Joint Risk & Stochastics and Financial Mathematics Seminar in 2021/2

Seminars are listed in reverse chronological order, most recent first

---

Thursday 24 March 2022 - Markus Fischer (University of Padova)

On correlated equilibria and mean field games
Mean field games are limit models for symmetric N-player games, as N tends to infinity, where the prelimit models are solved in terms of Nash equilibria. A generalization of the notion of Nash equilibrium, due to Robert Aumann (1974, 1987), is that of correlated equilibrium. In a simple discrete setting, we will discuss correlated equilibria for mean field games and their connection with the underlying N-player games. We first consider equilibria in restricted Markov open-loop strategies. In this case, N-player correlated equilibria are seen to converge to the mean field game limit and, conversely, correlated mean field game solutions allow one to construct approximate N-player correlated equilibria. We then discuss the problem of constructing approximate equilibria when deviating players have access to the aggregate system state. An explicit example of a correlated mean field game solution not of Nash-type will also be given.

Joint work with Ofelia Bonesini (University of Padua) and Luciano Campi (University of Milan "La Statale")

Thursday 10 March 2022 - Thomas Bernhardt (University of Manchester)

On the Continuity of the Root Barrier
Sharp lower bounds of VIX call options are famously connected to Root's solution of the Skorokhod embedding problem. Finding Root's solution for a given target measure involves finding a specific barrier function. Even though numerical schemes exist to find that barrier, those rely on apriori regularity assumptions like continuity and monotonicity, which have been notoriously challenging to prove. In this talk, we show that the barrier function in Root's solution to the Skorokhod embedding

problem is continuous and finite at every point where the target measure has no atom and its absolutely continuous part is locally bounded away from zero.

Thursday 24 February 2022 - Eyal Neumann (Imperial College London)

Trading with the Crowd
We formulate and solve a multi-player stochastic differential game between financial agents who seek to cost-efficiently liquidate their position in a risky asset in the presence of jointly aggregated transient price impact, along with taking into account a common general price predicting signal.

The unique Nash-equilibrium strategies reveal how each agent's liquidation policy adjusts the predictive trading signal to the aggregated transient price impact induced by all other agents. This unfolds a quantitative relation between trading signals and the order flow in crowded markets. We also formulate and solve the corresponding mean field game in the limit of infinitely many agents. We prove that the equilibrium trading speed and the value function of an agent in the finite N-player game converges to the corresponding trading speed and value function in the mean field game at rate O(N^{-2}). In addition, we prove that the mean field optimal strategy provides an approximate Nash-equilibrium for the finite-player game.

Thursday 10 February 2022 - Camilo Hernández (Imperial College London)

Type-I backward stochastic Volterra integral equations and their applications
We study extended type-I BSVIEs which are extensions of classic Backward Stochastic Differential Equations. The noticeable feature of extended type-I BSVIEs is the appearance of the *diagonal* processes of both elements of the solution in the generator. We will motivate them by a series of practical applications. In particular, they provide a rich framework to address the equilibrium approach to time-inconsistent control problems via either Bellman’s and Pontryagin’s principles and consequently time-inconsistent contract theory. We also present and discuss a type of flow property satisfied by this class of BSVIEs. This is a joint work with Dylan Possamaï.

Thursday 27 January 2022 - Haoyang Cao (Ecole Polytechnique)

Identifiability in inverse reinforcement learning
Inverse reinforcement learning attempts to reconstruct the reward function in a

Markov decision problem, using observations of agent actions. As already observed in Russell [1998] the problem is ill-posed, and the reward function is not identifiable, even under the presence of perfect information about optimal behavior. We provide a resolution to this non-identifiability for problems with entropy regularization. For a given environment, we fully characterize the reward functions leading to a given policy and demonstrate that, given demonstrations of actions for the same reward under two distinct discount factors, or under sufficiently different environments, the unobserved reward can be recovered up to a constant. We also give general necessary and sufficient conditions for reconstruction of time-homogeneous rewards on finite horizons, and for action-independent rewards, generalizing recent results of Kim et al. [2021] and Fu et al. [2018].

Thursday 2 December 2021 - Christian Furrer (University of Copenhagen)

Change of measure techniques for scaled insurance cash flows
Incidental policyholder behaviour, including free policy conversion and stochastic retirement, may have a significant impact on the market value of a life insurance contract; consequently, models should account for this. However, the inclusion of incidental policyholder behaviour typically leads to duration effects and thus an increase in computational complexity. In this talk, I show how change of measure techniques can be used to conveniently deal with this added layer of complication.

Thursday 18 November 2021 - Xiaoli Wei (TBSI)

Cooperative multi-agent reinforcement learning: A mean field perspective
Multi-agent reinforcement learning (MARL) has enjoyed substantial successes in many applications including real-time resource allocation, order matching for ride-sharing, and autonomous driving. Despite the empirical success of MARL, general theories behind MARL algorithms are less developed due to the intractability of interactions, complex information structure, and the curse of dimensionality. Instead of directly analyzing the multi-agent systems, mean-field theory provides a powerful approach to approximate the games under various notions of equilibria. Moreover, the analytical feasible framework of mean-field theory leads to efficient and tractable learning algorithms with theoretical guarantees.

In this talk, we will demonstrate how mean-field theory can contribute to analyzing a class of simultaneous-learning-and-decision-making problems under cooperation, with unknown rewards and dynamics. Moreover, we will show that the learning procedure can be further decentralized and scaled up if a network structure is specified. Our result lays the first theoretical foundation for the so-called "centralized training and decentralized execution" scheme, a widely used training scheme in the empirical works of cooperative MARL problems. This is based on joint work with Xin Guo (Berkeley), Haotian Gu (Berkeley) and Renyuan Xu (University of Southern California).

Tuesday 9 November 2021 - Mike Lipkin (NYU)

Time Scales in Finance - Introducing event-driven finance to graduate students
In classical finance statistical noise is introduced to drive stock dynamics and quantify option pricing. This means the stock sits in a "heat bath". But much of actual trading concerns itself with events, such as take-overs, earnings announcements, etc. Such an event introduces a singularity into the pricing kernel. And mesoscopic quantities, such as volatility, become irrelevant or misleading. Here we look at a spectrum of events and see how they manifest themselves in time and pricing.

Thursday 21 October - Philippe Casgrain (ETH Zürich)

Anytime-valid sequential testing for elicitable functionals via supermartingales
We consider the problem of testing statistical hypotheses and building confidence sequences for elicitable and identifiable functionals, a broad class of statistical quantities which are of particular interest in the field of quantitative risk management. Assuming a framework in which data is collected sequentially, where a user may choose to accept or reject a hypothesis at any point in time, we provide powerful distribution-free and anytime-valid testing methods which rely on controlled supermartingales. Leveraging tools from online convex optimization, we show that tests can be optimized to improve their statistical power, with asymptotic guarantees for rejecting false hypotheses. By "inverting the test", these methods are extended to the task of confidence sequence building. Lastly, we implement these techniques on a range of examples to demonstrate their effectiveness.

Thursday 7 October - Birgit Rudloff (Vienna University of Economics and Business)

Epic Math Battles: Nash vs. Pareto
Nash equilibria and Pareto optimization are two concepts when dealing with multiple

criteria. It is well known that the two concepts do not coincide. However, in this work we show that it is possible to characterize the set of all Nash equilibria for any non-cooperative game as the Pareto optimal solutions of a certain vector optimization problem. To accomplish this task, we enlarge the objective function and formulate a non-convex ordering cone under which Nash equilibria are Pareto efficient. We demonstrate these results, first, for shared constraint games in which a joint constraint is applied to all players in a non-cooperative game. This result is then extended to generalized Nash games, where we deduce two vector optimization problems providing necessary and sufficient conditions, respectively, for generalized Nash equilibria. Finally, we show that all prior results hold for vector- valued games as well. Multiple numerical examples are given and demonstrate the computational advantages of finding the set of Nash equilibria via our proposed vector optimization formulation in these cases.