TUMCAIR
TUM Center for Asset Management and Banking, Insurance, and Risk Management

15th International Workshop on Stochastic Models & Control

March 23 to 26, 2026
Tegernsee

Local Organizers

Christoph Knochenhauer (TUM)
Bettina Haas (TUM)

Program Committee

Nicole Bäuerle (Karlsruhe Institute of Technology)
Sören Christensen (University of Kiel)
Jörn Sass (RPTU Kaiserslautern-Landau)
Stefan Thonhauser (Graz University of Technology)
Ralf Wunderlich (Brandenburg University of Technology Cottbus-Senftenberg)

Program

Monday, March 23

17:30 – 19:00

Welcome Reception

NI BAR - The Rooftop Bar in Tegernsee

Head to the fifth floor of Caro & Selig Hotel to witness the highlight of the hotel: The roof terrace with the NI BAR, and a 360-degree panoramic view of Tegernsee and the Bavarian Alps. The NI BAR is named after Queen Caroline's favourite daughter.

Caro & Selig
Bahnhofstrasse 29
83684 Tegernsee

Program

Tuesday, March 24

08:45 - 09:00	Opening
09:00 - 09:40	Samuel Cohen
09:40 - 10:00	Jackson Hebner
10:00 - 10:20	Filippo de Feo
10:20 - 10:40	Christian Laudagé
10:40 - 11:10	Coffee Break
11:10 - 11:50	Eduardo Abi-Jaber
11:50 - 12:10	Alexander Kalinin
12:10 - 12:30	Laura Voß
12:30 - 12:50	Tobias Lausser
12:50 - 14:00	Lunch
14:00 - 14:40	Raul Tempone
14:40 - 15:00	Saifeddine Ben Naamia
15:00 - 15:20	Eya Ben Amar
15:20 - 15:50	Coffee Break
15:50 - 16:30	Michaela Hitz
16:30 - 16:50	Manuel Hasenbichler
16:50 - 17:10	Angel Moran Ledezma

Titles & Abstracts

09:00 – 09:40

Samuel Cohen University of Oxford

Neural Networks, PDEs and Control

Practical applications of stochastic control often depend on numerical computation of value functions and optimizing controls, typically through the solution of Hamilton-Jacobi-Bellman equations. The curse of dimensionality means that these equations are typically not numerically tractable and so, historically, many applications have depended on crude heuristic approaches for high dimensional systems. In this talk, we will consider recent developments in the use of neural networks as a class of function approximators, and their application to high dimensional control problems. We will explore how, when these methods can be applied, they work remarkably well, however they can also struggle in simple cases with low regularity due to their fundamental dependence on sampling in space. Based on joint work with Justin Sirignano, Deqing Jiang, and Jackson Hebner.

09:40 – 10:00

Jackson Hebner University of Oxford

Deep Hilbert Galerkin methods for PDEs on Hilbert spaces via derivative-informed operator learning with applications to stochastic optimal control of infinite-dimensional systems

Our previous research (joint with S. Cohen, F. de Feo, J. Sirignano) shows that Hilbert Neural Operators are able to approximate classical solutions of fully nonlinear second-order partial differential equations on Hilbert spaces, such as

Hamilton-Jacobi-Bellman and backwards Kolmogorov equations. Based on this result, we propose two actor-critic algorithms for solving Hilbert-valued HJB equations and two algorithms for solving Hilbert-valued backwards Kolmogorov equations. We then apply these algorithms to the control of the stochastic heat equation, a stochastic delay equation, and the stochastic Burgers equation. To the best of our knowledge, these algorithms are the first methods for solving PDEs directly on their whole Hilbert space domain.

10:00 – 10:20

Filippo de Feo TU Berlin

Derivative-informed Hilbert neural operators solve PDEs on Hilbert spaces and infinite-dimensional optimal control problems

We consider infinite-dimensional partial differential equations (PDEs) on separable Hilbert spaces with unbounded operators. These challenging equations arise in most applied sciences, e.g., as Kolmogorov PDEs and Hamilton-Jacobi-Bellman PDEs related to deterministic, stochastic, and controlled evolution equations (including PDEs and SPDEs, path-dependent deterministic and stochastic DEs, partially observed stochastic systems, and mean field systems) and functional differential equations. While a theoretical framework for these PDEs is well established in the literature, the development of numerical methods is an open area of research. In this talk, we provide a new theoretical analysis that rigorously justifies the development of Deep Galerkin methods, which will be presented in a companion talk. We start by parameterizing the solution of these PDEs via a Hilbert Neural Operator (HNO). We prove that HNOs can accurately represent classical solutions of these PDEs by showing new Universal Approximation Theorems for Frechet derivatives. Using these preliminary results, we show that HNOs approximately solve these PDEs. Finally, we consider optimal control problems of deterministic and stochastic evolution equations and we derive universal approximation results of optimal feedback controls in terms of our approximate solution HNO. Based on joint work Samuel Cohen, Jackson Hebner, and Justin Sirignano.

10:20 – 10:40

Christian Laudagé RPTU University Kaiserslautern-Landau

When risk defies order: On the limits of fractional stochastic dominance

Motivated by recent work on monotone additive statistics and questions regarding optimal risk sharing for return-based risk measures, we investigate the existence, structure, and applications of Meyer risk measures. Those are monetary risk measures consistent with so-called v-SD orders (based on a threshold utility function v) as refinement of second-order stochastic dominance (SSD). The test utilities defining the associated order are those at least as risk averse in absolute terms as v. The generality of v allows to subsume SSD and other examples from the literature. The structure of risk measures respecting the v-SD order is clarified by two types of representations. The existence of nontrivial examples is more subtle: for many choices of v outside the exponential (CARA) class, they do not exist. Additional properties like convexity or positive homogeneity further restrict admissible examples, even within the CARA class. We present impossibility theorems that demonstrate a deeper link between the axiomatic structure of monetary risk measures and SSD than previously acknowledged. If time permits, we discuss portfolio optimisation under a Meyer risk measure as objective. This is joint work with Felix-Benedikt Liebrich.

10:40 – 11:10 Coffee Break

11:10 – 11:50

Eduardo Abi Jaber Ecole Polytechnique

Hedging with Memory using Signatures

We show how signatures can be used for pricing and hedging derivatives in markets with memory using control-based approaches. Signatures provide a parsimonious and expressive representation of path dependence, allowing us to tackle inherently non-Markovian problems arising from stochastic volatility, path-dependent payoffs, and market frictions.

11:50 – 12:10

Alexander Kalinin LMU Munich

Stochastic Volterra equations with random functional coefficients in Banach spaces

We derive unique Banach-valued solutions to stochastic Volterra equations with random coefficients that may depend on pure chance and involve singular kernels. In particular, for controlled and distribution-dependent coefficients these solutions become strong, as a measurability analysis of the Wasserstein metric confirms. The presented novel approach is based on the proof that a stochastic Volterra integral admits a progressively measurable modification in a weak sense and on sharp moment estimates for non-negative product measurable processes.

12:10 – 12:30

Laura Voß Kiel University

Learning to Stop Multivariate Diffusions in Continuous Time via Nonparametric Estimation

We study optimal stopping for multivariate diffusion processes from a statistical perspective. Our approach is based on a structural representation of the value function as the minimum of a linear potential operator acting on non-negative drivers. This representation separates stochastic dynamics, payoff structure, and stopping geometry, and enables a direct statistical analysis of stopping rules via estimation of the value operator. The analysis proceeds in two steps. First, under full information, we establish a potential representation of the value function in terms of the martingale generator and an occupation measure formulation of optimal stopping. Second, in the data-driven setting, we construct nonparametric estimators of the value operator by estimating transition densities. A key feature of the analysis is that the resulting operator error bounds are uniform over a class of diffusion models and uniform over (t,x) in a prescribed domain, yielding genuinely nonparametric learning guarantees beyond a fixed-model setting.

12:30 – 12:50

Tobias Lausser Technische Universität München

Optimal investment under stochastic volatility models in discrete time

Continuous-time stochastic volatility models often yield elegant closed-form solutions for the problem of maximizing expected utility from terminal wealth. In practice, however, these solutions cannot be implemented exactly, and simple discretizations such as Euler–Maruyama schemes neither capture key stylized facts of financial return volatility nor preserve analytical tractability. In our contribution, we show that discrete-time stochastic volatility models based on GARCH dynamics provide a natural and tractable alternative. We derive approximate closed-form solutions to the expected utility maximization problem and identify a broad class of GARCH-type models for which such solutions exist. Our results establish a bridge between empirically successful GARCH volatility models and the classical continuous-time stochastic volatility literature based on diffusion processes.

12:50 – 14:00 Lunch Break

14:00 – 14:40

Raúl Tempone King Abdullah University

Belief-Space Stochastic Optimal Control with Discrete-Time Observations: Interlaced HJB and Pontryagin Principles

Continuous-time stochastic systems are controlled using intermittent sensor data, and in many applications, the sensing quality itself can be traded against control cost. This talk develops a unified view of continuous-time stochastic optimal control under partial observations available only at discrete times, combining continuous control of the dynamics with discrete-time design of the observation process (e.g., sensor/noise parameters). First, we lift the problem to the belief space: the conditional (filtering) distribution becomes the Markov state. Dynamic programming then yields interlaced Hamilton–Jacobi–Bellman (HJB) equations on this (infinite-dimensional) state space, coupled with Bayesian jump/impulse updates at observation times; in Gaussian-controlled models, invariance of the belief family yields finite-dimensional HJB equations for the mean and covariance, enabling computations. Second, we derive a Pontryagin Maximum Principle (PMP) on the belief space, accounting for continuous evolution between observations and for Bayes' rule jumps at observation instants. The resulting optimality system couples a forward filtering/Fokker–Planck equation with a backward adjoint equation and explicit jump conditions; under convexity and differentiability assumptions, the adjoint can be identified with the (flat) gradient of the value functional, linking the PMP and HJB viewpoints. We conclude with particle-based numerical schemes (particle filtering plus regression for the adjoint) and examples—linear–quadratic–Gaussian and nonlinear—that quantify how controllable noise levels reshape optimal strategies and highlight the benefits of active sensing.

14:40 – 15:00

Saifeddine Ben Naamia RWTH Aachen University

A Pontryagin Maximum Principle on the Belief Space for Continuous-Time Stochastic Optimal Control with Discrete Observations

This talk presents joint work with Christian Bayer, Saifeddine Ben Naamia, Erik von Schwerin, and Raúl Tempone on continuous-time stochastic optimal control under partial observations available only at discrete time instants. We formulate the problem on the space of beliefs, using the controller's posterior distribution as the state variable, and derive a Pontryagin maximum principle that accounts for continuous evolution between observations and Bayesian jump updates at observation times. A key result links the adjoint process to the gradient of the value functional on the belief space, connecting Pontryagin conditions with dynamic programming on the space of probability measures and yielding a predict–update structure related to the Zakai and Kushner-Stratonovich filtering equations. We also propose a particle-based numerical scheme for the coupled forward (filter) and backward (adjoint) system: particle filtering represents the belief, while regression approximates the adjoint, enabling computation of near-optimal controls under partial information. Numerical examples in linear and nonlinear settings illustrate the approach and highlight the benefits of actively controlling the observation process.

15:00 – 15:20

Eya Ben Amar King Abdullah University

Hierarchical Importance Sampling for Estimating Rare Event Probabilities in SDE Solutions

This study addresses the estimation of the complementary cumulative distribution function of the occupation time for a process governed by a stochastic differential equation. We focus on rare events in the right tail, where variance reduction is crucial for computational efficiency. Building on the connection between importance sampling (IS) and stochastic optimal control, we propose an optimal single-level IS (SLIS) estimator based on the solution of an auxiliary Hamilton–Jacobi–Bellman partial differential equation. The cost of solving this equation is included in the total computational work, and an optimal balance between preprocessing and sampling is derived. The approach is extended to a multilevel IS (MLIS) estimator, for which the single-level variance vanishes as the discretization level increases due to the zero-variance property of the optimal control. This leads to a necessary and sufficient condition under which MLIS outperforms SLIS. To ensure this condition, we introduce a smoothing of the drift and observable, together with a common-likelihood MLIS formulation that preserves variance decay. Numerical experiments on fade duration estimation confirm the theoretical results and demonstrate substantial efficiency gains over SLIS.

15:20 – 15:50 **Coffee Break**

15:50 – 16:30

Michaela Hitz Klagenfurt University

An optimal transport approach to quantifying model uncertainty of SDEs

A fundamental question in stochastic modelling is that of quantifying the effects of model uncertainty. In this context it is of interest to compute a distance between different stochastic models. A reasonable choice of distance is a modification of the Wasserstein distance on the space of probability measures called adapted Wasserstein distance, as it appears in bicausal optimal transport. We solve constrained optimal transport problems in which the marginal laws are given by the laws of solutions of stochastic differential equations (SDEs). We consider SDEs with irregular coefficients, making only minimal regularity assumptions. Numerical methods are employed as a theoretical tool to bound the adapted Wasserstein distance. This opens the door for computing the adapted Wasserstein distance in a simple way. We show that this method can be applied to quantifying model uncertainty in stochastic optimisation problems. Our approach successfully brings together optimal transport and numerical analysis of SDEs. This talk is based on joint work with Benjamin A. Robinson (University of Klagenfurt).

16:30 – 16:50

Manuel Hasenbichler Graz University of Technology

The Martingale Sinkhorn Algorithm

We develop a numerical method for the martingale analogue of the Benamou--Brenier optimal transport problem, which seeks a martingale interpolating two prescribed marginals which is closest to the Brownian motion. Recent contributions have established existence and uniqueness for the optimal martingale under finite second moment assumptions on the marginals, but numerical methods exist only in the one-dimensional setting. We introduce an iterative scheme, a martingale analogue of the celebrated Sinkhorn algorithm, and prove its convergence in arbitrary dimension under minimal assumptions. In particular, we show that convergence holds when the marginals have finite moments of order $p > 1$, thereby extending the known theory beyond the finite-second-moment regime. The proof relies on a strict descent property for the dual value of the martingale Benamou--Brenier problem. While the descent property admits a direct verification in the case of compactly supported marginals, obtaining uniform control on the iterates without assuming compact support is substantially more delicate and constitutes the main technical challenge.

16:50 – 17:10

Angel Moran Ledezma Karlsruhe Institute of Technology

Ultrametric Graphons and its application to hierarchical communities in graphs

We introduce the class of ultrametric graphons as a probabilistic framework for modeling random graphs with nested community structure. These graphons are generated from finite ultrametric trees endowed with a hierarchical partition of the unit interval, producing a block-constant kernel whose geometry encodes multiscale community organization. We prove that the associated graphon Laplacian admits an explicit closed-form description of its eigenvalues and eigenprojectors in terms of the underlying tree structure. Moreover, we show that for graphs sampled from the ultrametric graphon, the empirical spectrum converges with high probability to the deterministic spectrum of the limiting operator. This yields a precise characterization of the spectral scales associated with each hierarchical level. As applications, we demonstrate a phase transition in community detectability using spectral methods, where recovery becomes possible or impossible depending on the separation between hierarchical scales. We further establish almost sure convergence of the Moore–Penrose pseudoinverse of the sampled Laplacians toward its graphon counterpart, which allows us to derive limiting expressions for several quantities associated with random walks, including effective resistances, commute times, and diffusion distances. Finally, we investigate the impact of hierarchical community structure on the stability of epidemics in the SIS model, showing how multiscale organization modifies the epidemic threshold and influences metastable behavior.

Program

Wednesday, March 25

09:00 - 09:40	Roxana Dumitrescu
09:40 - 10:00	Philip Biegel
10:00 - 10:20	Sascha Desmettre
10:20 - 10:40	Lukas Eichhorn
10:40 - 11:10	Coffee Break
11:10 - 11:50	Christa Cuchiero
11:50 - 12:30	Sara Svaluto-Ferro
12:30 - 12:50	Jörn Sass
12:50 - 14:00	Lunch
14:00 - 17:00	Social Event
18:00 - 20:00	Workshop Dinner

Titles & Abstracts

09:00 – 09:40

Roxana Dumitrescu ENSAE, Institut Polytechnique de Paris

Fast-slow mean-field games with common noise

We propose a framework for approximating Nash equilibria in mean-field games (MFGs) with common noise based on a two-time-scale structure. In our model, the common noise is modeled by a fast variable evolving under ergodic dynamics. The framework applies to several classes of MFGs, including games with regular control and with optimal stopping. The main idea is to avoid solving the full MFG with common noise by approximating it with an “effective” MFG without common noise, whose coefficients are obtained by averaging with respect to the stationary measure of the fast-scale process. Starting from an equilibrium of the effective MFG, we construct an explicit ε-MFG equilibrium for the original game by introducing randomized control and stopping. To this end, we establish new existence results for MFG equilibria with randomized stopping. Our approach relies on convergence results for two-scale diffusions under various structural assumptions on the MFG, and we show that the time-scale separation parameter controls the error in the Nash equilibrium condition.

09:40 – 10:00

Philip Biegel RPTU Kaiserslautern-Landau

Expected Utility Spreads in a Black–Scholes Type Market with Drift Uncertainty

The literature advocates measuring model uncertainty for a payoff X by the width of its valuation interval across a family of models, i.e.

$$\mu_{\mathcal{Q}}(X) = \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[X] - \inf_{Q \in \mathcal{Q}} \mathbb{E}_Q[X].$$

In this talk we transfer this idea to continuous-time portfolio choice under drift ambiguity by introducing a utility spread

$$\Delta(\pi) = \sup_{\mu \in K} \mathbb{E}_\mu[U(X_T^\pi)] - \inf_{\mu \in K} \mathbb{E}_\mu[U(X_T^\pi)],$$

in a multivariate Black–Scholes market with drift uncertainty, where the drift ranges over an ellipsoidal set K. The first part of the talk focuses on structural properties of this spread. As a first result, under natural richness assumptions on the uncertainty set and non-degeneracy of the volatility, we show that $\Delta(\pi) = 0$ essentially forces investment exclusively in the risk-free asset. This motivates spread constraints as a quantitative control of uncertainty sensitivity. In the second part we discuss how $\Delta(\pi)$ can be incorporated into portfolio optimization, either as a primary objective (leading to "minimum ambiguity" allocations) or as a constraint that regularizes robust performance. We illustrate these ideas for different utility functions: log utility as a tractable benchmark with explicit solutions, and power utility which allows a richer analytical discussion. This is joint work with Jörn Saß.

10:00 – 10:20

Sascha Desmettre Johannes Kepler University Linz

Equilibrium investment under dynamic preference uncertainty

We study a continuous-time portfolio choice problem for an investor whose preferences are determined by an exogenous factor evolving as an Itô diffusion process. Uncertainty about future risk aversion leads to an objective that aggregates certainty equivalents across possible terminal preference states, generating inherent time-inconsistency. We address this problem using an equilibrium approach and derive an extended

Hamilton-Jacobi-Bellman system characterizing subgame-perfect investment strategies. In a tractable CRRA specification, the equilibrium portfolio admits a semi-explicit representation that decomposes into a myopic demand and a novel preference-hedging component that captures incentives to hedge against anticipated changes in risk aversion.

10:20 – 10:40

Lukas Eichhorn RPTU Kaiserslautern-Landau

Robustifying Expert Opinions for Portfolio Optimization with Partial Information

We study portfolio optimization under partial information in a mutivariate continuous-time financial market, where asset returns are driven by an unobservable stochastic drift. The drift is modeled as an Ornstein-Uhlenbeck process and is estimated by investors using both observed returns and external expert opinions. These expert opinions provide noisy signals about the current drift and are incorporated into the filtering framework via Kalman-style updates. We offer a base framework for integrating and robustifying expert opinions in continuous-time portfolio optimization and highlight the trade-off between informational efficiency and robustness under model uncertainty. In detail, recognizing that expert opinions may be biased or misspecified in practice, we introduce distributional uncertainty into the modeling of expert information. Specifically, we consider Wasserstein ambiguity sets around a nominal Gaussian distribution for the expert noise, allowing for potential systematic biases, deviations in higher moments and even different underlying distributions. The investor adopts a distributionally robust mean-square-error criterion and seeks an estimator that performs optimally against the worst-case distribution within the ambiguity set. Building on the distributionally robust estimation framework of Nguyen et al. (2021), we adapt their results to a filtering setting with expert opinions entering as noisy observations of the latent drift at deterministic time points. In contrast to the original setup, uncertainty is imposed solely on the distribution of the expert noise, while the prior distribution of the drift remains fixed and Gaussian. This leads to a reduced Wasserstein ambiguity set that is tailored to our model structure. We construct two conservative approximations of the distributionally robust estimation problem: one obtained by approximating the primal problem via restricting to affine estimators and enlarging the Wasserstein ambiguity set using the Gelbrich distance, and another derived from a conservative approximation of the corresponding dual minimax formulation by constraining to Gaussian distributions. For a Gaussian nominal distribution of the expert opinions, we show that these two approximations coincide and lead to the same optimal affine robust estimator yielding a minimax identity for the original robust mean-square-error estimation problem. Therefore the order of optimization can be interchanged: the investor's choice of the optimal estimator and nature's selection of the worst-case distribution within the Wasserstein ambiguity set commute. Finally, we analyze how the optimal robust estimator behaves as the Wasserstein radius increases, illustrating the resulting trade-offs between estimation accuracy and robustness.

10:40 – 11:10 Coffee Break

11:10 – 11:50

Christa Cuchiero Vienna University

Dynamic universal approximation and modeling with neural and signature SDEs

Motivated by generative AI inspired modeling, we investigate the mathematical foundations of neural and signature stochastic differential equations (SDEs). In these models, the coefficients of the SDE are parameterized either by neural networks or by functions of the path signature, leveraging their well-known universal approximation properties. While such classical universal approximation results concern static functions, our focus is on dynamic universal approximation at the level of the SDE solutions. We show that any standard SDE as well as path-dependent SDE can be approximated by a suitable neural or signature SDE, respectively. For signature SDEs, this requires the development of a dedicated well-posedness theory. Building on this foundation, we prove in particular that SDEs whose coefficients are entire functions of the signature serve as dynamic universal approximators for path-dependent SDEs. Since these signature SDEs exhibit an affine structure with respect to the lifted signature state, this result implies universality of affine processes within the class of Itô-processes. We also outline how a similar approach can be used to tackle generic path-dependent optimal control problems. The talk is based on joint works with Kristoffer Andersson, Tomas Carrondo, Eva Flonner, Alessandro Gnoatto, Paul Hager, Kurt Kevin, and Sara Svaluto-Ferro.

11:50 – 12:30
Sara Svaluto-Ferro University of Verona
Signature-based models: theory, calibration, and expansions
Signature methods provide a non-parametric approach to extracting features from trajectories, offering versatile applications in finance. The structure of signature components enables the use of advanced mathematical tools and the construction of highly general models capable of capturing diverse behaviors. In this talk, we introduce the concept of the signature and its key properties, illustrating its potential through

two financial applications. The first application examines the local-in-time expansion of a continuous-time process and its conditional moments, including the characteristic function. By leveraging the time-extended Itôsignature—composed of iterated integrals of deterministic and stochastic signals (time, multiple correlated Brownian motions, and compound Poisson processes)—we derive automated expansions to any order with explicit coefficients. This provides stochastic representations suitable for asymptotic analysis in the short-time limit. The second application focuses on a stochastic volatility model where volatility dynamics are described by linear functions of the (time-extended) signature of a primary process. When this process is of polynomial type, its truncated signature retains this structure, allowing for closed-form expressions of the squared VIX. By incorporating the Brownian motion driving the stock price, both the log-price and the squared VIX can be expressed linearly in terms of the signature of the augmented process, achieving highly accurate calibration results for SPX and VIX options. Joint work with Federico Bandi, Roberto Renò, Christa Cuchiero, Guido Gazzani, Janka Möller.

12:30 – 12:50

Jörn Saß RPTU Kaiserslautern-Landau

Regulation of Emission Trading Systems

Emission trading systems (ETS) provide an instrument to control greenhouse gas emissions. In an ETS the desired amount of emissions in a predefined time period is fixed in advance; and allowances for one ton of greenhouse gas emissions are handed out or auctioned to companies which underlie the system. Emissions, which are not covered by an allowance, are subject to a penalty at the end of the time period. Companies may trade allowances such that, ideally, emissions are reduced where this can be done at lowest costs. For a company, a continuous-time stochastic control problem arises by controlling the emissions via their abatement rate in order to minimize the costs that arise from these actions and from the penalty payments. The regulator of an ETS aims at the compliance with the given emission target. To check this, an additional stochastic differential equation (SDE) needs to be solved, in which the abatement rate enters as drift, resulting in a jump at terminal time. We show existence and uniqueness in this case and look at a simple model in which the Hamilton-Jacobi-Bellman equation and the SDE can be solved analytically. Then we extend the model by considering separate time periods in which the regulator may allow to transfer unused allowances to the subsequent time period (banking). We analyze two modeling approaches: First, a basic approach treating unused allowances as constant, which is analytically more tractable and computationally less costly. Second, a more realistic two-dimensional approach with stochastic dynamics for the unused allowances. Numerical results show that banking increases the probability of complying with the emissions cap. In combination with appropriate penalties, it essentially guarantees compliance and thus is a crucial policy for a regulator to improve the effectiveness and the reliability of an ETS. This is joint work with Ricarda Rosemann.

12:50 – 14:00 Lunch Break

Social Event

Meeting Point: Lobby Caro & Selig
Departure: 14:00

Event as booked:

or

Return: 17:00

Workshop Dinner

Meeting Point: Lobby Caro & Selig
Departure: 17:45

Program

Thursday, March 25

09:00 - 09:40	Mete Soner
09:40 - 10:00	Felix Höfer
10:00 - 10:20	Sebastian Höfer
10:20 - 10:40	Sebastian Zimper
10:40 - 11:10	Coffee Break
11:10 - 11:50	Anna Jaśkiewicz
11:50 - 12:10	Athanasios Vasileiadis
12:10 - 12:30	Florian Döttling
12:30 - 12:50	Alexander Schütt
12:50 - 14:00	Lunch
14:00 - 14:40	Jan Palczewski
14:40 - 15:00	Philip Le Borne
15:00 - 15:20	Sven Karbach
15:20 - 15:50	Coffee Break
15:50 - 16:30	Johannes Muhle-Karbe
16:30 - 16:50	Gemma Sedrakjan
16:50 - 17:10	Closing

Titles & Abstracts

09:00 - 09:40

Mete Soner Princeton University

Finite and Mean-Field Games with Discrete State Spaces

Differential games with a large or infinite number of players generally have multiple Nash equilibria. Their analysis and numerical resolution in high dimensions or in the Wasserstein space introduce challenging questions. When the underlying state space is discrete, however, several of these difficulties become more tractable. Leveraging this structure, one can establish general existence results without monotonicity assumptions and design efficient algorithms for computing equilibria. In this talk, I consider finite and mean-field games in these structures both in continuous and discrete time. Remarkably, in continuous time starting from a given distribution, finite player games always have unique Markov Perfect Equilibrium (MPE). Moreover, Picard type iterations converge to this Nash equilibrium. Some of these properties also hold in discrete time when the step size is small. Again in this set-up with general common noise, one can prove the existence of a MPE using the related forward-backward stochastic differential equations. This is joint work with Felix Hoefer of Princeton and Atilla Yilmaz of Temple.

09:40 – 10:00

Felix Hoefer Princeton University

Markov Perfect Equilibria in Discrete Finite-Player and Mean-Field Games

We study dynamic finite-player and mean-field stochastic games within the framework of Markov perfect equilibria (MPE).

Through the Kuramoto mean-field game, we investigate the emergence of multiple self-organizing equilibria with complex time-dependent dynamics. This is in contrast to continuous-time finite-player games, which admit unique MPE in the absence of monotonicity conditions. While discrete- time problems generally do not admit unique MPE, we show that uniqueness is remarkably recovered when the time steps are small. This result, established without relying on any monotonicity conditions, underscores the importance of inertia in dynamic games. This is joint work with Mete Soner (Princeton) and Atilla Yilmaz (Temple).

10:00 – 10:20

Sebastian Höfer Karlsruher Institut für Technologie

Continuous-Time Mean-Field Markov Decision Models

For many Markovian decision problems, it is reasonable to consider several statistically equal decision makers operating simultaneously on the same state space and interacting with each other (e.g., maintenance of identical machines in a production site, population of potentially infected persons). Depending on the model, the state transition and the profit of the individual may depend on the empirical distribution of the decision makers across the states. In the limiting case, as the number of decision makers tends to infinity, we show that the resulting mean-field model describes a classical deterministic control problem, for which the limit state process is characterized by a controlled ordinary differential equation. We show that an optimal control of the mean-field model yields an asymptotically optimal control for the model with N decision makers. In the end we discuss some applications.

10:20 – 10:40

Sebastian Zimper Freie Universität Berlin, Zuse Institue Berlin

Mean-field optimal control with stochastic leaders

We study the stochastic optimal control of a population of agents influenced by a single leader, for a finite number of agents as well as in the mean-field limit. Both the dynamics of the agents and the leader are driven by noise. The mean-field limit is described by a conditional McKean–Vlasov SDE coupled to the leader's dynamics, which is equivalent to a (nonlinear Fokker–Planck) PDE–SDE system. We generalize the connection between the finite-agent and mean-field optimal control problems, showing that, for multiplicative noise and stochastically evolving leaders, the optimal control of the finite system converges to that of the mean-field limit (when the state space is compact), in a Gamma-convergence result. Moreover, we present a numerical algorithm for approximating the optimal control of the limiting PDE–SDE system by first discretising the spatial component of the PDE and then adapting a gradient descent-based method for solving optimal control problems of finite systems. As an application, we study how agents can be brought to consensus by the intervention of a leader in a prototypical opinion dynamics system, the Hegselmann–Krause model. This is a joint work with Natasa Djurdjevac Conrad, Ana Djurdjevac, Carsten Hartmann and Christof Schütte.

10:40 – 11:10 Coffee Break

11:10 – 11:50

Anna Jaśkiewicz Wrocław University of Science Technology

Deterministic semi-Markov strategies in multiple objective Markov decision processes

This talk is based on the joint paper with A.S. Nowak and P. Więcek. I define the performance set in a Markov decision process with Borel state and action spaces and expected total reward criteria. The transition probabilities may have atoms but the initial state distribution is atomless. The main result, based on the Dvoretzky-Wald-Wolfowitz theorem from statistical decision theory, states that every performance vector determined by a general strategy can be obtained by using an appropriate deterministic semi-Markov strategy. An example shows that deterministic Markov strategies are not enough, when the transition probabilities are arbitrary and the initial state distribution is atomless. All papers in the literature proving the sufficiency of deterministic Markov or stationary strategies in Markov decision processes are based on the assumption that both transition probabilities and initial state distribution are atomless.

11:50 – 12:10

Athanasios Vasileiadis Karlsruhe Institute of Technology

Markov Decision Processes of the Third Kind: Learning Distributions by Policy Gradient Descent

Distributional Reinforcement Learning (DRL) has been a recent generalization of classical RL where instead of focusing on the conditional expectation of a (discounted) cumulative reward given the initial state, the classical value function of control literature, we are interested in the whole distribution of the terminal cumulative reward. Our goal in this talk is to develop a new Policy Gradient algorithm to find optimal controls that match this distribution of the terminal cumulative reward with a target one. We will present a probabilistic convergence proof and illustrate with numerical examples.

12:10 – 12:30

Florian Döttling Karlsruhe Institute of Technology (KIT)

A Hot Topic: Modeling Prosumer Heat Storage with a Markov Decision Process

Heat energy management plays an important role in the ongoing efforts to reduce carbon emissions. Due to the widespread use of solar collectors, many households are no longer just consumers of thermal energy but also producers (short: prosumers). A major challenge such prosumers face is the efficient use of the collected heat energy. Help is provided by technologies like long-term geothermal storage solutions and district heating networks. However, these more advanced systems raise questions about their optimal operation, for example, what amount of excess production should be stored long-term/sold vs. what amount should be kept on hand? This presentation aims to answer some of these questions by modeling prosumer heat storage with a Markov decision process, by deriving the structure of the optimal policy and by computing the optimal policy via learning algorithms. The talk is based on joint work with Nicole Bäuerle and funded by the Bundesministerium für Forschung, Technologie und Raumfahrt (BMFTR).

12:30 – 12:50

Alexander Schütt Technische Universität München

Deep Duality Methods for Constrained Optimal Portfolios

We present a deep learning-based approach for solving constrained optimal portfolio problems in high-dimensional markets. Our approach combines the duality framework of

Cvitanić and Karatzas (1992) with modern Physics-Informed Neural Networks (PINNs). The objective is to maximize the expected utility of terminal wealth, subject to convex constraints on the trading strategy. We solve the primal problem using a PINN to obtain an approximate optimal trading strategy. The solution to the corresponding dual problem provides a natural upper bound on the optimal performance under constraints. Comparing the performance of the learned trading strategy to this upper bound leads to an explicit and interpretable error estimate of its maximal deviation from the optimal strategy. Numerical experiments show the accuracy of this method in S&P 500-inspired 503-dimensional Black-Scholes and Heston markets, as well as in a rough Heston market model.

12:50 – 14:00 **Lunch Break**

14:00 – 14:40

Jan Palczewski Politechnika Wrocławska

When to Stop and Who Knows What? Zero-sum Optimal Stopping Games with Secrets

Imagine two players observing a Brownian motion on (0,1), each choosing when to stop it so as to optimise a payoff. Only one player possesses crucial private information – for instance, whether the process will eventually exit through 0 or 1. This simple setting already captures the key features of optimal stopping games with asymmetric information, including learning and the necessity of randomised strategies, and will serve as a guiding example for the talk. In the talk, I will present recent results on the existence of Nash equilibria (saddle points) of general zero-sum stopping games with asymmetric information, together with the martingale characterisation of equilibrium value processes and their links with differential equations and stochastic analysis. These results are part of a long-term research programme carried out with co-authors aimed at developing a systematic theory of Dynkin games beyond the symmetric information setting.

14:40 – 15:00

Philip Le Borne Kiel University

Learning to steer with Brownian noise

In this talk we consider an ergodic version of the bounded velocity follower problem, assuming that the decision maker lacks knowledge of the underlying system parameters and must learn them while simultaneously controlling. We propose algorithms based on moving empirical averages and develop a framework for integrating statistical methods with stochastic control theory. Our primary result is a logarithmic expected regret rate. Concluding we will compare our algorithm with standard deep learning results.

15:00 – 15:20

Sven Karbach University of Amsterdam

Semi-Static Variance-Optimal Hedging of Correlation Risk

We introduce variance-optimal hedging strategies for multi-asset contingent claims within the framework of multivariate stochastic covariance models. The proposed strategies are semi-static, combining continuous trading in the underlying assets with static positions in a set of auxiliary instruments. We solve the associated pricing and hedging problem via a multivariate Galtchouk–Kunita–Watanabe decomposition and show that the global mean–variance optimization naturally decouples into an inner dynamic and an outer static quadratic optimization problem. Moreover, we derive static replication formulas for multi-asset options, which provide natural candidates for the auxiliary hedging instruments. As an application, we consider covariance swaps and show that they can be hedged semi-statically by trading in the underlying assets and a suitable family of quanto options. For the broad class of affine stochastic covariance models, we obtain semi-closed-form expressions for optimal hedge ratios and prices. These analytical representations allow us to quantify the variance reduction attributable to the static component of the hedge and we demonstrate the tractability and practical relevance of the framework through numerical experiments for a range of exotic multi-asset derivatives, including covariance swaps.

15:20 – 15:50 Coffee Break

15:50 – 16:30

Johannes Muhle Karbe
Imperial College London

A unified theory of order flow, market impact and volatility

We propose a microstructural model for the order flow in financial markets that distinguishes between exogenous core orders and endogenous reaction flow, both modeled as Hawkes processes. This model has natural scaling limit that reconciles a number of robust yet apparently contradictory properties of empirical data: persistent signed order flow, rough trading volume and volatility, and power-law market impact. In our framework, all these quantities are pinned down by a single statistic H, which measures the persistence of the core flow. The signed flow converges to the sum of a fractional process with Hurst index H and a martingale, while traded volume is a rough process with Hurst index H-1/2. From these results, no-arbitrage constraints enable us to deduce that corresponding price volatility is rough, with Hurst parameter 2H-3/2, and that the price impact of trades follows a power law with exponent 2-2H. With H ~ 3/4, this model therefore is not only consistent with the square-root law of market impact, but also matches empirical estimates for signed order flow, unsigned trading volume and volatility remarkably well. (Joint work with Youssef O. Chahdi, Mathieu Rosenbaum, and Gregoire Szymanski)

16:30 – 16:50

Gemma Sedrakjan Technical University Berlin

Trading with the flow: Optimal execution and liquidity provision in a stylized limit order book model

We propose a stylized limit order book model in which price and liquidity dynamics are fully determined by the flow of limit and market orders. The presence of a strategic trading agent requires the imposition of conditions that preclude manipulative strategies, such as risk-free profitable round trips. Interestingly, these conditions turn out to reflect standard market practice and include, for example, the prohibition of wash trades. Next, we introduce a Poisson process to model the exogenous order flow, where order arrival rates depend on available liquidity, thereby creating a feedback mechanism that endows the system with resilience in supply and demand. We formulate the agent's expected utility maximisation problem, controlled by her decisions regarding liquidity provision and the use of market orders. To determine optimal investment and execution strategies, we derive the corresponding Hamilton–Jacobi–Bellman quasi-variational inequality. This framework allows us to investigate the time-varying use of limit and market orders. In particular, it gives rise to real-world market phenomenon such as market making, trend following and reactions volume imbalance.