Recent Advances on Quantitative Finance

The Hong Kong Polytechnic University

August 27-30, 2023

The aim and scope of the conference:
The meeting aims to feature the latest developments in the field of quantitative finance. We expect that the meeting will enhance interaction and cooperation among researchers worldwide working in this field and provide an ideal platform for local students and researchers to access cutting-edge research. We will focus on, but are not limited to, the following three topics:

(i) stochastic control in quantitative finance,
(ii) dynamic game and mean-field game in quantitative finance, and
(iii) machine learning and reinforcement learning in quantitative finance.

Organizers:

- Department of Applied Mathematics & Research Centre for Quantitative Finance, The Hong Kong Polytechnic University
- Centre for Financial Engineering, The Chinese University of Hong Kong
- Centre for Quantitative Finance, National University of Singapore

Local organizing committee:

- Min DAI (Chair), The Hong Kong Polytechnic University
- Kexin CHEN, The Hong Kong Polytechnic University
- Nan CHEN, The Chinese University of Hong Kong
- Guanxing FU, The Hong Kong Polytechnic University
- Xuedong HE, The Chinese University of Hong Kong
- Jianhui HUANG, The Hong Kong Polytechnic University
- Zhaoli JIANG, The Hong Kong Polytechnic University
- Xun LI, The Hong Kong Polytechnic University
- Zuoquan XU, The Hong Kong Polytechnic University
- Chen YANG, The Chinese University of Hong Kong
- Xiang YU, The Hong Kong Polytechnic University
- Chao ZHOU, National University of Singapore

On-campus Internet Access
Wi-Fi access is provided from 27–30 August 2023. Please use the login name and password listed below:

User ID: gstraqf

Password: August2023

Remark: Only SSID: PolyUWLAN can be used.

Schedule

Day 1 (Sunday, August 27, 2023)

08:00 - 8:50	Registration (**V322**)		
08:50 - 9:15	Opening Address (**V322**)		
Plenary Talks (**V322**)			
Chair: Min DAI			
09:15 – 10:00	Mete SONER	Synchronization in a Kuramoto Mean Field Game	
10:00 – 10:45	Xin GUO	MF-OMO: An Optimization Framework for Mean Field Games	
10:45 – 11:15	*Coffee Break*[1]		
Chair: Nan CHEN			
11:15 – 12:00	Steve KOU	Bitcoin Mining, Electricity Consumption, and Climate Damages	
12:00 – 12:45	Ulrich HORST	Mean-Field Liquidation Games with Market Drop-out	
12:45 – 13:45	*Lunch*[1] *(Communal Staff Restaurant, 4/F, Communal Building)*		
13:45 – 14:30	Free Discussion		
	Session 1 (**HJ303**)	Session 2 (**HJ304**)	Session 3 (**HJ305**)
	Chair: Kexin CHEN	Chair: Jianhui HUANG	Chair: Xun LI
14:30 – 14:55	Yue-Kuen KWOK	Xiaolu TAN	Hanqing JIN
14:55 – 15:20	Hoi Ying WONG	Xingye YUE	Shuoqing DENG
15:20 – 15:45	Ling QIN	Panpan ZHANG	Xun LI
15:45 – 16:10	Yang LIU	Huilin ZHANG	Ziqi ZHOU
16:10 – 16:40	*Coffee Break*[1]		
	Session 1 (**HJ303**)	Session 2 (**HJ304**)	Session 3 (**HJ305**)
	Chair: Zhaoli JIANG	Chair: Guanxing FU	Chair: Zuoquan XU
16:40 – 17:05	Ning CAI	Shuaijie QIAN	Yingda SONG
17:05 – 17:30	Henry CHIU	Cong QIN	Yuhong XU
17:30 – 17:55	Xiao WEI	Xiang YU	Zuoquan XU
17:55 – 18:20	Zhaoli JIANG	Shihao ZHU	Shuhui LIU
18:30 – 19:30	*Dinner*[2] *(Communal Staff Restaurant, 4/F, Communal Building)*		

1. Coffee Break & Lunch are provided to all registered attendees, including registered students, local organizers and student helpers.

2. Dinner is provided to plenary speakers, invited speakers, registered faculty and industry attendees (registered students are excluded), local organizers and student helpers.

Day 2 (Monday, August 28, 2023)

Plenary Talks (**V322**)		
Chair: Zuoquan XU		
09:00 – 09:45	Nizar TOUZI	Mean Field Game of Mutual Holding with Common Noise
09:45 – 10:30	Johannes MUHLE-KARBE	Pre-Hedging
10:30 – 11:00	*Coffee Break*[1]	
Chair: Xiang YU		
11:00 – 11:45	Zengjing CHEN	A Central Limit Theorem, Loss Aversion and Multi-Armed Bandits
11:45 – 12:30	Wolfgang Karl HÄRDLE	On Crypto-backed Loans
12:30 – 13:30	*Lunch*[1] *(Communal Staff Restaurant, 4/F, Communal Building)*	
13:30 – 18:00	Free Discussion	
18:00 – 20:00	*Conference Banquet*[2] *(V Cuisine, 2/F, Harbour Crystal Centre)*	

1. Coffee Break & Lunch are provided to all registered attendees, including registered students, local organizers and student helpers.
2. The Banquet is an exclusive event for plenary speakers, invited speakers, banquet-registered attendees, local faculty members, local organizers and student helpers.

Day 3 (Tuesday, August 29, 2023)

Plenary Talks (**V322**)			
Chair: Chao ZHOU			
09:00 – 09:45	Martin SCHWEIZER	A New Stochastic Fubini Theorem via Stochastic Integrals of Measure-valued Processes	
09:45 – 10:30	Huyên PHAM	Nonparametric Generative Modeling for Time Series via Schrödinger Bridge	
10:30 – 11:00	*Coffee Break*[1]		
Chair: Xuedong HE			
11:00 – 11:45	Paolo GUASONI	Lightning Network Economics: Channels and Topology	
11:45 – 12:30	Hao XING	The Dark Side of Circuit Breakers	
12:30 – 13:30	*Lunch*[1] *(Communal Staff Restaurant, 4/F, Communal Building)*		
13:30 – 14:30	Free Discussion		
	Session 1 (**HJ303**)	Session 2 (**HJ304**)	Session 3 (**HJ305**)
	Chair: Dohyun AHN	Chair: Yanwei JIA	Chair: Chen YANG
14:30 – 14:55	Simone SCOTTI	Nan CHEN	Phillip YAM
14:55 – 15:20	Jing XU	Xiaofei SHI	Nian YANG
15:20 – 15:45	Lu LIU	Yanwei JIA	Chen YANG
15:45 – 16:10	Dohyun AHN	Yufei ZHANG	Francis LIU
16:10 – 16:40	*Coffee Break*[1]		
	Session 1 (**HJ303**)	Session 2 (**HJ304**)	Session 3 (**HJ305**)
	Chair: Chao ZHOU	Chair: Wei JIANG	Chair: Lingfei LI
16:40 – 17:05	Bin LI	Marko WEBER	Xianhua PENG
17:05 – 17:30	Jiacheng ZHANG	Wenyuan LI	Lingfei LI
17:30 – 17:55	Anran HU	Wei JIANG	Wenhao XU
17:55 – 18:20	Yuwei WANG	Yinuo WANG	Bo WU
18:30-19:30	*Dinner*[2] *(Communal Staff Restaurant, 4/F, Communal Building)*		

1. Coffee Break & Lunch are provided to all registered attendees, including registered students, local organizers and student helpers.
2. Dinner is provided to plenary speakers, invited speakers, registered faculty and industry attendees (registered students are excluded), local organizers and student helpers.

Day 4 (Wednesday, August 30, 2023)

Plenary Talks (**V322**)			
Chair: Guanxing FU			
09:00 – 09:45	Thaleia ZARIPHOPOULOU	**Online via Zoom** Meeting ID: 974 4131 9609 Passcode: 847555 https://polyu.zoom.us/j/98621606279?pwd=bHA1VWxJR3ErV3FVMzAzT2lvODBVdz09	
9:45 – 10:15	*Coffee Break*[1]		
	Session 1 (**HJ303**)	Session 2 (**HJ304**)	Session (**HJ305**)
	Chair: Xiang Yu	Chair: Zhaoli JIANG	Chair: Kexin CHEN
10:15 – 10:40	Xiangyu CUI	Ruting WANG	Kexin CHEN
10:40 – 11:05	Yuchao DONG	Meng-Jou LUA	Hui BU
11:05 – 11:30	Xiaoli WEI	Zihao SONG	Jiacheng FAN
11:30 – 11:55	Boyu WANG	Yutian ZHOU	Yijie HUANG
12:00 – 13:00	*Lunch*[1] *(Communal Staff Restaurant, 4/F, Communal Building)*		
13:00 – 14:30	Free Discussion		
Plenary Talks (**V322**)			
Chair: Xun LI			
14:30 – 15:15	Xunyu ZHOU	Continuous-Time Reinforcement Learning	
15:15 – 16:00	Julien PRAT	Contagion in Decentralized Lending Protocols	
16:00 – 16:10	Closing Address (**V322**)		

1. Coffee Break & Lunch are provided to all registered attendees, including registered students, local organizers and student helpers.

Plenary Talks

A Central Limit Theorem, Loss Aversion and Multi-Armed Bandits

Zengjing CHEN

School of Mathematics and Zhongtai Securities Institute for Financial Studies, Shandong University

This paper studies a multi-armed bandit problem where the decision-maker is loss averse, in particular she is risk averse in the domain of gains and risk loving in the domain of losses. The focus is on large horizons. Consequences of loss aversion for asymptotic (large horizon) properties are derived in a number of analytical results. The analysis is based on a new central limit theorem for a set of measures under which conditional variances can vary in a largely unstructured history-dependent way subject only to the restriction that they lie in a fixed interval.

Lightning Network Economics: Channels and Topology

Paolo GUASONI

School of Mathematical Sciences, Dublin City University

Designed to address Bitcoin's scalability challenge, the Lightning Network (LN) is a protocol allowing two parties to secure bitcoin payments and escrow holdings between them. In a lightning channel, each party commits collateral towards future payments to the counterparty and payments are cryptographically secured updates of collaterals. First, we identify conditions for two parties to optimally establish a channel, find explicit formulas for channel costs and optimal collaterals, and derive the implied reduction in congestion of the blockchain. Then we obtain necessary conditions for cost-minimizing topologies and bounds on the cost of the optimal topology, showing the unusual circumstances in which it is a hub that connects all other nodes.

MF-OMO: An Optimization Framework for Mean Field Games

Xin GUO

Department of Industrial Engineering and Operations Research, University of California, Berkeley

The literature on theory and computation of mean-field games (MFGs) has grown exponentially since its inception. In this talk, we present MF-OMO (Mean-Field Occupation Measure Optimization), a mathematical framework to analyze MFGs. This approach is not limited to contractive or monotone settings, nor will it need an a priori assumption of the uniqueness of the Nash equilibrium (NE). MF-OMO reformulates the problem of finding NE solutions in MFGs as a single optimization problem. This formulation thus allows for directly utilizing various optimization tools, algorithms and solvers to find NE solutions of MFGs. We also provide convergence guarantees for finding (multiple) NE solutions using popular algorithms such as projected gradient descent. For MFGs with linear rewards and mean-field independent dynamics, solving MF-OMO can be reduced to solving a finite number of linear programs, hence solved in finite time.

On Crypto-backed Loans

Wolfgang Karl HÄRDLE

Institut für Statistik und Ökonometrie, Humboldt-Universität zu Berlin

The development of distributed ledger technology renovates the form of lending and borrowing: borrowers and lenders can now transact without traditional intermediaries, e.g. banks. Instead, the borrowing and lending process is performed and governed by pre-agreed, immutable protocols and applications deployed on blockchains. The lending process heavily reply on crypto collateral - crypto-backed loans. This paper introduces a mathematical model to price the loan position. The model justifies the motivation of borrowers and allows us to determine the arbitrage-free borrowing premium known as the fair borrowing premium. We study the characteristics of the loan position and fair borrowing premium, showing that the crypto-back loan term structure must be contango.

Mean-Field Liquidation Games with Market Drop-out

Ulrich HORST

Department of Mathematics, and School of Business and Economics, Humboldt- Universität zu Berlin

We consider a novel class of portfolio liquidation games with market drop-out ("absorption"). More precisely, we consider mean-field and finite player liquidation games where a player drops out of the market when her position hits zero. In particular round-trips are not admissible. This can be viewed as a no statistical arbitrage condition. In a model with only sellers we prove that the absorption condition is equivalent to a short selling constraint. We prove that equilibria (both in the mean-field and the finite player game) are given as solutions to a non-linear higher- order integral equation with endogenous terminal condition. We prove the existence of a unique solution to the integral equation from which we obtain the existence of a unique equilibrium in the MFG and the existence of a unique equilibrium in the N-player game. We establish the convergence of the equilibria in the finite player games to the obtained mean-field equilibrium and illustrate the impact of the drop-out constraint on equilibrium trading rates.

Bitcoin Mining, Electricity Consumption, and Climate Damages

Steven KOU

Department of Finance, Questrom School of Business, Boston University

We propose a tractable dynamic industry equilibrium model for Bitcoin mining by incorporating miners' endogenous exit and entry with technology innovation. A key modeling step is to track the average operating costs rather than the exact operating costs, thus overcoming the difficulty of strong path dependency incurred by the interaction among endogenous exit, entry, and technology innovation. The model can capture and predict co-movements of miners' electricity consumption, computing power, and mining revenue. The model also predicts that Bitcoin mining is not long-term sustainable regarding climate damages.

Pre-Hedging

Johannes MUHLE-KARBE

Department of Mathematics, Imperial College London

This paper studies a dealer that pre-hedges anticipated potential trades and analyses how this impacts the client's overall execution outcomes. We show that pre-hedging can benefit both parties: improved risk management enables the dealer to charge reduced spreads that more than offset any adverse impact the pre hedging activity has on the execution price. However, when a dealer pre-hedges too aggressively, this can be detrimental to the client. This result is robust to a setting where competing dealers simultaneously pre-hedge. Any counter-productive pre-hedge activity is mitigated with uncertainty about the timing of the potential trade.

Nonparametric Generative Modeling for Time Series via Schrödinger Bridge

Huyên PHAM

Laboratoire de Probabilités, Statistique et Modèlisation, Université Paris Cité

We propose a novel generative model for time series based on Schr¨odinger bridge (SB) approach. This consists in the entropic interpolation via optimal transport between a reference probability measure on path space and a target measure consistent with the joint data distribution of the time series. The solution is characterized by a stochastic differential equation on finite horizon with a path-dependent drift function, hence respecting the temporal dynamics of the time series distribution. We can estimate the drift function from data samples either by kernel regression methods or with LSTM neural networks, and the simulation of the SB diffusion yields new synthetic data samples of the time series. The performance of our generative model is evaluated through a series of numerical experiments. First, we test with a toy autoregressive model, a GARCH Model, and the example of fractional Brownian motion, and measure the accuracy of our algorithm with marginal and temporal dependencies metrics. Finally, we use our SB generated synthetic samples for the application to deep hedging on real-data sets. Joint work with Mohamed Hamdouche and Pierre Henry-Labord`ere.

Contagion in Decentralized Lending Protocols

Julien PRAT

CREST, Ecole Polytechnique

We study financial contagion in Compound V2, a decentralized lending protocol deployed on the Ethereum blockchain. We explain how to construct the balance sheets of Compound's liquidity pools and use our methodology to characterize the financial network. Our analysis reveals that most users either borrow stablecoins or engage in liquidity mining. We then study the resilience of Compound V2 through a series of stress tests, identifying the pools that are most likely to set off a cascade of defaults

A New Stochastic Fubini Theorem via Stochastic Integrals of Measure-valued Processes

Martin SCHWEIZER

Department of Mathematics, ETH

Classical stochastic Fubini theorems start from a fixed semimartingale S, say, and a family of integrands for S which are parametrized by a parameter z, say, from some parameter space Z. There is a (non-random) measure μ, say, on Z, and the stochastic Fubini theorem then says that integration with respect to μ and stochastic integration with respect to S can be interchanged. In other words, we can either 1) first integrate the parametrized integrands with respect to μ and then stochastically integrate the resulting process with respect to S, or we can 2) first stochastically integrate, for each z, the corresponding integrand with respect to S and then integrate the result with respect to μ – and both double integrals yield the same result. What happens now if we replace the fixed measure μ by a stochastic kernel from the predictable sigma- field to Z? Approach 1) still makes sense, but how about 2)? And can we still get some kind of stochastic Fubini theorem? We show that we can, but we need to define for that a stochastic integral, with respect to S, of suitable measure-valued processes. The origin of this question comes from a (still open) question in mathematical finance. There are also some connections to a class of Volterra-type semimartingales. Based on joint work with Tahir Choulli and with Jiaming Chen.

Synchronization in a Kuramoto Mean Field Game

Mete SONER

Department of Operations Research and Financial Engineering, Princeton University

Originally motivated by systems of chemical and biological oscillators, the classical Kuramoto model has found an amazing range of applications from neuroscience to Josephson junctions in superconductors, and has become a key mathematical model to describe self organization in complex systems. These autonomous oscillators are coupled through a nonlinear interaction term which plays a central role in the long term behavior of the system. While the system is unsynchronized when this term is not sufficiently strong, fascinatingly, they exhibit an abrupt transition to a full synchronization above a critical value of the interaction parameter. We explore this system in the mean field formalism. We treat the system of oscillators as an infinite particle system, but instead of positing the dynamics of the particles, we let the individual particles determine endogenously their behaviors by minimizing a cost functional and eventually, settling in a Nash equilibrium. The mean field game also exhibits a bifurcation from unsynchronization to self-organization. This approach has found interesting applications including circadian rhythms and jet-lag recovery. This is joint work with Rene Carmona and Quentin Cormier of Princeton University.

Mean Field Game of Mutual Holding with Common Noise

Nizar TOUZI

Centre de Mathḿatiques Appliquées, École Polytechnique

We consider the mean field game modeling of optimal cross-holding within a population of investors endowed with some idiosyncratic risk process. In this talk, we consider the extension to the situation where the individual risks are correlated through some common noise. Unlike the uncorrelated case, we recover here the trade off between average returns and risks, which is the main novel difficulty to solve the problem. Our main finding is that the mean field no- arbitrage condition imposes some structure restrictions on the model, which turn out to play a crucial role. Under these conditions, we provide a quasi-explicit solution for the mean field game of cross holding.

The Dark Side of Circuit Breakers

Hao XING

Questrom School of Business, Boston University

Market-wide trading halts, also called circuit breakers, have been widely adopted as part of the stock market architecture, in the hope of stabilizing the market during dramatic price declines. We develop an intertemporal equilibrium model to examine how circuit breakers impact market behavior and welfare. We show that a circuit breaker tends to lower the overall level of the stock price and significantly alters its dynamics. In particular, as the price approaches the circuit breaker, its volatility rises drastically, accelerating the chance of triggering the circuit breaker – the so-called "magnet effect"; in addition, returns exhibit increasing negative skewness and positive drift, while trading activity spikes up. Our empirical analysis finds supportive evidence for the model's predictions. Moreover, we show that a circuit breaker can affect the overall welfare either negatively or positively, depending on the relative significance of investors' trading motives for risk sharing vs. irrational speculation. This is a joint work with Hui Chen, Anton Petukhov, and Jiang Wang.

TBA

Thaleia ZARIPHOPOULOU

Department of Mathematics, The University of Texas at Austin

Continuous-Time Reinforcement Learning

Xunyu ZHOU

Department of Industrial Engineering and Operations Research, Columbia University

I will give an overview of the recent development on reinforcement learning in continuous time. I'll highlight three main points: 1) there are fundamental theoretical questions in machine learning in general, and in reinforcement learning in particular, that are begged for answers, more than just applying neural networks and feeding into known algorithms; 2) answering these questions demands fundamentally and conceptually different thinkings than in the classical model-based paradigm; and yet 3) the mathematical techniques that are employed to answer the questions are still largely based on what we (namely researchers in stochastic analysis, stochastic control and mathematical finance) are good at, namely, stochastic analysis, stochastic control and differential equations.

Research Talks

Risk-Sensitive Model Selection in Finance: A Simulation Perspective

Dohyun AHN

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

We investigate the problem of risk-sensitive model selection, which aims to identify the "least risky" model among a fixed number of stochastic models. This problem commonly arises in financial practices, such as selecting hedging or investment strategies, and is typically addressed using simulation tools due to the complexity of characterizing each model's risk in closed form. However, due to the rarity of extreme outcomes and lack of information on their distributions, conventional simulation approaches to this problem exhibit significant performance degradation. To circumvent this issue, we leverage extreme value theory to learn the tail behavior of the outcomes and use it to compare and rank models based on large deviations theory. These results enable us to develop a sequential simulation budget allocation rule that asymptotically maximizes the likelihood of identifying the least risky model as the sample size grows large. Our proposed approach demonstrates substantial improvement over the conventional methods in various numerical experiments.

A Framework for Machine Learning Methods to Evaluate the Information Value of News in Predicting Defaults of Corporate Bond Issuers

Hui BU

School of Economics and Management, Beihang University

This study investigates the information value and predictive power of news headlines for predicting first-time defaults of corporate bond issuers in the Chinese market. It explores the integration of domain knowledge-derived topics with cause-effect logic into a newly proposed credit risk dictionary and constructs news variables to enhance early warning capabilities. A novel machine learning framework is proposed to evaluate the incremental information, predictability, and early warning abilities of these variables. The findings demonstrate that news headlines contain firm- specific incremental information beyond traditional financial ratios and economic variables. Moreover, news variables incorporating topic information exhibit predictive power, although the effectiveness depends on the choice of the forecasting model. Machine learning models that incorporate news variables outperform those relying solely on financial ratios, with decision tree (DT) and support vector machine (SVM) models showing superior forecasting performance. Additionally, news variables demonstrate an early warning ability with forecasting horizons of 3 months or longer, providing valuable insights for credit risk management practices. The study emphasizes the importance of incorporating cause-effect domain knowledge-derived topics, which offer more informative insights than individual words in text analysis. These findings enhance our understanding of news-driven credit risk evaluation and hold implications for financial practitioners and policymakers.

Data Collection and Machine Learning with Privacy Preservation

Ning CAI

Thrust of Financial Technology, Hong Kong University of Science and Technology (Guangzhou)

Data contain valuable information that could play important roles in decision making, but many data are sensitive in areas such as finance, economics, political science, and life science. How to collect and analyze data while still protecting individual privacy has attracted increasing attention from researchers, practitioners, and policymakers. We propose a privacy-preserving data collection algorithm that allows a central administration to recover the whole data set exactly while still preserving individual privacy. Moreover, we demonstrate the applications of our algorithm to machine learning. This is joint work with Siyi Wang.

Irreversible Consumption Habit under Ambiguity: Singular Control and Optimal G-stopping Time

Kexin CHEN

Department of Applied Mathematics, The Hong Kong Polytechnic University

Consider a robust utility maximization with irreversible consumption habit, where an agent concerned about model ambiguity does not tolerate any decline in his/her consumption and at the same time faces a disutility (i.e., adjustment cost) from consumption increase. While the optimization is a robust analogue of singular control problems over a class of consumption- investment strategies and a set of probability measures, it is a new formulation with the non- dominated probability measures on the diffusion process of underlying assets, in addition to the singular controls with adjustment cost. This paper provides a novel connection between the singular controls in the optimization and optimal G-stopping times on a G-expectation space, using a duality theory. This connection enables us to derive the robust consumption strategy as a running maximum of stochastic boundary which is characterized by a free boundary arising from the optimal G-stopping times. The duality, relying on the arguments of reflected G-BSDEs, is attained through verifying the first-order optimality conditions for the singular control, the budget constraint equation for the robust strategies, and the worst-case realization under the non-dominated measures.

A Two Time-Scale Evolutionary Game Theoretic Approach to Multi-Agent Reinforcement Learning

Nan CHEN

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

We propose a two-time scale evolutionary game theoretic approach to solving multiagent reinforcement learning problems. Different from the existent literature that requires to solve Nash equilibrium strategies, exactly or approximately, in each period of learning, the new approach incorporates three innovative designs. First, inspired by the literature of mean-field games, we represent each agent by a population of pure-action-playing individuals and use the empirical distribution of actions in each population to approximate the agent's mixed policy. Second, we propose a simple pairwise comparison protocol to update policies. This enables us to avoid the computationally expensive step of finding the exact equilibrium at each state. Third, we operate the evolution of population state and the updating of Q-values in two different time scales to ensure the convergence. The new approach provably converges to epsilon-Nash equilibria of multiagent reinforcement learning problems without imposing the global optima or saddle point conditions, two restrictive assumptions that are typically needed in the literature. The numerical experiments show that our algorithm is computationally efficient for such challenging setups as systems with multiple equilibria and systems of more than two agents. This talk is based on a joint work with Chengli Ren.

A Model-free Approach to Continuous-time Finance

Henry CHIU

Department of Mathematics, Imperial College London

We present a path-wise approach to continuous-time finance based on causal functional calculus. Our framework does not rely on any probabilistic concept. We introduce a definition of continuous-time self-financing portfolios, which does not rely on any integration concept and show that the value of a self-financing portfolio belongs to a class of non-anticipative functionals, which are path-wise analogs of martingales. We show that if the set of market scenarios is generic in the sense of being stable under certain operations, such self-financing strategies do not give rise to arbitrage. We then consider the problem of hedging a path-dependent payoff across a generic set of scenarios. Applying the transition principle of Rufus Isaacs in differential games, we obtain a path-wise dynamic programming principle for the super-hedging cost. We show that the super- hedging cost is characterized as the solution of a path-dependent equation. For the Asian option, we obtain an explicit solution.

Discrete-Time Mean Variance Strategy based on Reinforcement Learning

Xiangyu CUI

School of Statistics and Management, Shanghai University of Finance and Economics

We have considered a discrete-time mean-variance model based on reinforcement learning. Com- pared to its continuous-time counterpart of Wang and Zhou (2020), the discrete-time model makes more general assumptions about the asset distribution. By using entropy to measure the cost of exploration, we derive the optimal investment strategy, which is also Gaussian type. Additionally, we design a corresponding reinforcement learning algorithm. Simulation experiments and empirical analysis indicate that the discrete-time model exhibits better applicability when analyzing real-world data compared to the continuous-time model.

On Optimal Time-consistent Equilibrium Stopping under Aggregation of Diverse Discount Rates

Shuoqing DENG

Department of Mathematics, Hong Kong University of Science and Technology

We revisit the collective decision making for a group under diverse discount rates. In the context of optimal stopping, we propose a smooth aggregation preference to incorporate all heterogeneous discount rates and an attitude function reflecting the aggregation weight similar to the smooth ambiguity preference in Klibanoff et al. (2005) when handling the model uncertainty. The optimal stopping problem renders to be time inconsistent, for which we develop an iterative approach using consistent planning and characterize all time-consistent equilibrium stopping polices as fixed points of an operator in the setting of one-dimensional diffusion processes. More importantly, we provide some sufficient conditions on both the underlying models and the attitude function such that the smallest equilibrium attains the optimal equilibrium and the attitude function is equivalent to the linear aggregation rule as of diversity neutral. Joint work with Xiang Yu and Jiacheng Zhang.

Randomized Optimal Stopping Problem in Continuous time and Reinforcement Learning Algorithm

Yuchao DONG

School of Mathematical Sciences, Tongji University

In this paper, we study the optimal stopping problem in the so-called exploratory framework, in which the agent takes actions randomly conditioning on current state and an entropy-regularized term is added to the reward functional. Such a transformation reduces the optimal stopping problem to a standard optimal control problem. For the American put option model, we derive the related HJB equation and prove its solvability. Furthermore, we give a convergence rate of policy iteration and compare our solution to the classical American put option problem. Our results indicate a balance between the convergence rate and bias in the choice of the temperature constant. Based on the theoretical analysis, a reinforcement learning algorithm is designed and numerical results are demonstrated for several models.

Asset Pricing with α-MEU Model

Jiacheng FAN

Department of Applied Mathematics, The Hong Kong Polytechnic University

In this paper, we study an asset pricing problem in continuous time where the representative agent can trade the stock, risk free asset and human capital continuously in time with an objective of maximizing her preference value of consumption. We assume the agents' preferences are represented by the α-maximin expected utility model. It is, however, known in the literature that when applied to dynamic decision problem, α-maximin model lead to time inconsistency. To address this issue, we consider intra-personal equilibrium for the agent, which is a rational strategy of the agent choosing action today by anticipating her actions in the future. We then define the market equilibrium to be the set of asset prices under which the intra-personal equilibrium taken by the representative agent clears the market. We prove the existence and uniqueness of the market equilibrium and characterizes it by the solution to a second-order ordinary differential equation. Finally, we compute comparative statics and study the effect of ambiguity attitude on asset prices.

Optimization Frameworks and Sensitivity Analysis of Stackelberg Mean-Field Games

Anran HU

Mathematical Institute, The University of Oxford

In this talk we present a class of discrete-time finite-time-horizon Stackelberg mean-field games, with one leader and an infinite number of identical and indistinguishable followers. In this game, the objective of the leader is to maximize her reward considering the worst-case cost over all possible E-Nash equilibria among followers. A new analytical paradigm is established by showing the equivalence between this Stackelberg mean-field game and a minimax optimization problem. This optimization framework facilitates studying both analytically and numerically the set of Nash equilibria for the game, and leads to the sensitivity and the robustness analysis of the game value. In particular, when there is model uncertainty, the game value for the leader suffers non-vanishing sub-optimality as the perturbed model converges to the true model. In order to obtain a near-optimal solution, the leader needs to be more pessimistic with anticipation of model errors and adopts a relaxed version of the original Stackelberg game.

An Extended Merton Problem with Relaxed Benchmark Tracking

Yijie HUANG

School of Mathematical Sciences, University of Science and Technology of China, and Department of Applied Mathematics, The Hong Kong Polytechnic University

This paper studies a Merton's optimal portfolio and consumption problem in an extended formulation incorporating the tracking of a benchmark process described by a geometric Brownian motion. We consider a relaxed tracking formulation such that that the wealth process compensated by a fictitious capital injection outperforms the external 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 with reference to the benchmark. By introducing an auxiliary state process with reflection, we formulate and tackle an equivalent stochastic control problem by means of the dual transform and probabilistic representation, where the dual PDE can be solved explicitly. On the strength of the closed-form results, we can derive and verify the feedback optimal control in the semi- analytical form for the primal control problem, allowing us to observe and discuss some new and interesting financial implications on portfolio and consumption decision making induced by the additional risk-taking in capital injection and the goal of tracking. Joint work with Lijun Bo and Xiang Yu.

Learning Equilibrium Mean-Variance Strategy

Yanwei JIA

Department of Industrial Engineering and Operations Research, Columbia University and Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong

We study a dynamic mean-variance portfolio optimization problem under the reinforcement learning framework, where an entropy regularizer is introduced to induce exploration. Due to the time-inconsistency involved in a mean-variance criterion, we aim to learn an equilibrium policy. Under an incomplete market setting, we obtain a semi-analytical, exploratory, equilibrium mean- variance policy that turns out to follow a Gaussian distribution. We then focus on a Gaussian mean return model and propose a reinforcement learning algorithm to find the equilibrium policy. Thanks to a thoroughly designed policy iteration procedure in our algorithm, we prove the convergence of our algorithm under mild conditions, despite that dynamic programming principle and the usual policy improvement theorem failing to hold for an equilibrium policy. Numerical experiments are given to demonstrate our algorithm. The design and implementation of our reinforcement learning algorithm apply to a general market setup.

Nonlinear Dependence and Households' Portfolio Decisions over the Life Cycle

Wei JIANG

Department of Industrial Engineering and Decision Analytics,
Hong Kong University of Science and Technology

This talk uncovers the nonlinear relationship between earning risk and stock returns, as measured by the between-squares correlation. By incorporating this between-squares correlation into a life-cycle model, we demonstrate that it lowers households' participation rate and generates moderate risky asset holdings. We identify two pathways through which the between-squares correlation affects portfolio choices: the skewness and kurtosis channels. The extent to which these channels dominate each other depends on the level of between-squares correlation, leading to a nonlinear relationship between this variable and household decisions. Our empirical studies support the model's predictions. Moreover, we find that ignoring between-squares correlations leads to substantial welfare loss and contributes to increasing wealth inequality.

Strategic Investment under Uncertainty with First-and Second-Mover Advantages

Zhaoli JIANG

Department of Applied Mathematics, The Hong Kong Polytechnic University

Standard strategic real-option models predict that firms exercise their growth options too soon because they fear losing the advantage of being the first mover (Fudenberg and Tirole, 1985; Dixit and Pindyck, 1994; Grenadier, 1996). In contrast, we show that firms often enter too late when the second mover has a cost-saving advantage over the first mover. In the unique sym- metric equilibrium, firms enter probabilistically when market demand is sufficiently high (mixed entry strategies), causing inefficient entry delays similar to the war of attrition. Moreover, the interaction between first-mover and second-mover advantages fundamentally changes firm equi- librium strategies and real-option values, producing highly nonlinear and non-monotonic entry predictions. Technically, we show that the symmetric equilibrium solution features five regions. In addition to the option-value-of-waiting and rushing-to-enter regions in standard models with the first-mover advantage only, three new regions arise: a waiting-to-be-Follower region and two disconnected probabilistic-entry regions.

Optimal Stopping without Time Consistency

Hanqing JIN

Mathematical Institute, University of Oxford

We study a continuous time dynamic optimal stopping problem with a flow of preferences, which can be in non-expectation form and can depend on both the current time and state of the system in general. We will define a solution to the problem by the rationality of the agent, and compare it with other solutions appeared in literature.

Analytic Solvability and Efficient Simulation Schemes for Pricing under General Stochastic Volatility Models

Yue-Kuen KWOK

Thrust of Financial Technology, Hong Kong University of Science and Technology (Guangzhou)

A thorough investigation on analytic solvability of general stochastic volatility models with jumps is presented. By an effective measure change technique, we develop an efficient analytic approach of deriving the conditional moment generating functions that are essential for exact simulation schemes. Further development of elegant analytic formulas provides enhancement of computational efficiency to the extent that pricing of discrete exotic variance derivatives only requires a few dozen CPU seconds with numerical accuracy up to 3 to 4 significant figures. The bottleneck of the time-consuming Fourier inversion algorithm in traditional exact simulation schemes is resolved, thus making our new exact simulation schemes to be the most effective numerical approach in pricing discrete exotic derivatives when compared with other numerical algorithms in the literature.

Risk-Sharing Pricing of Variable Annuities within a Principal-Agent Framework

Bin LI

Department of Statistics and Actuarial Science, University of Waterloo

We propose a new risk-sharing pricing approach of variable annuities within a principal-agent framework where an insurer (principal) is the contract provider and a policyholder (agent) is the follower having the surrender option. While the risk neutral pricing approach adopted in the existing literature leads to significantly higher fees and more frequent surrendering than market observations, this new risk-sharing pricing approach reconciles the misalignment between theoretical results and market observations.

Deep Learning for Enhanced Index Tracking

Lingfei LI

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

We develop a novel deep learning method for the enhanced index tracking problem, which aims to outperform an index while effectively controlling the tracking error. We generate a dynamic trading policy from a neural network that accepts a set of features as inputs. We design four blocks in the neural network architecture to handle different types of features, including regimes of the index and stocks, their short-term characteristics, and the current allocation. Outputs from the blocks are integrated into the final output that changes the portfolio allocation. We test our model on the S&P 500 index in an empirical study based on real market data. Out- of-sample results reveal the importance of different features and demonstrate the ability of our method in obtaining impressive excess returns while effectively controlling the tracking error, downside risk, and transaction costs.

Optimal Defined-contribution Pension Management with Financial and Mortality Risks

Wenyuan LI

Department of Statistics and Actuarial Science, The University of Hong Kong

This talk studies optimal defined-contribution (DC) pension management under stochastic interest rates and expected inflation. Besides financial risk, we consider the risk of pre-retirement death and introduce life insurance to the pension account as an option to manage this risk. We formulate this pension management problem as a random horizon utility maximization problem and derive its explicit solution under the constant relative risk aversion (CRRA) utility. We calibrate our model to the U.S. data and demonstrate that the pension member's demand for life insurance has a hump-shaped pattern with age and a U-shaped pattern with the real interest rate and expected inflation. The optimal pension account balance in the model resembles a variable annuity with death benefits that are endogenously determined and depend on various factors such as age, mortality, account balance, future contributions, preferences, and market conditions. The study suggests that offering variable annuities with more flexible death benefits in the DC account could better cater to the bequest demands of its members.

Optimal Consumption with Loss Aversion and Reference to Past Spending Maximum

Xun LI

Department of Applied Mathematics, The Hong Kong Polytechnic University

This paper studies an optimal consumption problem for a loss-averse agent with reference to past consumption maximum. To account for loss aversion on relative consumption, an S-shaped utility is adopted that measures the difference between the non-negative consumption rate and a fraction of the historical spending peak. We consider the concave envelope of the utility with respect to consumption, allowing us to focus on an auxiliary HJB variational inequality on the strength of concavification principle and dynamic programming arguments. By applying the dual transform and smooth-fit conditions, the auxiliary HJB variational inequality is solved in closed-form piecewisely and some thresholds of the wealth variable are obtained. The optimal consumption and investment control can be derived in the piecewise feedback form. The rigorous verification proofs on optimality and concavification principle are provided. Some numerical sensitivity analysis and financial implications are also presented. Joint work with Xiang Yu and Qinyi Zhang.

On Crypto Traders' Preferences toward Jumps – A Measure Change Approach

Francis LIU

Department of Business and Economics, Humboldt-Universität zu Berlin

Preferences of crypto traders occasionally, if not always, surprise us. When Bitcoin price was at its peak, the deep ITM Bitcoin call options were traded at 150% IV, i.e. traders were willing to pay significant premiums based on their belief in an imminent and substantial upward price movement. On the other hand, crypto prices, so as crypto traders' confidence, plummet sharply upon negative news arrivals, e.g. FTX scandal. A joint influence of traders' sentiments, preferences, and news induces clusters of positive and negative jumps in price. We develop a pricing model based on the Hawkes process, designated to capture the cluster of jumps. The model deliberately separates the effects of positive and negative jumps by two Hawkes processes, allowing us to examine the self- and cross-excitement of positive and negative jumps in the price process. This model setup also allows us to unveil traders' preferences toward positive and negative jumps by comparing the historical and risk-neutral measures. We study crypto traders' preferences unveiled by the model.

Price Forecasting of CSI 300 Stock Index Futures Using PCA-Autoformer

Lu LIU

International Institute of Big Data in Finance, Beijing Normal University

In the financial market, Transformer series models always exhibit powerful predictions due to their ability to capture the non-linear characteristics of the markets. In this talk, we propose a PCA-Autoformer model to predict the price of CSI 300 stock index futures by integrating Principal Components Analysis (PCA) with Autoformer. We select 5 principal components using PCA among 19 indicators for the Autoformer model. Empirical tests show (1) PCA processing could increase the prediction accuracy for Transformer, Informer, and Autoformer, and PCA-Autoformer outperforms others; (2) Using radical trading strategy with a lower trading threshold, PCA-Autoformer works well with an impressive annualized return of 24.58% and an annual Sharpe ratio of 1.3761.

Bang–bang Control for a Class of Optimal Stochastic Control Problems with Symmetric Cost Functional

Shuhui LIU

School of Mathematics, Shandong University

This talk applies the method of backward stochastic differential equations (BSDEs) to study the bang–bang optimal stochastic control problem, where the optimal control is of the feedback form and the final cost functional is given by a symmetric function. In addition to obtaining the existence of the optimal control, we also give the explicit representation of the optimal control and the optimal value function of the stochastic control problem by the explicit solution of nonlinear BSDEs with symmetric terminal condition.

A Framework for Measures of Risk under Uncertainty

Yang LIU

School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen

A risk analyst assesses potential financial losses based on multiple sources of information. Often, the assessment does not only depend on the specification of the loss random variable but also various economic scenarios. Motivated by this observation, we design a unified axiomatic framework for risk evaluation principles which quantifies jointly a loss random variable and a set of plausible probabilities. We call such an evaluation principle a generalized risk measure. We present a series of relevant theoretical results. The worst-case, coherent, and robust generalized risk measures are characterized via different sets of intuitive axioms. We establish the equivalence between a few natural forms of law invariance in our framework, and the technical subtlety therein reveals a sharp contrast between our framework and the traditional one. More- over, coherence and strong law invariance are derived from a combination of other conditions, which provides additional support for coherent risk measures such as Expected Shortfall over Value- at-Risk, a relevant issue for risk management practice.

Spectral Risk for Digital Assets

Meng-Jou LUA

Department of Finance, Asia University

Digital assets (DAs) are a unique asset class that presents investors with opportunities and risks that are contingent upon their particular characteristics such as volatility, type, and profile, among other factors. Among DAs, cryptocurrencies (CCs) have emerged as the most liquid asset class, holding this distinction for almost a decade. However, while CCs offer a high level of liquidity, investors must be aware of the potential risks and rewards associated with investing in this asset class, and should conduct a thorough evaluation before making any investment decisions. Our study examines the risk profile of CCs through portfolio analysis, utilizing Spectral Risk Measures (SRMs) as the commonly applied method. In this study, we investigate the application of SRMs in assessing the risk structure of CC portfolios, and their alignment with investors' risk preferences. We employ SRMs to evaluate the CC index CRIX and portfolios constructed from the most liquid 10 CCs from the Blockchain Research Center (BRC), optimizing different SRMs. Our empirical findings suggest that various optimal portfolio allocations can be formulated to meet the unique risk appetites of individual investors. All Quantlets (macros, code snippets) are available via quantlet.com and instructive educational element are available on quantinar.com.

Reinforcement Learning for Financial Index Tracking

Xianhua PENG

HSBC Business School, Peking University

We propose a novel reinforcement learning based approach for financial index tracking. The approach better captures the joint market dynamics of a large number of stocks through learning from a much longer time period of data than existing approaches, which only utilize a short time period of data right before the portfolio rebalancing time. We also propose to solve the portfolio rebalancing equation using a Banach fixed point iteration, which allows to accurately take into account the transaction cost specified as nonlinear functions of transaction volume used in practice. To effectively train the reinforcement learning agent, we develop a training strategy that addresses data limitation issues. Empirical results demonstrate that the proposed approach outperforms existing methods in terms of tracking accuracy and incurs relatively low transaction costs. This is a joint work with Chenyin Gong and Xuedong He.

Non-Concave Utility Maximization with Transaction Costs

Shuaijie QIAN

Department of Mathematics, Hong Kong University of Science and Technology

This talk studies a finite-horizon portfolio selection problem with non-concave terminal utility and proportional transaction costs. The commonly used concavification principle for terminal value is no longer valid here, and we establish a proper theoretical characterization of this problem. We first give the asymptotic terminal behavior of the value function, which implies any transaction close to maturity only provides a marginal contribution to the utility. After that, the theoretical foundation is established in terms of a novel definition of the viscosity solution incorporating our asymptotic terminal condition. Via numerical analyses, we find that the introduction of transaction costs into non-concave utility maximization problems can prevent the portfolio from unbounded leverage and make a large short position in stock optimal despite a positive risk premium and symmetric transaction costs.

Dynamic Trading with Realization Utility

Cong QIN

Center for Financial Engineering, Soochow University

An investor receives utility bursts from realizing gains and losses at the individual-stock level (Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013) and dynamically allocates his mental budget between risky and risk-free assets at the trading-account level. Using savings, he reduces his stockholdings and is more willing to realize losses. Using leverage, he increases his stockholdings beyond his mental budget and is more reluctant to realize losses. While lever- age strengthens the disposition effect, introducing leverage constraints mitigates it. Our model predicts that investors with stocks in deep losses sell them either immediately or after stocks rebound a little.

Subjective Valuation of Defined Contribution Plans

Ling QIN

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

In the United States, many employees can save for retirement by participating in Defined Contribution Plans (DCPs). We propose a tractable framework for calculating the employee's subjective value of the participation option. The framework allows us to separate multiple benefits of participation, including the gains from deferring ordinary income taxes, from receiving employer's matching, and from the capital gains tax and diversification benefit of investing with the pension accounts. We find that suboptimal asset location will significantly reduce the option value, but suboptimal capital gains tax-timing will not. Moreover, we show that participation can lead to a higher life-time consumption level.

Optimal Ratcheting Dividends Policy with Resets

Simone SCOTTI

Department of Economics and Management, University of Pisa

We study an optimal dividend problem under ratcheting and bankruptcy constraints. Firms face a tradeoff between the negative signal of a dividend reduction and a potential future bankruptcy. In contrast to previous works, the dividend rate is allowed to decrease after incurring a reputation cost in accordance with a large empirical literature on the topic. We provide an explicit sufficient and necessary condition

for the finiteness of the value function and prove the differentiability of the value function, and its characterisation via viscosity solutions and the so-called super-contact property. The value function is given by the solution of a two-dimensional PDE combining optimal stopping, singular and impulse controls. We show that the problem is equivalent to a one-dimensional setup simplifying both mathematical and financial analysis. Finally, we detail the numerical study of the model and we provide a sensitivity analysis with respect to the model parameters. This is a joint work with Alexandre Roch

Deep Learning Algorithms for Hedging with Frictions

Xiaofei SHI

Department of Statistical Sciences, University of Toronto

This work studies the deep learning-based numerical algorithms for optimal hedging problems in markets with general convex transaction costs. Our main focus is on how these algorithms scale with the length of the trading time horizon. Based on the comparison results of the FBSDE solver by Han, Jentzen, and E (2018) and the Deep Hedging algorithm by Buehler, Gonon, Teichmann, and Wood (2019), we propose a Stable-Transfer Hedging (ST-Hedging) algorithm, to aggregate the convenience of the leading-order approximation formulas and the accuracy of the deep learning-based algorithms. Our ST-Hedging algorithm achieves the same state-of-the- art performance in short and moderately long time horizon as FBSDE solver and Deep Hedging, and generalize well to long time horizon when previous algorithms become suboptimal. With the transfer learning technique, ST-Hedging drastically reduce the training time, and shows great scalability to high-dimensional settings. This opens up new possibilities in model-based deep learning algorithms in economics, finance, and operational research, which takes advantage of the domain expert knowledge and the accuracy of the learning-based methods.

Risk-free Rate Caplets Pricing

Yingda SONG

Antai College of Economics and Management, Shanghai Jiao Tong University

The benchmark rate reform has significant impact on fixed income derivatives markets, and poses challenges in pricing risk-free rate (RFR) instruments. We propose a unified framework for pricing RFR caplets, including forward-looking, barrier, and backward-looking caplets, un- der general Markov short rate models. The framework offers a comprehensive and efficient approach to accurately price RFR caplets across various short rate models. Numerical experiments demonstrate the efficiency and accuracy of our pricing method under popular short rate models including the Vasicek model, CIR model, and jump-extended CEV model. We also discuss the potential extension of our method to other RFR derivatives.

Designing Structured Products on Traded Accounts

Zhihao SONG

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

Structured products on traded accounts are financial products whose payoff depends on a trading account managed by the product buyer. In this talk, we study how to design these products from the perspectives of both buyers and sellers, and our proposed model features both the transaction costs in the trading account and the buyer's risk aversion. The buyer's problem is formulated as a singular stochastic control problem with a non-concave objective, and the seller's contract design strategy is characterized as a Stackelberg game. Via numerical analyses, we find that the price that the risk-averse buyer is willing to pay is much smaller than the risk- neutral price reported in the literature, and the buyer's optimal trading strategy is significantly more complicated. For the seller's perspective, it is optimal to offer a contract to less risk-averse buyers at a time with lower market volatility.

A C^1-Itô's formula for flows of semimartingale distributions

Xiaolu TAN

Department of Mathematics, The Chinese University of Hong Kong

We provide an Itô's formula for C^1-functionals on the flows of conditional marginal distributions of a continuous semimartingale. This is based on the notion of the weak Dirichlet process, and extends the

C^1-Itô's formula in Gozzi and Russo (2006) to functionals which depend also on marginal distributions of semimartingales. As a first application, we study a class of McKean- Vlasov optimal control problems, and establish a verification theorem by requiring only C^1- regularity of its value function, which is equivalently the (viscosity) solution of the associated HJB master equation. Joint work with Bruno Bouchard and Jixin Wang.

Reinforcement Learning for Continuous-Time Optimal Execution: Actor-Critic Algorithm and Error Analysis

Boyu WANG

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

We propose an actor-critic reinforcement learning (RL) algorithm for the optimal execution problem. We consider the celebrated Almgren-Chriss model in continuous time and formulate a relaxed stochastic control problem for execution under an entropy regularized mean-quadratic variation objective. We obtain in closed form the optimal value function and the optimal feedback policy, which is Gaussian. We then utilize these analytical results to parametrize our value function and control policy for RL. While standard actor-critic RL algorithms perform policy evaluation update and policy gradient update alternately, we introduce a recalibration step in addition to these two updates, which turns out to be critical for convergence. We develop a finite-time error analysis of our algorithm and show that it converges linearly under suitable conditions on the learning rates. We test our algorithm in three different types of market simulators built on the Almgren-Chriss model, historical data of order flow, and a stochastic model of limit order books. Empirical results demonstrate the advantages of our algorithm over the classical control method and a deep learning based RL algorithm.

Financial Risk Meter for Cryptocurrencies

Ruting WANG

Business School, Sun Yat-sen University

The FRM Financial Risk Meter is a technique based on quantile LASSO regression, which is able to identify individual risk characteristics in a network topology. We explain why the FRM index, being an average of penalisation parameters, proves to be a good estimator for variation in the market risk, and even displays predictive ability for it. To exemplify, we focus on the case of cryptocurrency market, which to this day lacks well-developed risk measures. As a particularly highly connected asset class, cryptocurrencies call for a measure that accounts for links and mutual dependencies, especially with respect to tail events; and the conditional quantile based FRM goes a long way towards filling the niche.

Machine Learning in Forecasting Implied Volatility

Yinuo WANG

Department of Financial and Actuarial Mathematics, Xi'an Jiaotong-Liverpool University

Implied volatility is a crucial metric in option pricing and has garnered significant attention from scholars and investors alike due to its utility as a gauge of market value. Therefore, accurately forecasting implied volatility is crucial for implementing effective trading strategies, mitigating market risks, and gaining insights into financial market trends. This research utilizes machine learning methods for implied volatility forecasting to better under- stand option theory-based implied volatility. First, our critical innovation lies in proposing a novel framework, LSTM- DNN, wherein the Long Short-Term Memory (LSTM) model can capture the hidden state from historical time series of implied volatility and generate a temporal feature. The feature is then fed into a Deep Neural Network (DNN), along with other characteristics, to accomplish the prediction task. Second, after considering the time-varying nature of volatility, as well as the "smile" and "term structure" aspects of implied volatility, a new feature matrix is established that incorporates moneyness, time-to-maturity, and historical sequences of varying lengths. Further- more, we aim to validate the predictive performance of LSTM-DNN through a bench- mark experiment. Instead of extracting a temporal feature, the DNN and XGBoost models are directly fed with all historical sequences to predict implied volatility. Moreover, we select ETF50 and SP500 as proxies for the Chinese and American financial markets, respectively, to examine the efficacy of machine learning models in both emerging and mature economies. The empirical results show that the new framework has the best performance in the Chinese market. Other general patterns observed in both markets indicate that the new framework and the traditional DNN model exhibit superior performance compared to the XGBoost model, while long-term historical sequences are utilized for more accurate feature prediction. Additionally, the implied volatility prediction efficacy for at-the-money option contracts is optimal and improves with longer time-to-maturity. We ultimately devise a straightforward trading strategy to demonstrate that the predicted outcomes can generate profitable results. Our study

provides a complete spectrum of the effects of machine learning models and historical states on implied volatility prediction.

Predictable Relative Forward Preferences: Multi-agent and Mean-field Games for Portfolio Management

Yuwei WANG

Department of Statistics, University of Warwick, and Department of Information Systems and Management Engineering, Southern University of Science and Technology

We study portfolio management in a competitive environment where the agents appraise their investment performance via predictable relative forward performance processes. The agents are concerned with their wealth relative to the performance of other peer fund managers in the market. In a binomial model, we explicitly construct such forward processes and the associated forward Nash and mean field equilibria under exponential performance criteria and establish the connection between the games for finite and infinite populations.

Holding stocks, trading bonds

Marko WEBER

Department of Mathematics, National University of Singapore

In an economy with random growth, several long-lived agents with heterogeneous risk aversions, time-preferences, and personal income streams make consumption and investment decisions, trading stocks and a consol bond, and borrowing from and lending to each other. We find in closed form equilibrium stock prices, interest rates, consumption, and trading policies. Agents do not trade stocks, although their returns are time-varying and predictable. Agents dynamically trade the consol bond in response to growth shocks, as to hedge their effect on interest rates, dividends, and personal incomes. Static fund separation holds if agents have also access to a linear bond and two additional hedges for dividend and growth shocks. Such additional assets can be dynamically replicated with the stock and the consol bond. No representative agent exists.

Computing Risk Measures of Variable Annuities in Lévy Models

Xiao WEI

School of Insurance, Central University of Finance and Economics

We propose an efficient numerical method for calculating the Value-at-Risk for variable annuities in Lévy models. In the proposed approach, the probability density of the net liabilities is approximated using the theory of frames and Riesz bases. The key element of the numerical method is a new algorithm for calculating the integral of the exponential L´evy process, approximated by a discrete sum whose

expectation coincides with the expected value of the desired integral. Numerical experiments on the application of the developed method for the Black-Scholes and CGMY models clearly demonstrate its high accuracy and speed.

Continuous Time q-learning for McKean-Vlasov Control Problems

Xiaoli WEI

Institute for Advanced Study in Mathematics, Harbin Institute of Technology

This talk studies the q-learning, recently coined as the continuous-time counterpart of Q-learning by Jia and Zhou (2022c), for continuous time Mckean-Vlasov control problems in the setting of entropy-regularized reinforcement learning. In contrast to the single agent's control problem in Jia and Zhou (2022c), the mean-field interaction of agents render the definition of q-function more subtle, for which we reveal that two distinct q-functions naturally arise: (i) the integrated q-function (denoted by q) as the first-order approximation of the integrated Q-function that can be learnt by a weak martingale condition involving test policies; and (ii) the essential q-function (denoted by q_e) that is employed in the policy improvement iterations. We show that two q- functions are related via an integral representation under all test policies. Based on the weak martingale condition of the integrated q-function and our proposed searching method of test policies, some model-free offline and online learning algorithms are devised. In two financial applications, one in LQ control framework and one beyond LQ control framework, we can obtain the exact parameterization of the value function and two q-functions and illustrate our algorithms with simulation experiments. Joint work with Xiang Yu.

Deep Impulse Control: Application to Interest Rate Intervention

Hoi Ying WONG

Department of Statistics, The Chinese University of Hong Kong

We propose a deep learning framework for impulse control problems involving multivariate stochastic processes, which can be controllable or uncontrollable. We then apply it to estimate central bank interventions on the (controllable) interest rate to stabilize the (uncontrollable) inflation rate, where the two rates are correlated and cointegrated. This method is useful for small banks with high exposure to Treasury securities to predict and stress-test their potential losses from central bank interventions. This is a joint work with Bowen Jia.

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

Bo WU

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

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.

Expected Stock Return and Executive Stock Options: A Dynamic Analysis

Jing XU

School of Finance, Renmin University

I study optimal exercise of executive stock options (ESOs) when the underlying stock's expected return is time varying, allowing the executive to gradually exercise the ESOs over time. Compared to constant expected return, time-varying expected return induces an incentive to speculate on future variations of the expected return by holding onto the ESOs longer, and the effect is stronger if the executive is restricted from directly investing in the underlying stock. When the stock's expected return and past return are correlated, a hedging effect also significantly affects the optimal exercising process: with a positive (negative) correlation, the executive desires low (high) exposure to the stock and thus exercises her ESOs early (lately). Consequently, depending on the stock return characteristics, return predictability may either increase or decrease the firm's cost of issuing ESOs. Furthermore, I show that allowing for gradual option exercise is important when the underlying stock return exhibits a momentum effect.

Regret Bounds for Markov Decision Processes with Recursive Optimized Certainty Equivalents

Wenhao XU

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

The optimized certainty equivalent (OCE) is a family of risk measures that cover important examples such as entropic risk, conditional value-at-risk and mean-variance models. In this paper, we propose a new episodic risk-sensitive reinforcement learning formulation based on tabular Markov decision processes with recursive OCEs. We design an efficient learning algorithm for this problem based on value iteration and upper confidence bound. We derive an upper bound on the regret of the proposed algorithm, and also establish a minimax lower bound. Our bounds show that the regret rate achieved by our proposed algorithm has optimal dependence on the number of episodes and the number of actions.

Dynamic Programming Principle for a Controlled FBSDE System and Associated Extended HJB Equation

Yuhong XU

School of Mathematical Sciences, Soochow University

This talk investigates the dynamic programming principle for a general stochastic control problem in which the state processes are described by a forward-backward stochastic differential equation (FBSDE). Using the method of *S*-topology, we show that there exists an optimal control for the value function. Then a dynamic programming principle is established. As a consequence, an extended Hamilton-Jacobi-Bellman (HJB) equation is derived. The existence and uniqueness of both smooth

solution and a new type of viscosity solution are investigated for this extended HJB equation. Compared with the extant researches on stochastic maximum principle, this paper is the first normal work on partial differential equation (PDE) method for a controlled FBSDE system.

Stochastic LQ Control with Regime Switching and Random Coefficients

Zuoquan XU

Department of Applied Mathematics, The Hong Kong Polytechnic University

We introduce our recent progress on stochastic linear-quadratic optimal control problems with regime switching, random coefficients, and cone control constraint. The results are applied to study mean variance and expected utility maximization problems.

Taylor's Law of Fluctuation Scaling for Semivariances and Higher Moments of Heavy-tailed Data

Phillip YAM

Department of Statistics, The Chinese University of Hong Kong

We introduce Taylor's law and its generalization for variance of light-tailed distributions to many sample statistics of heavy-tailed distributions with infinite mean, namely, as sample size increases, the sample upper and lower semivariances, sample higher moments, skewness, and kurtosis of a random sample from the law increase asymptotically in direct proportion to a power of sample mean. These and additional scaling laws characterize the asymptotic behavior of commonly used measures of the risk-adjusted performance of investments, such as Sortino and Sharpe ratios, when returns follow heavy-tailed nonnegative distributions. Such power-law scaling relationships are known in ecology as Taylor's law and in physics as fluctuation scaling. We find the asymptotic distribution and moments of number of observations exceeding the sample mean, and finally propose estimators of tail index based on these scaling laws and the number of observations exceeding the sample mean and compare these estimators with some prior estimators.

Patience Is a Virtue: Optimal Investment in the Presence of Limit Order Book

Chen YANG

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

We study an optimal investment problem of a CARA investor trading in a limited order book (LOB) market. The model synergizes three key features of market microstructure: the bid-ask spread, the market depth, and a finite market resilience. We develop explicit characterization on the investor's optimal trading strategy under a Bachelier model for the dynamic of the fundamental value of asset, and we derive asymptotic optimal trading strategies in closed-form in the extension with return-predicting signals. Our theoretical and numerical results unveil the importance of patience in trading, for an investor who seeks to strike a balance among several competing goals including achieving the optimal risk exposure, incorporating predictive signals, and minimizing trading costs. Contributing to the existing literature, our model also helps to quantify significant impacts of the market resilience on the trading decisions. This is a joint work with Nan Chen, Min Dai, and Qiheng Ding.

Optimal Order Execution Using After-hour Fixed-price Trading

Nian YANG

Department of Finance and Insurance, Nanjing University

Recently, several exchanges have introduced the after-hour fixed-price trading (hereinafter referred to as the fixed-price session) such as Nasdaq's Extended Trading Close launched in 2022. Although the fixed trading price eliminates temporary price impact, it is subject to execution risk, and empirical evidence also suggests that the disclosed imbalance of the fixed-price orders would generate permanent price impact on future asset prices. In this paper, we present a novel order execution model that captures these features. Our model considers a trader first liquidates the asset position in the fixed-price session, and then transfers the remaining position to the regular session. Using the principle of dynamic programming, we derive explicit optimal trading strategies to the liquidation problem. These strategies are robust across different model settings such as partial fills of fixed-price orders and time-varying permanent price impact, and they rationalize the "opportunistic crossing" strategy used in practice. Additionally, the introduction of the fixed-price session can reduce the trader's execution cost.

Stochastic Control Problems with State-reflections Arising from Relaxed Benchmark Tracking

Xiang YU

Department of Applied Mathematics, The Hong Kong Polytechnic University

We study a nonstandard stochastic control problem motivated by the optimal consumption with wealth tracking of a benchmark process. The benchmark process is modelled by a geometric Brownian motion plus the running maximum of a drifted Brownian motion, indicating the growing property of the benchmark in the long run. We consider a relaxed tracking formulation using capital injection such that the wealth compensated by the injected capital dominates the benchmark at all times. The stochastic control problem is to maximize the expected utility on consumption deducted by the cost of the capital injection under the dynamic floor constraint. By introducing two auxiliary state processes with reflections, an equivalent auxiliary control problem is formulated and studied such that the singular control of capital injection and the floor constraint can be hidden. To tackle the HJB equation with two Neumann boundary conditions, we establish the existence of a unique classical solution to the dual PDE in a separation form using some novel probabilistic representations and a new decomposition-homogenization technique. The proof of the verification theorem on the optimal feedback control can be carried out by some technical stochastic flow analysis of the dual reflected processes and estimations of the optimal control.

Trinomial Tree Method for G-expectation and Sampling for G-normal Random Variable after Measurement

Xingye YUE

School of Mathematical Sciences, Soochow University

Given a 1-D G-normal random variable (RV) X, a trinomial tree method is proposed to approximate the G-Expectation $\mathrm{E}[\varphi(X)]$ for any test function φ. The method is stable and convergent. There is no need to worry about the boundary condition. Furthermore, the whole numerical process actually yields the samples to estimate the expectation $\mathrm{E}[\varphi(X)]$. As we know, direct sampling for a G-normal RV X is infeasible due to the uncertainty of its distribution. But for a 'measurement' $\mathrm{E}[\varphi(X)]$, the sampling is feasible. This is just like a measurement on quantum state: Before measurement, the state is uncertain and unknown, after measurement, the state is definite.

Stochastic Optimal Control with Rough Drivers

Huilin ZHANG

Research Center of Mathematics and Interdisciplinary Sciences, Shandong University, and Department of Mathematics, Humboldt University Berlin

In this talk, I would like to introduce the study of optimal control for systems with rough drivers. The stochastic optimal control for such systems is motivated by the pathwise stochastic control, which is applied to describe the case that an investor has his/her own personal information about the market. Besides, such systems also concern those with singular drivers. I will talk about the stochastic maximum principle and the dynamical programming for such systems. A linear quadratic case is also involved if time permits. The talk is based on joint work with U. Horst, M. Grillo, P. Friz and K. Le.

Major-Minor Mean Field Singular Control and Related Skorokhod Problem

Jiacheng ZHANG

Department of Industrial Engineering and Operations Research Department, University of California, Berkeley

In this work, we aim at characterize the singular control of a conditional McKean-Vlasov dynamic in the context of major-minor setting. We study the related Skorokhod problem of the free boundary problem on the space of probability measure and the regularity of the value function via displacement convexity and semi-concavity.

Zero-sum Linear-Quadratic Stochastic Differential Games with Non-Markovian Regime Switching

Panpan ZHANG

School of Control Science and Engineering, Shandong University

We study the two-player zero-sum linear-quadratic stochastic differential games in a regime switching model. The controlled inhomogeneous system coefficients depend on the underlying noises, so it is a non-Markovian regime switching model. Based on a new kind of multidimensional indefinite stochastic Riccati equation and a multidimensional linear backward stochastic differential equation with unbounded coefficients, we provide optimal feedback control-strategy pairs for the two players in a closed-loop form. We also obtain the corresponding optimal feedback control-strategy pairs for homogeneous systems under closed convex cone control constraints. Finally, these results are applied to portfolio selection problems with different short- selling prohibition constraints in a regime switching market with random coefficients.

Statistical Learning with Sublinear Regret of Propagator Models

Yufei ZHANG

Department of Statistics, London School of Economics and Political Science

We consider a class of learning problems in which an agent liquidates a risky asset while creating both transient price impact driven by an unknown convolution propagator and linear temporary price impact with an unknown parameter. We characterize the trader's performance as maximization of a revenue-risk functional, where the trader also exploits available information on a price predicting signal. We present a trading algorithm that alternates be- tween exploration and exploitation phases and achieves sublinear regrets with high probability. For the exploration phase we propose a novel approach for non-parametric estimation of the price impact kernel by observing only the visible price process and derive sharp bounds on the convergence rate, which are characterised by the singularity of the propagator. These kernel estimation methods extend existing methods from the area of Tikhonov regularisation for inverse problems and are of independent interest. The bound on the regret in the exploitation phase is obtained by deriving stability results for the optimizer and value function of the associated class of infinite-dimensional stochastic control problems. As a complementary result we propose a regression-based algorithm to estimate the conditional expectation of non-Markovian signals and derive its convergence rate.

Design Good Liquidity Pools on Automated Market Makers

Yutian ZHOU

Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong

Automated Market Makers (AMMs) are a popular type of decentralized exchanges in which users trade tokens with each other directly and automatically through a liquidity pool and a fixed pricing function. The Liquidity Providers (LPs) contribute to the liquidity pool by supplying tokens to the pool, and in return, they earn transaction fees from users who trade through the pool. We proposed a model of optimal liquidity provision in which the risk-averse LP decides the number of tokens she would like to supply to the pool and trade in an open market and the amount of consumption in multiple periods. We derived the liquidity provider's optimal strategy by dynamic programming and proved the existence and uniqueness of the solution. Also, we studied the impact of the transaction fees and AMM pricing formula on the LP's optimal decision. In the class of constant mean pricing formula, we numerically found the optimal liquidity pool that maximizes the LP's utility.

Hybrid Public Pension under Loss Aversion

Ziqi ZHOU

Lingnan College, Sun Yat-sen University

We propose an optimal design of hybrid pension plans that considers loss aversion utility for a public pension fund, with a reference point that updates with the agent's consumption history. The public pension fund is a combination of pay-as-you-go (PAYG) and fully-funded (FF) schemes, financed by the payroll tax of our representative agent. We formulated the optimal consumption-investment problem before retirement, handled the S-shaped utility. We obtained closed-form value functions for the problem and characterized the optimal consumption and, most importantly, determine the optimal mix of the two pension schemes and examine the marginal effect of model parameters.

Consumption Decision, Portfolio Choice and Healthcare Irreversible Investment

Shihao ZHU

Center for Mathematical Economics (IMW), Bielefeld University

We propose a tractable dynamic framework for the joint determination of optimal consumption, portfolio choice, and healthcare irreversible investment. Our model is based on a Merton's portfolio and consumption problem, where, in addition, the agent can choose the time at which undertaking a costly lump sum health investment decision. Health depreciates with age and directly affects the agent's mortality force, so that investment into healthcare reduces the agent's mortality risk. The resulting optimization problem is formulated as a stochastic control-stopping problem with a random time-horizon and state-variables given by the agent's wealth and health capital. We transform this problem into its dual version, which is now a two-dimensional optimal stopping problem with interconnected dynamics and finite time-horizon. Regularity of the optimal stopping value function is derived and the related free boundary surface is proved to be Lipschitz continuous and it is characterized as the unique solution to a nonlinear integral equation. In the original coordinates, the agent thus invests into healthcare whenever her wealth exceeds an age- and health-dependent transformed version of the optimal stopping boundary.