Program and Abstracts

ZiF Focus Group

March 2026

ZiF

UNIVERSITÄT BIELEFELD

Information

ZiF Focus Group

Center for Interdisciplinary Research (ZiF), Bielefeld

March 02 – 13, 2026

Location:
Zentrum für interdisziplinäre Forschung
Universität Bielefeld
Methoden 1
33615 Bielefeld
Germany

Please note that all sessions will be held in the **Long Table Room**, with the exception of the Media Training Event on Friday the 13th, which will take place in the Nexus (old library).

Visit the ZiF website for travel information and a detailed map:
https://www.uni-bielefeld.de/einrichtungen/zif/about/contact/

Quantifying and Managing Uncertainty (QUAMU) advances economics, mathematics, and statistics by developing innovative methods to address emerging challenges. These methods have relevance well beyond economics, extending to fields such as data science, sports, psychology, and biology, where quantifying and managing uncertainty is critical.

Organizers:

Frank Riedel (Bielefeld University) – frank.riedel@uni-bielefeld.de

Maren Schmeck (Bielefeld University) – maren.schmeck@uni-bielefeld.de

Gerrit Bauch (Bielefeld University) – gerrit.bauch@uni-bielefeld.de

Visit our conference website:
https://www.uni-bielefeld.de/forschung/profil/fokusbereiche/quamu/focusgroup/

Week 1 (March 2–6, 2026)

Monday, March 2nd QUAMU Development (internal)	Tuesday, March 3rd Health and Insurance	Wednesday, March 4th Data Science and OR	Thursday, March 5th Sports	Friday, March 6th Teambuilding Event (internal)
	9.30 am – 11.00 am 9.30 – 10.15: An Chen 10.15 – 11.00: Steven Vanduffel	**9.30 am – 10.30 am** 9.30 – 10.30: Eyke Hüllermeier	**9.30 am – 11.00 am** 9.30 – 10.15: Raphael Flepp 10.15 – 11.00: Carl Singleton	**10.00 am – 11.00 am** Teambuilding Event - Part I
10.00 am – 11.00 am QUAMU General Meeting		**10.30 am – 11.00 am** Coffee Break		
11.00 am – 11.30 am Coffee Break	**11.00 am – 11.30 am** Coffee Break	**11.00 am – 1.00 pm** 11.00 – 11.40: Marília Barandas 11.40 – 12.20: Paul-Christian Bürkner 12.20 – 1.00: Bernhard Kasberger	**11.00 am – 11.30 am** Coffee Break	**11.00 am – 11.30 am** Coffee Break
11.30 am – 1.00 pm QUAMU General Meeting	**11.30 am – 1.00 pm** 11.30 – 12.15: Ondřej Zíka 12.15 – 1.00: Vanessa Brown		**11.30 am – 1.00 pm** 11.30 – 12.15: Rouven Michels 12.15 – 1.00: Christian Deutscher	**11.30 am – 1.00 pm** Teambuilding Event - Part II
1.00 pm – 2.00 pm Lunch Break	**1.00 pm – 2.00 pm** Lunch Break	**1.00 pm – 2.00 pm** Lunch Break	**1.00 pm – 2.00 pm** Lunch Break	**1.00 pm – 2.00 pm** Lunch Break
	2.00 pm – 3.30 pm 2.00 – 2.45: Stéphane Loisel	**2.00 pm – 3.00 pm** 2.00 – 3.00: Peter Grünwald	**2.00 pm – 3.30 pm** Panel Discussion	
	2.45 pm – 3.30 pm Panel Discussion	**3.00 pm – 4.00 pm** Panel Discussion		
	3.30 pm – 4.00 pm Coffee Break		**3.30 pm – 4.00 pm** Coffee Break	
	4.00 pm – 6.00 pm Research Groups	**4.00 pm – 4.30 pm** Coffee Break	**4.00 pm – 6.00 pm** Research Groups	
		4.30 pm – 6.00 pm Research Groups		
	6.00 pm – 8.00 pm Dinner	**6.00 pm – 8.00 pm** Joint Dinner with other Focus Groups	**6.00 pm – 8.00 pm** Dinner	

Week 2 (March 9–13, 2026)

Monday, March 9th Numerics & Learning	Tuesday, March 10th Social & Pol. Dynamics	Wednesday, March 11th QUAMU General Meeting (internal)	Thursday, March 12th ZiF Resident Group (internal)	Friday, March 13th Media Training (internal)
9.30 am – 11.00 am 9.30 – 10.15: Stefan Ankirchner 10.15 – 11.00: Yufei Zhang	**9.30 am – 11.00 am** 9.30 – 10.15: Jurek Preker 10.15 – 11.00: Gergely Horváth	**10.00 am – 11.00 am** QUAMU General Meeting	**10.00 am – 11.00 am** ZiF Resident Group	**9.00 am – 9.15 am** Welcome **9.15 am – 10.00 am** Science communication and content strategy (Jörg Heeren)
11.00 am – 11.30 am Coffee Break	**11.00 am – 11.30 am** Coffee Break	**11.00 am – 11.30 am** Coffee Break	**11.00 am – 11.30 am** Coffee Break	**10.00 am – 11.00 am** Media Training - Part I (Xenia Raufeisen)
11.30 am – 1.00 pm 11.30 – 12.15: Xinyu Li 12.15 – 1.00: Tsiry Randrianasolo	**11.30 am – 1.00 pm** 11.30 – 12.15: Manuel Förster 12.15 – 1.00: Ron Peretz	**11.30 am – 1.00 pm** QUAMU General Meeting	**11.30 am – 1.00 pm** ZiF Resident Group	**11.00 am – 11.30 am** Coffee Break
1.00 pm – 2.00 pm Lunch Break	**1.00 pm – 2.00 pm** Lunch Break	**1.00 pm – 2.00 pm** Lunch Break	**1.00 pm – 2.00 pm** Lunch Break	**11:30 am – 1:30 pm** Media Training - Part II - IV (Xenia Raufeisen)
2.00 pm – 2.45 pm 2.00 – 2.45: Jodi Dianetti	**2.00 pm – 3.30 pm** Panel Discussion			**1:30 pm – 2:30 pm** Lunch Break
2.45 pm – 3.30 pm Free Discussion				
3.30 pm – 4.00 pm Coffee Break	**3.30 pm – 4.00 pm** Coffee Break			
4.00 pm – 6.00 pm Research Groups	**4.00 pm – 6.00 pm** Research Groups			
6.00 pm – 8.00 pm Dinner	**6.00 pm – 8.00 pm** Dinner			

ABSTRACTS

Stefan Ankirchner

Friedrich-Schiller-Universität Jena

Learning optimal search strategies

Abstract:

We explore the question of how to learn an optimal search strategy within the example of a parking problem where parking opportunities arrive according to an unknown inhomogeneous Poisson process. The optimal policy is a threshold-type stopping rule characterized by an indifference position. We propose an algorithm that learns this threshold by estimating the integrated jump intensity rather than the intensity function itself. We show that our algorithm achieves a logarithmic regret growth, uniformly over a broad class of environments. Moreover, we prove a logarithmic minimax regret lower bound, establishing the growth optimality of the proposed approach. The talk is based on joint work with Maximilian Thiel.

Maília Barandas

Fraunhofer Portugal AICOS

Uncertainty Quantification Across the Machine Learning Lifecycle

Abstract:

Uncertainty quantification (UQ) plays a central role in reliable machine learning, yet in practice it is frequently handled at the level of individual models rather than embedded in ML methodology. This talk frames UQ as a structural element of the iterative machine learning lifecycle, spanning model development, evaluation, deployment, and human interaction. The presentation introduces a practical uncertainty library designed to compute a diverse set of predictive uncertainty metrics and to generate lightweight audit reports. Uncertainty is then discussed from three complementary perspectives. First, a method-oriented perspective, where uncertainty informs model design, training procedures, and evaluation strategies. Second, an application-oriented perspective, highlighting how uncertainty-aware models improve robustness, risk management, and decision quality in real-world deployments. Finally, a human-oriented perspective, focusing on how uncertainty shapes trust, interpretability, and interaction between humans and machine learning systems. Overall, the talk treats UQ as a lifecycle-level variable rather than a model-specific add-on.

Vanessa Brown

The University of Texas at Austin

Using the explore/exploit dilemma to understand the role of uncertainty in anxious avoidance

Abstract:

The explore/exploit dilemma refers to the need to balance exploring uncertain, potentially rewarding options versus exploiting known good ones. Whether the value-maximizing choice is to explore or exploit depends on many contextual factors, including the reward distribution, choice horizon, and type of uncertainty present in the environment. I will first present work showing that humans adaptively scale exploration based on cognitive demands, specifically working memory. Next, I will connect exploration under uncertainty to avoidance in anxiety, and discuss results showing that people with anxiety fail to adaptively scale exploration based on the current context. These findings are in line with theories suggesting that anxiety is related to difficulties learning the correct uncertainty statistics of the environment, rather than competing theories like loss aversion.

Paul-Christian Bürkner

TU Dortmund

Bayes in the age of machine learning

Abstract:

Bayesian inference is one of the prime methods for uncertainty quantification. It also allows to incorporate prior knowledge into statistical and machine learning models, thereby providing additional information sources independent of the data at hand. Due to these features, coupled with algorithmic and computational innovations, Bayes has seen a huge uptick in popularity in recent decades. Yet, as the age of machine learning progresses, the place of Bayes in the landscape of data-based learning methods seems to be threatened yet again. Increasingly complex deep learning models learn from immense amounts of data to perform tasks previously thought impossible to achieve by a machine. Due to the sheer size of these models and data, standard Bayesian inference is neither possible nor sensible in these contexts. Will Bayes thus become obsolete? In my talk, I will address this question highlighting where Bayes will continue to play important roles, even becoming more important and useful in the future.

An Chen

Ulm University

Ambiguity about Future States: Insurance and Care Decisions

Abstract:

Economic decisions about consumption, insurance, and care are often made under deep uncertainty about future states rather than well-defined risks. This talk examines how ambiguity—uncertainty about probabilities themselves—shapes behavior in insurance and care-related decisions. I first study optimal contingent consumption under smooth ambiguity, showing that ambiguity aversion induces endogenous shifts in subjective beliefs and effective intertemporal trade-offs. These mechanisms provide new insights into insurance puzzles such as underinsurance and low annuitization. I then analyze long-term care insurance and informal family care in a parent–child framework where future care needs are uncertain and health outcomes are endogenous. Unlike deterministic settings, ambiguity activates new incentive channels, reducing informal care provision while increasing demand for long-term care insurance. The results highlight the central role of ambiguity about future states in shaping insurance demand and care arrangements.

Jodi Dianetti

University of Rome Tor Vergata.

Reinforcement learning for exploratory optimal stopping: A singular control formulation

Abstract:

In this talk we discuss continuous-time and state-space optimal stopping problems from a reinforcement learning perspective. We begin by formulating the stopping problem using randomized stopping times, where the decision maker's control is represented by the probability of stopping within a given time–specifically, a bounded, non-decreasing, càdlàg control process. To encourage exploration and facilitate learning, we introduce a regularized version of the problem by penalizing it with the cumulative residual entropy of the randomized stopping time. The regularized problem takes the form of an (n+1)-dimensional degenerate singular stochastic control with finite-fuel. We address this through the dynamic programming principle, which enables us to identify the unique optimal exploratory strategy. For the specific case of a real option problem, we derive a semi-explicit solution to the regularized problem, allowing us to assess the impact of entropy regularization and analyze the vanishing entropy limit. Finally, we propose a reinforcement learning algorithm based on policy iteration. We show both policy improvement and policy convergence results for our proposed algorithm. This talk is based on a joint project together with Giorgio Ferrari and Renyuan Xu.

Christian Deutscher
Department for Sports Science, Bielefeld University

Uncertainty on the Pitch and in Prices: Qualification Thresholds and Goal Anticipation in European Football

Abstract:

Uncertainty is intrinsic to sports: randomness in match events, schedules, and tie-breaks means qualification targets can't be captured by a single points "cut-off," but should be treated as season- and format-dependent probability distributions. A similar uncertainty shows up in betting markets, which may react to play rather than predict it. We study how uncertainty is quantified and priced in football in two settings. First, under UEFA's new 36-team UCL league phase (from 2024/25), we simulate seasons using a bivariate Dixon–Coles model with Elo-based strengths and a draw-rate adjustment, estimating point-threshold distributions for the top eight and for places 9–24. Second, using 1 Hz Bundesliga odds and stakes around the first goal, regression and state-space models show strong real-time adjustment but no systematic pre-goal signal, consistent with strong-form efficiency. Overall, uncertainty is a central feature shaping club strategy under new formats and constraining anticipation even in information-rich betting markets.

Raphael Flepp
University of Zurich

Suspense and Surprise in European Football

Abstract:

Based on the premise that outcome uncertainty enhances entertainment value and is desirable for sports fans, league organizers frequently redesign competitions, and policymakers often justify market interventions such as salary caps, draft systems, or revenue redistribution schemes. However, prior research has failed to provide consistent evidence that outcome uncertainty is a key driver of stadium attendance or TV viewership, particularly in European football. As a result, monitoring outcome uncertainty over time and treating it as a primary policy target appears misleading. Match-level suspense and surprise, which capture the entertainment utility created by competitive balance and outcome uncertainty for sports spectators, are promising alternative policy targets. These measures are valuable because they provide a more accurate understanding of how match attractiveness related to competitive balance evolves across clubs and seasons. Through simulations, we derive a benchmark range for suspense and surprise based on a perfectly balanced match before analyzing over 25,000 men's matches (2010/11–2023/24) and 725 women's matches (2023/24) from Europe's top football leagues. Our findings reveal that an average match generates lower suspense than the benchmark range, particularly for top teams, while surprise values largely align with the benchmark. Moreover, we observe nuanced trends over time in men's football and highlight notable differences across leagues and clubs in both men's and women's competitions. These insights enhance our understanding of how match attractiveness arises from competitive balance and carry important policy implications.

Manuel Förster

Center for Mathematical Economics (IMW), Bielefeld University

From prejudice to racial profiling and back: A naive intuitive statistician's curse

Abstract:

A designer conducts random searches to detect criminals, and may condition the search probability on individuals' appearance. She updates her belief about the distribution of criminals across appearances using her search results, but incorrectly takes her sample distribution for the population distribution. In equilibrium she employs optimal search probabilities given her belief, and her belief is consistent with her findings. We show that she will be discriminating an appearance if and only if she overestimates the probability of this appearance's being criminal. Notably, the ranking of two appearances' being criminal may be reversed in equilibrium. Moreover, in a linear model, we provide conditions for uniqueness of the equilibrium and show how tightening the budget may worsen the situation of those most discriminated against. Finally, we demonstrate that our model can explain data from New York City, while a rational benchmark cannot. In particular, our results suggest that a policy change in 2013 has increased the aversion against being discriminating.

Peter Grünwald
CWI in Amsterdam

E is the new P

Abstract:

How much evidence do the data give us about one hypothesis versus another? The standard way to measure evidence is still the p-value, despite a myriad of problems surrounding it. In this talk I will provide a gentle introduction to the e-value (wikipedia), a recently popularized notion of evidence which overcomes some of these issues. E-values were only given a name as recently as 2019, but since then, interest in them has exploded with Oberwolfach and BIRS workshops and 100s of papers, many in top journals such as the Annals of Statistics. Crucially, e-values allow for effortless testing under optional continuation of data collection and combination of data from different sources while keeping Type-I error control - something that practitioners yearn for, yet is simply impossible to do with the classical approaches. Relatedly, they resolve some paradoxes involving p-values and counterfactuals and the likelihood principle. In simple cases, e-values coincide with Bayes factors. But if the null is composite or nonparametric, or an alternative cannot be explicitly formulated, e-values and Bayes factors become distinct and e-processes can be seen as a type of betting strategies leading to anytime-valid confidence intervals that remain valid under continuous monitoring and optional stopping. In parametric settings they tend to be strictly wider than, hence consistent with Bayesian credible intervals. This led to the development of the e-posterior, an analogue to the Bayesian posterior that *gets wider rather than wrong* if the prior is chosen badly.

Gergely Horváth
Duke Kunshan University

The Fragility of Reputation under Manipulated Information Co-authors: Jaromir Kovarik

Abstract:

Reputation is a central mechanism for sustaining cooperation in human societies, yet–unlike in typical theories and most experimental setups–real-world reputational information about others' past behavior is often noisy, delayed, or deliberately manipulated, raising the question of whether cooperation can survive when reputational signals are unreliable. We experimentally study how cooperation evolves when individuals can strategically manipulate their public image and whether allowing others to verify such information can counteract manipulation and restore cooperation. As a benchmark, we replicate a standard repeated network cooperation game in which participants observe one another's past actions and freely choose partners for future interactions. This baseline produces stable, high levels of cooperation driven by conditional association and reputation-based reciprocity. When participants can manipulate information about their past behavior at a cost, manipulation becomes widespread and cooperation collapses. Introducing the option to verify others' reputations does not prevent this decline: verification is infrequent, poorly targeted, and insufficient to restore reliable reputational information. Consequently, cooperation disappears even when truth-checking is possible. These findings expose the vulnerability of reputation-based cooperation to strategic distortion and highlight the crucial role of trustworthy, verifiable reputation systems. This fragility is magnified in the modern digital era, where profit-driven communication technologies and social media make distortion, selective presentation, and manipulation of information cheap and widespread. Such dynamics foster the spread of disinformation, including fake news, climate change denial, and vaccine refusal, undermining collective efforts to address global challenges such as climate change, resource depletion, and pandemics.

Eyke Hüllermeier

Ludwig Maximilian University of Munich

Uncertainty Quantification in Machine Learning: From Aleatoric to Epistemic

---

Bernhard Kasberger

Johannes Kepler University Linz

Algorithmic cooperation

Abstract:

Reinforcement learning algorithms play an increasingly important role in economic situations. These situations are often strategic, where the artificial intelligence may or may not be cooperative. We compare human and algorithmic cooperation rates in the infinitely repeated two-player prisoner's dilemma and study which strategies they choose to cooperate and punish deviations. Through a sequence of computational Q-learning and human-player experiments, we find that our Q-learning algorithms tend to cooperate less than humans, particularly when cooperation is risky or not incentive-compatible. Algorithms often use different strategies than humans, leading to different on and off-path behavior.

Xinyu Li
University of Oxford

Learning Path-Dependent McKean Vlasov Control

Abstract:

We study numerical approximation and learning algorithms for path-dependent McKean-Vlasov (MKV) stochastic control problems with open-loop controls. We introduce an Euler discretisation with piecewise constant open-loop controls. Under suitable regularity and convexity assumptions, we show that the original continuous-time control problem can be approximated by the discrete-time Euler scheme with an error of order $(\Delta t)^{1/4}$. We further establish a value function equivalence between piecewise constant open-loop controls and controls adapted to a smaller filtration generated by discrete Brownian motion increments at each time step. A propagation of chaos result is then established for the associated particle system. Building on this structure, we design a neural network-based algorithm whose control parameterisation depends on the state distribution, is path-dependent, and is given in feedback form with respect to the Brownian motion increments.

Stéphane Loisel
Laboratoire LIRSA, Paris

Construction of climate change impact scenarios in insurance, end of assurability and impact of prevention

Abstract:

Following initial pilot stress test exercises proposed by insurance and banking regulators, insurance and reinsurance companies face the challenge to design climate change impact scenarios up to time horizons ranging from 2050 to 2100, much beyond the traditional 3 to 5-year view usually provided in their Own Risk and Solvency Assessment. We provide general recommendations for the construction of such scenarios and illustrate them in the context of health and mortality risks of a Metropolitan France insurer. We also explain how prevention can play a great role to mitigate both life and non-life impacts of climate change, and how satellite data offers promising inputs for some climate risk management. We analyze the risk that some risks become uninsurable and identify optimal prevention strategies in several contexts.

The talk is based on joint papers with coauthors including A. Stephan (Guy Carpenter, Paris) and R. Vigneron (Univ. of Lyon and ACTIONS research chair), E. Dal Moro (SI Re), H. Albrecher (Lausanne), and J. Trufin (ULB).

Rouven Michels
TU Dortmund University

Using Hidden Markov Models for Decision-Making under Uncertainty in Sports Analytics

Abstract:

In sports, players and coaches continuously make tactical decisions under uncertainty. This talk presents statistical approaches to quantify such decision-making processes using Hidden Markov Models (HMMs). Three applications demonstrate how tactical structures can be inferred from observed tracking and event data and how these insights enable predictions of future game situations. Specifically, we model offensive play-calling and defensive coverage schemes in American football as well as latent playing styles in soccer. The results highlight how probabilistic models like HMMs help to understand strategic behaviour under uncertainty – a challenge that translates beyond sports analytics.

Jurek Preker
Center for Mathematical Economics (IMW), Bielefeld University

Strategic Information Selection

Abstract:

An agent who receives a stream of messages that contain information about an unknown state of the world might ignore some of them. This might be because of cognitive overload, but it might also be strategically: information that contradicts one's belief increases cognitive dissonance. A decision maker who experiences disutility in this way wants to be sure to make the right decision—rather than make the right decision. This means that her utility is belief-based. In our model we consider a Markovian agent who observes a stream of noisy messages that stops in each round with some positive probability. Her anticipatory utility function is defined by a functional equation that takes all possible future beliefs after a message into account. We show that she will always update her belief after a confirming message, but will ignore contradicting messages, if her belief is too strong. The relevant threshold solely depends on the messages' precision or—in case the precision is unknown and she has maxmin preferences—on the lowest precision she deems possible. Such an agent has strictly higher anticipatory utility than an agent who uses every message to update. However, the latter has a higher chance to choose the correct outcome in the end. In a population of strategic agents, who only differ in their initial beliefs, polarization is inevitable. This project is joint work with Professor Dr. Dominik Karos.

Tsiry Randrianasolo

Bielefeld University

Numerical approximation of a Dynkin game with asymmetric information

Abstract:

Dynkin games are zero-sum optimal stopping games where two players choose stopping times to optimize a payoff. The value of such games solves a Hamilton– Jacobi–Bellman equation with two obstacles. In the asymmetric information case studied by Christine Grün (2013), one player knows the game configura- tion while the other only knows it up to a certain probability, leading to an additional convexity constraint and a third obstacle in the associated PDE. In this joint work with Ľubomír Baňas and Giorgio Ferrari, we propose a fully discrete numerical scheme to approximate the value function of this game. The value function depends on time, space, and a probability vector in a simplex. We first introduce a semi-discrete scheme in time and probability that preserves convexity and prove its convergence to the viscosity solution using the Barles–Souganidis framework. We then construct an implementable fully discrete scheme using feedforward neural networks for regression. Under suitable conditions on the meshgrids and the neural network, we obtain uniform convergence on compacts. Numerical experiments illustrate the performance of the method and comparisons with Semi-Lagrangian schemes.

Carl Singleton

University of Stirling

Are inefficient odds within in-play markets profitable for bookmakers?

Abstract:

In-play sports betting markets have experienced a significant increase in economic relevance in recent years, yet academic research on their dynamics remains limited. This study examines whether the apparent inefficiencies found within popular in-play betting odds contribute directly to bookmaker profitability. Using high-frequency data that includes the total amounts bet by customers, from three consecutive seasons of the German Bundesliga, we analyse bookmaker price-setting strategies. By comparing bookmaker odds with an efficient dynamic forecasting model, we characterise the profit function for final-result in-play markets.

Ron Peretz
Bar-Ilan University

Authority Measure for Opinion Dynamics

Abstract:

We propose a formal framework for measuring individual authority in opinion dynamics. Our approach is based on an axiomatic foundation in the tradition of coalitional game theory. We show that the unique solution concept satisfying our axioms generalizes the authority distribution notion of Hu and Shapley 2003. The framework is grounded in the dynamics of opinion exchange rather than network topology. However, for any fixed opinion dynamics, our authority concept induces a corresponding network centrality measure for each network topology. Both existing and new centrality measures can be obtained in this way.

Steven Vanduffel
Brussels University

Optimal Life-Contingent Payoffs: A Peer-to-Peer Solution

Abstract:

We consider an investor with a general state-dependent utility, in that she uses different utility functions for the states of survival and death. We explicitly derive her optimal life-contingent payoff, assuming that a provider charges actuarially fair premiums for mortality risk. In real markets, however, a risk loading applies, and the optimal life-contingent payoff can no longer be purchased directly. Instead, we propose a peer-to-peer solution that approximates the optimal life-contingent payoff via a pool, which delivers a proxy payout to each pool member. We show that the proxy payout converges almost surely to the optimal life-contingent payoff and performs very well in realistic scenarios where pools are finite. In particular, our solution demonstrates the potential for developing a capital market that can effectively manage longevity risk.

Yufei Zhang
Imperial College London

The alpha-Potential Game Paradigm: Theory, Algorithms, and Applications

Abstract:

Designing and analyzing non-cooperative multi-agent systems that interact within shared dynamic environments is a central challenge across many established and emerging applications, including autonomous driving, production management, and e-commerce. A key objective in these systems is to identify Nash equilibria, where no agent can benefit by unilaterally deviating from her strategy. However, computing such equilibria is generally intractable unless specific structural properties of the interactions can be leveraged.

In this talk, we provide an overview of a recently developed framework for dynamic N-player games, called alpha-potential games. This approach extends classical potential games to dynamic settings by capturing how individual incentives approximately align with a global alpha-potential function, up to an error level alpha. Within this framework, the problem of computing approximate Nash equilibria reduces to a stochastic control problem for the alpha-potential function, significantly simplifying both analysis and computation. The parameter alpha also reveals important structural properties of the game, such as the population size, the strength of player interactions, and the degree of heterogeneity across agents. We complement these theoretical insights with numerical experiments based on policy gradient methods, illustrating how the alpha-potential framework enables efficient equilibrium computation in large-scale, dynamic multi-agent systems.

Ondřej Zíka

University College Dublin

How our brains make sense of the uncertain world: The neuroscience and psychology of uncertainty

Abstract:

Uncertainty is ubiquitous and subjective. From perceptual ambiguity to uncertainty over social interactions, our brains scramble to correctly estimate, combine and act on different sources and types of uncertainty. Unsurprisingly, inability to cope with uncertainty, a phenomenon known as "intolerance of uncertainty", has been identified as one of the key risks driving multiple mental health issues. In this talk I will introduce how areas of psychology and neuroscience have grappled with the topic of uncertainty, introducing the key experimental and analytic frameworks. I will also focus on reconciling different naming conventions and link them to more broadly used terminology. I will then move on to introducing key findings about how the brain computes and represents uncertainty. Finally, I will discuss uncertainty in the context of mental health, highlighting my ongoing work on structural inference in internalizing psychopathology.
