Workshop:

YOUNG RESEARCHERS IN STOCHASTIC CONTROL AND GAMES

Bielefeld, December 14 – December 15, 2023

SFB 1283 IMW UNIVERSITÄT BIELEFELD

Contents

About

Aim

The workshop *Young Researchers in Stochastic Control and Games* aims to bring together early career researchers working in the field of stochastic control.

The main topics of the workshop comprise the modeling and the analysis of optimal decision problems in a stochastic environment, new trends relating stochastic control and machine learning, and theoretical and numerical perspectives on multi agents decision problems.

Sponsor

The workshop *Young Researchers in Stochastic Control and Games* is funded by the the CRC1283 *Taming uncertainty and profiting from randomness and low regularity in analysis, stochastics and their application* (https://www.sfb1283.uni-bielefeld.de/Pages/home).

Organizers

The workshop is organized by Jodi Dianetti (jodi.dianetti@uni-bielefeld.de) with the help of the Center for Mathematical Economics (https://www.uni-bielefeld.de/zwe/imw/).

Timetable

Thursday, 14.12.20023

09:00–09:20	**Registration: room V2-210-216**	
09:20–09:30	**Welcome Address**	
9:30-10:10	**Greco**	Stochastic control meets Sinkhorn on a Schrödinger bridge
10:10-10:50	**Tzouanas**	Ergodic mean-field games of singular control with regime-switching
10:50–11:20	**Coffee: room V10-151**	
11:20–12:00	**Denkert**	Extended mean-field games with multi-dimensional singular controls
12:00-12:40	**Cannerozzi**	Coarse correlated equilibria in mean field games
12:40–14:00	**Lunch**	
14:00-14:40	**Kassing**	Stochastic modified flows, mean-field limits and dynamics of stochastic gradient descent
14:40-15:20	**Perko**	Unlocking optimal batch size schedules using continuous–time control and perturbation theory
15:20–16:00	**Coffee: room V10-151**	
16:00–16:40	**Merkel**	Optimal adaptive control with separable drift uncertainty
16:40–17:20	**Milazzo**	Learning-by-doing for irreversible investment under incomplete information
17:20-18:00	**Wang**	Bayesian sequential least-squares estimation in discrete time
20:00	**Social Dinner at *The Bernstein Bielefeld***	

Friday, 15.12.20023

9:30-10:10	**Dexheimer**	Data-driven optimal stopping: A pure exploration analysis
10:10-10:50	**Zhu**	Optimal retirement choice under age-dependent force of mortality
10:50–11:20	**Coffee: room V10-151**	
11:20–12:00	**Bovo**	A saddle point in stopper vs. singular-controller games with free boundaries
12:00-12:40	**Neumann**	Markovian randomized equilibria for general Markovian Dynkin games in discrete time
12:40–14:00	**Lunch**	
14:00-14:40	**Eichinger**	Meanfield control problems: turnpike property and propagation of chaos
14:40-15:20	**Martini**	Fourier expansion of mean field control problems
15:20–15:30	**Concluding Remarks**	

List of Abstracts

A saddle point in stopper vs. singular-controller games with free boundaries

Andrea Bovo

In this work, we look for a saddle point for a class of Stopper vs. Singular-Controller games over a finite-time horizon. In this game the controller (the minimiser) can choose a control from the class of singular controls, whereas the stopper (the maximiser) can choose the time at which the game ends. We prove that the game admits a value function, and it is related to a variational inequality with two constraints – an obstacle constraint and a gradient constraint. These constraints lead to two continuous free-boundaries which divide the whole 'space' in three regions. We prove that the value function is continuously differentiable in both time and space, and the mixed derivative and the second order spatial derivative are continuous across the free boundary associated to the controller. Finally, we obtain an optimal strategy for both the controller and the stopper player.

Coarse correlated equilibria in mean field games

Federico Cannerozzi

In the context of continuous-time stochastic mean field games (MFGs), we introduce a generalization of MFG solution, called coarse correlated solution, which can be seen as the mean field game analogue of a coarse correlated equilibrium. The latter is a generalization of Nash equilibrium for stochastic games, which allows for correlation between the strategies of non-cooperative players. Our notion of solution can be justified by showing that approximate N-player correlated equilibria can be constructed starting from a correlated solution to the mean field game, and existence can be proved by means of a minimax theorem. If time allows, an application to an abatement game between greenhouse gas emitters will be presented, in which coarse correlated solutions lead both to greater abated quantities and higher payoffs than the usual MFG solution. The talk is based upon joint works with L. Campi (University of Milan "La Statale"), M. Fischer (University of Padua) and F. Cartellier (ENSAE Paris).

Extended mean-field games with multi-dimensional singular controls

Robert Denkert

We consider a class of extended mean-field games (MFG) involving multi-dimensional singular controls where the impact of the singular control both on the state-process and the reward functional is not necessarily linear. Our key idea is linking singular controls to parametrisations which carry the missing information on how jumps of the singular control are executed. By first considering the bounded velocity case and then proving an approximation result, we show the existence of Nash equilibria to the MFG in the general singular control case. This presentation is

based on joint work with Ulrich Horst.

Data-driven optimal stopping: A pure exploration analysis

Niklas Dexheimer

We investigate data-driven optimal stopping from a pure exploration perspective and derive non-asymptotic results on the simple regret of the proposed estimator of the optimal stopping barrier. Additionally, uniform and non-asymptotic PAC-bounds and minimax lower bounds for the simple regret, matching the upper bounds, are provided. All results are shown under general conditions on the payoff functions and also under more refined assumptions, mimicking the Tsybakov noise condition used in binary classification, which leads to an improved rate of convergence. Furthermore, we also investigate how our results transfer to the cumulated regret.

Meanfield control problems: turnpike property and propagation of chaos

Katharina Eichinger

In this talk we provide an analysis for a class of meanfield 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 meanfield 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.

Stochastic control meets Sinkhorn on a Schrödinger bridge

Giacomo Greco

The Schrödinger bridge problem deals with finding the most likely evolution of a cloud of independent particles conditionally to the observation of the particles densities at an initial and final time. Remarkably this problem can be equivalently stated as an entropic regularization of the Kantorovich-OT problem and its solution can be efficiently computed via an iterative procedure known as Sinkhorn's algorithm. This made this algorithm and the Schrödinger bridge problem extremely popular in the ML community, particularly for their use in generative modelling, and because of that the theoretical study of the convergence rate of Sinkhorn's algorithm is still a very active area of research. Particularly, not much is known for unbounded marginals in unbounded domains. In this talk, after a brief introduction, 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 equation of value functions obtained from stochastic optimal control problems. Based on G. Greco, M. Noble, G. Conforti, A. Durmus - Proceedings of Thirty Sixth Conference on Learning Theory, PMLR 195:716-746, 2023.

Stochastic modified flows, mean-field limits and dynamics of stochastic gradient descent

Sebastian Kassing

We propose new limiting dynamics for stochastic gradient descent in the small learning rate regime called stochastic modified flows. These SDEs are driven by a cylindrical Brownian motion and improve the so-called stochastic modified equations by having regular diffusion coefficients and by matching the multi-point statistics. As a second contribution, we introduce distribution dependent stochastic modified flows which we prove to describe the fluctuating limiting dynamics of stochastic gradient descent in the small learning rate - infinite width scaling regime. This is joint work with Benjamin Gess and Vitalii Konarovskyi.

Fourier expansion of mean field control problems

Mattia Martini

The goal of this talk is to present a novel approach for approximating mean field control problems without relying on particle-based methods. Our primary tool is Fourier analysis, and for simplicity, we confine our discussion to the periodic setting. First, we introduce a family of approximated control problems that exhibit good properties in terms of computability of Fourier coefficients. Subsequently, our main result is that, under suitable regularity conditions, the optimal control, trajectory, and value function from the original problem converge to their counterparts within the approximation. Noticeably, we show that in the smooth setting our method yields a convergence rate directly proportional to the data's regularity. This convergence rate is faster than the one achieved by the usual particles approach, offering a more efficient alternative. Furthermore, our technique also provide an explicit method for constructing an approximate optimal control along with its corresponding trajectory. This talk is based on a joint work with François Delarue.

Optimal adaptive control with separable drift uncertainty

Alexander Merkel

We consider a problem of stochastic optimal control with separable drift uncertainty in strong formulation on a finite horizon. The drift coefficient of the state Y^u is multiplicatively influenced by an unknown random variable λ, while admissible controls u are required to be adapted to the observation filtration. Choosing a control actively influences the state and information acquisition simultaneously and comes with a learning effect. The problem, initially non-Markovian, is embedded into a higher-dimensional Markovian, full information control problem with control-dependent filtration and noise. To that problem, we apply the stochastic Perron method to characterize the value function as the unique viscosity solution to the HJB equation, explicitly construct ε-optimal controls and show that the values of strong and weak formulations agree. Numerical illustrations show a significant difference between the adaptive control and the certainty equivalence control. Joint work with Christoph Knochenhauer and Samuel Cohen.

Learning-by-doing for irreversible investment under incomplete information

Alessandro Milazzo

We study a problem of irreversible investment for a decision-maker who aims to invest into a project with unknown profitability. We introduce a notion of “learning-by-doing”: by increasing the investment into the project more information is acquired and, thus, the learning rate of the unknown profitability is improved. We formulate this problem as a singular control problem with incomplete information, which has a two-dimensional sufficient statistics comprised of the current level of investment together with the conditional probability that the profitability takes the larger value. Under some conditions on the learning procetodure, we show that the optimal investment strategy is to gradually invest into the project so that the sufficient statistics reflects along a monotone boundary. This boundary can be characterised as the solution of a differential problem.

Markovian randomized equilibria for general Markovian Dynkin games in discrete time

Berenice Neumann

We study a general formulation of the classical two-player Dynkin game in a Markovian discrete time setting. We show that an appropriate class of mixed, i.e., randomized, strategies in this context are Markovian randomized stopping times, which correspond to stopping at any given state with a state-dependent probability. One main result is an explicit characterization of Wald-Bellman type for Nash equilibria based on this notion of randomization. In particular, this provides a novel characterization for randomized equilibria for the zero-sum game, which we use, e.g., to establish a new condition for the existence and construction of pure equilibria, to obtain necessary and sufficient conditions for the non-existence of pure strategy equilibria, and to construct an explicit example with a unique mixed, but no pure equilibrium. We also provide existence and characterization results for the symmetric specification of our game. Finally, we establish existence of a characterizable equilibrium in Markovian randomized stopping times for the general game formulation under the assumption that the state space is countable.

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

Stefan Perko

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.

Ergodic mean-field games of singular control with regime-switching

Ioannis Tzouanas

This paper studies a class of stationary mean-field games of singular stochastic control with regime-switching. The representative agent adjusts the dynamics of a Markov-modulated Itô-diffusion via a two-sided singular stochastic control and faces a long-time-average expected profit criterion. The mean-field interaction is of scalar type and it is given through the stationary distribution of the population. Via a constructive approach, we prove the existence and uniqueness of the stationary mean-field equilibrium. Furthermore, we show that this realizes a symmetric ε-Nash equilibrium for a suitable ergodic N-player game with singular controls. The proof hinges on the characterization of the optimal solution to the representative player's ergodic singular stochastic control problem with regime switching, which is of independent interest and appears here for the first time. This talk based on the joint work with Jodi Dianetti and Giorgio Ferrari.

Bayesian sequential least-squares estimation in discrete time

Yuqiong Wang

We study the sequential estimation problem regarding an unknown parameter in an exponential family in the presence of observation cost. The problem is formulated using the mean squared error within a Bayesian setting, and it is further integrated into a Markovian framework. Building upon this, we provide sufficient conditions for the monotonicity of the value function in space. We further establish the concentration property of the posterior distribution of the unknown parameter and discuss its implication on the time-monotonicity of the value function along with the structure of the continuation regions. Aside from the obvious statistical applications, I will briefly comment on further applications of control problems with learning features.

Optimal retirement choice under age-dependent force of mortality

Shihao Zhu

This paper examines the retirement decision, optimal investment, and consumption strategies under an age-dependent force of mortality. We formulate the optimization problem as a combined stochastic control and optimal stopping problem with a random time horizon, featuring three state variables: wealth, labor income, and force of mortality. To address this problem, we transform it into its dual form, which is a finite time horizon, three-dimensional degenerate optimal stopping problem with interconnected dynamics. We establish the existence of an optimal retirement boundary that splits the state space into continuation and stopping regions. Regularity of the optimal stopping value function is derived and the boundary is proved to be Lipschitz continuous, and it is characterized as the unique solution to a nonlinear integral equation, which we compute numerically. In the original coordinates, the agent thus retires whenever her wealth exceeds an age-, labor income- and mortality-dependent transformed version of the optimal stopping boundary. We also provide numerical illustrations of the optimal strategies, including the sensitivities of the optimal retirement boundary concerning the relevant model's parameters.

List of Participants

Andrea Bovo	University of Turin
Matteo Buttarazzi	University Rome La Sapienza
Federico Cannerozzi	University of Milan
Felix Dammann	Bielefeld University
Robert Denkert	Humboldt University
Jodi Dianetti	Bielefeld University
Niklas Dexheimer	Aarhus University
Katharina Eichinger	École Polytechnique
Giorgio Ferrari	Bielefeld University
Fabian Fuchs	Bielefeld University
Giacomo Greco	Eindhoven University of Technology
Sebastian Kassing	Bielefeld University
Annika Kemper	Bielefeld University
Bingbing Li	Bielefeld University
Mattia Martini	Université Côte d'Azur
Alexander Merkel	Technische Universität Berlin
Alessandro Milazzo	University of Turin
Max Nendel	Bielefeld University
Anna Pajola	University of Milan
Berenice Anne Neumann	Trier University
Umberto Pappalettera	Bielefeld University
Stefan Perko	Friedrich-Schiller-University Jena
Frank Riedel	Bielefeld University
Alessandro Sgarabottolo	Bielefeld University
Maren Diane Schmeck	Bielefeld University
Ioannis-Paraskevas Tzouanas	Bielefeld University
Yuqiong Wang	Uppsala University
Shihao Zhu	Bielefeld University

Useful Information

Talks will be held in the room **V2-210-216** of the **Main Building (UHG)** at Bielefeld University (see below). The room is situated on the second floor and it is accessible from the main hall.

Coffee breaks will be offered in the room **V10-151**.

The **lunches** are not offered. Participants can find canteens (the Westend or the Mensa), restaurants and coffees inside the Main Building or in the nearby.

The **conference dinner** will be held at *The Bernstein Bielefeld*, Niederwall 2, 33602 Bielefeld.

How to get to Bielefeld University?

The workshop will take place at the room **V2-210-216** in the Main Building (UHG) at Bielefeld University, Universitätsstraße 25, 33615 Bielefeld, Germany.

To get there, you can follow the following instruction:
1. Take the **tram** "Stadtbahnlinie 4" (red line) in direction "Lohmannshof";
2. Get off at "Universität", and take the staircase to the bridge which leads from the tram station to the campus (ca. 5 minutes).