THE 2023 ALPINE GAME THEORY SYMPOSIUM

CONFERENCE PROGRAM
AND BOOK OF ABSTRACTS

Plenary Speakers

Galit Ashkenazi-Golan	LSE
Itai Ashlagi	Stanford University
Yakov Babichenko	Technion – Israel Institute of Technology
John Birge	University of Chicago
Simina Brânzei	Purdue University
Mario Bravo	Universidad Adolfo Ibañez
Jean-Paul Carvalho	University of Oxford
Nicolò Cesa-Bianchi	Università degli Studi di Milano & Poliecnico di Milano
Julien Combe	CREST, IP Paris & École Polytechnique
Gabriele Farina	Meta AI
Stéphane Gaubert	Inria & CMAP, École Polytechnique
Olivier Gossner	CNRS/CREST & LSE
Josef Hofbauer	University of Vienna
Maryam Kamgapour	EPFL
Frederic Koessler	HEC Paris & CNRS/GREGHEC
David Leslie	Lancaster University
Patrick Loiseau	Inria
Heinrich Nax	University of Zurich
Abraham Neyman	Hebrew University of Jerusalem
Miquel Oliu-Barton	Paris-Dauphine University
Vianney Perchet	ENSAE & Criteo AI Lab
Georgios Piliouras	Singapore University of Technology and Design
Marek Pycia	University of Zurich
Catherine Rainer	University of Western Brittany
Jérôme Renault	Toulouse School of Economics
Rahul Savani	University of Liverpool & Alan Turing Institute
Marco Scarsini	LUISS Guido Carli University
Eran Shmaya	Stony Brook University
Eilon Solan	Tel-Aviv University
Sylvain Sorin	Sorbonne University
Tristan Tomala	HEC Paris
Xavier Venel	LUISS Guido Carli University
Nicolas Vieille	HEC Paris
Guillame Vigeral	Paris-Dauphine University
Yannick Viossat	Paris-Dauphine University
Bernhard von Stengel	LSE
Peyton Young	University of Oxford
Bruno Zilliotto	CNRS / CMM

Conference Program

Time	Mon, June 26	Tue, June 27	Wed, June 28	Thu, June 29	Fri, June 30
09:00	Registration		**Birge**		
09:30	**Neyman**	**Venel**	**Brânzei**	**Gaubert**	**Hofbauer**
10:00	**Rainer**	**Vigeral**	**Shmaya**	**von Stengel**	**Viossat**
10:30	Break	Break	Break	Break	Break
11:00	**Savani**	**Sorin**	**Babichenko**	**Bravo**	**Gossner**
11:30	**Solan**	**Ashkenazi-Golan**	**Loiseau**	**Farina**	**Koessler**
12:00	**Ziliotto**	**Piliouras**	**Tomala**	**Renault**	**Oliu-Barton**
12:30	Lunch	Lunch		Lunch	Lunch
14:00	**Cesa-Bianchi**	**Flash Talks**		**Ashlagi**	
14:30	**Perchet**			**Combe**	
15:00	**Vieille**			**Pycia**	
15:30	Break	Break	**Hike in the Alps**	Break	
16:00	**Scarsini**			**Carvalho**	**Symposium Ends**
16:30	**Kamgarpour**			**Nax**	
17:00	**Leslie**	**Posters & Drinks**		**Young**	
17:30	———			———	
19:30				**O2 Dinner**	

Welcome from the Organizing Committee

It is our great pleasure to welcome you to Grenoble for the 2023 Alpine Game Theory Symposium! We hope you will enjoy both the exciting technical program as well as Grenoble and the French Alps in June.

This meeting is a follow-up to the Game Theory Symposium that took place in Paris in 2018. The purpose of this meeting is to bring together students, researchers, and practitioners that are active in the broad area of game theory, and to create a fertile forum for presenting research results, exchanging ideas, and initiating collaborations. As organizers, we hope there will be more to come in the future, and that the "Alpine" designation will initiate a series of events, in any of the countries bordering the Alps.

This symposium would not have been possible without the generous support of our sponsors, as listed on the cover of this booklet: the French National Research Agency (ANR) that provided generous financial support through the bilateral grants ALIAS; the French National Center for Scientific Research (CNRS) and its special interest groups on game theory (GdR JEMMA) and optimization (GdR MOA); Inria; the Grenoble-Alpes Multidisciplinary Institute in Artificial Intelligence (MIAI); the Laboratoire d'Informatique de Grenoble (LIG); and, last but not least, the administrative staff of our institutions for their valuable time and support in organizing this event: Sophie Azzaro, Sabine Persico, and Annie Simon.

The 2023 Alpine Game Theory Sympoisum is bringing together ca. 40 plenary speakers and 25 flash talk / poster presenters from all over the globe. Their contributions have been organized into a series of plenary sessions and a session of rapid-fire presentations followed by an open-ended poster session. All talks and presentations will be taking place at the IMAG building of the Université Grenoble-Alpes (UGA); we are likewise thankful to the executive board of the UGA for its hospitality and support.

On behalf of the organizing committee, we wish you an exciting conference and a great time in Grenoble!

General Co-Chairs

Panayotis Mertikopoulos	(CNRS)
Bary Pradelski	(CNRS)

Junior Chair

Simon Jantschgi	(Univ. Zurich / Univ. Grenoble-Alpes)

Scientific Committee

Nicolas Gast	(Inria)
Bruno Gaujal	(Inria)
Jérôme Malick	(CNRS)
Panayotis Mertikopoulos	(CNRS)
Bary Pradelski	(CNRS)

The Symposium at a Glance

Day 1 – Monday, June 26 2
Morning Session 1 (*09:30 – 10:30*) 2
Morning Session 2 (*11:00 – 12:30*) 2
Afternoon Session 1 (*14:00 – 15:30*) 3
Afternoon Session 2 (*16:00 – 17:30*) 4

Day 2 – Tuesday, June 27 6
Morning Session 1 (*09:30 – 10:30*) 6
Morning Session 2 (*11:00 – 12:30*) 6
Flash Talks, Posters & Drinks (*14:00*) 7

Day 3 – Wednesday, June 28 10
Morning Session 1 (*09:00 – 10:30*) 10
Morning Session 2 (*11:00 – 12:30*) 10
Social Event (*13:00*) 11

Day 4 – Thursday, June 29 12
Morning Session 1 (*09:30 – 10:30*) 12
Morning Session 2 (*11:00 – 12:30*) 12
Afternoon Session 1 (*14:00 – 15:30*) 13
Afternoon Session 2 (*16:00 – 17:30*) 14
Conference Dinner (*19:30*) 15

Day 5 – Friday, June 30 16
Morning Session 1 (*09:30 – 10:30*) 16
Morning Session 2 (*11:00 – 12:30*) 16

Day 1 – Monday, June 26

Morning Session 1 (09:30 – 10:30)

Abraham Neyman **09:30 – 10:00**

Equilibrium of Stochastic Games with lim*-average Payoffs*

We consider non-zero-sum stochastic games and define a lim-average ε-equilibrium, $\varepsilon > 0$. The definition implies that any lim-average ε-equilibrium is an ε-equilibrium in the game where the payoff to each player is either the lim sup or the lim inf, as $n \to \infty$, of the average of his payoffs in the first n stages of the game. We prove that a stochastic game with finitely many players and states and compact convex and semialgebraic action sets and multilinear payoff and transition functions has, for every $\varepsilon > 0$, a lim-average ε-equilibrium. A corollary of the main result is that any finite stochastic game with observable mixed actions has, for every $\varepsilon > 0$, a lim-average ε-equilibrium.

Catherine Rainer **10:00 – 10:30**

Continuous-Time Information Provision

A principal privately observes the evolution of a continuous-time Markov chain and sends messages to an agent over time. The agent doesn't observe the Markov chain. Her payoff depends on her own action and the values taken by the Markov chain, while the payoff of the principal depends only on the action of the agent. Assuming that the agent plays its best replay against the signalling strategy of the principal, this strategy is chosen accordingly, to optimise the payment of the principal. Inspired by the literature on zero-sum games with asymmetric information, it is possible to characterize the optimal payoff of the principal as the solution of a couple of variational inequalities. In a second step we analyse the case where the optimal strategy of the principal can be described in terms of a piecewise deterministic Markov process (PDMP). The characterization of the optimal strategy of the principal and the associated payoff in terms of PDMP's provides an explicite description of both. As an example, we study the continuous time counterpart of a model introduced independently by Renault, Solan & Vieille (2017) and Ely (2017).

Morning Session 2 (11:00 – 12:30)

Rahul Savani **11:00 – 11:30**

Ordinal Potential-Based Player Rating

A two-player symmetric zero-sum game is transitive if for any pure strategies x, y, z, if x is better than y, and y is better than z, then x is better than z. It was recently observed that the Elo rating fails at preserving transitive relations among strategies and therefore cannot correctly extract the transitive component of a game. Our first contribution is to show that the Elo rating actually does preserve transitivity when computed in the right space.

Precisely, using a suitable invertible mapping φ, we first apply φ to the game, then compute Elo ratings, then go back to the original space by applying φ^{-1}. We provide a characterization of transitive games as a weak variant of ordinal potential games with additively separable potential functions. Leveraging this insight, we introduce the concept of transitivity order, the minimum number of invertible mappings required to transform the payoff of a transitive game into (differences of) its potential function. The transitivity order is a tool to classify transitive games, with Elo games being an example of transitive games of order one. Most real-world games have both transitive and non-transitive (cyclic) components, and we use our analysis of transitivity to extract the transitive (potential) component of an arbitrary game. We link transitivity to the known concept of sign-rank: transitive games have sign-rank two; arbitrary games may have higher sign-rank. Using a neural network-based architecture, we learn a decomposition of an arbitrary game into transitive and cyclic components that prioritises capturing the sign pattern of the game. In particular, a transitive game always has just one component in its decomposition, the potential component. We provide a comprehensive evaluation of our methodology using both toy examples and empirical data from real-world games.

Eilon Solan **11:30 – 12:00**

On Equilibria in Games in Continuous Time

I will discuss the proper notion of strategies in games in continuous time, and equilibrium existence and characterization in both repeated games and stopping games in continuous time.

Bruno Ziliotto **12:00 – 12:30**

Blackwell's Approachability with Time-Dependent Payoff Functions and Metrics: Application to the Big Match

We define an extension of Blackwell's approachability where the payoff function and the metric are time-dependent. We establish a general guarantee for the natural extension to this framework of Blackwell's algorithm. In the case where the target set is an orthant, we present a choice of time-dependent metric which yields different convergence speeds for each coordinate of the average vector-valued payoff. As an application, an epsilon-uniformly optimal strategy for Player I in the Big Match is given by Blackwell's algorithm in a well-chosen auxiliary approachability problem.

Afternoon Session 1 (14:00 – 15:30)

Nicolò Cesa-Bianchi **14:00 – 14:30**

Online Learning, Bandits, and Digital Markets

The explosive growth of online markets has created complex ecosystems of algorithmic agents. To optimize their revenue, agents need to understand how the market works, and to do so they often resort to strategies that learn from past observations. In this talk, we describe some recent results characterizing the strengths and limitations of sequential decision-making approaches applied to various problems arising in digital markets. The analysis sheds light on how the learning rates depend on the interplay between the form of the revenue function and the feedback provided during the learning process.

Vianney Perchet **14:30 – 15:00**

Decentralized Learning in Online Queuing Systems

Motivated by packet routing in computer networks, online queuing systems are composed of queues receiving packets at different rates. Repeatedly, they send packets to servers, each of them treating only at most one packet at a time. In the centralized case, the number of accumulated packets remains bounded (i.e., the system is *stable*) as long as the ratio between service rates and arrival rates is larger than 1. In the decentralized case, individual no-regret strategies ensure stability when this ratio is larger than 2. Yet, myopically minimizing regret disregards the long-term effects due to the carryover of packets to further rounds. On the other hand, minimizing long-term costs leads to stable Nash equilibria as soon as the ratio exceeds $e/(e-1)$. Stability with decentralized learning strategies with a ratio below 2 was a major remaining question. We first argue that for ratios up to 2, cooperation is required for stability of learning strategies, as selfish minimization of policy regret, a *patient* notion of regret, might indeed still be unstable in this case. We, therefore, consider cooperative queues and propose the first learning decentralized algorithm guaranteeing the stability of the system as long as the ratio of rates is larger than 1, thus reaching performances comparable to centralized strategies.

Nicolas Vieille **15:00 – 15:30**

Blackwell Optimality in Robust MDPs

Robust MDPs are a generalization of MDPs that share similarities with stochastic games, but were studied independently. In this paper we provide results on the existence of Blackwell optimal strategies in robust MDPs, using tools that were developed for stochastic games.

Afternoon Session 2 (16:00 – 17:30)

Simina Brânzei **16:00 – 16:30**

Online Learning in Multi-unit Auctions

We consider repeated multi-unit auctions with uniform pricing, which are widely used in practice for allocating goods such as carbon licenses. In each round, K identical units of a good are sold to a group of buyers that have valuations with diminishing marginal returns. The buyers submit bids for the units, and then a price p is set per unit so that all the units are sold. We consider two variants of the auction, where the price is set to the K-th highest bid and $(K+1)$-th highest bid, respectively. We analyze the properties of this auction in both the offline and online settings. In the offline setting, we consider the problem that one player i is facing: given access to a data set that contains the bids submitted by competitors in past auctions, find a bid vector that maximizes player i's cumulative utility on the data set. We design a polynomial time algorithm for this problem, by showing it is equivalent to finding a maximum-weight path on a carefully constructed directed acyclic graph. In the online setting, the players run learning algorithms to update their bids as they participate in the auction over time. Based on our offline algorithm, we design efficient online learning algorithms for bidding. The algorithms have sublinear regret, under both full information and bandit feedback structures. We complement our online learning algorithms with regret lower bounds. Finally, we analyze the quality of the equilibria in the worst case through the lens of the core solution

concept in the game among the bidders. We show that the $(K+1)$-th price format is susceptible to collusion among the bidders; meanwhile, the K-th price format does not have this issue.

Maryam Kamgarpour **16:30 – 17:00**

No-Regret Learning in Games with Correlated Payoffs

We consider the problem of learning to play a repeated multi-agent game with an unknown reward function. Single player online learning algorithms attain strong regret bounds when provided with full information feedback, which unfortunately is unavailable in many real-world scenarios. Bandit feedback alone, i.e., observing outcomes only for the selected action, yields substantially worse performance. We consider a natural model where, besides a noisy measurement of the obtained reward, the player can also observe the opponents' actions. This feedback model, together with a regularity assumption on the reward function, allows us to estimate the player's reward function. We propose a novel confidence-bound based bandit algorithm and derive kernel-dependent regret bounds that are comparable to the known bounds in the full information setting. We extend our framework to different class of games, including contextual games and multi-agent reinforcement learning.

David Leslie **17:00 – 17:30**

Learning in Stochastic Games

I will present recent results on fictitious-play-like processes in stochastic games. This extension of normal form games introduces stochastic state transitions, as in a Markov decision process, and presents interesting challenges to learning agents. If the agents are allowed to calculate full stochastic game best responses to opponent strategies then the analysis essentially reduces to a (very) large normal form game, but if we assume more boundedly-rational agents that can simple take single-stage best responses then things become challenging. As well as theoretical results in this more boundedly-rational setting, I will demonstrate some applications to a smuggling game, in which a border agent moves around a border in an attempt to minimise the flow of contraband across the border.

Day 2 – Tuesday, June 27

Morning Session 1 (09:30 – 10:30)

Gabriele Farina **09:30 – 10:00**

Kernelized Multiplicative Weights for 0/1*-Polyhedral Games: Bridging the Gap Between Learning in Extensive-Form and Normal-Form Games*

While extensive-form games (EFGs) can be converted into normal-form games (NFGs), doing so comes at the cost of an exponential blowup of the strategy space. So, progress on NFGs and EFGs has historically followed separate tracks, with the EFG community often having to catch up with advances (e.g., last-iterate convergence and predictive regret bounds) from the larger NFG community. In this talk I will show that the Optimistic Multiplicative Weights Update (OMWU) algorithm – the premier learning algorithm for NFGs – can be simulated on the normal-form equivalent of an EFG in linear time per iteration in the game tree size using a kernel trick. The resulting algorithm, Kernelized OMWU (KOMWU), applies more broadly to all convex games whose strategy space is a polytope with 0/1 integral vertices, as long as the kernel can be evaluated efficiently. In the particular case of EFGs, KOMWU closes several standing gaps between NFG and EFG learning, by enabling direct, black-box transfer to EFGs of desirable properties of learning dynamics that were so far known to be achievable only in NFGs. Specifically, KOMWU gives the first algorithm that guarantees at the same time last-iterate convergence, lower dependence on the size of the game tree than all prior algorithms, and $\widetilde{\mathcal{O}}(1)$ regret when followed by all players.

Sylvain Sorin **10:00 – 10:30**

Some Advances on No-Regret Dynamics

We present a unifying approach to no-regret algorithms (in continuous and discrete time) including projected gradient, mirror descent and dual averaging. We describe the link with variational inequalities and games and offer some directions of research on alternative first-order dynamics.

Morning Session 2 (11:00 – 12:30)

Marco Scarsini **11:00 – 11:30**

Best-Response Dynamics in Two-Person Random Games with Correlated Payoffs

We consider finite two-player normal form games with random payoffs. Player A's payoffs are i.i.d. from a uniform distribution. Given $p \in [0,1]$, for any action profile, player B's payoff coincides with player A's payoff with probability p and is i.i.d. from the same uniform distribution with probability $1-p$. This model interpolates the model of i.i.d. random payoff used in most of the literature and the model of random potential games. First we study the number of pure Nash equilibria in the above class of games. Then we show that, for any positive p, asymptotically

in the number of available actions, best-response dynamics reaches a pure Nash equilibrium with high probability.

Xavier Venel **11:30 – 12:00**

0-1 *Laws for a Control Problem with Random Action Sets*

In many control problems there is only limited information about the actions that will be available at future stages. We introduce a framework where the controller sequentially chooses actions $a_0, a_1, \ldots$, one at a time. Her goal is to maximize the probability that the infinite sequence $a_0, a_1, \ldots$, is an element of a subset G of the set of infinite sequence of natural numbers. The set G is assumed to be a Borel tail set. The controller's choices are restricted: having taken a sequence h_t of actions prior to stage t, she must choose a decision at stage t from a restricted set $A(h_t)$ of natural number. The set $A(h_t)$ is chosen randomly from a distribution p_t, independently over all time periods and past histories. We consider several information structures defined by how far into the future the controller knows what actions will be available. In the special case where all the action sets are singletons (and thus the controller is a dummy), Kolmogorov's 0-1 law says that the probability for the goal to be reached is 0 or 1. We construct a number of counterexamples to show that in general the value of the control problem can be strictly between 0 and 1, and derive several sufficient conditions for the 0-1 "law" to hold.

Guillaume Vigeral **12:00 – 12:30**

Structure of the Sets of Nash Equilibria of Finite Games: Applications to the Complexity of Some Decision Problems in Game Theory

The set of Nash equilibrium payoffs of a one-shot finite game is always nonempty, compact and semialgebraic. For 3 players or more, the reverse also holds: given E a subset of $\mathbb{R}^N$ that is nonempty, compact and semialgebraic, one constructs a finite N-player game such that E is its set of equilibrium payoffs. Related results also hold when one considers only games with integral payoffs, or when the focus is on equilibria instead of equilibrium payoffs. We apply this to understand the complexity class of some natural decision problems on finite games.

Flash Talks, Posters & Drinks (14:00)

Flash Talks (*14:00 –15:30*) | *Poster Presentations* (*16:00 – 19:00*) 14:00 – 19:00

This session will consist of ca. 25 very short "flash" presentations (2-3 minutes each), followed by an open-ended poster session. Drinks and some light finger food will be served during the poster session (which, weather permitting, may take place outdoors).

The list of flash/poster presentations is as follows:

1. **Luc Attia**
   "*Percolation games*"
2. **Omar Boufous**
   "*Constrained Correlated Equilibria*"

3. **Matthew Darlington**
   "*A Stochastic Game Framework for Patrolling a Border*"
4. **Nikos Karagiannis-Axypolitidis**
   "*An Inspection Game where a Penalty Convention between the Players Determines the Strategic Advantage of the Inspector*"
5. **Davide Legacci**
   "*Hodge Decomposition for Continuous Games*"
6. **Emilien Macault**
   "*Keep the Deadline: Cooperation in Revision Games with Frequency-Dependent Payoffs*"
7. **Sayan Mukherjee**
   "*Equilibrium Selection via Stochastic Evolution in Continuum Potential Games*"
8. **Ivan Novikov**
   "*Zero-Sum Repeated Games with Stage Duration and Public Signals*"
9. **Waïss Azizian**
   "*The Last-Iterate Convergence Rate of Optimistic Mirror Descent in Stochastic Variational Inequalities*"
10. **Victor Boone**
    "*The Long-Run Behavior of Regularized Learning via Asymptotic Pseudotrajectories in Payoff Space*"
11. **Maurizio D'Andrea**
    "*Stochastic Multiplicative Weights Algorithms with Constant Step-size*"
12. **Etienne Montbrun**
    "*Certified Multi-Fidelity Zeroth-order Optimization*"
13. **Felipe Garrido-Lucero**
    "*DataSet Shapley: A Shapley Value Proxy for Efficient Data-set Valuation*"
14. **Yu-Guan Hsieh**
    "*No-Regret Learning in Games with Noisy Feedback: Faster Rates and Adaptivity via Learning Rate Separation*"
15. **Mohammad Reza Karimi**
    "*SDEs and Games via Dynamical Systems*"
16. **Edward Plumb**
    "*Convergence of Policy Gradient Methods in Repeated Games*"
17. **Hedi Hadiji & Sarah Sachs**
    "*Towards Characterizing the First-order Query Complexity of Learning (Approximate) Nash Equilibria in Zero-sum Matrix Games*"
18. **Aditya Aradhye**
    "*Sender-receiver Games with Bounded Recall*"
19. **Thomas Asikis**
    "*Using controlled experimental data to predict market structure*"
20. **Cesare Carissimo**
    "*Dynamic Effects of Recommendations on Users*"

21. **Roberto Colomboni**
    "*A Regret Analysis of Bilateral Trade*"
22. **Luca Damonte**
    "*Targeting Interventions for Polarization Minimization in Opinion Dynamics*"
23. **Simon Finster**
    "*Selling Multiple Complements with Packaging Costs*"
24. **Margarita Kirneva**
    "*Informing to Divert Attention*"
25. **Simon Mauras**
    "*Auction with Private Interdependent Valuations*"
26. **Raimundo Saona**
    "*Prophet Inequalities: Separating Random Order from Order Selection*"
27. **Yevgeny Tsodikovich**
    "*Minimal Contagious Sets in Innovation Diffusion Networks*"

Day 3 – Wednesday, June 28

Morning Session 1 (09:00 – 10:30)

John Birge **09:00 – 09:30**

Learning and Information Aggregation in Dynamic Games

In many settings with partially informed agents, agents' actions and noisy random outcomes lead to full information aggregation as agents have repeated interactions. When the learning of other agents cannot be directly observed by a focal agent, however, that agent may have an incentive to choose actions that reduce the other agents' learning as well as the focal agent's learning. This result can yield situations in which information is not aggregated and all agents benefit relative to the full information equilibrium. The talk will describe examples of such behavior including settings in which information of an exogenous state can be attained by all agents but uncertainty over the diffusion of information about that state can lead to cooperative outcomes not possible otherwise.

Georgios Piliouras **09:30 – 10:00**

Towards a Unified Theory of Learning in Games: Equilibration, Optimization, Recurrence & Chaos

We examine some classic questions in game theory and online learning. How do standard regret-minimizing learning dynamics such as multiplicative weights update, online gradient descent, and follow-the-regularized-leader behave when applied in games? The standard textbook answer to this question is that these dynamics "converge" in a time-average sense to weak notions of equilibria. We discuss why such results can be non-informative in practice and instead focus on understanding the day-to-day behavior. In the seminal class of zero-sum games, continuous-time dynamics "cycle" around the equilibrium, whereas naive discretizations (e.g., simultaneous gradient descent ascent) result in unstable and chaotic behavior. On the other hand, alternating gradient descent ascent leads again to cyclic/recurrent behavior. We offer a unified perspective to understand these results combining tools from optimization theory (regret analysis), dynamical systems (chaos theory) and physics (energy conservation) and discuss extensions of these ideas and directions for future work.

Eran Shmaya **10:00 – 10:30**

Optimal Queuing Regimes

Morning Session 2 (11:00 – 12:30)

Yakov Babichenko **11:00 – 11:30**

Quantile Shares

Fair allocation of indivisible goods is considered. A share is a mapping from individual's utility function to the reals. A share is feasible if for every tuple of utilities there exists an allocation that allocates every agent a bundle

that exceeds her share value. We suggest a new notion of shares that is based on ordinal comparisons. For every $0 < q < 1$ we consider the q-quantile share that equals the q-quantile of the distribution of agent's value under the random allocation that allocates each item uniformly at random. We focus on the question: *For which values of q is the q-quantile share feasible?* Previously suggested notions of shares are typically infeasible. We provide surprisingly positive result on the feasibility of quantile shares. For all monotone valuations the $(1/2e)$-quantile share is feasible assuming that the Erdős Matching Conjecture is true. This bound is relatively tight: The $(1/e)$-quantile share is infeasible. We also provide feasibility results that do not rely on conjectures for special classes of valuations such as additive, unit demand, and matroid rank functions valuations.

Patrick Loiseau **11:30 – 12:00**

Statistical Discrimination in Stable Matching

We study statistical discrimination in matching, where multiple decision-makers are simultaneously facing selection problems from the same pool of candidates. We propose a model where decision-makers observe different, but correlated estimates of each candidate's quality. The candidate population consists of several groups that represent gender, ethnicity, or other attributes. The correlation differs across groups and may, for example, result from noisy estimates of candidates' latent qualities, a weighting of common and decision-maker specific evaluations, or different admission criteria of each decision maker. We show that lower correlation (e.g., resulting from higher estimation noise) for one of the groups worsens the outcome for all groups, thus leading to efficiency loss. Further, the probability that a candidate is assigned to their first choice is independent of their group. In contrast, the probability that a candidate is assigned at all depends on their group, and – against common intuition – the group that is subjected to lower correlation is better off. The resulting inequality reveals a novel source of statistical discrimination.

Tristan Tomala **12:00 – 12:30**

Mechanism Design with Restricted Communication

We consider a Sender-Receiver environment where the sender is informed of states and the receiver chooses actions. There is a communication channel between them consisting of sets of input/output messages and a fixed transition probability. The receiver commits to the mechanism which selects distribution of actions and possibly monetary transfers, contingent on output messages. We give characterizations of the joint distributions which can be implemented by communication over the channel, given the incentives of the sender. We consider both one-shot problems and series of i.i.d. problems. We show that for series of problems, linking decisions together achieves more outcomes than monetary transfers.

Social Event (13:00)

Hiking in the Alps 13:00 –18:00

If you have signed up for the hike, you will receive a lunchbox and buses will take us to the beautiful Belledonne mountain range. Starting from the small ski resort of Chamrousse we will hike to the lac Achard, a beautiful mountain lake. The hike is considered easy and covers ca. 200 meters of elevation and a 5km return. Please use suitable footwear!

Day 4 – Thursday, June 29

Morning Session 1 (09:30 – 10:30)

Stéphane Gaubert **09:30 – 10:00**

Solving Irreducible Stochastic Mean-Payoff Games and Entropy Games by Relative Krasnoselskii-Mann Iteration

We analyse an algorithm solving stochastic mean-payoff games, combining the ideas of relative value iteration and of Krasnoselskii-Mann damping. We derive parameterized complexity bounds for several classes of games satisfying irreducibility conditions. We show in particular that an ε-approximation of the value of an irreducible concurrent stochastic game can be computed in a number of iterations in $\mathcal{O}(|\log(\varepsilon)|)$ where the constant in the $\mathcal{O}(\cdot)$ is explicit, depending on the smallest non-zero transition probabilities. This should be compared with a bound in $\mathcal{O}(\varepsilon^{-1}|\log(\varepsilon)|)$ obtained by Chatterjee and Ibsen-Jensen (ICALP 2014) for the same class of games, and to a $\mathcal{O}(\varepsilon^{-1})$ bound by Allamigeon, Gaubert, Katz and Skomra (ICALP 2022) for turn-based games. We also establish parameterized complexity bounds for entropy games, a class of matrix multiplication games introduced by Asarin, Cervelle, Degorre, Dima, Horn and Kozyakin, in which the payoff is given by a topological entropy. We derive these results by methods of variational analysis, establishing contraction properties of the relative Krasnoselskii-Mann iteration with respect to Hilbert's semi-norm.

Bernhard von Stengel **10:00 – 10:30**

Multi-Agent Learning in a Pricing Game

This project proposes a framework, under development, of using machine learning to study larger and more realistic game-theoretic models. In a classic dynamic pricing game due to Selten, firms compete by repeatedly setting a price for a product, where higher prices lead to higher short-term but lower long-term profits. The classic subgame-perfect equilibrium is very competitive, whereas strategic experiments show a tendency of agents to collude. We want to employ machine learning of pricing strategies, with the learning environment given by a mixed equilibrium of existing strategies. The equilibrium is selected with the Harsanyi-Selten tracing procedure, which works fast. Successfully learned strategies are added to the pool and a new equilibrium is computed. This is akin to double-oracle learning in zero-sum games, except that the game is not zero-sum and the resulting equilibrium depends on the learning history. The approach is modular rather than a large simulation, which should allow a better study of the relevant features of the underlying game, its learning mechanisms, and the employed equilibrium concept.

Morning Session 2 (11:00 – 12:30)

Galit Ashkenezi-Golan **11:00 – 11:30**

Projected Gradient Learning and the Folk Theorem

In the projected gradient method the players (independently) employ the learning method to estimate each a direction for improving payoffs Typically, convergence of the strategy profile to the Nash equilibria of the one

shot game is studied, and it is known that a strict Nash equilibrium has a basin of attraction (that is, if the initial strategy profile is sufficiently close to the Nash equilibrium, it converges to it). However, this learning method assumes the ability to repeat playing the game so that improvements can be made. We show that using finite memory strategies, any feasible payoff strictly pareto dominating a pure minmax payoff has a basin of attraction. Thus, if the learning is by repeating a game, the richness of equilibrium payoffs guaranteed by the folk theorem is recovered by considering strategies that condition on finite memory.

Mario Bravo **11:30 – 12:00**

Stochastic Iteration for Fixed Points of Nonexpansive Maps: Error Bounds and Applications

We study a stochastically perturbed version of a well-known iteration for computing fixed points of nonexpansive maps in finite-dimensional spaces. We discuss sufficient conditions on the random noise and stepsizes that guarantee almost sure convergence, and derive non-asymptotic error bounds for the fixed-point residuals. Our main result concerns the case of a martingale difference noise with variances that can grow unbounded. Next, we show that our result can be used to study a known reinforcement learning algorithm for average reward Markov decision processes, a much less studied type of algorithm than its discounted reward counterpart. Finally, time permitting, we will discuss how our general approach leads to explicit non-asymptotic error rates with computable constants, in the case where the noise sequence has a bounded variance.

Jérôme Renault **12:00 – 12:30**

Optimistic Gradient Descent Ascent in Bilinear Games

We study the convergence of Optimistic Gradient Descent Ascent (OGDA) in unconstrained bilinear games. For zero-sum games, we clarify and extend earlier results by proving the exponential convergence of OGDA to a saddle-point, and provide the exact ratio of convergence as a function of the step size. Then we introduce OGDA for general-sum games, and show that in many cases, either OGDA converges to a Nash equilibrium, or the payoffs for both players converge to infinity. We also show how to increase drastically the speed of convergence of a zero-sum problem, by introducing a general-sum game using the Moore-Penrose inverse of the original payoff matrix. Hence, general-sum games can be used here to optimally improve algorithms designed for min-max problems. We finally illustrate our results on a stylized example of Generative Adversarial Network.

Afternoon Session 1 (14:00 – 15:30)

Itai Ashlagi **14:00 – 14:30**

Price Discovery in Waiting Lists: A Connection to Stochastic Gradient Descent

Waiting lists allocate items by offering agents a choice among items with associated waiting times. These waiting times serve as prices that are determined endogenously and adjust according to the stochastic arrivals and departures of agents. We study the allocative efficiency under such dynamically adjusting prices by drawing a connection between this price adjustment process and the stochastic gradient descent optimization algorithm. We show that the loss due to price fluctuations is bounded by the granularity of price changes. Additional conditions allow us to identify markets where the loss is close to the bound or exponentially small. Our results show that a simple price

adjustment heuristic can perform well, but may be slow to adjust to changes in demand because of a trade-off between the speed of adaptation and loss from price fluctuations.

Julien Combe **14:30 – 15:00**

Market Design for Distributional Objectives in (Re)Assignment: An Application to Improve the Distribution of Teachers in Schools

Centralized (re)assignment of workers to jobs is increasingly common in public and private sectors. These markets often suffer from distributional problems. To alleviate these, we propose two new strategy-proof (re)assignment mechanisms. While they both improve individual and distributional welfare over the status quo, one achieves two-sided efficiency and the other achieves a novel fairness property. We quantify the performance of these mechanisms in teacher (re)assignment where unequal distribution of experienced teachers in schools is a widespread concern. Using French data, we show that our efficient mechanism reduces the teacher experience gap across regions more effectively than benchmarks, including the current mechanism, while also effectively increasing teacher welfare. As an interesting finding, while our fairness-based mechanism is very effective in reducing teacher experience gap, it prevents the mobility of tenured teachers, which is a detrimental teacher welfare indicator.

Marek Pycia **15:00 – 15:30**

Evaluating with Statistics: Which Outcome Measures Differentiate Among Matching Mechanisms?

The selection of mechanisms to allocate school seats in public school districts can be highly contentious. At the same time the standard statistics of student outcomes calculated from districts' data are very similar for many mechanisms. This paper contributes to the debate on mechanism selection by explaining the similarity puzzle as being driven by the invariance properties of the standard outcome statistics: outcome measures are approximately similar if and only if they are approximately anonymous.

Afternoon Session 2 (16:00 – 17:30)

Jean-Paul Carvalho **16:00 – 16:30**

Radicalization

To analyze the mechanisms behind radicalization and the design of counter-radicalization policies, we study a population (or subpopulation) known as the identity group that derives meaning from participation in identity-based activities. A forward-looking organization provides a platform (i.e., club) for these activities. By dynamically tuning its membership requirements, the organization determines both current participation and the future share of radicals. The warning sign for radicalization is cultural purification, i.e., the screening out of moderates and exclusive recruitment of radicals. While this shrinks the club, it puts it on a growth path along which it becomes larger and more extreme over time. Conventional counter-radicalization policies can backfire and fuel this process of radicalization. The organization can itself boost radicalization through outreach and by inducing discrimination against group members. The radicalization mechanisms we identify can be disabled by mild anti-radical messaging and informational interventions that eliminate stereotypes.

Heinrich Nax **16:30 – 17:00**

Learning in Buyer-Seller Markets: Experimental Evidence

Convergence dynamics toward Walrasian competitive equilibrium in experimental continuous bid/ask double auctions is a finding of central importance in experimental economics and behavioral finance. First recorded by Smith (1962), and contrary to the common expectations in the field based on earlier results by Chamberlin (1948), DA markets do not fail to converge even under purely private information. This finding has been replicated and reproduced in numerous studies since, and convergence typically fails when markets are strongly skewed towards one of the two market sides. In this talk, we dig into the learning models that might explain the convergence based on evidence from a large-scale online reproduction of the canonical experiments on the DA, and discuss implications for behavioral market design.

Peyton Young **17:00 – 17:30**

Stochastic Stability in Theory and Practice

Stochastic stability theory predicts which equilibria are selected when players deviate from rational behavior with small probability. Although the theory is now well-developed it needs to be brought to bear on empirical cases. I discuss recent experiments and field data in which the magnitude of players' deviations can be estimated, and the observed dynamics compared with theoretical predictions.

Conference Dinner (19:30)

Dinner at the Bastille 19:30

The conference dinner will take place at the O2 Restaurant, which is located at the top of the Bastille. You can reach the restaurant either by foot (a steep, ca. 1 hour and 300 meters elevation, walk) or by cable car – affectionately known as "*les bulles*".

Day 5 – Friday, June 30

Morning Session 1 (09:30 – 10:30)

Josef Hofbauer **09:30 – 10:00**

Evolutionary Dynamics in Costly-Signaling Games

For some simple signaling games (with 2 states of nature, 2 signals, and 2 actions) we determine all equilibrium components and their index and study the replicator and best response dynamics.

Yannick Viossat **10:00 – 10:30**

Survival of Dominated Strategies under Evolutionary Game Dynamics

This talk will review past and recent results around elimination or survival of dominated strategies under evolutionary game dynamics, insisting on underlying intuitions, and discuss a couple of open questions.

Morning Session 2 (11:00 – 12:30)

Olivier Gossner **11:00 – 11:30**

Rationalizable Distributions in Games with Incomplete Information: A SCAMP Characterization

We study (interim correlated) rationalizability in a game with incomplete information. We characterize the recursive set of possible rationalizable hierarchies through a finite automaton, and provide a revelation principle that characterizes the distributions over these hierarchies that arise from any common prior. We show that a simple and finitely parametrized class of information structures, Stationary Common Automaton Markov Priors (SCAMP), is sufficient to generate every outcome distribution induced by general common prior information structures. Using this result, we characterize the set of rationalizable distributions as a convex polyhedron.

Frédéric Koessler **11:30 – 12:00**

Convergence and Correlation in Large Games

We consider multi-population Bayesian games with a large number of players. Each player aims at minimizing a cost function that depends on her own action, the distribution of players' actions in all populations, and an unknown state parameter. We study the nonatomic limit versions of these games. We introduce the concept of Bayes correlated Wardrop equilibrium, which extends the concept of Bayes correlated equilibrium to nonatomic games. We prove that Bayes correlated Wardrop equilibria are limits of action flows induced by Bayes correlated equilibria of the game with a large finite set of small players. For nonatomic games with complete information admitting a convex potential, we prove that the set of correlated and of coarse correlated Wardrop equilibria, are reduced to the set of probability distributions over Wardrop equilibria, and that all equilibrium outcomes have

the same costs. We get the following consequences. First, all flow distributions of (coarse) correlated equilibria in convex potential games with finitely many players converge to Wardrop equilibria when the weight of each player tends to zero. Second, for any sequence of flows satisfying no-regret property, its empirical distribution converges to the set of distributions over Wardrop equilibria and the average cost converges to the unique Wardrop cost.

Miquel Oliu-Barton **12:00 – 12:30**

Value-Positivity for Matrix Games

Matrix games are the most basic problem in Game Theory, but robustness to small perturbations is not yet fully understood. A perturbed matrix game is one where the entries depend on a parameter which varies smoothly around zero. We introduce two new concepts: (*a*) value-positivity if, for every sufficiently small error, there is a strategy that guarantees the value of the error-free matrix game; and (*b*) uniform value-positivity if there exists a fixed strategy that guarantees, for every sufficiently small error, the value of the error-free matrix game. While the first concept captures the dependency of optimal strategies to small perturbations, the second naturally arises where the data is uncertain and a strategy is sought which remains optimal despite that uncertainty. In this paper, we provide explicit polynomial-time algorithms to solve these two problems for any polynomially perturbed matrix game. For (*a*) we further provide a functional form for the error-dependent optimal strategy. Last, we translate our results into robust solutions for LPs.

Conference Venue

The conference will take place at the main amphitheater of the IMAG building at the Saint-Martin d'Hères campus of Université Grenoble-Alpes. The address is

Bâtiment IMAG
700 Av. Centrale
38400 Saint-Martin-d'Hères

and it can be found on Google Maps at `https://goo.gl/maps/97sNT8zNwisKcbUB9` (see)

The IMAG building is easily accessible from downtown via tram, Lines B and C; you can get off at either Gabriel Fauré or Bibliothèques Universitaires (see map above). You should expect a tram ride of approx. 15 minutes from downtown, plus a 5-minute walk from Gabriel Fauré or Bibliothèques Universitaires.