation

Tianyang Nie

We investigate an indefinite linear-quadratic partially observed large population system with common noise, where both the state-average and control-average are considered. All weighting matrices in the cost functional can be indefinite. We obtain the decentralized optimal strategies by the Hamiltonian approach and demonstrate the well-posedness of Hamiltonian system by virtue of relaxed compensator. The related Consistency Condition and the feedback form of decentralized optimal strategies are derived. Moreover, we prove that the decentralized optimal strategies are ε-Nash equilibrium by using the relaxed compensator.

Path Dependent Trading Strategies Using Signatures and Kernels

Futter Owen

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

Criticality of HJB equations in stochastic control problems

Jongjin Park

We investigate the relationship between Hamilton-Jacobi-Bellman (HJB) equations and ergodic-type elliptic eigenvalue problems. In particular, we establish the existence and uniqueness of the associated critical eigenpair and analyze the asymptotic behavior of HJB equation solutions over time. Specifically, we prove that the solution of the HJB equation converges to the critical eigenfunction. Moreover, our work provides a mathematical perspective on the interplay between stochastic control and nonlinear elliptic equations. As an application, we formulate the expected utility maximization problem using a weak formulation of Markovian control and demonstrate that the optimal expected utility corresponds to the solution of a specific HJB equation, interpreted in the weak sense of stochastic differential equations. Our results contribute to the understanding of the long-term behavior of optimal portfolio strategies in stochastic environments. This research was conducted in collaboration with Hyungbin Park.

A BSDE approach to designing funding rates for perpetual swaps in cryptocurrency markets

Hyungbin Park

In cryptocurrency markets, a significant challenge for perpetual swap issuers is ensuring that

the perpetual swap price remains aligned with the underlying asset value. This paper addresses this problem by exploring the relationship between funding rates and perpetual swap prices. Given specific funding rates, we uniquely determine the price and replicating portfolio of perpetual swaps through an arbitrage argument. Our findings indicate that by appropriately designing funding rates, the perpetual swap can be pegged to the underlying asset value. Additionally, we provide approximate funding rates for practical applications and assess the difference between the original funding rates and these approximate rates. To achieve these results, our study utilizes path-dependent infinite-horizon backward stochastic differential equations in conjunction with arbitrage pricing theory. Our key results are derived by establishing the existence and uniqueness of solutions to these BSDEs and developing the corresponding Feynman-Kac formula.

A Risk Sensitive Contract-unified Reinforcement Learning Approach for Option Hedging

Xianhua Peng

We propose a new risk sensitive reinforcement learning approach for the dynamic hedging of options. The approach focuses on the minimization of the tail risk of the final P&L of the seller of an option. Different from most existing reinforcement learning approaches that require a parametric model of the underlying asset, our approach can learn the optimal hedging strategy directly from the historical market data without specifying a parametric model; in addition, the learned optimal hedging strategy is contract-unified, i.e., it applies to different options contracts with different initial underlying prices, strike prices, and maturities. Our approach extends existing reinforcement learning methods by learning the tail risk measures of the final hedging P&L and the optimal hedging strategy at the same time. We carry out comprehensive empirical study to show that, in the out-of-sample tests, the proposed reinforcement learning hedging strategy can obtain statistically significantly lower tail risk and higher mean of the final P&L than delta hedging methods.

A target approach to Stackelberg games

Dylan Possamai

In this talk, we provide a general approach to reformulating any continuous-time stochastic Stackelberg differential game under closed-loop strategies as a single-level optimisation problem with target constraints. More precisely, we consider a Stackelberg game in which the leader and the follower can both control the drift and the volatility of a stochastic output process, in order to maximise their respective expected utility. The aim is to characterise the Stackelberg equilibrium when the players adopt “closed-loop strategies”, i.e. their decisions are based solely on the historical information of the output process, excluding especially any direct dependence on the underlying driving noise, often unobservable in real-world applications. We first show that, by considering the-second-order-backward stochastic differential equation associated with the continuation utility of the follower as a controlled state variable for the leader, the latter’s unconventional optimisation problem can be reformulated as a more standard stochastic control

problem with stochastic target constraints. Thereafter, adapting the methodology developed by Soner and Touzi or Bouchard, Élie, and Imbert, the optimal strategies, as well as the corresponding value of the Stackelberg equilibrium, can be characterised through the solution of a well-specified system of Hamilton–Jacobi–Bellman equations. For a more comprehensive in- sight, we illustrate our approach through a simple example, facilitating both theoretical and numerical detailed comparisons with the solutions under different information structures studied in the literature. This is a joint work with Camilo Hern´andez, Nicol´as Hern´andez Santib´an˜ez, and Emma Hubert.

Quantile-based Distributional Reinforcement Learning under Cumulative Prospect Theory and its Dynamic Optimality Characterization

Chi Seng Pun

Distributional reinforcement learning (RL) emerges as a powerful tool for modeling risk-sensitive sequential decisions, where leveraging distribution functions in place of scalar value functions has allowed for the flexible incorporation of risk measures. However, due to the inherent time inconsistency (TIC) in the use of numerous risk measures in sequential decision making, the nature of controls under distributional RL has remained a mystery. For its use in the risk-sensitive problems in management science, this paper seeks to fill the research gap by building on the cumulative prospect theory (CPT)-based analysis of human gambling behavior and the emergence of three policy classes under TIC: precommitment, equilibrium, and dynamically optimal. We focus on the prevailing quantile-based distributional RL (QDRL) for CPT risk measures. Our theoretical results extend some results from the risk-insensitive QDRL theory to CPT prediction, from which we derive the characterization of QDRL control as an approximate equilibrium of an intrapersonal game. We empirically demonstrate the efficacy of our CPT QDRL algorithm in approaching the equilibrium. Finally, by further exploring the economic interpretation of the three policy classes in their handling of TIC, we devise some metrics and instances relevant for driving interesting patterns of interactions between these policies, including when and how the equilibrium may be more desirable than the precommitment.

Existence of optimal contract for principal-agent problem with quadratic cost function

Shuaijie Qian

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

the agent's optimal effort through numerical analysis.

Periodic Evaluation with Non-Concave Utility

Cong Qin

A fund manager's performance is often evaluated annually and compared with a benchmark, such as a market index. In addition, the manager may be subject to trading constraints, such as limited use of leverage, no short-selling, and a forced liquidation clause. We formulate this as a periodic evaluation problem with a non-concave utility, a stochastic reference point, and trading constraints. The value function is characterized as the unique fixed point of a Hamilton-Jacobi-Bellman equation with periodic terminal and boundary conditions, which must be imposed carefully due to possible discontinuities at the terminal time and/or on the liquidation boundary. We find that adding a stochastic reference point and/or trading constraints can offer distinct economic insights and implications, consistent with existing empirical findings.

Mean-Field Games with Common Poissonian Noise

Joshue Ricalde

Mean-Field Games arise as limit problems of symmetric stochastic differential games and their equilibria, where the coupling between players is due only to statistics of the population. Later, it was proposed to strengthen this coupling by adding a common Gaussian process to the individual dynamics. These are the so-called Mean-Field Games with Common Noise. In this talk we propose adding a pure jump martingale process instead of a Brownian motion for the Common Noise. Furthermore, we consider the case where the stochastic intensity is driven by the statistical characteristics of the population as well, and derive the corresponding Stochastic Hamiltonian System along with a stochastic version of Pontryagin's Maximum Principle.

When BSDEs Meet Semimartingales

Marco Rodrigues

We consider backward stochastic differential equations (BSDEs) and reflected BSDEs in a general framework that allows for a unified treatment of certain discrete-time and continuous-time control problems with random time horizons. We discuss well-posedness results for BSDEs and reflected BSDEs with optional obstacle processes, assuming appropriately weighted square-integrable data. Furthermore, we explore related second-order BSDEs, focusing on their structure, well-posedness, and relevance within this broader framework. This talk is based on joint work with Dylan Possamaï and Alexandros Saplaouras.

Time-Inconsistent Stochastic Control: State Dependency

Mateo Rodriguez Polo

We address the challenge of Time-Inconsistent Stochastic Control within a continuous-time framework. The focus lies in uncovering a probabilistic representation in the form of a system of Backward Stochastic Differential Equations (BSDEs) that encapsulate the equilibrium value function in the case where the present state affecting the target functional triggers the inconsistency. Furthermore, we present applications that illustrate this theory using classical examples from the Stochastic Control literature.

Dependence-Based Distortion Risk Measures for Insurance Portfolio Analysis

Aniq Rohmawati

This work examines dependence-based distortion risk measures for assessing risks in insurance portfolios by incorporating dependence structures into the distortion risk framework. Distortion risk measures transform probability distributions to assess risk beyond traditional expectation-based methods. However, they are often applied to marginal distributions, which may overlook the role of dependencies in multi-risk environments. To address this limitation, this work introduces dependence modeling through copulas, allowing for a structured representation of joint risk behavior while preserving the marginal properties of individual risks. The conditional threshold is adjusted to account for tail dependence, improving the accuracy of risk evaluation for interconnected losses. The proposed approach is applied to insurance portfolio data, demonstrating the role of dependence-aware distortion risk measures in providing a more detailed assessment of risk concentration and extreme event behavior. These findings contribute to extending the applicability of distortion risk measures in portfolio risk management by highlighting the impact of dependence structures on overall risk exposure.

Time-Inconsistent Stochastic Games with Mean-Variance Preferences

Chiara Rossato

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

The Impact of Risk Transfers on Systemic Risk in Banking Networks

Fabian Sandmeier

We develop a framework to analyze the impact of risk transfers on systemic risk in banking networks. An initial network is constructed from overlapping portfolios with country-specific liquid- and illiquid-asset exposures. Risk transfers are then artificially imposed via credit default swaps and interest rate swaps. By choosing the share of illiquid asset exposures subject

to risk transfers, we obtain several networks upon which we apply macroprudential stress tests. Across the resulting networks, we are able to quantify both solvency and liquidity risks that emerge as the share of risk transfers increases.

Stability of backward propagation of chaos

Alexandros Saplaouras

It will initially be considered the asymptotic behavior of the solution of a mean-field system of Backward Stochastic Differential Equations with Jumps (BSDEs), as the multitude of the system equations grows to infinity, to independent and identically distributed (IID) solutions of McKean–Vlasov BSDEs. This property is known in the literature as backward propagation of chaos. Afterwards, it will be provided the suitable framework for the stability of the aforementioned property to hold. In other words, assuming a sequence of mean-field systems of BSDEs which propagate chaos, then their solutions, as the multitude of the system equations grows to infinity, approximates an IID sequence of solutions of the limiting McKean–Vlasov BSDE. The generality of the framework allows to incorporate either discrete-time or continuous-time approximating mean-field BSDE systems.

Dynamic Monotone Mean-variance Portfolio Optimisation with Independent Returns

Martin Schweizer

Mean-variance portfolio choice, while popular, has some serious drawbacks: it can lead to time-inconsistent problems, and it is not monotone. We consider instead its modification to monotone mean-variance (MMV) and want to find a dynamic portfolio strategy which maximises the MMV criterion for final wealth on a finite horizon. Assuming only that the underlying semimartingale asset price model has independent returns, we are able to provide a complete and explicit solution. The only assumption we need is a weak local absence-of-arbitrage condition, and we can show that our results are sharp. This is based on joint work with Ales Cerny and Johannes Ruf.

Non-concave distributionally robust stochastic control in a discrete time finite horizon setting

Julian Sester

In this article we present a general framework for non-concave distributionally robust stochastic control problems in a discrete time finite horizon setting. Our framework allows to consider a variety of different path-dependent ambiguity sets of probability measures comprising, as a natural example, the ambiguity set defined via Wasserstein-balls around path-dependent reference measures, as well as parametric classes of probability distributions. We establish a dynamic programming principle which allows to derive both optimal control and worst-case measure by solving recursively a sequence of one-step optimization problems. As a concrete application, we study the robust hedging problem of a financial derivative under an asymmetric

(and non-convex) loss function accounting for different preferences of sell- and buy side when it comes to the hedging of financial derivatives. As our entirely data-driven ambiguity set of probability measures, we consider Wasserstein-balls around the empirical measure derived from real financial data. We demonstrate that during adverse scenarios such as a financial crisis, our robust approach outperforms typical model-based hedging strategies such as the classical Delta hedging strategy as well as the hedging strategy obtained in the non-robust setting with respect to the empirical measure and therefore overcomes the problem of model misspecification in such critical periods.

Two-sided sticky short rates with jumps

Cheolmin Shin

In this paper, we study the sticky short rate structure whose process moves within the upper and lower bounds of a target rate. We formulate its dynamics using a system of stochastic differential equations (SDEs) and investigate the existence, uniqueness in law, and Feynman Kac formula of its weak solution. This formulation accommodates two arbitrary stickiness parameters, which characterize the speed at which the short rate departs from the left and right boundaries. To model jumps in the target rate, we incorporate a Hidden Markov chain whose transition probability matrix is governed by a stochastic process representing economic variables. As an application, we derive bond prices under this short rate model by using a solution of a parabolic partial differential equation (PDE) with boundary conditions. Additionally, we develop numerical methods for parameter calibration and estimation of bond prices. This research was conducted in collaboration with Dong-Hyun Ahn and Hyungbin Park.

The Effects of Income and HARA Utility on Optimal Consumption, Investment, and Retirement Decisions

Yong Hyun Shin

This study examines an infinite-horizon optimal consumption, investment, and retirement decision problem, in which an agent maximizes the expected utility of consumption based on a Hyperbolic Absolute Risk Aversion (HARA) utility function. The model distinguishes between pre-retirement and post-retirement phases, characterized by differing income levels Y_i $(i = 1,2)$ and baseline lump sums L_i $(i = 1,2)$ in the HARA utility function, respectively. Specifically, we assume that the agent receives a higher income but has a lower baseline lump sum before retirement, whereas after retirement, the agent receives a lower income but has a higher baseline lump sum, that is, $Y_1 > Y_2 \geq 0$ and $L_2 > L_1 \geq 0$. The decision to retire is modeled as a one-time, irreversible choice. To analyze this problem, we employ the martingale approach combined with a dual method. The key contributions of this study are: (1) deriving an explicit condition for the optimal retirement timing in terms of income and utility parameters, and (2) demonstrating the monotonic relationships between optimal consumption, portfolio allocation, and retirement timing with respect to income and the baseline lump sum in the HARA utility function.

Representation theorem of temporal copula families and their corresponding Markov processes

Liulei Sun

This paper develops a unified framework for characterizing the temporal dependence structure of Markov processes via the copula method. We first establish a representation theorem for temporal copulas of homogeneous Markov processes, and propose constructive methods for temporal copulas of general Markov processes. In both cases, intuitive probabilistic structures for the corresponding Markov processes are provided. Building on the representation theorem, we (i) introduce an explicit Markov martingale construction, with applications including the extensively studied fake Brownian motion, and (ii) generalize existing results on the β-mixing property for stationary Markov chains and extend them to both continuous time and non-stationary settings. Numerical simulations confirm the theoretical results, highlighting the importance of fitting the copula method into the framework of Markov processes.

Unbiased simulation of Asian options

Xiaolu Tan

We provide an extension of the unbiased simulation method for SDEs in Henry Labordère et al. [Ann Appl Probab. 27:6 (2017) 1-37] to the path-dependent case. In the setting, both the payoff and the SDE coefficient functions depend on the (weighted) average of the process, so that the expected value is solution to a linear path-dependent PDE.

Path dependent modelling in Finance with a view towards large language models

Josef Teichmann

We provide a continuous time version of Takens theorem for stochastic differential equations. This explains on the one hand the universality of path dependent modelling, which appears for instance in large language models or recently in Mathematical Finance, and, on the other hand, paves a path towards regression-based scenario generators.

Portfolio Selection in Contests

Alex Tse

In an investment contest with incomplete information, a finite number of agents dynamically trade assets with idiosyncratic risk and are rewarded based on the relative ranking of their terminal portfolio values. We explicitly characterize a symmetric Nash equilibrium of the contest and rigorously verify its uniqueness. The connection between the reward structure and

the agents' portfolio strategies is examined. A top-heavy payout rule results in an equilibrium portfolio return distribution with high positive skewness, which suffers from a large likelihood of poor performance. Risky asset holding increases when competition intensifies in a winner-takes-all contest.

Stochastic Graphon Games with Memory

Sturmius Tuschmann

We study finite-player dynamic stochastic games with heterogeneous interactions and non-Markovian linear-quadratic objective functionals. We derive the Nash equilibrium explicitly by converting the first-order conditions into a coupled system of stochastic Fredholm equations, which we solve in terms of operator resolvents. When the agents' interactions are modeled by a weighted graph, we formulate the corresponding non-Markovian continuum-agent game, where interactions are modeled by a graphon. We also derive the Nash equilibrium of the graphon game explicitly by first reducing the first-order conditions to an infinite-dimensional coupled system of stochastic Fredholm equations, then decoupling it using the spectral decomposition of the graphon operator, and finally solving it in terms of operator resolvents. Moreover, we show that the Nash equilibria of finite-player games on graphs converge to those of the graphon game as the number of agents increases. This holds both when a given graph sequence converges to the graphon in the cut norm and when the graph sequence is sampled from the graphon. We also bound the convergence rate, which depends on the cut norm in the former case and on the sampling method in the latter. Finally, we apply our results to various stochastic games with heterogeneous interactions, including systemic risk models with delays and stochastic network games.

Kernel Conditional Factor Models

Urban Ulrych

Factor models are widely employed in finance to capture the relationship between asset returns and their underlying factors. Traditionally, these models assume a linear relationship in learning factor loadings. This paper enhances factor models by introducing non-linearity through low-rank kernel functions, offering a flexible, non-parametric representation of complex, non-linear relationships between factors, returns, and asset characteristics. We utilize a reproducing kernel Hilbert space (RKHS) with the associated reproducing kernel as a hypothesis space for modeling the factor loadings, while cross-sectional ridge regression is used to directly learn the factor portfolios. This approach extends existing methods by incorporating non-linear dependence on characteristics, regularization for more factors, and additional characteristics such as industries. Empirical analysis shows that the proposed non-linear learning framework significantly outperforms traditional linear models in terms of out-of-sample performance, as measured by explained variation and optimal factor portfolio performance. This is a joint work with Pierre Collin-Dufresne and Damir Filipović.

A Tail-Driven Nonparametric Estimator for State Price Densities

Yating Wan

This paper proposes and implements a novel nonparametric method for estimating the state price density (SPD) over the entire state space including the tails. This SPD estimator achieves shape consistency at the tails in theory, which is unattainable by other existing nonparametric methods. Monte Carlo simulations demonstrate the accuracy and robustness of our method. In particular, our estimator accurately captures the risk-neutral tail distribution, which is often underestimated by existing alternative methods. In an empirical analysis based on S&P 500 options data, we evaluate the performance of our SPD estimation method out of sample and demonstrate that the estimates can serve as effective indicators for market conditions. Combining these two perspectives, we suggest that our SPD estimator renders a valuable tool for market prediction, with potential applications in asset and risk management practices.

PreFER: Interactive Robo-Advisor with Scoring Mechanism

Yuwei Wang

Instead of asking a client to specify her risk preference or learning it from her investment choice, we propose an inverse reinforcement learning (IRL) framework to learn her risk preference or the reward function by scoring. Specifically, the robo-advisor requests the client to score unadopted investment advice and extracts information from adopted ones for non-interactive periods. We develop the IRL through discrete-time Predictable Forward Exploratory Reward (PreFER) processes, where the exploration is regularized by Tsallis entropy. By interpreting the score as the acceptance probability of an advice, the preference learning becomes an inverse problem of finding the exploratory investment distribution of the client, given investment distribution recommended by the robo-advisor and an acceptance probability in the context of acceptance-rejection method proposed by Von Neumann. Demonstrations are made for both classes of CARA and CRRA utilities. We prove that the density function of the optimal exploratory control attains maximum at the classic optimal strategy in the absence of exploration. We establish a one-to-one correspondence between the observed score and inferred risk aversion parameter. As long as the scores are consistent in ordering, the biasness of scores does not affect the identification of the client's risk aversion for a sufficient large number of interactions. The PreFER process further predicts the risk preference at the next time point from the one just learned, leading to an aggregation of learning power.

Convergence of Policy Iteration for Entropy-Regularized Stochastic Control Problems

Zhenhua Wang

We investigate the Policy Iteration Algorithm(PIA) for the so-called exploratory HJB equations. This type of equation is first proposed by Wang, Zariphopoulou and Zhou (J. Mach. Learn. Res., 21, 2020). First, we study the convergence of the PIA for entropy-regularized stochastic control problems on an infinite time horizon. We consider the convergences of PIA with large

discountings for the cases when small control or no controls on the diffusion terms. Then, we extend the PIA framework for entropy-regularized control-stopping mixed problems. The first part of this presentation is based on a joint work with Hung Tran, and Yuming Paul Zhang.

Continuous-time q-learning algorithms for mean-field control problems

Xiaoli Wei

We study the continuous-time q-learning for continuous time mean-field control problems from two perspectives: the central controller, who can observe the population distribution, and the representative agent, who cannot. To this end, we introduce distinct q-functions: integrated q-function (Iq-function) for the central controller and integrated q-function in decoupled form (decoupled Iq-function) for the representative agent. We establish their (weak) martingale characterization together with the associated value functions. Additionally, we explore their relationships with the essential q function used in policy improvement iterations. Furthermore, we develop corresponding learning algorithms and present financial examples to validate their effectiveness.

Reinforcement learning without a market simulator: Learning the stopping intensity

Hoi Ying Wong

Reinforcement learning (RL) has been introduced to continuous-time finance with various applications, including portfolio management, robo-advising, and optimal stopping problems. Unlike stochastic control problems, which assume complete knowledge about the environment, RL obtains an optimal policy with incomplete knowledge about the environment and/or the reward function. In an unknown environment (or model), RL has to strike a balance between exploitation (optimal decision under an assumed model) and exploration, which allows the agent to interact with the environment in uncharted territory. In the continuous-time finance literature, this interaction is often achieved by constructing a market simulator of the environment. However, the construction of the market simulator is tricky in practice. In this talk, I will introduce a new concept to implement RL without a market simulator. The key is to use alternative data to signal the interaction between the agent and the market (environment). There are three potential sources of data: publicly available data, service-generated data, and proprietary data. While the talk will briefly mention some examples, the focus will be on learning the surrender intensity of variable annuities (VAs) managed by insurance companies. We will explain how a company can use its proprietary historical surrender data to learn about surrender decisions in the form of surrender intensity. Unlike financial products, the underlying fund of a VA, which is the investment portfolio of the firm, often only updates its value on a monthly basis. In addition, there are no financial derivatives on the fund for calibrating the process of the fund. We show that the surrender data can be used to learn the surrender intensity and implied volatility using RL. Numerical illustrations will be provided. (This is a joint work with Ling Wang).

Reinforcement Learning for Continuous-Time Mean-Variance Portfolio Selection in a Regime-Switching Market

Bo Wu

We propose a reinforcement learning (RL) approach to solve the continuous-time mean variance portfolio selection problem in a regime-switching market, where the market regime is unobservable. To encourage exploration for learning, we formulate an exploratory stochastic control problem with an entropy-regularized mean-variance objective. We obtain semianalytical representations of the optimal value function and optimal policy, which involve unknown solutions to two linear parabolic partial differential equations (PDEs). We utilize these representations to parametrize the value function and policy for learning with the unknown solutions to the PDEs approximated based on polynomials. We develop an actor-critic RL algorithm to learn the optimal policy through interactions with the market environment. The algorithm carries out filtering to obtain the belief probability of the market regime and performs policy evaluation and policy gradient updates alternately. Empirical results demonstrate the advantages of our RL algorithm in relatively long-term investment problems over the classical control approach and an RL algorithm developed for the continuous-time mean-variance problem without considering regime switches.

The Mechanism of Transaction Price

Lan Wu

金融市场的交易（高频）价格是市场参与者多方博弈的结果，在逐笔成交或撤单的颗粒度下，价格的生成是确定的，但却无法建立交易订单与价格（变化）的传导机制模型。本报告将分享我们组近年在交易价格形成机制方面的研究尝试。具体包括：1）周期性 Hawkes 订单行为下的理论价格模型。2）市场主力交易者的周期性交易行为如何通过博弈传导为价格的特征。3）交易者的个体信息与价格的关系。

Choquet Rating Criteria and Risk Consistency

Chenxi Xia

Credit ratings are widely used by investors as a screening device. We incorporate a notion of risk aversion that favors prudent investment decisions into credit rating criteria. Specifically, we are interested in a class of rating criteria generated by Choquet integrals. First, we show that the corresponding Choquet integral is unique when presented in a standard form, and, as such, these criteria can be appropriately referred to as "Choquet rating criteria". Second, we explore risk-aversion related properties designed to reflect investors' risk preferences concerning pooling effect consistency, increasing convex orders and economic scenarios during the asset pooling phase of structuring a deal. We characterize all Choquet rating criteria that are consistent with such properties and provide new examples. We also offer several new results about distortion risk measures for broader interests.

SOFR Futures Pricing under Affine Term Structure Model with Scheduled Jumps

Wencan XIA

Affine term structure models (ATSM) are a class of established models for risk-free rates (RFRs). It is well known that the Secured Overnight Financing Rate (SOFR, the overnight RFR rate index for USD) fluctuates around the Federal fund rate, and the latter is reset routinely by the Federal Open Market Committee (FOMC) and thus follows a jump process. In this paper, we adopt the ATSM with jumps to price 3-month and 1-month SOFR futures and then derive the convexity adjustment formulae between the SOFR futures rates and backward-looking SOFR term rates. We also develop price formulae for major SOFR derivatives, including mid-curve options on SOFR futures and SOFR swaptions. The major efforts of this paper, however, is put on the estimation of a three-dimensional Cox-Ingersoll-Ross (CIR) processes with jumps, a subclass of the ATSM with jumps, as a model for the SOFR dynamics, using techniques consisting of maximum likelihood estimation and extended Kalman filter. The estimated model is applied to derive the SOFR term rates through the convexity adjustment formulae, and SOFR discount curve is then constructed based on inputs of both SOFR term rates and SOFR swap rates.

A Study of Algorithmic Collusion in Multi-dealer-to-client Platforms

Wei Xiong

The rise of algorithmic trading has revolutionized market microstructure, and a typical example is multi-dealer-to-client (MD2C) platforms where multiple market makers compete for order flow. While algorithmic market making enhances efficiency, it also raises concerns regarding the potential for algorithmic collusion — where market makers can learn to set bid-ask spreads at supra-competitive levels, even without explicit coordination. In this study, we develop a game-theoretic framework to model competition among market makers on MD2C platforms and analyze their pricing strategies through decentralized multi-agent reinforcement learning (MARL). We first establish a notion of Competitive Nash Equilibrium in the MD2C setting, where dealers compete under a 'winner-takes-all' market share mechanism. We then introduce the notion of collusion that induces a formal reward-punishment scheme. To investigate the interactions of market making algorithms, we model the market makers' learning process through a decentralized multi-agent reinforcement learning algorithm, where market makers optimize their quoting strategies using policy gradient methods. We show that, even in the absence of price information sharing, under the specific 'winner-takes-all' mechanism, market prices may converge to levels which are similar to a collusion situation, resulting in 'tacit collusion'. We conduct extensive simulations across different market configurations, varying the number of dealers and tick size constraints. Our results show that, in concentrated markets with a small number of dealers, algorithmic strategies frequently converge to collusive outcomes, characterized by consistently maintained supra-competitive spreads. We show that the 'tacit collusion' phenomenon is robust under different market configurations. Increasing the number of market makers mitigates the level of 'tacit collusion', while a too crowded market reverts back to supra-competitive spreads. Furthermore, we observe that smaller tick sizes do

not tend to provide extra improvement in market competitiveness, with the learned spreads above Competitive Nash Equilibrium. We then discuss the implications of our research for market regulators, on optimal execution rules and market transparency. While current market regulations primarily focus on explicit collusion, our study highlights the potential risks posed by autonomous learning agents that inadvertently generate collusive pricing patterns. Market design choices, such as optimal execution rules and transparency enhancements, can play an important role in mitigating such risks. Our research contributes to the growing literature on the intersection of artificial intelligence, market microstructure, and financial regulation, emphasizing the need for regulatory oversight to address unintended consequences of autonomous market making algorithms.

Pricing and Hedging Strategies for Cross-Currency Equity Protection Swaps

Huansang Xu

In this paper, we explore the pricing and hedging strategies for an innovative insurance product called the equity protection swap (EPS). Notably, we focus on the application of EPSs involving cross-currency reference portfolios, reflecting the realities of investor asset diversification across different economies. The research examines key considerations regarding exchange rate fluctuations, pricing and hedging frameworks, in order to satisfy dynamic requirements from EPS buyers. We differentiate between two hedging paradigms: one where domestic and foreign equities are treated separately using two EPS products and another that integrates total returns across currencies. Through detailed analysis, we propose various hedging strategies with consideration of different types of returns nominal, effective, and quanto for EPS products in both separate and aggregated contexts. The aggregated hedging portfolios contain basket options with cross-currency underlying asset, which only exists in the OTC market, thus we further consider a superhedging strategy using single asset European options for aggregated returns. A numerical study assesses hedging costs and performance metrics associated with these hedging strategies, illuminating practical implications for EPS providers and investors engaged in international markets. We further employ Monte Carlo simulations for the basket option pricing, together with two other approximation methods—geometric averaging and moment matching. This work contributes to enhancing fair pricing mechanisms and risk management strategies in the evolving landscape of cross-currency financial derivatives.

Optimal portfolio selection with VaR and portfolio insurance constraints under the rank-dependent expected utility theory

Zuoquan Xu

We explore two optimal portfolio selection problems for an investor with rank-dependent utility preferences, focusing on managing risk exposure. The first problem incorporates a single Value-at-Risk (VaR) constraint, while the second includes joint VaR and portfolio insurance constraints. These models extend existing frameworks under expected utility theory, offering a more comprehensive approach to risk management. Using the martingale method, quantile

formulation, and relaxation techniques, we derive explicit optimal solutions. Notably, we identify a specific condition under which the VaR constraint becomes effective. Numerical analysis is conducted to validate the theoretical findings, providing deeper financial insights. Our results demonstrate that, in adverse market conditions, the optimal investment strategy significantly reduces risk compared to models with no constraints or only a single constraint. This work is a collaborative effort with Hui Mi from Nanjing Normal University.

Volatility-Managed Volatility Trading

Aoxiang Yang

We develop volatility risk premium timing strategies that trade two assets: a volatility asset and a risk-free asset. We first analyze a benchmark portfolio that sells a constant weight of volatility assets each month. Then, we show that a volatility-managed portfolio, which reduces selling volatility assets during periods of heightened volatility, considerably enhances long-run performance. Our findings are robust across variance swaps, VIX futures, and S&P 500 straddles, and even in the presence of transaction costs. An ex-post study indicates that timing portfolios yield positive alpha and reduce exposure relative to constant weight portfolios, mostly during volatility-spike periods rather than stable periods. Our findings help differentiate asset pricing theories on risk-return relations in the volatility asset market.

Portfolio Selection with Time-Varying Taxation

Chen Yang

The capital gains tax rate has fluctuated significantly over time, leading to substantial changes in investors' optimal strategies, as documented by the empirical studies. This paper proposes a novel continuous-time portfolio selection framework with a time-varying capital gains tax rate. Featuring differential tax rate announcement time and implementation time, our framework is able to capture the investors' anticipation over a potential future tax rate change before its announcement, as well as their reaction to an announced tax change yet to be implemented. The optimal investment strategy embodies the interaction between the time-varying tax rate and the lock-in and diversification effects proposed in the existing literature. Furthermore, our findings provide theoretical support for the permanent and transitory effects of tax rate changes documented in the empirical studies. The strength of the transitory effect depends on the size of the tax rate change, and the tax rate uncertainty mostly affects the transitory effect and has a negligible impact on the permanent effect. Moreover, the permanent effect vanishes under a zero interest rate while the transitory effect persists. This talk is based on a joint work with Xianhao Zhu.

A fast and universal framework for option pricing within 3D models

Nian Yang

This paper develops a unified framework for pricing options in three-dimensional (3D) models via a recursive numerical integration method based on small-time expansions of transition densities, such as Hermite polynomial based expansions. Specifically, we first extend smalltime expansions to longer horizons and analyze error propagation in our recursive algorithm, linking local transition density errors to global errors and determining the minimal expansion order for stability. Then we use anisotropic numerical integration to enable 3D model pricing on consumer-grade PCs, reducing storage needs from 10^{13} to 10^{9} float32 elements (10 GB), making the method feasible. In numerical experiments, we apply this framework to two classes of 3D models: one with two stochastic volatility factors (SV2F) and another with a stochastic interest rate and a volatility factor (SVSI), both including non-affine processes. Our CPU/GPU implementation ensures computational efficiency across a wide range of options, including European, Bermudan, discretely monitored barrier, lookback, and Asian options. Under a typical 1-year SV2F model, the GPU prices European, Bermudan, and Barrier options in 1.5 minutes, compared to 41 minutes on the CPU. For lookback and Asian options, the GPU takes 3 minutes, whereas the CPU requires 700 minutes (11.7 hours).

An Extended Merton Problem with Relaxed Benchmark Tracking

Xiang Yu

This paper studies a Merton's optimal portfolio and consumption problem in an extended formulation by incorporating the benchmark tracking on the wealth process. We consider a relaxed tracking formulation such that the wealth process compensated by a fictitious capital injection outperforms the benchmark at all times. The fund manager aims to maximize the expected utility of consumption deducted by the cost of the capital injection, where the latter term can also be regarded as the expected largest shortfall of the wealth with reference to the benchmark. By considering an auxiliary state process, we formulate an equivalent stochastic control problem with state reflections at zero. For general utility functions and Itô benchmark process, we rigorously develop a convex duality theorem, new to the literature, to the stochastic control problem with state reflections in which our dual process also exhibits reflections from above at a constant barrier. For power utility and geometric Brownian motion benchmark process, we can further derive the optimal portfolio and consumption in feedback form using the new duality theorem, allowing us to discuss some interesting financial implications induced by the additional risk-taking from the capital injection and the goal of tracking.

Random risk measures on the order statistics of correlated risks

Xin Zang

The random risk measure has been studied for extending the traditional deterministic risk measure by incorporating available information. This paper investigates the random risk measure based on order statistics (RRM-OS), i.e., the linear combination of order statistics (L-statistics) from observations correlated by a copula. Based on an integral representation and in a random axiomatic system, we show that a random risk measure can be represented as the

RRM-OS under some axiomatic conditions. Thus a connection between the risk measures and the L-statistics is established, and the copula method is introduced into the study of random risk measures. For the RRM-OS under the conditional independent structure, some limit properties are discussed, including the strong law of large numbers and the central limit theorem. The numerical experiments reveal the connections and differences between the conditional independent and traditional independent structure. The empirical analysis show that the estimator of VaR under the conditional independent assumption is more effective than the estimator under the traditional independent assumption.

Analytical solvability and exact simulation in models with affine stochastic volatility and Lévy jumps

Pingping Zeng

We investigate analytical solvability of models with affine stochastic volatility (SV) and Lévy jumps by deriving a unified formula for the conditional moment generating function of the log-asset price and providing the condition under which this new formula is explicit. The results lay a foundation for a range of valuation, calibration, and econometric problems. We then combine our theoretical results, the Hilbert transform method, various interpolation techniques, with the dimension reduction technique to propose unified simulation schemes for solvable models with affine SV and Lévy jumps. In contrast to traditional exact simulation methods, our approach is applicable to a broad class of models, maintains good accuracy, and enables efficient pricing of discretely monitored path-dependent derivatives. We analyze various sources of errors arising from the simulation approach and present error bounds. Finally, extensive numerical results demonstrate that our method is highly accurate, efficient, simple to implement, and widely applicable.

Trading Volume Alpha

Chao Zhang

We model the economic benefits of predicting individual stock trading volume. We translate volume forecasts into a component of expected trading costs and analyze the value of our predictions through a portfolio framework. By recasting the volume prediction problem into a portfolio optimization problem that trades off tracking error versus net-of-cost performance, we quantify volume predictions into economic outcomes. In addition, incorporating the economic loss function directly into a machine learning algorithm yields even better out-of-sample performance. While volume is only one component of what drives trading costs, it is highly predictable, readily available, and we show that the economic benefits of predicting individual stock volume are as large as those from stock return predictability.

Pricing American Parisian Options under General Time-Inhomogeneous Markov Models

Gongqiu Zhang

This paper develops general approaches for pricing various types of American-style Parisian options (down-in/-out, perpetual/finite-maturity) with general payoff functions based on continuous-time Markov chain (CTMC) approximation under general 1D time inhomogeneous Markov models. For the down-in types, by conditioning on the Parisian stopping time, we reduce the pricing problem to that of a series of vanilla American options with different maturities and their prices integrated with the distribution function of the Parisian stopping time yield the American Parisian down-in option price. This facilitates an efficient application of CTMC approximation to obtain the approximate option price by calculating the required quantities. For the perpetual down-in cases under time homogeneous models, significant computational cost can be reduced. The down-out cases are more complicated, for which we use the state augmentation approach to record the excursion duration and then the approximate option price is obtained by solving a series of variational inequalities recursively with the Lemke's pivoting method. We show the convergence of CTMC approximation for all the types of American Parisian options under general time-inhomogeneous Markov models, and the accuracy and efficiency of our algorithms are confirmed with extensive numerical experiments.

Profitability of collusive sandwich attack in automated market maker exchanges

Junhuan Zhang

In decentralized exchanges (DEXs), sandwich attackers may collude with validators through private memory pools. We investigate the collusive behaviors between sandwich attackers and colluding validators in terms of a new game-theoretic market microstructure model of automated market maker based decentralized exchange. We also explore the impacts of this collusion on the other agents including arbitrageurs employing the two-point (TA) arbitrage strategy, liquidity providers using the balancing swap fees and impermanent loss strategy, swap traders using the moving average (MA) and zero intelligence (ZI) strategies, and ordinary validators with honest strategy. The deep reinforcement learning is used to optimize the collusive trading strategy of sandwich attackers. We validate our model using daily closing prices of USDT/BTC from January 1, 2023 to December 31, 2023. The results show that: 1) the collusive sandwich attack strategy is effective and profitable; 2) colluding validators receive larger rewards than those from only verifying transactions; 3) the collusion does not affect the rewards of the ordinary validators, but reduces the rewards of the liquidity providers; 4) rewards for the TA arbitrageurs are diminished as a result of collusion; 5) MA and ZI swap traders consistently suffer losses.

No Sparsity in Asset Pricing: Evidence from a Generic Statistical Test

Lingxiao Zhao

We present a novel test to determine sparsity in characteristic-based factor models. Applying the test to industry and pseudo-random portfolios, we reject the null hypothesis that fewer than ten factors are sufficient to explain returns, and show that at least forty factors are needed for the various sample periods examined. We find that dense models outperform sparse ones in

both pricing and investing. Testing with tree-based portfolios also indicates no sparsity. Our results suggest that most existing factor models, which have fewer than six factors, are questionable, and that future research on such low-dimensional models is unlikely to be fruitful.

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

Ruixun Zhang

Financial scenario simulation is essential for risk management and portfolio optimization, yet it remains challenging especially in high-dimensional and small data settings common in finance. We propose a diffusion factor model that integrates latent factor structure into generative diffusion processes, bridging econometrics with modern generative AI to address the challenges of the curse of dimensionality and data scarcity in financial simulation. By exploiting the low-dimensional factor structure inherent in asset returns, we decompose the score function--a key component in diffusion models--using time-varying orthogonal projections, and this decomposition is incorporated into the design of neural network architectures. We derive rigorous statistical guarantees, establishing nonasymptotic error bounds for both score estimation at $O(d^{\frac{2}{5}}\ n^{\frac{-2}{k+5}})$ and generated distribution at $O(d^{\frac{4}{5}}\ n^{\frac{-1}{2(k+5)}})$, primarily driven by the intrinsic factor dimension k rather than the number of assets d, surpassing the dimension-dependent limits in the classical nonparametric statistics literature and making the framework viable for markets with thousands of assets. Numerical studies confirm superior performance in latent subspace recovery under small data regimes. Empirical analysis demonstrates the economic significance of our framework in constructing mean-variance optimal portfolios and factor portfolios. This work presents the first theoretical integration of factor structure with diffusion models, offering a principled approach for high-dimensional financial simulation with limited data.

α-Potential Games: A New Paradigm for N-player Dynamic Games

Yufei Zhang

Static potential games, pioneered by Monderer and Shapley (1996), are non-cooperative games in which there exists an auxiliary function called static potential function, so that any player's change in utility function upon unilaterally deviating from her policy can be evaluated through the change in the value of this potential function. The introduction of the potential function is powerful as it simplifies the otherwise challenging task of finding Nash equilibria for non-cooperative games: maximizers of potential functions lead to the game's Nash equilibria. In this talk, we propose an analogous and new framework called α-potential game for dynamic N-player games, with the potential function in the static setting replaced by an α-potential function. We present an analytical characterization of α-potential functions for any dynamic game. For stochastic differential games in which the state dynamic is a controlled diffusion, α is explicitly identified in terms of the model coefficients. We further show the α-Nash equilibrium of the stochastic game can be constructed through an associated conditional

McKean-Vlasov control problem. To illustrate our findings, we examine a linear-quadratic game on a graph, where α captures asymmetric interactions and player heterogeneity beyond the mean-field paradigm.

Convergence of PPGM for Linear Quasi-Convex Problems

Harry Zheng

We prove convergence of the proximal policy gradient method (PPGM) for a class of constrained stochastic control problems with control in both the drift and diffusion of the state process. The problem requires either the running or terminal cost to be strongly convex in either state or control, but the other functions may be non-convex and non-differentiable. We provide sufficient conditions under which the control iterates converge linearly to the optimal control, via deriving representations and estimates of the solution to adjoint BSDEs. We introduce numerical algorithms for implementing this method using deep learning and ODE based techniques, allowing us to solve stochastic control problem with high accuracy and scalability with dimension. Numerical examples are provided to demonstrate the accuracy and verify convergence of the algorithm.

Convergence analysis on the particle systems with centralized control

Chao Zhou

This paper deals with the optimization problem of a class of controlled N-particle systems. We establish the regularity results, which is uniform in N , on the HJB equations corresponding to the N-particle system. The uniform regularity results are obtained by the stochastic maximum principle and the analysis on a Riccati type BSDE. Using the uniform regularity results, we show the convergence of value function and optimal control as the number N of particles tends to infinity, where the convergence rates are also given. This is based on a joint work with Huafu Liao, Alpar Meszaros and Chenchen Mou.

Liquidity, exploration and inference-based learning to reduce transaction cost

Sixian Zhuang

Reinforcement learning (RL) for trading has been the fast-growing field in mathematical finance since the success of machine learning algorithms in many practical problems. Unlike the classic stochastic control approach, which outputs an optimal policy for a given stochastic model, RL enables the investor not only to optimize the reward with a model (exploitation) but also to interact with the market, or equivalently the environment, through generating alternative policies (exploration). The success of RL relies on striking the best balance between exploitation and exploration. In this research, we show that exploration is unavoidable due to the ever-existing liquidity issue in trading. The optimized policy cannot be executed perfectly because the random liquidity spread makes the realized policy randomized, causing the investor

to encounter natural and passive exploration. We propose calibrating the temperature parameter, or the level of exploration, in RL through liquidity data. Then, we propose an inference-based learning approach to form investment policies which reduce transaction costs due to liquidity. Our proposal involves a novel relaxed stochastic control problem with regularization that combines Shannon entropy and KL-divergence relative to the optimal exploratory policy. We derive an explicit solution to the relaxed stochastic control problem. Simulation demonstrates how this approach effectively reduces transaction costs and outperforms classic and exploratory optimal policies.

Multiverse Equivalent Expectation Measures for Computing Moments of Contingent Claim Returns

Xiaoyang Zhuo

This paper introduces Multiverse Equivalent Expectation Measures (MEEMs) for deriving analytical solutions to the variance, covariance, and higher-order moments and co-moments of contingent claim returns over a finite horizon. We show that the solutions for the Mth-order moment or co-moments require constructing a MEEM within an expanded probability product space of M identical marginal probability spaces, representing $M - 1$ parallel universes in addition to the original universe. The state variables evolve identically across all universes up to the horizon date but diverge independently thereafter. Using different classes of MEEMs, we offer a comprehensive analytical framework for deriving analytical solutions of higher-order moments and co-moments of returns on financial derivatives and fixed income securities. This unified framework significantly broadens the application of econometric methods to asset pricing and portfolio management for contingent claims.

Numerical approximation of RBSDEs via regularization

Yihan Zou

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