BOOK OF ABSTRACTS

Stochastic Games and Randomised Strategies

8-10 September 2025

Andi Bodnariu

Time-inconsistent singular control problems: Reflection and Absolutely continuous controls with exploding rates

We consider the game-theoretic approach to a time-inconsistent singular stochastic control problem for a one-dimensional diffusion. The time-inconsistency is due to the presence of a non-exponential discount function. We find equilibria in a novel class of (mild) controls that allow for explosion of an absolutely continuous control rate such that an inaccessible boundary is created for the controlled process. We provide a verification theorem giving sufficient as well as necessary conditions for equilibria in terms of a set of variational inequalities. We apply the theory in the study of an inventory problem. We show that for some parameter values the problem has a (strong) equilibrium in the form of Skorohod reflection. However, for some other parameter values there is no strong equilibrium, but we do find an equilibrium control given by an exploding rate that creates an inaccessible boundary for the controlled process.

Andrea Bovo

An optimal control problem arising in dynamic contracting with limited liability

We study a stochastic optimal control problem motivated by an application to a model of dynamic contracting (Principal-Agent problem) with limited liability. The problem is set on a finite-time horizon T and the principal designs a contract by controlling drift and diffusion of a stochastic differential equation living on $[\underline{x}, \infty)$. The solution of the SDE represents the current value of the contract and the agent can walk away from the contract if such value falls below $\underline{x}$, because of limited liability.

The original problem is connected to an auxiliary optimal control problem via a time-change. In the latter problem, the diffusive dynamics is uncontrolled whereas the time-dynamics is controlled. We show that the value function of this problem is the unique smooth solution of an associated HJB equation and we construct an optimal control. Then we transfer these results to the original problem by undoing the time-change. The methods of proof are almost entirely based on probabilistic estimates.

Galen Cao

Logarithmic regret in the ergodic Avellaneda–Stoikov market making model

We analyse the regret arising from learning the price sensitivity parameter κ of liquidity takers in the ergodic version of the Avellaneda–Stoikov market making model. We show that a learning algorithm based on a maximum-likelihood estimator for the parameter achieves the regret upper bound of order $\ln^2 T$ in expectation. To obtain the result we need two key ingredients. The first is the twice differentiability of the ergodic constant under the misspecified parameter in the Hamilton–Jacobi–Bellman (HJB) equation with respect to κ, which leads to a second–order performance gap. The second is the learning rate of the regularised maximum-likelihood estimator which is

obtained from concentration inequalities for Bernoulli signals. Numerical experiments confirm the convergence and the robustness of the proposed algorithm.

Yukun Cao

Exit Game with Asymmetric Information

We study a two-player, nonzero-sum exit game with asymmetric information. Each player privately observes a type determining their terminal payoff; while both remain active, they earn a duopoly running payoff, and if one exits, the other receives a monopoly running payoff until stopping. We extend Kwon–Palczewski (2024) by allowing a richer private type space under which the associated single-player monopoly benchmark problem may possibly admit a finite optimal stopping time. Our main results are: (i) the existence of a symmetric Nash equilibrium and showing that the belief process is continuous on $(0, \infty)$ with at most one jump at $t = 0$; and (ii) a timing comparison – the player who exits first exits strictly before the stopping time in the corresponding single-player monopoly problem.

Piotr Chlebicki

A game formulation of the (SSC) de Finetti problem with randomized strategies

We study a game formulation of resource extraction of a common good under one-dimensional diffusive dynamics with singular stochastic control up to absorption at 0, implying a trade-off between profitable resource extraction and sustainability. Unsurprisingly, immediate extraction of all available resources is an equilibrium. A main result is that we characterise and prove existence of several other, non-trivial, equilibria, that do not result in immediate absorption, but instead are attained with both players extracting resources according to a state-dependent rate. The assumption which guarantees this result states, roughly, that the drift coefficient of the uncontrolled state process grows sufficiently fast. We also study a generalization of the game, by allowing the instantaneous singular stochastic control reward to be state-dependent. In this game we again characterise and prove existence of several non-trivial equilibria. In particular, a main novelty is that we find equilibria where the state process is controlled with its own local time such that we have either reflection points or points where the control process increases in a singular manner (skew points).

Sören Christensen

From Randomised Stopping to Generation – On Time-Reversal in Diffusion Models

Randomisation is not only central to game theory, but also arises in many other areas of stochastic modeling. In this talk, we focus on randomised stopping times and show how they naturally arise in the theory of time-reversed Markov processes. We then explore their application to modern generative modeling.

We begin with a theoretical overview of time-reversed stochastic processes and the connection to Doob's h-transform. In particular, we discuss the connection to stochastic control and highlight the natural emergence of Markovian randomised stopping times in this context.

Building on these ideas, we introduce a new type of diffusion-based generative model. These models generate data by gradually removing randomness from a known random sample until a sample from a complex target distribution emerges. Unlike conventional models in this class, our approach employs time-homogeneous dynamics in both directions (adding and removing noise) and

uses a randomised stopping rule rather than a fixed number of steps. This leads to a flexible and efficient generative procedure that adapts to the data's intrinsic structure.

Our analysis reveals that this approach is particularly well suited for data concentrated near low-dimensional structures (manifolds), where the stopping rule reduces to a first-hitting time.

The talk is based on joint work with Claudia Strauch and Lukas Trottner (https://arxiv.org/abs/2501.19373).

Tiziano De Angelis

Nash equilibria for dividend distribution with competition

We construct Nash equilibria in feedback form for a class of two-person stochastic games of singular control with absorption, arising from a stylized model for corporate finance. More precisely, the paper focuses on a strategic dynamic game in which two financially-constrained firms operate in the same market. The firms distribute dividends and are faced with default risk. The strategic interaction arises from the fact that if one firm defaults, the other one becomes a monopolist and increases its profitability. The firms choose their dividend distribution policies from a class of randomised strategies and we identify two types of equilibria, depending on the firms' initial endowments. In both situations the optimal strategies and the equilibrium payoffs are found explicitly.

Joint work with Fabien Gensbittel and Stephane Villeneuve (Toulouse School of Economics)

Jean-Paul Décamps

The War of Attrition under Uncertainty: Theory and Robust Testable Implications

We study a symmetric-information war of attrition in which the players' payoffs depend on exogenous market conditions that evolve according to a homogeneous linear diffusion. Using that a Markov strategy can be represented as a stopping region along with an intensity measure of stopping, we fully characterize mixed-strategy Markov-perfect equilibria through a variational system for the players' value functions. When players are asymmetric, in any such equilibrium each player randomizes at a discrete set of thresholds for market conditions. As a result, players may alternatively find themselves in a position of strength or weakness on the equilibrium path. Delayed concessions occur because a player currently in a position of weakness can hope for market conditions to reverse in his favor. In the standard duopoly model of exit under uncertainty, the firms' stock prices and their return volatilities comove negatively over the attrition region and exhibit patterns documented by technical analysis.

Marina Di Giacinto

Executions in competition under Erlang kernel

We consider the problem of multiple order executions in competition as differential games and investigate their associated equilibria in two regimes: Stackelberg and Nash games. Price impact during execution comprises permanent, transient, with an Erlang-type delay kernel, and temporary or slippage components. Unlike purely decaying kernels such as exponential or power-law types, the Erlang-type kernel offers the flexibility to specify a peak impact from past trades at a finite lag. The resulting price dynamics remain Markovian in an extended, yet finite-dimensional, state space. This structure enables the use of the dynamic programming principle, and equilibria in the two differential games are obtained by solving systems of Riccati equations. Numerical experiments are conducted to illustrate the theoretical results.

Henri Doerks

How much better is sequential analysis than methods with fixed sample size?

While many statistical procedures rely on a fixed sample size, sequential methods allow for dynamic decision-making as data is collected. In this way, sequential analysis reduces the average number of observations required to achieve a given power of the test – but by how much? To answer this, we restrict our attention to the sequential probability ratio test for the drift of a Brownian motion. In this setting, we give precise bounds for the average reduction in sample size needed to obtain a certain power. More specifically, we show that the sequential test decreases the average sample size by at least 36% and by at most 75%. Moreover, the reduction is monotone increasing in the power of the test, so the advantage of a sequential test in comparison with a fixed sample size test is larger if a high power is sought-after.

Erik Ekström

Optimal search problems: when to follow the leader

We consider a three-dimensional process where two coordinates are standard Brownian motions and the remaining (unknown) coordinate is a Brownian motion with drift. For this process we study a search problem: if only one coordinate can be observed at each point in time, what strategy should be used to detect the drifting coordinate as accurately as possible within a given time-frame?

Fabien Gensbittel

Mixed Markov-Perfect Equilibria in the Continuous-Time War of Attrition

We prove the existence of a Markov-perfect equilibrium in randomized stopping times for a model of the war of attrition in which the underlying state variable follows a homogenous linear diffusion. The proof uses the fact that the space of Markovian randomized stopping times can be topologized as a compact absolute retract, which in turn enables us to use a powerful fixed-point theorem by Eilenberg and Montgomery. We illustrate our results with an example of a war of attrition that admits a mixed-strategy Markov-perfect equilibrium but no pure-strategy Markov-perfect equilibrium.

Pavel Gapeev

Discounted nonzero-sum optimal stopping games in diffusion-type models with maxima and minima

We present closed-form solutions to some nonzero-sum optimal stopping games of two players whose rewards depend on the running maximum or minimum processes of a generalised geometric Brownian motion and have exponential negative or positive discounting rates. It is assumed that the coefficients of the model depend on the current values of the running maximum and minimum of the diffusion-type process. The optimal stopping times forming Nash equilibria are shown to be the first times at which the underlying process hits certain boundaries depending on the running values the associated maximum and maximum processes. The proof is based on the reduction of the original games to the equivalent coupled free-boundary problems and the solution of the latter problems by means of the smooth-fit and normal-reflection conditions. We show that the optimal stopping boundaries are partially determined as solutions of either the appropriate systems of arithmetic equations or solutions of the appropriate first-order nonlinear ordinary differential equations. The

obtained results are related to the rational valuation of the perpetual lookback game options with fixed and floating strikes in the Black-Merton-Scholes model.

David Hobson

Non-zero sum games, convertible bonds and manipulation

We consider a non-zero sum game in the form of a convertible bond and use the Dayanik-Karatzas approach to reduce the optimal stopping problems to a geometric problem. Then we investigate the form of the solution. We discuss the form of Nash equilibria, (non)-uniqueness, pre-commitment and equilibria, and the failure of the order condition. Joint work with Edward Wang and Gechun Liang.

Anna Jaśkiewicz

Markov perfect equilibria in Markov decision processes with quasi-hyperbolic discounting

In this talk I present some results on the existence of Markov perfect equilibria for a large class of Markov decision processes with Borel state spaces under quasi-hyperbolic discounting. This type of discounting nicely models human behaviour, which is time-inconsistent in the long run. The decision maker has preferences changing in time. Therefore, the standard approach based on the Bellman optimality principle fails. I shall consider two cases: the first case, when the decision maker is risk-neutral and the second case when the decision maker is risk-averse.

Damian Jelito

Time-inconsistency in selected long-run and risk-sensitive control problems

In the context of stochastic control, time-inconsistency arises when a decision-maker's preferences change over time in such a way that a strategy deemed optimal at time $t = 0$ is no longer optimal when re-evaluated at $t > 0$. This phenomenon is typically linked to non-exponential discounting, endogenous habit formation, or non-linear optimality criteria.

In this talk, we consider two frameworks where time-inconsistency may appear. In the first, we examine the long-run impulse control problem under non-exponential discounting. Somewhat surprisingly, despite the non-standard form of discounting, we show that under mild conditions the problem remains time-consistent and is, in fact, equivalent to a simpler undiscounted formulation. In the second framework, we turn to the risk-sensitive optimal stopping problem with exponential discounting. Here, time-inconsistency can arise, and we analyze the corresponding time-consistent equilibria.

The results are based on joint work with Łukasz Stettner and ongoing research with Jan Palczewski.

Kristofer Lindensjö

Markovian randomized stopping times and stopping games

One of the most well-known results in classical game theory is that randomized strategies facilitate existence of Nash equilibria. This talk is about recent research regarding the role of randomized strategies for stopping games based on Markov processes. In particular, we are interested in defining Markovian randomized stopping strategies that allow for equilibrium existence as well as characterization similar to how non-randomized solutions to Markovian stopping problems can be characterized using e.g., variational inequalities (continuous time) and Wald-Bellman equations (discrete time). In particular:

We study a general formulation of the zero-sum Dynkin (stopping) game with an underlying one-dimensional diffusion, and present a complete equilibrium solution (existence and characterization) relying on the notion of Markovian randomized stopping times.

We study a general formulation of the classical two-player Dynkin game in a discrete time Markovian setting, and give an explicit equilibrium characterization of Wald-Bellman-type using (discrete time) Markovian randomized stopping times. In particular, we derive a novel characterization of randomized equilibria in zero-sum Dynkin games, which we use to (i) establish the existence and explicit construction of Markovian randomized equilibria, (ii) provide necessary and sufficient conditions for the non-existence of pure strategy equilibria, and (iii) construct an example that admits a unique randomized equilibrium but no pure one.

Based on joint works with Andi Bodnariu (Stockholm), Berenice Neumann (Trier), and Sören Christensen (Kiel).

Alessandro Milazzo

A journey through zero-sum singular-controller vs. stopper games

Zero-sum stochastic games involving a controller and a stopper were first introduced by Maitra and Sudderth in 1996, inspiring further research in the area. Since then the contributions to this type of games have grown steadily. This talk focuses on a specific class of such games, where the controller employs singular control strategies, and presents the results that have recently been obtained. These include the existence of a value for the game, the characterisation of the optimal strategies for the players (also under constrained control directions or degenerate diffusions) and the global regularity of the value function. These findings are achieved using a blend of analytical and probabilistic techniques.

Jan Palczewski

Equilibria in zero-sum stopping games with asymmetric information

Zero-sum stopping games (Dynkin games) are a natural generalisation of optimal stopping problems with a well-developed theory for construction of players' strategies. However, when players have access to different information flows (filtrations), little is known. A few specific games have been solved, but the existence of an equilibrium in randomised strategies was only established in 2022. In this talk, I will present our progress towards a general martingale theory for the characterisation of value processes and equilibrium strategies. Based on joint work with Tiziano De Angelis and Jacob Smith.

Jan-Henrik Steg

Strategic Irreversible Investment

This paper studies oligopolistic irreversible investment with closed-loop strategies. These permit fully dynamic interactions that result in much richer strategic behavior than previous studies with open-loop strategies allow. The tradeoff between preemption incentives and the option value of waiting becomes distinctly visible. Strategies that depend on present capital stocks enable credible reactions that deter from excessive preemption and support positive option values in equilibrium. Simpler strategies lead into a "preemption trap" with perfectly competitive outcome and zero net present values. To obtain these results, a novel concept of Markov perfect equilibrium is developed that copes with optimal investment taking the form of singular control.

Topias Tolonen-Weckström

Learning from the past in an irreversible investment problem

We consider an irreversible investment problem under incomplete information, where the investor is able to exercise multiple rights to invest into a project. Upon each investment, the investor acquires previously hidden information from the project's past. This phenomenon is modelled by allowing each investment induce a diffusion, realizations of which the investor uses to update their knowledge about the project ("learning from the past") in addition to observing a traditional observation process. We set up the model as an N-step optimal stopping problem with incomplete information, where each step collapses into a corresponding free-boundary problem. We present the optimal investing strategy as a sequence of stopping boundaries for each investment step. The boundaries are then solved numerically from their respective smooth fit conditions. When discussing the numerically solved boundaries, we provide some comparative statistics and discuss the optimal strategy when N tends to infinity.

Xavier Venel

Two-Player Zero-Sum Discrete-time Stochastic Games with a More Informed Player

We will present the general framework of two-player zero-sum discrete-time stochastic games with asymmetric information. The focus will be on the situation where one player is more informed than the other. Recent contributions have provided positive results on the existence of equilibrium in this framework. In particular, I will discuss the work of Venel, Gensbittel and Oliu-Barton (2014) on a more informed controller, and the work of Venel and Li (2015) on recursive games. We will see that the type of results that one can expect depends on the type of assumption.

Chaorui Wang

A Measure-Valued HJB Perspective on Bayesian Optimal Adaptive Control

We consider a Bayesian adaptive optimal stochastic control problem where a hidden static signal has a non-separable influence on the drift of a noisy observation. Being allowed to control the specific form of this dependence, we aim at optimising a cost functional depending on the posterior distribution of the hidden signal. Expressing the dynamics of this posterior distribution in the observation filtration, we embed our problem into a genuinely infinite-dimensional stochastic control problem featuring so-called measure-valued martingales. We address this problem by use

of viscosity theory and approximation arguments. Specifically, we show equivalence to a corresponding weak formulation, characterise the optimal value of the problem in terms of the unique continuous viscosity solution of an associated HJB equation, and construct a piecewise constant and arbitrarily-close-to-optimal control to our main problem of study. (Joint work with Alexander Cox and Sigrid Källblad).