Junior Researchers in Stochastic Optimal Control

31st August – 1st September 2023

Contents

1. Welcome

It is our great pleasure to welcome you to the Junior Researchers in Stochastic Optimal Control workshop in Berlin. We hope that you will enjoy a pleasant and productive stay.

Venue and further information

The workshop will take place at TECHNISCHE UNIVERSITÄT (TU) BERLIN, Straße des 17. Juni 135, 10623 Berlin in room H 0111, on the ground floor of the main building:

`https://osm.org/go/0MZu3QNP?m=&relation=3348239`.

All lectures and talks will be held in person. Keynote lectures are scheduled for 60 minutes including open discussion. Contributed talks are scheduled for 15 minutes plus 5 minutes for open discussion.

The social dinner will take place in the Indian restaurant Amrit Mitte, Oranienburger Str. 45, 10117 Berlin:

`https://www.openstreetmap.org/node/279302123`,

on Thursday starting at 7 p.m.

Conference organisers

Peter Bank
Sebastian Ertel
Sascha Gaudlitz
Yifan Jiang
Hannes Kern
Laura Körber
Alexander Merkel
Shyam Popat
Emanuel Rapsch
Thomas Wagenhofer

Supporting Institutions

WIAS

Stochastic Analysis in Interaction

TECHNISCHE UNIVERSITÄT BERLIN

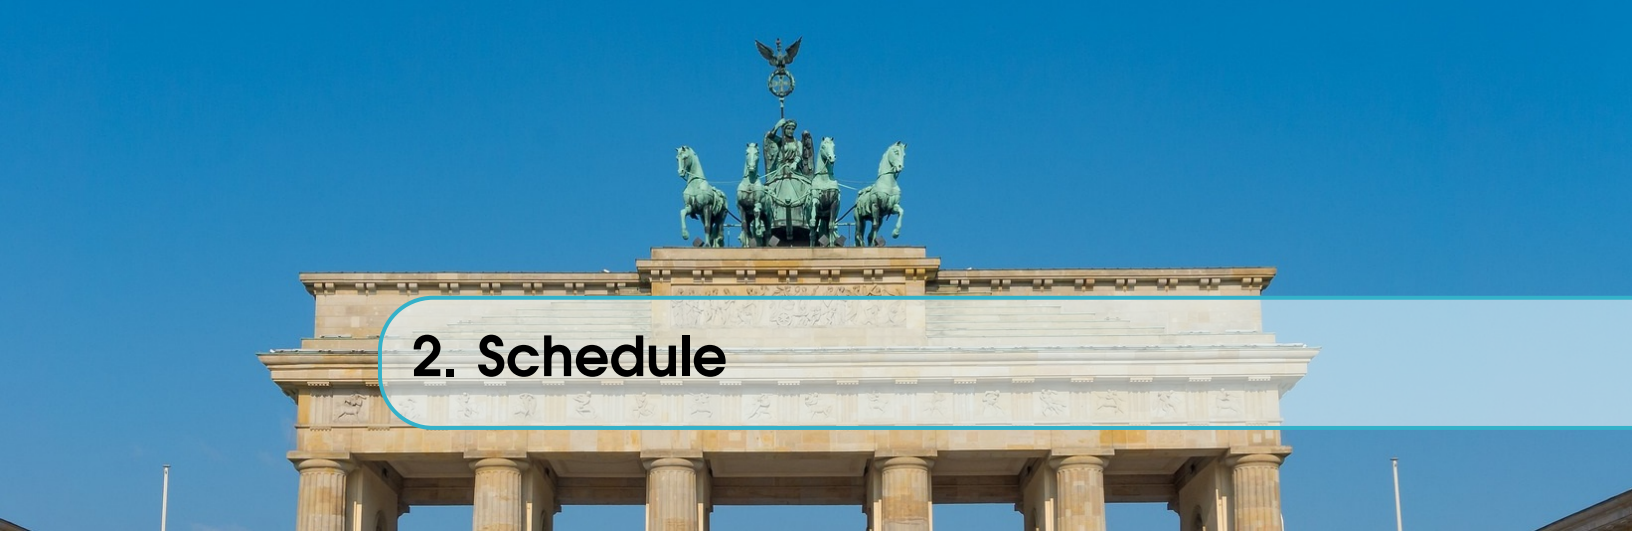

2. Schedule

Thursday, 31st August

Location: **TU Berlin**, Straße des 17. Juni 135, 10623 Berlin
`https://osm.org/go/0MZu3QNP?m=&relation=3348239`
Room: H 0111 (Main Building, Ground Floor)

08:20–09:00	**Registration**		
09:00–09:10	**Opening Remarks**		
09:10–10:10	PETER TANKOV (ENSAE Paris)	DECARBONIZATION OF LARGE FINANCIAL MARKETS	8
10:10–10:40	**Coffee Break**		
10:40–11:00	Redouane Silvente (ENSAE Paris)	*Optimal control of storage and short-term price formation in electricity market*	13
11:00–11:20	Katharina Eichinger (École Polytechnique)	*On McKean-Vlasov control problems: turnpike property and propagation of chaos*	9
11:20–11:40	Fanny Cartellier (ENSAE Paris)	*Mean field coarse correlated equilibria in linear quadratic games and application to an abatement game*	9
11:40–12:00	Berenice Anne Neumann (Trier University)	*Finite state mean field games with common shocks*	12
12:00-13:30	**Lunch Break**		
13:30–14:30	DAVID SISKA (University of Edinburgh)	GRADIENT FLOWS ARISING IN STOCHASTIC CONTROL PROBLEMS	8
14:30-14:50	**Coffee Break**		
14:50–15:10	Giacomo Greco (Eindhoven University of Technology)	*Exponential convergence of Sinkhorn's algorithm: a stochastic optimal control approach*	10
15:10–15:30	Stefan Perko (Friedrich Schiller University Jena)	*Unlocking optimal batch size schedules using continuous-time control and perturbation theory*	12
15:30–15:50	Thibaut Bourdais (ENSTA Paris)	*An entropy penalized approach for stochastic control problems*	9
15:50–16:10	Jonathan Tam (University of Oxford)	*Controlling the information flow: trade-off between cost and observations*	13
16:10–16:30	**Coffee Break**		
16:30–17:30	JOHANNES MUHLE-KARBE (Imperial College London)	PRE-HEDGING	7

Friday, 1st September

Location: **TU Berlin**, Straße des 17. Juni 135, 10623 Berlin
`https://osm.org/go/0MZu3QNP?m=&relation=3348239`
Room: H 0111 (Main Building, Ground Floor)

09:00–10:00	ROXANA DUMITRESCU (King's College London)	ENERGY TRANSITION: A MEAN-FIELD GAME APPROACH	7
10:00–10:40	**Coffee Break**		
10:40-11:00	Chiara Rossato (ETH Zürich)	*Golden parachutes under the threat of accidents*	12
11:00–11:20	Daniel Krŝek (ETH Zürich)	*Relaxed principal-agent problem*	11
11:20–11:40	Topias Tolonen (Uppsala University)	*Hiring and firing – a signaling game*	13
11:40–12:00	Maike Klein (Kiel University)	*On the gain of collaboration*	11
12:00–13:30	**Lunch Break**		
13:30–13:50	David Itkin (Imperial College London)	*Are linear strategies nearly optimal when trading with superlinear frictions?*	11
13:50–14:10	Tolulope Latunde (University of Edinburgh)	*Guarding investments amidst turbulence: sensitivity analysis of an optimal control model for the management of capital assets in the face of debt crisis*	11
14:10–14:30	Justin Gwee (London School of Economics)	*The explicit solution to a risk-sensitive ergodic singular stochastic control problem*	10
14:30–14:50	**Coffee Break**		
14:50–15:10	Filippo de Feo (Politecnico di Milano)	*Optimal control of stochastic delay differential equations*	10
15:10–15:30	Alessandro Milazzo (University of Torino)	*A detection problem with a monotone observation rate*	12
15:30–15:50	Marco Rodrigues (ETH Zürich)	*Reflections on BSDEs*	13
15:50–16:15	**Coffee Break**		
16:15–17:15	GIANMARIO TESSITORE (University Milano-Bicocca)	CONTROL OF TWO-SCALE EVOLUTIONARY SYSTEMS BY BSDES	8

3. Keynote Lectures and Abstracts

3.1 Energy transition: a mean-field game approach

Roxana Dumitrescu, King's College London

Friday
09:00-10:00

In this lecture, I will present two models developed in the context of energy transition - from both supply and demand sides - using as a mathematical tool the theory on mean-field games. I will first present a model for the industry dynamics in the electricity market, based on mean-field games of optimal stopping. In this model, there are two types of agents: the renewable producers and the conventional producers. The renewable producers choose the optimal moment to build new renewable plants, and the conventional producers choose the optimal moment to exit the market. The agents interact through the market price, determined by matching the aggregate supply of the two types of producers with an exogenous demand function. Using a relaxed formulation of optimal stopping mean-field games, we prove the existence of a Nash equilibrium and the uniqueness of the equilibrium price process. An empirical example, inspired by the UK electricity market, is presented. In the second part of the talk, we focus on demand-side strategies to deal with the increase of intermittent renewable energy. We consider an energy system with N consumers who are linked by a Demand Side Management (DSM) contract, i.e. they agreed to diminish, at random times, their aggregated power consumption by a predefined volume during a predefined duration. Their failure to deliver the service is penalised via the difference between the sum of the N power consumption and the contracted target. We are led to analyse a non-zero sum stochastic game with N players, where the interaction takes place through a cost which involves a delay induced by the duration included in the DSM contract. The asymptotic formulation can be written in terms of Mean-Field Game (MFG) with random jump time penalty and interaction on the control. We prove the existence of an equilibria, which admits a semi-explicit representation in the linear-quadratic case, and present several numerical illustrations. The first part is based on a joint work with R. Aid (Univ. Dauphine) and P. Tankov (Ensae), and the second part on a joint work with C. Alasseur (EDF), L. Campi (Univ of Milan) and J. Zeng (King's College).

3.2 Pre-Hedging

Johannes Muhle-Karbe, Imperial College London

Thursday
16:30-17:30

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. This is joint work with Roel Oomen.

3.3 Gradient flows arising in stochastic control problems

David Siska, University of Edinburgh

Thursday
13:30-14:30

Iterative approximation methods arise naturally in the study of stochastic control problems as these problems typically elude analytical solutions. I will discuss how many such iterative schemes (policy iteration due to Bellman, the method of successive approximations, and policy gradient from RL) give rise (in the limit) to continuous time gradient flows. With the appropriate notion of convexity and differentiability, one can then study their convergence by adapting well-developed techniques from optimization. In particular, I'll present the Fisher-Rao flow and its mirror-descent dual formulation: 1) for the modified method of successive approximations in the case of a relaxed and entropy regularized finite-time-horizon stochastic control problem, and 2) for the policy gradient algorithm for a relaxed and entropy regularized infinite-time-horizon Markov decision problem. This is joint work with James-Michael Leahy (Imperial), Bekzhan Kerimkulov (formerly University of Edinburgh), Deven Sethi (University of Edinburgh), Lukasz Szpruch (University of Edinburgh), and Yufei Zhang (LSE).

3.4 Decarbonization of large financial markets

Peter Tankov, ENSAE Paris

Thursday
09:10-10:10

We build a model of a financial market where a large number of firms determine their dynamic emission strategies under climate transition risk in the presence of both green-minded and neutral investors. The firms aim to achieve a trade-off between financial and environmental performance, while interacting through the stochastic discount factor, determined in equilibrium by the investors' allocations. We formalize the problem in the setting of mean-field games and prove the existence and uniqueness of a Nash equilibrium for firms. We then present a convergent numerical algorithm for computing this equilibrium and illustrate the impact of climate transition risk and the presence of green-minded investors on the market decarbonization dynamics and share prices. We show that uncertainty about future climate risks and policies leads to higher overall emissions and higher spreads between share prices of green and brown companies. This effect is partially reversed in the presence of environmentally concerned investors, whose impact on the cost of capital spurs companies to reduce emissions. However, if future climate policies are uncertain, even a large fraction of green-minded investors is unable to bring down the emission curve: clear and predictable climate policies are an essential ingredient to allow green investors to decarbonize the economy. This is joint work with Pierre Lavigne.

3.5 Control of two-scale evolutionary systems by BSDEs

Gianmario Tessitore, University Milano-Bicocca

Friday
16:15-17:15

The purpose of this talk is to present a review of the results regarding the control in infinite dimensional stochastic systems which evolve according to two different time scales. The first objective is to represent the limit of the value function as a solution to an HJB equation or as the value function of a reduced control problem. The techniques used are primarily related to the study of backward stochastic equations. Finally, some of the obtained results are applied to a model, related to climate evolution, where the slow evolution follows the climate scale and the fast evolution follows the human scale. This last part is a work in progress with F. Flandoli, G. Guatteri, and U. Pappalettera.

4. Contributed Talks and Abstracts

4.1 An entropy penalized approach for stochastic control problems

Thibaut Bourdais, ENSTA Paris

Thursday
15:30-15:50

We propose an alternative technique to dynamic programming for solving stochastic control problems. We consider a weak formulation that is written as an optimization problem on the space of probabilities. We then propose a regularized version of this problem obtained by splitting the optimization variables and penalizing the entropy between the two probabilities to be optimized. We show that the regularized problem provides a good approximation of the original problem when the weight of the entropy regularization term is large enough. Moreover, the regularized problem has the advantage of giving rise to optimization problems that are easy to solve in each of the two optimization variables when the other is fixed. We take advantage of this property to propose an alternating optimization algorithm whose convergence to the infimum of the regularized problem is shown. The relevance of this approach is illustrated by solving a high-dimensional stochastic control problem aimed at controlling consumption in electrical systems.

4.2 Mean field coarse correlated equilibria in linear quadratic games and application to an abatement game

Fanny Cartellier, ENSAE Paris

Thursday
11:20-11:40

Coarse correlated equilibria are an alternative to Nash equilibria which have first been introduced in Moulin et Vial (1978). They include a correlation device which can be interpreted as a mediator recommending strategies to the players, which makes it particularly relevant in a context of market failure. We develop a methodology to compute mean-field coarse correlated equilibria (CCEs) in a linear-quadratic framework. We discuss the specifications of the objective function under which CCEs outperform Nash equilibria (NEs) in terms of both social utility and control levels. We show that the mean field limit CCEs we found allow to build approximate CCEs in N-player settings. We apply this computation methodology to some relevant models, in particular to a CO2 abatement game between countries (Barrett 1994, in a slightly modified version). We show that in the latter CCEs allow to reach higher abatement levels than the NE, with higher global utility. This is a joint work with Luciano Campi and Federico Cannerozzi.

4.3 On McKean-Vlasov control problems: turnpike property and propagation of chaos

Katharina Eichinger, École Polytechnique

Thursday
11:00-11:20

In this talk we provide an analysis for a class of McKean-Vlasov control problems. Our technique is based on coupling by reflection adapted to controlled processes allowing us to treat dynamics consisting of a drift component only being strongly decreasing outside a large ball, and possibly an additional sufficiently small mean-field interaction. With this we prove a uniform in time Lipschitz estimate of the value function and its (measure) derivatives under suitable assumptions. This enables of to prove existence and uniqueness of the ergodic problem and the exponential turnpike property. We also prove uniform in time quantitative propagation of chaos of the dynamics of the corresponding N-agent control problem towards its meanfield counterpart.This is based on a joint work in progress with Alekos Cecchin, Giovanni Conforti and Alain Durmus.

4.4 Optimal control of stochastic delay differential equations

Filippo de Feo, Politecnico di Milano

Friday 14:50-15:10

In this talk we will discuss an optimal control problem for stochastic differential delay equations. We will only consider the case with delays in the state. We will show how to rewrite the problem in a suitable infinite-dimensional Hilbert space. Then using the dynamic programming approach we will characterize the value function of the problem as the unique viscosity solution of an infinite dimensional Hamilton-Jacobi-Bellman (HJB) equation. We will discuss partial $C^{1,\alpha}$-regularity of the value function. This regularity result is particularly interesting since it permits to construct a candidate optimal feedback control. This is the first partial regularity result for solutions of fully non-linear second order HJB equations on Hilbert spaces with unbounded operators. In the final part of the talk we will also present a strategy which permits to show the optimality of the candidate feedback control. This is the first optimal synthesis result for stochastic optimal control problems in Hilbert spaces with unbounded operators by means of viscosity solutions of HJB equations. This is a joint project with S. Federico and A. Święch.

4.5 Exponential convergence of Sinkhorn's algorithm: a stochastic optimal control approach

Giacomo Greco, Eindhoven University of Technology

Thursday 14:50-15:10

Sinkhorn's algorithm is an iterative fitting procedure that allows to quickly compute the unique solution of the Entropic Optimal Transport problem (EOT), i.e. the entropic regularization of the Kantorovich-OT problem. Despite its wide use in machine learning and computational optimal transport, the theoretical study of its convergence rate is still a very active area of research. In this talk, after a brief introduction to Sinkhorn's algorithm and EOT, we are going to prove the exponential convergence of the algorithm by relying on the connection between Sinkhorn's iterates and the evolution along the Hamilton-Jacobi-Bellman equations of value functions obtained from stochastic optimal control problems. Based on `https://arxiv.org/pdf/2304.06549.pdf` (accepted to COLT2023)

4.6 The explicit solution to a risk-sensitive ergodic singular stochastic control problem

Justin Gwee, London School of Economics

Friday 14:10-14:30

We consider a two-sided singular stochastic control problem with a risk-sensitive ergodic criterion. In particular, we consider a stochastic system whose uncontrolled dynamics are modelled by a linear diffusion. The control that can be applied to this system is modelled by an additive finite variation process. The objective of the control problem is to minimise a risk-sensitive long-term average criterion that penalises deviations of the controlled process from a nominal point as well as the expenditure of control effort. Such a problem has been motivated by applications in target tracking and in portfolio selection with transaction costs. We derive the complete solution to the problem under general assumptions by relating a suitable Sturm-Liouville problem with the free-boundary problem associated with the control problem's Hamilton-Jacobi-Bellman equation. In the limit as the risk-sensitivity parameter tends to zero, we recover the solution to the risk neutral case. Furthermore, we show that the optimal long-term average cost converges to the upper value of a deterministic zero sum differential game in the totally risk-sensitive limit. The talk is based on joint work with Mihail Zervos.

4.7 Are linear strategies nearly optimal when trading with superlinear frictions?

David Itkin, Imperial College London

Friday 13:30-13:50

We consider an investor with mean-variance preferences who trades on a noisy signal in the presence of transaction costs. When the costs are quadratic in the trade size, the celebrated results of Garleanu & Pedersen (2013) establish that the explicit optimal strategy is a linear feedback function of the current signal and position size. However, empirical evidence suggests that for many assets costs are superlinear, but not quadratic. In this case no explicit solution is available and the optimization problem can be computationally intensive to solve. Motivated by the quadratic cost case we consider a tractable one parameter class of linear strategies. We show that under realistic choices of the parameters the best strategy in this class performs nearly as well as the true optimum, which we compute with a brute force numerical method. Our result gives a simple and practical rule of thumb that can be efficiently implemented and yields nearly optimal performance. This is joint work with Xavier Brokmann, Johannes Muhle-Karbe and Peter Schmidt.

4.8 On the gain of collaboration

Maike Klein, Kiel University

Friday 11:40-12:00

We consider two firms with endowment processes given by continuous diffusions driven by independent Brownian motions. Both, the drift and diffusion rates are controlled, which can be interpreted as a collaboration of the firms. The aggregate endowment is assumed to be a Brownian motion with constant drift and diffusion rate. We show that the maximal joint survival probability depends only on the aggregate risk-adjusted return and on the maximal risk-adjusted return that can be implemented in each firm. Here the risk-adjusted return is understood as the drift rate divided by the squared diffusion rate. Moreover, we quantify the gain of collaboration in terms of the relative increase of the maximal joint survival probability due to a collaboration. This talk is based on a joint work with Stefan Ankirchner and Robert Hesse (both University of Jena).

4.9 Relaxed principal-agent problem

Daniel Krŝek, ETH Zürich

Friday 11:00-11:20

We study a principal-agent problem with a lump-sum payment on a finite-time horizon. Extending the dynamic programming approach in Cvitanić, Possamaï and Touzi (2017), which reduces the principal's problem to a standard optimal control problem, we consider problems involving constraints on the optimal contract. Since such a stochastic target problem ceases to be tractable with the standard PDE approach, we introduce a framework, in which the agent is allowed to choose measure-valued controls. This in turn gives more freedom to the principal, when she chooses the contract, and allows us to show existence of an optimal contract in problems with fairly general constraints.

4.10 Guarding investments amidst turbulence: sensitivity analysis of an optimal control model for the management of capital assets in the face of debt crisis

Tolulope Latunde, University of Edinburgh

Friday 13:50-14:10

Controlled asset management is a critical aspect influencing decision-making in investments. A significant problem on which investors base their judgements in developed asset allocation and during a debt crisis is determining the optimal timing and quantity of investments to maximize expected net asset values while minimizing risks for long-term capital asset management. Recognizing this enables investors to make rational decisions. This study intends to give a control policy to assist investors in making decisions for asset management sustainability, as well as to provide a greater knowledge required of the problems and opportunities connected with investment decisions during times of financial insecurity, such as debt crises. This can be approached by employing an asset-liability optimal control model to address real-life economic issues of a country and hence carrying out the sensitivity analysis by examining the impact of varying debt crisis scenarios on the optimal control parameters and investment decisions. The result shed light on the robustness of the model and provides insights into the adjustments required to safeguard investments during periods of financial instability.

4.11 A detection problem with a monotone observation rate

Alessandro Milazzo, University of Torino

Friday
15:10-15:30

In quickest detection problems the goal is to detect, as soon as possible, a random "disorder" time that affects the dynamics of an underlying process. We study a quickest detection problem where the observation rate of the underlying process can be increased at any time for higher precision in the detection, but at an observation cost that grows linearly in the observation rate. This leads to a problem of combined control, filtering and stopping. We show that the problem has a two-dimensional sufficient statistic comprised of the current observation rate together with the conditional probability that disorder has happened. Moreover, the problem is shown to have a semi-explicit solution, where for some parameter values it is too costly to exert control at all, whereas for other parameter values the optimal behaviour is to increase the observation rate in such a way that the sufficient statistic stays below a certain boundary until the optimal stopping time. In the latter case, we characterise the reflection boundary and the optimal stopping strategy with the help of appropriate smooth fit conditions.

4.12 Finite state mean field games with common shocks

Berenice Anne Neumann, Trier University

Thursday
11:40-12:00

We present a new framework for mean field games with finite state space and common noise, where the common noise is given through shocks that occur at random times. We first analyse the game for up to n shock times in which case we are able to characterize mean field equilibria through a family of parametrized and coupled forward-backward systems and prove existence of solutions to these systems for a small time horizon. Thereafter, we show that for the case of an unbounded number of shocks the equilibria of the game restricted to n shocks are approximate mean field equilibria.

4.13 Unlocking optimal batch size schedules using continuous-time control and perturbation theory

Stefan Perko, Friedrich Schiller University Jena

Thursday
15:10-15:30

Stochastic Gradient Descent (SGD) and its variants are almost universally used to train neural networks and to fit a variety of other parametric models. An important hyperparameter in this context is the batch size, which determines how many samples are processed before an update of the parameters occurs. Previous studies have demonstrated the benefits of using variable batch sizes. In this work, we will theoretically derive optimal batch size schedules for SGD and similar algorithms, up to an error that is quadratic in the learning rate. To achieve this, we approximate the discrete process of parameter updates using a family of stochastic differential equations indexed by the learning rate. To better handle the state-dependent diffusion coefficient, we further expand the solution of this family into a series with respect to the learning rate. Using this setup, we derive a continuous-time optimal batch size schedule for a large family of diffusion coefficients and then apply the results in the setting of linear regression.

4.14 Golden parachutes under the threat of accidents

Chiara Rossato, ETH Zürich

Friday
10:40-11:00

Based on recent work by Possamaï and Touzi, we consider an extension of Sannikov's principal-agent problem by letting the agent control the drift and the jump intensity of the output process. We investigate whether the problem exhibits a golden parachute, that is, whether there is a scenario in which the agent retires and receives a continuous stream of payments or the agent receives a lump-sum payment as compensation for the termination of the contract by the principal. With the introduction of the face-lifted utility, we can study the two cases simultaneously in a different principal-agent problem that we reduce to a standard mixed control-stopping problem. This is based on joint work with Dylan Possamaï.

4.15 Reflections on BSDEs

Marco Rodrigues, ETH Zürich

Friday
15:30-15:50

We consider backward stochastic differential equations (BSDEs) and reflected BSDEs in the genereality that will allow a unified study of certain discrete-time and continuous-time control problems on random time horizons. We provide well-posedness results for the BSDEs and reflected BSDEs with optional obstacle process in case of appropriately weighted square-integrable data. This is based on joint work with Dylan Possamaï.

4.16 Optimal control of storage and short-term price formation in electricity market

Redouane Silvente, ENSAE Paris

Thursday
10:40-11:00

We develop a model for the storage dynamics in the electricity market. We consider that the storage facilities behave as price takers in this market. Given an exogenous price process, a single storage player solves a maximisation problem representing the best earning strategy for the asset owner. Using a stochastic optimal control approach, we find an explicit expression for the optimal strategy of a Pumped Hydroelectric Energy Storage (PHES). With this knowledge, the focus shifts to viewing the electricity market as a platform for interaction between energy demand, price taker storage players, renewable energy producers bidding always their full, stochastic, capacity, and conventional producers who have a supply function supposedly known. We prove the existence and uniqueness of the price process resulting from a market equilibrium between energy supply and demand under some conditions on the storage facilities. We then illustrate through multiple numerical examples the concrete influence of storage in terms of the variation of the short-term electricity price.

4.17 Controlling the information flow: trade-off between cost and observations

Jonathan Tam, University of Oxford

Thursday
15:50-16:10

This talk revolves around aspects of actively controlling the information flow within the framework of Markov decision processes. This models situations which require a trade-off between quality of information and cost. We discuss the problem reformulation in terms of a partially observable MDP, and examine the use of penalty methods as a numerical approach for the resulting system of quasi-variational inequalities obtained from dynamic programming. In the associated mean-field game setup, we define the appropriate notion of a mean-field Nash equilibrium on an augmented problem, and employ entropy regularisation to demonstrate a contraction result for an iterative scheme to compute approximate Nash equilibria.

4.18 Hiring and firing – a signaling game

Topias Tolonen, Uppsala University

Friday
11:20-11:40

We study a strategic game between an employer and a potential employee, where the employee has private information regarding their production capacity. At the initial stage, the employee communicates a salary claim, after which their true production capacity is gradually revealed to the employer as the unknown drift of a Brownian motion representing the revenues generated by the employee. Subsequently, the employer has the possibility to reject the salary claim or choose a time to fire the employee in case the estimated production capacity understeps the salary. We set-up a semi-separating Perfect Bayesian Equilibrium in which the employee provides a mixed salary claim, and the employer uses a conditional sequential probability for the high production capacity to calculate an stopping time to use as an optimal strategy. In addition, we briefly discuss alternative equilibria and showcase the robustness of our set-up by discussing few immediate extensions, such as introducing a firing cost, uncertainty for employee regarding their type, and additional grading as an interview phase. Submission is based on a joint work together with professor Erik Ekström.

5. Participants

Gevorg Adamyan
Peter Bank
Dirk Becherer
Thibaut Bourdais
Simon Breneis
Umur Can Kaya
Federico Cannerozzi
Fanny Cartellier
Robert Denkert
Roxana Dumitrescu
Katharina Eichinger
Stefan Engelhardt
Sebastian Ertel
Filippo de Feo
Lukas Gräfner
Maalvladedon Ganet Some
Sascha Gaudlitz
Giacomo Greco
Justin Gwee
Paul Hager
Franz Haniel
Stefanie Hesse
Robert Hesse
David Itkin
Dave Jacobi
Yifan Jiang
Likai Jiao
Adam Christopher Jones
Hannes Kern
Maike Klein
Konrad Kleinberg
Laura Körber
Fritz Krause
Daniel Krŝek
Tolulope Latunde
Olexandr Lednyts'kyy
Ibrahim Mbouandi Njiasse
Alexander Merkel
Alessandro Milazzo
Johannes Muhle-Karbe
Berenice Anne Neumann
Ajla Nurkanovic
Maria Olympia Tsianni
Luca Pelizzari
Stefan Perko
Nicolas Perkowski
Shyam Popat
Alexandra Quitmann
Laurena Ramadani
Emanuel Rapsch
Hans Reimann
Marco Rodrigues
Chiara Rossato
Rahama Sani Abdullahi
Gemma Sedrakjan
Redouane Silvente
David Siska
Lukasz Sliwinski
Wilhelm Stannat
Janine Steck
Jan Szalankiewicz
Sato Takashi
Jonathan Tam
Peter Tankov
Gianmario Tessitore
Topias Tolonen
Theresa Traxler
Alexander Vogler
Thomas Wagenhofer
Julian Wendt
Ruben Wiedemann
César Zarco-Romero
Sebastian Zimper