RECENT DEVELOPMENTS IN MEAN-FIELD GAME, MACHINE LEARNING AND QUANTITATIVE FINANCE

21-24 May 2019

Tuan Chau, Vietnam

Jointly organized by:

Ministry of Education and Training
Vietnam Institute for Advanced Study in Mathematics

THE HONG KONG POLYTECHNIC UNIVERSITY
香港理工大學

Centre for Quantitative Finance
Faculty of Science

Organizing Committee

- **PHAM** Xuan Huyên (University Paris Diderot and JVN Institute, Ho-Chi-Minh City)
- **HA** Le Minh (VIASM, Hanoi)
- **YU** Xiang (Hong Kong Polytechnic University)
- **ZHOU** Chao (National University of Singapore, Singapore)

Sponsors

Ministry of Education and Training
Vietnam Institute for Advanced Study in Mathematics

Invited Speakers

ANKIRCHNER Stefan (Jena University)

CAI Ning (Hong Kong University of Science and Technology, Hong Kong)

CHASSAGNEUX Jean-François (Paris Diderot university, France)

CHEN Ying (NUS, Singapore)

DJEHICHE Boualem (KTH Stockholm, Sweden)

GUO Xin (University of California, Berkeley, USA)

HO Tu Bao (VIASM and JVN Institute, Vietnam)

JAIMUNGAL Sebastian (University of Toronto, Canada)

JIN Yong Jimmy (Hong Kong Polytechnic University, Hong Kong)

LUDKOVSKI Mike (University of California Santa Barbara, USA)

LY VATH Vathana (ENSIIE, France)

NADTOCHIY Sergey (IIT Illinois, USA)

NGUYEN Quang (JVN Institute, HCM, Vietnam)

NUTZ Marcel (Columbia University, USA)

SIRBU Mihai (University of Texas at Austin, USA)

TAN Xiaolu (Paris Dauphine university, France)

XING Hao (Boston University, USA and LSE, UK)

ZHENG Harry (Imperial College London, United Kingdom)

Programme Overview

Tuesday 21 May 2019	Wednesday 22 May 2019	Thursday 23 May 2019	Friday 24 May 2019
09:00 – 09:15 Opening Address		**Excursion to Halong Bay**	
09:15 – 10:00 Tu Bao HO	**09:00 – 09:45** Sebastian JAIMUNGAL		**09:00 – 09:45** Marcel NUTZ
10:00 – 10:45 Ying CHEN	**09:45 – 10:30** Ning CAI		**09:45 – 10:30** Vathana LY VATH
10:45 – 11:15 Tea break	**10:30 – 11:00** Tea Break		**10:30 – 11:00** Tea Break
11:15 – 12:00 Boualem DJEHICHE	**11:00 – 11:45** Stefan ANKIRCHNER		**11:00 – 11:45** Xiaolu TAN
12:00 – 12:45 Sergey NADTOCHIY	**11:45 – 12:30** Jean-François CHASSAGNEUX		**11:45 – 12:30** Harry ZHENG
12:45 – 15:00 Lunch + free time	**12:30 – 15:00** Lunch + free time		**Closing and Lunch**
15:00 – 15:45 Hao XING	**15:00 – 15:45** Mihai SIRBU		
15:45 – 16:30 Jimmy Yong JIN	**15:45 – 16:30** Quang NGUYEN		
16:30 – 17:00 Tea Break	**16:30 – 17:00** Tea Break		
17:00 – 17:45 Xin GUO	**17:00 – 17:45** Mike LUDKOVSKI	**Conference Dinner**	

Daily Schedule

Tuesday ,	21 May 2019
TIME	**ACTIVITY**
09:00 – 09:15	Opening Address
09:15 – 10:00	**Tu Bao HO** VIASM, Hanoi and JVN Institute, Ho-Chi-Minh City, Vietnam *Learning and Recommending Treatments Using Electronic Medical Records*
10:00 –10:45	**Ying CHEN** National University of Singapore, Singapore *Topic Sentiment Asset Pricing with DNN Supervised Learning*
10:45 – 11:15	Tea break
11:15 – 12:00	**Boualem DJEHICHE** KTH Stockhom, Sweden *Mean-field type modeling and control in pedestrian crowd dynamics*
12:00 – 12:45	**Sergey NADTOCHIY** Illinois Institute of Technology, Chicago, USA *Mean Field Systems for solving Free-Boundary problems: from Finance to Neuroscience and Physics*
12:45 – 15:00	Lunch + free time
15:00 – 15:45	**Hao XING** Boston University, USA and LSE, UK *Rational inattention and dynamic discrete choice*
15:45 – 16:30	**Jimmy Yong JIN** Hong Kong Polytechnic University *Robust Two-Funds and Three Funds Rules: A Cross-Validation Approach*
16:30 – 17:00	Tea Break
17:00 – 17:45	**Xin GUO** University of California, Berkeley, USA *Learning Mean Field Games*

Daily Schedule

Wednesday , 22 May 2019	
TIME	**ACTIVITY**
09:00 – 9:45	**Sebastian JAIMUNGAL** University of Toronto, Canada *Deep Q-Learning for Nash Equilibria: Nash-DQN*
09:45 –10:30	**Ning CAI** Hong Kong University of Science and Technology *Regime Classification and Stock Loan Valuation*
10:30 – 11:00	Tea break
11:00 – 11:45	**Stefan ANKIRCHNER** Jena University, Germany Solving fully coupled FBSDEs by minimizing a directly calculable error functional
11:45 – 12:30	**Jean-François CHASSAGNEUX** Université Paris Diderot, France *Weak error expansion for Mean Field SDE*
12:30 – 15:00	Lunch + free time
15:00 – 15:45	**Mihai SIRBU** University of Texas at Austin, USA *Optimal investment and consumption with labor income or liability streams in incomplete markets*
15:45 – 16:30	**Quang NGUYEN** John Von Neumann Institute, Ho-Chi-Minh City, Vietnam *Study of financial market systemic risk using a complex network approach*
16:30 – 17:00	Tea Break
17:00 – 17:45	**Mike LUDKOVSKI** University of California Santa Barbara, USA *A Machine Learning Approach to Adaptive Robust Utility Maximization and Hedging*

Daily Schedule

Friday, 24 May 2019	
TIME	**ACTIVITY**
09:00 – 09:45	**Marcel NUTZ** Columbia University, New York, USA *Fine Properties of the Optimal Skorokhod Embedding Problem*
09:45 –10:30	**Vathana LY VATH** ENSIIE, Evry, France *Growth Opportunities and Dividend Control Problem under Financing Constraints*
10:30 – 11:00	Tea break
11:00 – 11:45	**Xiaolu TAN** University Paris Dauphine, France *Mean-Field games with branching*
11:45 – 12:30	**Harry ZHENG** Imperial College, London, UK *Constrained time-inconsistent utility deviation-risk optimization*
12:30 – 14:00	Lunch

Abstracts

Stefan ANKIRCHNER, Iena University

Solving fully coupled FBSDEs by minimizing a directly calculable error functional

In this talk we present a new scheme for approximating solutions of forward-backward stochastic differential equations (FBSDEs). The scheme is particularly suitable for fully coupled FBSDEs. The approximation relies on a piecewise in time approximation by minimizing an error functional that measures how well a process triplet satisfies the FBSDE. The error functional is minimized in a finite-dimensional linear space based on iterated integrals. We provide sufficient conditions for the approximations to converge at the rate 1/2. The talk is based on joint work with Alexander Fromm.

Ning CAI, Hong Kong University of Science and Technology

Regime Classification and Stock Loan Valuation

For traditional perpetual American put options under regime-switching models, the risk-free interest rates are typically assumed to be positive and optimal stopping usually can occur in any regime. However, this may not hold true when the risk-free interest rates are allowed to be equal to zero (the interest rate may drop to zero in reality, e.g., in certain periods in Japan); there may exist "continuation regimes" within which optimal stopping can never occur, i.e., within which stopping is never optimal. In this paper, we develop a unified, fixed point approach to determining all continuation regimes explicitly for the pricing problem of perpetual American put options under general regime-switching exponential Levy models with any finite numbers of regimes and general Levy types, where the risk-free interest rate in each regime is nonnegative and the discounted stock price with the dividends reinvested is not necessarily a martingale. As an important application of this regime classification result, we provide a unified framework for the valuation of infinite maturity stock loans under general regime-switching exponential Levy models with any finite numbers of regimes and general Levy types, which can be formulated as the pricing problem of perpetual American call options with negative interest rates (see Xia and Zhou 2007 for the pioneering research of stock loan valuation under the geometric Brownian motion model). Applying this unified approach yields analytical solutions to the infinite maturity stock loan prices under both the general exponential Levy models without regime switching and the regime-switching phase-type jump diffusion models with any finite numbers of regimes and with or without dividends, which include the related results of Xia and Zhou (2007), Zhang and Zhou (2009), and Cai and Sun (2014) as special cases.

Jean François CHASSAGNEUX, University Paris Diderot

Weak error expansion for Mean Field SDE

In this work, we study the weak approximation error by particle system of Mean Field SDE. We prove an expansion of this error in terms of the number of particle. Our strategy of proof follows the approach of Talay-Tubaro for weak approximation of SDE by an Euler Scheme. We thus consider a PDE on the Wasserstein space (called the Master Equation in mean-field games literature) and, relying on smoothness properties of the solution, obtain our expansion. We also prove the required smoothness properties under sufficient conditions on the coefficient function.
This is a joint work with L. Szpruch and A. Tse.

Ying CHEN, National University of Singapore

Topic Sentiment Asset Pricing with DNN Supervised Learning

We develop an innovative deep neural network (DNN) supervised learning approach to extracting insightful topic sentiments from analyst reports at the sentence level and incorporating this qualitative knowledge in asset pricing and portfolio construction. The topic sentiment analysis is performed on 113,043 Japanese analyst reports and the topic sentiment asset pricing model delivers superior predictive power on stock returns with adjusted R2 increasing from 1.6% (benchmark model without sentiment) to 14.0% (in-sample) and 13.4% (out-of-sample). We find that topics reflecting the subjective opinions of analysts have greater impact than topics of objective facts and justification of the quantitative measures. This is a joint work with Hitoshi Iwasaki.

Boualem DJEHICHE, KTH Stockholm

Mean-field type modeling and control in pedestrian crowd dynamics.

We will review some recent results on modeling and control of pedestrian crowds. We first consider a model which allows for nonlocal crowd aversion and arbitrarily but finitely many interacting crowds. We then consider a mean-field model for the movement of tagged pedestrians, distinguishable from a surrounding crowd, with a targeted final destination. The tagged pedestrians move through a dynamic crowd, interacting with it while optimizing their path. The model includes distribution-dependent effects like congestion and crowd aversion. Finally, we focus on modeling and controlling pedestrian motion in confined domains, where interaction with solid obstacles like walls and pillars is considered. This amounts to control a particular class sticky SDEs of mean-field type.

Xin GUO, UC Berkeley

Learning MFGs

Motivated by the Ad auction problem for advertisers, we consider the general problem of simultaneous learning and decision-making in a stochastic game setting with a large population. We formulate this type of games with unknown rewards and unknown dynamics as a generalized mean-field-game (GMFG), with incorporation of action distributions. We first analyze the existence of the solution to this GMFG, and show that naively combining Q-learning with the three-step fixed- point approach in classical MFGs yields unstable algorithms. We then propose an alternating approximating Q-learning algorithm and establish its rate of convergece. The numerical performance of this new algorithm on the repeated Ad auction problem shows superior computational efficiency.

Tu Bao HO, VIASM and JVN Institute

Learning and Recommending Treatments Using Electronic Medical Records

Treatment is a vital factor which directly affects patient health status. In recent years, the fast development of electronic medical records has provided valuable resources for developing data mining methods to address different healthcare issues, especially learning and recommending treatments. However, most of the related studies are limited in exploiting various patient information, incorporating medical domain knowledge, capturing different kinds of treatment patterns as well as interpreting the recommendation mechanism. In this paper, we propose novel methods of learning and recommending treatments to overcome the above drawbacks. Our methods adopt a mixed-variate restrict Boltzmann machine to represent different kinds of patient records. To handle the varying-length of treatments, we incorporate medical domain knowledge, which is drug indication, to capture significant changes in prescription indication and split treatment records into periods more flexibly. Our treatment learning approach organizes sequences of prescription drugs in a new representation, namely regimen trees that reveals many more different kinds of treatment patterns. More interestingly, we propose two neighbor-based treatment recommendation methods to suggest new patients' treatments that are not only derived from common treatments of patient groups but also personalized based on treatments of neighbor patients. The experimental evaluation was conducted on three acute disease cohorts extracted from the MIMIC III database. The obtained results show that the proposed methods are able to provide different kinds of treatment patterns and achieve competitive efficacy but better interpretability compared to related works. Joint work with Hung Hoang.

Sébastian JAIMUNGAL, University of Toronto

Deep Q-Learning for Nash Equilibria: Nash-DQN

Model-free learning for multi-agent stochastic games is an active area of research. Existing reinforcement learning algorithms, however, are often restricted to zero-sum games, and are applicable only in small state-action spaces or other simplified settings. Here, we develop a new data efficient Deep-Q-learning methodology for model-free learning of Nash equilibria for general-sum stochastic games. The algorithm uses a local linear-quadratic expansion of the stochastic game, which leads to analytically solvable optimal actions. The expansion is parametrized by deep neural networks to give it sufficient flexibility to learn the environment without the need to experience all state-action pairs. We study symmetry properties of the algorithm stemming from label-invariant stochastic games and as a proof of concept, apply our algorithm to learning optimal trading strategies in competitive electronic markets. https://arxiv.org/abs/1904.10554

Jimmy Yong JIN, Hong Kong Polytechnic University

Robust Two-Funds and Three Funds Rules: A Cross-Validation Approach

Kan and Zhou (2007) theoretically and empirically analyze the estimation risk in Markowitz mean-variance portfolio theory and conclude that the standard plug-in method replacing the parameters by the sample / moment estimators performs poorly when the estimation window is not long enough. Further they propose their two-funds ($\widehat{w}(c) = \frac{c}{\gamma}\hat{\Sigma}^{-1}\hat{\mu}$) and three-funds rules ($\widehat{w} = \frac{1}{\gamma}(d\hat{\Sigma}^{-1}\hat{\mu} + e\,\hat{\Sigma}^{-1}1_N)$), which optimally combine the risk-free assets, the sample tangency portfolio and the sample global minimum-variance portfolio and effectively reduce the estimation risk. However, another estimation risk exists in their forms of two-funds and three-funds rules – the parameters (i.e., c, d and e) involved in the rules actually are functions of the parameters in the portfolio theory assumptions (the true mean vector and variance-covariance matrix) – therefore how to improve the estimation of the parameters in the two-funds and three-funds rules becomes another challenge. In this talk, I propose an alternative approach to estimate the two-funds and three funds rules by applying the trimmed mean approach in repeated cross-validation practice. A comprehensive empirical study with several different weighting schemes across several well diversified portfolio universes are further conducted to compare the out-of-sample performance.

Mike Ludkovski, UCSB

A Machine Learning Approach to Adaptive Robust Utility Maximization and Hedging

We investigate the adaptive robust control framework for portfolio optimization and loss-based hedging under drift and volatility uncertainty. Adaptive robust problems offer many advantages but require handling a nested optimization problem (infimum over market measures, supremum over the control) at each instance. Moreover, the underlying Bellman equations are intrinsically multi-dimensional. As our major contribution, we propose a novel machine learning approach that solves for the local saddle-point at a chosen set of inputs and then uses a nonparametric (Gaussian Process) regression to obtain a functional representation of the value function. Our algorithm resembles control randomization and regression Monte Carlo techniques but also brings multiple innovations, including adaptive experimental design, separate surrogates for optimal control and the local worst-case measure, and computational speed-ups for the sup-inf optimization. Thanks to the new scheme we are able to consider settings that have been previously computationally intractable and provide several new financial insights regarding the interplay between learning and trading under unknown market parameters. We further compare adaptive robust strategies to myopic-adaptive and static robust alternatives. This is joint work with Tao Chen (UCSB).

Vathana LY VATH, ENSIEE

Growth Opportunities and Dividend Control Problem under Financing Constraints

We consider the problem of determining an optimal control on the dividend under debt constraints and investment opportunities in an economy with business cycles. We allow the company to accept or reject investment opportunities arriving at random times with random sizes. Investment may be done through debt financing, therefore changing the firm's indebtedness and impacting its capital structure and risk profile. We formulate this problem as a bi-dimensional singular control problem under regime switching in presence of jumps. We use a viscosity solution approach to get qualitative descriptions of the solution and further show that the bi-dimensional problem could be reduced by considering the precessus representing the leverage ratio. The reduced value function and the optimal policy may be completely characterized.

Joint work with Simone Scotti and Luca Presiozo

Sergey NADTOCHIY, Illinois Institute of Technology

Mean Field Systems for solving Free-Boundary problems: from Finance to Neuroscience and Physics.

This talk is concerned with the recent advances in the study of particle systems with threshold-type interaction. Such systems appear naturally in the mathematical models for neuron cells, supercooled liquids, computer and credit networks, etc. In particular, our recent work shows how to use these particle systems in order to solve the supercooled Stefan PDE. The latter is a free-boundary problem which describes the process of freezing of a supercooled liquid, and whose well-posedness has been an open question for the last 50 years. We prove the well-posedness and provide a complete description of such solutions in one space dimension, by relating the temperature distribution in the liquid to the regularity of the interface between the solid and liquid states. These results provide a methodology for establishing the well-posedness and describing the structure of other mean-filed systems with threshold-type interaction (and of the associated free-boundary problems). This is a joint work with F. Delarue and M. Shkolnikov.

Quang NGUYEN, JVN Institute

Study of financial market systemic risk using a complex network approach

For the last 2 decades, there has been a great interest in the study of complex systems that cover from engineering system such as the internet, the power-grid, to biological or social ones, including financial market. Profiting from such advances, the systemic risk of the financial market, the risk that the whole market fluctuates in an unpredictable way or even dysfunctions, can be evaluated using a complex network approach. In this work, we show that financial networks share various properties of a complex system, as well as other social networks. And we measure the robustness of the network constructed from financial data against a loss of a proportion of its components. This robustness could be used to indirectly deduce the probability of a structural collapse. We use a theoretical approach that is based on percolation theory, and simulation approach in parallel. Practical applications of this study are numerous. For example, one can answer questions such as: when the market can be considered under great systemic risk? Which company is critical when the crisis starts to spread out (that are needed to be protected beforehand)? The same analysis can be applied to other social networks that we mention in our work: an email network, a co-authorship network and a bitcoin trusting network. Inverse question is equally relevant: how can we spread the news through a network; or how do we break an unwanted network (fraudulent network, contaminated network).

Marcel NUTZ, Columbia University

Fine Properties of the Optimal Skorokhod Embedding Problem

We study the problem of stopping a Brownian motion at a given distribution ν while optimizing a reward function that depends on the (possibly randomized) stopping time and the Brownian motion. Our first result establishes that the set $T(\nu)$ of stopping times embedding ν is weakly dense in the set $R(\nu)$ of randomized embeddings. In particular, the optimal Skorokhod embedding problem over $T(\nu)$ has the same value as the relaxed one over $R(\nu)$ when the reward function is semicontinuous, which parallels a fundamental result about Monge maps and Kantorovich couplings in optimal transport. A second part studies the dual optimization in the sense of linear programming. While existence of a dual solution failed in previous formulations, we introduce a relaxation of the dual problem and establish existence of solutions as well as absence of a duality gap, even for irregular reward functions. This leads to a monotonicity principle which complements the key theorem of Beiglbock, Cox and Huesmann. These results can be applied to characterize the geometry of optimal embeddings through a variational condition. (Joint work with Mathias Beiglbock and Florian Stebegg)

Mihai SIRBU, U Texas, Austin

Optimal investment and consumption with labor income or liability streams in incomplete markets

We consider the problem of optimal investment and consumption for an investor endowed with a stream of income/liabilities, where the constraint that wealth be positive is imposed not only at terminal times, but at all times. We consider a general semi-martingale market which is incomplete (even without the constraints). The dual problem is a singular control-problem related to the martingale measure for the unconstrained problem. In order to deal with the general incomplete case we parametrize the labor income stream by a time-dependent multiplier q, in the spirit of Hugonnier-Kramkov. We can show that the primal and dual value functions are conjugate and optimizers exist. Under appropriate conditions we show the the primal value function is differentiable (with respect to q).
Based on joint work with Oleksii Mostovyi.

Xiaolu TAN, Paris Dauphine University

Mean-Field games with branching

Mean field games are concerned with the limit of large-population stochastic differential games where the agents interact through their empirical distribution. In the classical setting, the number of players is large but fixed throughout the game. However, in various applications, such as population dynamics or economic growth, the number of players can vary across time which may lead to different Nash equilibria. For this reason, we introduce a branching mechanism in the population of agents and obtain a new formulation of mean field games. We then study the problem with both PDE and probabilistic approach. Joint work with Julien Claisse and Zhenjie Ren.

Hao XING, Boston University and LSE

Rational inattention and dynamic discrete choice

We adopt the posterior-based approach to study dynamic discrete choice problems with rational inattention. We show that the optimal solution for the Shannon entropy case is characterized by a system of equations that resembles the dynamic logit rule. We propose an efficient algorithm to solve this system and apply our model to explain phenomena such as status quo bias, confirmation bias, and belief polarization. We also study the dynamics of consideration sets. Unlike the choice-based approach, our approach applies to general uniformly posterior-separable information cost functions. A key condition for our approach to work in dynamic models is the convexity of the difference between the discounted (generalized) entropy of the prior beliefs about the future states and the entropy of the current posterior. This is a joint work with Jianjun Miao.

Harry ZHENG, Imperial College London

Constrained time-inconsistent utility deviation-risk optimization

We propose a unified utility deviation-risk model which covers both utility maximization and mean-variance analysis as special cases. We derive the time-consistent HJB equation for the equilibrium value function and recover the known results in the literature and go beyond with several examples, including mean-variance model with stochastic volatility dependent risk aversion, utility deviation-risk model with state dependent risk aversion and control constraint, and constrained portfolio selection model. The numerical tests show that the utility and deviation-risk have significant impact on the equilibrium control strategy and the distribution of the terminal wealth. (Joint work with Jiawen Gu and Shijing Si).
