2ND LEEDS CONFERENCE ON STOCHASTIC CONTROL AND GAMES UNDER AMBIGUITY

LEEDS, 8TH–9TH APRIL 2019

List of Abstracts

1. Invited talks (alphabetical order)

Speaker: René Aid, Université Paris-Dauphine
Title: Optimal electricity demand response contracting with responsiveness incentives
Abstract: The interaction between the producer and the consumer is modeled as a Principal-Agent problem, thus accounting for the moral hazard underlying demand response programs. The producer, facing the limited flexibility of production, pays an appropriate incentive compensation in order to encourages the consumer to reduce his average consumption and to enhance his responsiveness. We provide closed–form solution for the optimal contract in the case of linear energy valuation. Without responsiveness incentive, this solution decomposes into a fixed premium for enrolment and a proportional price for the energy consumed, in agreement with previously observed demand response contracts. The responsiveness incentive induces a new component in the contract with payment rate on the consumption quadratic variation. Furthermore, in both cases, the components of the premium exhibit a dependence on the duration of the demand response event. In particular, the fixed component is negative for sufficiently long events. The calibration of our model to publicly available data of a large scale demand response experiment predicts a significant increase of responsiveness under our optimal contract, a significant increase of the producer satisfaction, and a significant decrease of the consumption volatility. The stability of our explicit optimal contract is justified by appropriate sensitivity analysis. Joint work with Dylan Possamaï and Nizar Touzi.

Speaker: Erik Bardoux, London School of Economics
Title: Optimality of doubly reflected Lévy processes in singular control
Abstract: We consider a class of two-sided singular control problems. A controller either increases or decreases a given spectrally negative Levy process so as to minimize the total costs comprising of the running and control costs where the latter is proportional to the size of control. We provide a sufficient condition for the optimality of a double barrier strategy, and in particular show that it holds when the running cost function is convex. Using the fluctuation theory of doubly reflected Levy processes, we express concisely the optimal strategy as well as the value function using the scale function. Numerical examples are provided to confirm the analytical results. Based on joint work with Kazutoshi Yamazaki

Speaker: Daniel Bartl, University of Vienna
Title: Wasserstein and adapted Wasserstein distances in mathematical finance
Abstract: In this talk we model uncertainty through neighborhoods in

Wasserstein distance within a one-period framework. After a short discussion on the choice of distance, we show (semi-)explicit formulas for some robust risk measures. We then conduct a sensitivity analysis (of e.g. utility maximization) and finally study a scaling limit in continuous time of Wasserstein neighborhoods. If time permits, we elaborate why Wasserstein distances are not suited for a general multi-period analysis and introduce an adapted modification. Based on joint works with J.Backhoff, M.Beiglboeck, S.Drapeau, M.Eder, M.Kupper, J.Obloj, L.Tangpi, J.Wiesel.

Speaker: Erhan Bayraktar, University of Michigan
Title: On the asymptotic optimality of the comb strategy for prediction with expert advice
Abstract: For the problem of prediction with expert advice in the adversarial setting with geometric stopping, we compute the exact leading order expansion for the long time behavior of the value function. Then, we use this expansion to prove that as conjectured in Gravin, Peres and Sivan [12], the comb strategies are indeed asymptotically optimal for the adversary in the case of 4 experts. Joint work with Ibrahim Ekren and Yili Zhang.

Speaker: Luciano Campi, London School of Economics
Title: N-player games and mean-field games with smooth dependence on past absorptions
Abstract: Mean-field games with absorption is a class of games, that have been introduced in Campi and Fischer (2018) and that can be viewed as natural limits of symmetric stochastic differential games with a large number of players who, interacting through a mean-field, leave the game as soon as their private states hit some given boundary. In this talk, we push the study of such games further, extending their scope along two main directions. First, a direct dependence on past absorptions has been introduced in the drift of players' state dynamics. Second, the boundedness of coefficients and costs has been considerably relaxed including drift and costs with linear growth. Therefore, the mean-field interaction among the players takes place in two ways: via the empirical sub-probability measure of the surviving players and through a process representing the fraction of past absorptions over time. Moreover, relaxing the boundedness of the coefficients allows for more realistic dynamics for players' private states. We prove existence of solutions of the mean-field game in strict as well as relaxed feedback form. Finally, we show that such solutions induce approximate Nash equilibria for the N-player game with vanishing error in the mean-field limit as $N \to \infty$. This is based on a joint paper with M. Ghio and G. Livieri (SNS, Pisa).

Speaker: Erik Ekström, University of Uppsala
Title: Nonzero-sum Dynkin games with uncertain competition

Abstract: We consider an optimal stopping game (Dynkin game) with uncertainty about the existence of competitors. This is natural for example when considering real options that are freely available to everyone, and where different actors do not want to publicly reveal their interest in a certain business opportunity. Naturally, hiding of information is important, and we therefore consider the strategies that are randomized stopping strategies. We show that there exists a symmetric equilibrium in randomized stopping times. Moreover, the equilibrium is described explicitly in terms of the corresponding one-player game. This is joint work with Tiziano De Angelis.

Speaker: Sigrid Källblad, Vienna University of Technology
Title: Stochastic control of measure-valued martingales with applications to robust pricing and Skorokhod embedding problems
Abstract: We consider a stochastic control problem where the controlled processes are (probability) measure-valued martingales (MVMs). We consider the problem in a weak formulation where the MVMs appear as weak solutions to certain (controlled) SDEs for which we prove existence of solutions. We then show that our control problem satisfies the Dynamic Programming Principle and that the associated value function is a viscosity solution in a certain sense to a corresponding HJB equation. A key motivation for the study of control problems featuring MVMs is that a number of interesting probabilistic problems can be formulated as optimisation problems over such processes; we illustrate this by applying our results to optimal Skorokhod embedding problems as well as robust pricing problems. The talk is based on joint work with A. Cox, M. Larsson and S. Svaluto.

Speaker: Jan Obłój, University of Oxford
Title: Information (data-driven) approach to (robust) pricing and hedging
Abstract: I introduce the robust approach to pricing and hedging which does not start with an priori probability measure but is instead data driven. The framework is designed to interpolate between model-independent and model-specific settings and to allow to address and quantify model risk. I explain briefly how classical fundamental notions and theorems in quantitative finance extend to the robust setting. I then focus on a simple two-dimensional study case of pricing and hedging a spread option, introducing suitable numerical methods and presenting concrete examples. I use vanilla option prices, together with agent-prescribed bounds on key market characteristics, to drive the interval of no-arbitrage prices and the associated hedging strategies. The setting can be seen as a constrained variant of the classical optimal transportation problem and comes with a natural pricing-hedging duality. I discuss numerical methods based on discretization and LP implementation but subsequently focus on a deep NN optimization. At the end I will outline some of the higher-dimensional challenges for such

methods as well as way to coherently combine option price data with past time series data in one estimation procedure.

Speaker: Catherine Rainer, University of Brest
Title: Solving two-state Markov games with incomplete information on one side
Abstract: We study the optimal use of information in Markov games with incomplete information on one side and two states. We provide a finite-stage algorithm for calculating the limit value as the gap between stages goes to 0, and an optimal strategy for the informed player in the limiting game in continuous time. This limiting strategy induces an ϵ-optimal strategy for the informed player, provided the gap between stages is small. This is a common work with Galit Ashkenazy-Golan and Eilon Solan, University of Tel Aviv, Israel.

Speaker: Neofytos Rodosthenous, Queen Mary University of London
Title: Optimal control of the debt ratio in a regime-switching economy
Abstract: We study the problem of a government wishing to control the country's debt-to-GDP ratio. The debt-to-GDP ratio evolves stochastically and the interest on debt is affected by an N-state continuous-time Markov chain, representing the country's credit ratings. The debt-to-GDP ratio can be reduced through fiscal interventions or increased by public investments. The government aims to choose a policy minimising the total expected cost of having debt and fiscal interventions counterbalanced by the gain from public investments. The problem is modelled by a bounded-variation stochastic control problem, that we explicitly solve through the analysis of an associated Dynkin game. This is joint work with Giorgio Ferrari.

Speaker: Ronnie Sircar, Princeton University
Title: Oligopoly Mean Field Games & Energy Production
Abstract: We discuss oligopoly games with a continuum of players that have mean field structure. These may be of Bertrand (price setting) or Cournot (quantity setting) type and may apply to analysis of consumer goods or energy markets respectively. Key advantages over finite player nonzero sum differential games are analytical and numerical tractability of the associated PDEs. Models for energy markets with competition between producers with heterogeneous costs (fossils vs. renewables) are presented as motivation. Sources of uncertainty in the stochastic version of the problem include controlled random discovery of reserves, and uncertain demand environments.

Speaker: Stéphane Villeneuve, Toulouse School of Economics
Title: On a monotone dynamic approach of optimal stopping problems for

Continuous-time Markov chains
Abstract: This paper is concerned with the solution of the optimal stopping problem associated to the valuation of American Options driven by continuous time Markov chains with discrete state spaces. The value function of an American option in this setting is characterized by the limit of a monotone sequence of values associated to hitting times strategies. The main advantage of this monotone approach, is that it works with minimal assumptions on the continuous- time Markov chain and the payoff function.

2. Contributed talks (alphabetical order)

Speaker: Katia Colaneri, University of Leeds
Title: A class of recursive optimal stopping problems with applications to stock trading in the dark pool
Abstract: In this paper we introduce and solve a class of optimal stopping problems of recursive type. In particular, the stopping payoff depends directly on the value function of the problem itself. In a multi-dimensional Markovian setting we show that the problem is well posed, in the sense that the value is indeed the unique solution to a fixed point problem in a suitable space of continuous functions, an! d an optimal stopping time exists. We then apply our class of problems to a model for stock trading in the dark pool and we determine the optimal stopping rule in that case.

Speaker: Davide De Santis, London School of Economics
Title: Nonzero-sum stochastic differential games between an impulse controller and a stopper
Abstract: We study a two-player nonzero-sum stochastic differential game where one player controls the state variable via additive impulses while the other player can stop the game at any time. The main goal of this work is characterize Nash equilibria through a verification theorem, which identifies a new system of quasi-variational inequalities whose solution gives equilibrium payoffs with the correspondent strategies. Moreover, we apply the verification theorem to a game with a one-dimensional state variable, evolving as a scaled Brownian motion, and with linear payoff and costs for both players. Two types of Nash equilibrium are fully characterized, i.e. semi-explicit expressions for the equilibrium strategies and associated payoffs are provided. Both equilibria are of threshold type: in one equilibrium the intervention regions of the players are separated, while in the other one they can overlap producing a situation where the first player induces her competitor to stop the game. Finally, we prove some asymptotic results with respect to the intervention costs.

Speaker: Dominykas Norgilas, University of Warwick
Title: Pricing and hedging of American puts under model uncertainty
Abstract: In a two-period setting we derive the model-independent upper bound of the American put option. The model associated with the highest price of the American put is the extended left-curtain martingale coupling. Moreover we derive the cheapest superhedge.

Speaker: Benjamin A. Robinson, University of Bath
Title: Optimal control of martingales in a radially symmetric environment
Abstract: We study a stochastic control problem where the controlled process is a d-dimensional martingale constrained to have unit quadratic variation, and the cost function is radially symmetric and continuous. We give a characterisation of the value function, showing that it satisfies a smooth fit condition. We see that the problem can be reduced to a one-dimensional control problem with two behaviour regimes. Under one of these regimes, the controlled process has zero local time, and so the usual rationale for smooth fit breaks down. However, we show that the problem still exhibits smooth fit, by providing an explicit construction of the value function. We prove optimality by showing that our construction gives the unique viscosity solution to the associated d-dimensional Hamilton-Jacobi-Bellman equation. We demonstrate, by means of an example, that smooth fit can fail for a discontinuous radially symmetric cost function. We will also see that the theory of viscosity solutions is insufficient to study the problem in this case. This is joint work with Alexander Cox.

Speaker: Haodong Sun, University of Warwick
Title: Dynkin games with Poisson random intervention times
Abstract: This talk introduces a new class of Dynkin games, where the two players are allowed to make their stopping decisions at a sequence of exogenous Poisson arrival times. The value function and the associated optimal stopping strategy are characterized by the solution of a backward stochastic differential equation. The talk further provides a replication strategy for the game, and applies the model to study the optimal conversion and calling strategies of convertible bonds. Our recent work on the Dynkin games with dependent point processes will also been covered in the talk.

Speaker: Yuqiong Wang, University of Uppsala
Title: Optimal stopping with discrete costly observations
Abstract: We study stopping problems where each observation of the underlying is associated with a positive cost for the optimizer. The discrete cost of an observation forces a discrete structure upon the problem, which therefore has similarities with stopping theory in discrete time. However, along with finding an optimal stopping time, the optimizer additionally faces

the problem of sequentially distributing the observation times in an efficient way. Adopting a Markovian formulation of the problem, we show that the problem is closely related to fixed points of a certain functional operator. In fact, the value function can be characterized in terms of a fixed point of this operator, and the optimal strategy of observation times and a stopping time can be described explicitly. This is joint work with Erik Ekström.

Speaker: Diego Zabaljauregui, London School of Economics
Title: Optimal market making under partial information with general intensities
Abstract: Starting from the Avellaneda-Stoikov framework, we consider a market maker who wants to optimally set bid/ask quotes over a finite time interval, to maximise her expected utility. The intensities of the orders she receives depend not only on the spreads she quotes, but also on unobservable factors modelled by a hidden Markov chain. We tackle this stochastic control problem under partial information with a model that unifies and generalises many existing ones, combining several risk metrics and constraints, and using general decreasing intensity functionals. We use stochastic filtering, control and piecewise-deterministic Markov processes theory, to reduce the dimensionality of the problem and characterise the reduced value function as the unique continuous viscosity solution of its dynamic programming equation. We then solve the analogous full information problem and compare the results numerically through a concrete example. We show that the optimal full information spreads are biased when the exact market regime is unknown, and the MM needs to adjust for 'regime risk' in terms of liquidity volatility and sensitivity to regime changes. This effect becomes higher, the longer the waiting time in between orders

3. Posters (alphabetical order)

Speaker: Jason Anquandah, University of Leeds
Title: Optimal Stopping in a Simple Model of Unemployment Insurance
Abstract: Managing unemployment is one of the key issues in social policies. Unemployment insurance schemes are designed to cushion the financial and morale blow of loss of job but also to encourage the unemployed to seek new jobs more pro-actively due to the continuous reduction of benefit payments. A simple model of unemployment insurance is proposed with a focus on optimality of the individual's entry to the scheme. The corresponding optimal stopping problem is solved, and its similarity and differences with the perpetual American call option are discussed. Some examples in this direction are worked out.

Speaker: Cheng Cai, University of Leeds
Title: Optimal hedging for American options with a single transaction
Abstract: In the Black-Scholes model, an option seller constructs a self-financing stock-bond portfolio in order to Delta hedge a short position in a perpetual American put option. In contrast to the continuous trading of the Black-Scholes model, the option seller can only rebalance her portfolio **once** before the time $\tau_a \wedge \tau_b$ at which the underlying stock price S_t leaves an interval (a, b). Here the lower end-point a is the optimal exercise price of the put option and $b(> a)$ is chosen arbitrarily. The goal is to determine the optimal time to rebalance the portfolio and the optimal hedge ratio (stock holding after the trade) that minimize the variance of the so-called tracking error (at time $\tau_a \wedge \tau_b$). First, we formulate the optimal hedging problem for a fixed initial stock holding as a one dimensional optimal stopping problem. This is solved by constructing three different free boundary problems depending on possible parameter choices. Second, we study analytically how the stopping boundaries move in response to variations in the initial stock holding. Finally, we obtain an equation that must be satisfied by an optimal initial stock holding.

Speaker: Nikita Merkulov, University of Leeds
Title: An optimal stopping game with asymmetric information
Abstract: We study an optimal stopping game of two players. A scenario - a "state of the world" in which the game is played - is chosen according to a probability distribution known to both players. Immediately after the game starts, one of the players receives the information about the scenario, therefore knowing exactly the structure of the game. They then need to use a randomized stopping strategy in order to hide the information from the other player and use their informational advantage.

We restrict ourselves to a special case when only the terminal time payoff depends on the random scenario. We discuss alternative representations of the game: in particular, usage of randomized strategies gives rise to an equivalent game of optimal controls. Ultimately, we provide the conditions necessary for the game to have a value and an ϵ-Nash equilibrium.

Speaker: Shilei Niu, Xi'an Jiaotong-Liverpool University
Title: Implications of Electricity Price Regimes on Hydroelectric Power Plant Valuation
Abstract: In a Markov regime switching model we investigate how the introduction of a third regime and its associated jump size for electricity prices affect the hydro power plant's value with and without ramping rate restrictions. Specifically, we propose the jump size elasticity of value to measure the value of price jumps to the hydro operator. Our numerical experiments show that the plant value is positively affected by the jump size, but, to some extent, the impact depends on the restrictiveness of ramping

constraints. We also provide further evidence that the hydro plant faces a larger negative impact from ramping restrictions when the expected price variation is increased due to an increase in the jump size. This makes it desirable to adjust the water release rate more substantially and frequently. (Joint work with T. Wirjanto)

Speaker: Johannes Wiesel, University of Oxford
Title: TBA
Abstract: TBA