Abstracts

Games on networks

1 Nizar Allouch

University of Kent, UK
Aggregation in Networks

Abstract

In this paper, we show that a concept of aggregation can hold in large network games with linear best replies. Breaking up large networks into smaller subnetworks, which can be replaced by representative players, leads to a coarse-grained description of strategic interactions.

This method of summarizing complex strategic interactions by simple ones can be applied to compute all Nash equilibria for the special network structure of cograph. A key finding is that a stable Nash equilibrium of the large network game can be decomposed into a collection of Nash equilibria of subnetwork games. Thereby, we establish a systematic relationship between player's position in a subnetwork and his equilibrium action in the large network game.

2 Dario Bauso

University of Groningen, Netherlands
Cascading failures: dynamics, stability and control

Abstract

In this talk, we try to understand and mitigate cascading failures in a financial network from a systems and control theoretic perspective. We shall view cross-holdings as feedback, external assets as inputs and failure penalties as static nonlinearities. We provide sufficient milder and stronger conditions for the system to be a positive one, and study equilibrium points and stability. Stability means absence of cascades and convergence of market values to a constant value. We provide a constructive method for control design to obtain stabilizing market investments in the form of feedback-feedforward control inputs. Time permitting, we shall look at the stochastic extension with noisy market values, which will provide insights into the probability of failure of each company during the transient and at steady-state.

3 John R Birge

The University of Chicago Booth School of Business, USA
Prediction Games and Market Mechanisms

Abstract

Prediction markets can aggregate information from a group of heterogeneously informed agents but well-informed agents may have an incentive to manipulate the market and to extract surplus at the expense of the other players and a market maker if any. In some cases, a market maker may, however, be able to design a mechanism that prevents manipulation but can still encourage participation. This talk will describe conditions that lead to such outcomes and general situations in which information is aggregated and a market maker also has an incentive for participation.

4 Peter Caines

McGill University, Canada
Mean Field Games on Large Networks

Abstract

Graphon Mean Field Games (GMFGs) constitute generalizations of Mean Field Games to the situation where agents are concentrated in subpopulations located at the nodes of a large network. The solution to the corresponding GMFG equations determines the Nash equilibrium of the entire population distributed over the network, and an epsilon-Nash property relates infinite population equilibria on infinite networks to finite population equilibria on finite networks. For infinite networks embedded in Euclidean space (described via graphexons) the locations of maximum Nash value necessarily correspond to nodes with stationarity GMFG values; in initial examples, this is shown to correspond to maximal nodal degree. Current work is aimed at extending this theory to large sparse networks.

5 Giacomo Como

Politecnico of Torino, Italy
On a network centrality optimization game

Abstract

We study a network formation game whereby n players, identified with the nodes of the network to be formed, chose where to wire their outgoing links in order to maximize their Bonacich centrality. Specifically, the action of every player i consists in the wiring of a predetermined number d_i of directed out-links, and her utility is her own Bonacich centrality in the network resulting from the actions of all players. We show that this is an ordinal potential game and that the best response correspondence always exhibits a local structure in that it is never convenient for a node to link to other nodes that are at incoming distance more than d_i from her. We then study the equilibria of this game determining necessary conditions for a graph to be a (strict, recurrent) Nash equilibrium. Moreover, in the special cases where for all players $ \deg_i=1 $ or $ \deg_i=2 $, we provide a complete classification of the set of (strict, recurrent) Nash equilibria. This analysis shows in particular that the considered formation mechanism leads to the emergence of undirected and disconnected or loosely connected networks.

Joint work with Maria Castaldo, Costanza Catalano, and Fabio Fagnani

6 Gabriele Dragotto

Princeton University, USA
Integer Programming Games: Do You Really Need Them?

Abstract

This talk deals with Integer Programming Games (IPGs), i.e., simultaneous non-cooperative games where each player solves a parametrized mixed-integer optimization problem. The large class of IPGs encompasses, among others, any finite game and provides a flexible framework to model games that cannot be efficiently represented in normal form. We restrict our attention to applying IPGs to model and solve (weighted) congestion games. First, we focus on a direct task, i.e., on computing Nash equilibria whenever the players' utilities and strategies are common information; we show that \pagebreak modeling these games as IPGs enables us to compute, enumerate and select the Nash equilibria. Second, we focus on an inverse task, i.e., on recovering the missing parameters of the players' utilities and strategy sets; in this case, we present a novel algorithm that can efficiently learn the missing parameters starting from historical data.

7 Ezzat Elokda

ETH Zürich, Switzerland

Dynamic population games for the modeling of complex socio-economic systems

Abstract

First, we present dynamic population games: a novel formulation of large population games with individual agent dynamics. Using the running example of epidemic propagation, we introduce the components of our formulation: the individual state dynamics (e.g., the infection state) are affected by the agent's own action (e.g., how much they engage in social interactions) and the states and actions of those they interact with (e.g., whether they interact with an infected neighbour). We guarantee the existence of a stationary Nash equilibrium, which is a condition where all agents are playing optimally and the distribution of states and actions in the population is stationary. In contrast to similar works on anonymous sequential games and mean-field games, we show that stationary Nash equilibria can be reduced to standard Nash equilibria in classical (static) population games. This allows adopting the well studied notions of evolutionary dynamics and stable games to our dynamic setting. In our epidemic example, it allows modelling the evolution of the population's social behaviors on a similar timescale as the disease propagation, thereby providing a principled tool to investigate the often counter-intuitive effects of public health measures.

Second, we present a dynamic population game model of a self-contained karma economy in which population agents use dedicated 'karma' tokens to repeatedly trade access to shared resources (e.g., roads, energy, bandwidth). Our model provides valuable insights on the equilibrium behaviours under the self-contained economy and serves as a versatile design tool to achieve long-term fairness and efficiency objectives. In particular, we demonstrate that the karma economy leads to highly efficient resource allocations without resorting to the possibly problematic monetary pricing of the resource, making it an attractive alternative in many of present day's equity-sensitive domains.

8 Selman Erol

Carnegie Mellon University, USA

Regulating Clearing in Networks

Abstract

Recent regulations in the U.S. and Europe incentivize the use of central counterparty clearing houses (CCP) to clear derivatives, arguably to create a less complex and more transparent interbank network that is less prone to financial instabilities. We construct a network model with endogenous exposures and show that the core and the periphery react asymmetrically to these regulations. The core values opacity more and adopts clearing less. Consequently, bilaterally netted exposures to the core increase. The regulation also makes the CCP more exposed to the core than the periphery was pre-regulation. This endogenous network reaction to the regulation creates the unanticipated effect of reducing financial stability through more frequent coordination failures that start at the core and spread to the periphery and the CCP. A novel dataset on U.S. counterparty exposures, before and after the regulations, confirm the model's testable implications.

9 Paolo Frasca

Université Grenoble Alpes, France

Potential deterioration in transportation network efficiency due to route recommendations

Abstract

On-line navigation devices are widely used by drivers to obtain recommendations on the fastest route to their destination. Despite their wide usage, their impact on the efficiency of the transportation network is not fully understood. In this work, we propose a road traffic model to assess the effects of real-time route

recommendations on an origin-destination pair connected by two possible routes. We suppose that a certain percentage of the users constituting this demand has access to a navigation application, which suggests the route with the shortest travel time. The model aims to highlight potential drawbacks arising when the percentage of routed users is too high. After a comprehensive stability analysis determining the asymptotic behaviour of the system, we show that an excessive use of routing apps leads to deterioration of the network efficiency, corresponding both to the decrease of performance measures and to failure to satisfy user demand. Joint work with Tommaso Toso and Alain Kibangou (University of Grenoble, France).

10 David Grimsman

Brigham Young University, USA

Information, Order, and Utility in Submodular Games

Abstract

Submodular maximization is a well-known problem with a plethora of applications across many disciplines. Since it is NP-Hard, there has been much work devoted to finding approximations to the optimal solution with some level of optimality guarantee. A recent thread of research has been devoted to exploring how these problems can be formulated and solved in a distributed way – specifically as a game where many agents choose actions with the goal of maximizing the quality of the overall action profile. This talk will explore how different network structures of information sharing and utility design among the agents affects the optimality guarantees of the resulting solution.

11 Ziv Hellman

Bar-Ilan University, Israel

Mentors and Recombinators: Multi-Dimensional Social Learning

Abstract

We study games in which the set of strategies is multi-dimensional, and new agents might learn various strategic dimensions from different mentors. We introduce a new family of dynamics, the recombinator dynamics, which is characterised by a single parameter, the recombination rate r in [0,1]. The case of r=0 coincides with the standard replicator dynamics. The opposite case of r=1 corresponds to a setup in which each new agent learns each new strategic dimension from a different mentor, and combines these dimensions into her adopted strategy. We fully characterise stationary states and stable states under these dynamics, and we show that they predict novel behaviour in various applications.

(Joint work with Srinivas Arigapudi, Omer Edhan, and Yuval Heller)

12 Julien Hendrickx

U C Louvain, Belgium

Open Multi-Agents Models with Evolving Composition

Abstract

Almost all theoretical on the evolution of multi-agent systems assume a fixed composition: the set of agents in the system remain the same all along the process considered. By contrast, many phenomena amenable to multi-agent modelling involve frequent arrivals and departures. Similarly, scalability and robustness to agent losses are often cited as advantages of multi-agent solutions.

We consider open multi-agent systems, where agents keep arriving or leaving during the process considered. We discuss the general challenges to analyze such systems - for example, arrivals and departures keep creating perturbations and forbid any classical convergence – and propose a descriptor-based analysis. We then turn to the design challenges: correction mechanisms designed to cope with a small number of arrivals or departures may indeed fail when these events keep taking place. We analyze some

design strategies, focusing in particular on in particular on averaging, decentralized estimation, and decentralized optimization, and present some fundamental performance limitations in open systems.

13 Ryan Chong Luck Kor

National University of Singapore, Singapore
Welfare and Distributional Effects of Joint Intervention in networks

Abstract

We study a planner's optimal interventions in both the standalone marginal utilities of players on a network and weights on the links that connect players. The welfare-maximizing joint intervention exhibits the following properties: (a) when the planner's budget is moderate (so that optimal interventions are interior), the change in weight on any link connecting a pair of players is proportional to the product of eigen-centralities of the pair; (b) when the budget is sufficiently large, the optimal network takes a simple form: It is either a complete network under strategic complements or a complete balanced bipartite network under strategic substitutes. We show that the welfare effect of joint intervention is shaped by the principal eigenvalue, while the distributional effect is captured by the dispersion of the corresponding eigenvectors, i.e., the eigen-centralities. Comparing joint intervention in our setting with single intervention solely on the standalone marginal utilities, as studied by Galeotti et al. (2020), we find that joint intervention always yields a higher aggregate welfare, but may lead to greater inequality, which highlights a possible trade-off between the welfare and distributional impacts of joint intervention.

14 Suraj Malladi

Cornell University, USA
Rational Disagreement and the Fragility of Social Learning

Abstract

We examine how agents learn when information from original sources only reaches them after noisy relay. In the presence of random mutation of message content and transmission failures, there is a sharp threshold such that a receiver fully learns if and only if they have access to more chains than the threshold number and they perfectly understand the noise process. However, even slight uncertainty over the relative rates of mutations makes learning from long chains impossible, no matter how many distinct sources information emerges from. The identification failure is that an agent cannot distinguish uncertainty about the state from uncertainty about the noise in communication. This result rationalizes long-run disagreement: even agents with a common prior and access to an arbitrary number of primary sources can end up with different beliefs if their network positions place them at different distances relative to primary sources.

15 Dario Paccagnan

Imperial College London, UK
The Unintended Consequences of Optimizing the Worst-case Equilibrium Performance

Abstract

Much research within the field of game theory has focused on the design of mechanisms that optimize the performance of either worst-case or best-case equilibria. Examples are plenty and range from auction theory to incentive design in congestion games. Naturally, both approaches are well motivated as bounds on the worst-case equilibrium performance hold universally, i.e. require no behavioral assumption, while best-case equilibria often arise as the result of natural game dynamics. However, these metrics are often studied in isolation, raising the following central question: do mechanisms that improve the worst-case performance have any unintended consequences on the best-case one? In this work we answer the

previous question in the affirmative and fully characterize the trade-off between price of anarchy and price of stability in congestion games with incentives.

16 Nicolò Pagan

University of Zurich, Switzerland

Strategic Coalition Formation among Social Media Content Creators

Abstract

As social media platforms continue to grow in popularity, they are not only becoming a hub for connecting with others but are also increasingly being leveraged as a means of making a living. The Creator Economy has grown rapidly in the past decade, with over 50 million Content Creators (CCs) and a market size of over 104 billion dollars, but fair opportunities and remuneration are not always provided by the social media platforms on which they post their content. To closely study the phenomenon of the rise of social media influencers on today's most popular platforms, we recently proposed a meritocratic network formation model where followers continuously search for the best user-generated content (UGC) provider. This research provided valuable insights into how users form networks based on UGC and highlighted the importance of content quality in social network formation. Further extensions analyzed the impact of the recommender systems on the individual fairness of the CCs, also in the presence of biases in the recommendation as well as of users' homophilic behavior. Building on this prior work, here we investigate to what extent it is beneficial for Content Creators to strategically form coalitions, and whether these coalitions are more efficient when the characteristics of the strategic users are substantially different. Furthermore, we are interested in examining how CCs from minority groups may face unique challenges in this context, and how these challenges can be addressed through interventions that promote equity and fairness on social media platforms.

17 Francesca Parise

Cornell University, USA

Network games with large populations: non-uniqueness and higher-order interactions

Abstract

Understanding the role of network interactions is fundamental for improving efficiency, resilience and welfare of in many social-economic settings. The large size of these systems (e.g., involving billions of users in the case of social platforms) however introduces some challenges from the perspective of a planner that aims at regulating interactions. In fact in many cases, collecting exact network data is either too costly or impossible due to privacy concerns. In these cases however it might be feasible for the planner to collect statistical information about agents' interactions that can be used to infer a random graph model. A key question is then whether knowledge of such a random graph model is sufficient to infer relevant features of the realized network or to control a dynamical process evolving over it. This question has been addressed in a number of recent works by focusing on network games with unique Nash equilibrium and on random graph models in which each link is formed independently. Yet, many relevant networks may exhibit multiple-equilibria and/or higher-order structure that cannot be captured by such models. In this talk we will present novel results in these directions. Specifically, we will present a convergence theory for graphon games with multiple equilibria and concentration inequality results for subgraph generated models (which can be used to capture higher-order interactions).

18 Lacra Pavel

University of Toronto, Canada

Games over Networks: What issues does the network bring in?

Abstract

In this talk, we consider decision-making over networks under the umbrella of game theory. In the classical setting of game theory, each player is assumed to interact with all others and to have complete information. The presence of the network brings in several new issues: complexity of agents' networked interaction, local/partial information, delayed/asynchronous communication, the curse of dimensionality. What's needed is a unified framework for games on/over/in networks - a framework that allows both the analysis of strategic networked interactions, and the design of nearest-neighbour agent-interaction/learning rules. These rules should rely on locally available information, minimize superfluous communication, and by using them, agents should reach an optimal collective state, such as (generalized) Nash equilibrium. In this talk we review some of our group's work towards getting there, focused on operator theoretic and control theoretic approaches.

19 Paolo Pin

Universita degli Studi di Siena, Italy

Revealing information -- or not -- in trading with asymmetric information

Abstract

We propose a micro-founded model of trading, with ex-ante asymmetric information.

We analyze under which conditions the informed trader may want to share her information for free with other traders.

We find that, despite the strictly competitive setup and conventional wisdom, there is a unique separating equilibrium in which the informed trader reveals some signals and conceals others.

This implies that the price is not fully revealing of the aggregate information in the market, even if traders are risk neutral.

We find results in equilibrium when there is a network between traders, so that not any trader can inform any other trader: the interval for which a trader will reveal or not depends on her centrality in the network.

20 Ketan Savla

University of Southern California, USA

Information Design for Non-atomic Games: Computation, Repeated Setting, and Experiment

Abstract

Bayesian information design is a compelling framework for generating persuasive recommendations. Feasibility of a recommendation policy is often characterized by the "obedience constraint" according to which following recommendation is, in a posteriori expectation, no worse than any other action. The probabilistic recommendation policy has continuous support for non-atomic games. This makes it computationally challenging to handle the obedience constraint for optimal design, and further questions the practical validity of Bayesian computation in decision-making for non-atomic setting. We address these issues in the context of routing games. For polynomial link latency functions, we provide an exact finite-dimensional formulation of the obedience constraint. We also study a repeated setting, in which the likelihood of an agent following a recommendation in a stage is proportional to the population-wide average regret from previous stages, e.g., provided by a review aggregator platform. We present results

from a human subject experiment on validity of this "learning model" and establish its convergence to Bayes correlated equilibrium under obedient recommendation policy.

21 Muhammed Sayin

Bilkent University, Türkiye
Learning in Stochastic Games

Abstract

Reinforcement learning (RL) has been the backbone of many frontier artificial intelligence (AI) applications, such as game playing and autonomous driving, by addressing how intelligent and autonomous systems should engage with an unknown dynamic environment. The progress and interest in AI are now transforming social systems with human decision-makers, such as (consumer/financial) markets and road traffic, into socio-technical systems with AI-powered decision-makers. However, self-interested AI can undermine the social systems designed and regulated for humans. We are delving into the uncharted territory of AI-AI and AI-human interactions. The new grand challenge is to predict and control the implications of AI selfishness in AI-X interactions with systematic guarantees. Hence, there is now a critical need to study self-interested AI dynamics in complex and dynamic environments through the lens of game theory.

In this talk, I will present the recent steps we have taken toward the foundation of how self-interested AI would and should interact with others by bridging the gap between game theory and practice in AI-X interactions. I will specifically focus on stochastic games to model the interactions in complex and dynamic environments since they are commonly used in multi-agent reinforcement learning. I will present new learning dynamics converging almost surely to equilibrium in important classes of stochastic games. The results can also be generalized to the cases where (i) agents do not know the model of the environment, (ii) do not observe opponent actions, (iii) can adopt different learning rates, and (iv) can be selective about which equilibrium they will reach for efficiency. The key idea is to use the power of approximation thanks to the robustness of learning dynamics to perturbations. I will conclude my talk with several remarks on possible future research directions for the framework presented.

22 Marc Schroder

Maastricht University, Netherlands
Negative Prices in Network Pricing Games}

Abstract

In a Stackelberg network pricing game a leader sets prices for a given subset of edges so as to maximize revenue, after which one or multiple followers choose a shortest path. Our main results study the impact of the seemingly innocent assumption that prices are non-negative. First, we characterize the class of graphs in which negative prices cannot increase the revenue. Second, we characterize the difference in revenue between allowing negative prices and assuming non-negative prices.

23 Jeff Shamma

University of Illinois at Urbana-Champaign, USA
Higher order uncoupled dynamics do not lead to Nash equilibrium–except when they do

Abstract

One line of investigation in the "learning in games" framework concerns whether learning dynamics may or may not lead to Nash equilibrium. There is a multitude of specific cases of learning dynamics/game combinations that illustrate outcomes ranging from convergence to limit cycles to chaotic behaviour. A natural characteristic of these learning dynamics is that they are "uncoupled", i.e., the learning dynamics of

an agent cannot depend explicitly on the utility functions of other agents. Prior work constructed a specific anti-coordination game for which no learning dynamics can converge to the (unique mixed strategy) Nash equilibrium of this game. The class of admissible uncoupled dynamics considered therein had a restricted order associated with the dimension of the game. Subsequent work showed that relaxing this restriction, i.e., allowing for uncoupled but higher order dynamics, can overcome non-convergence for the very same anti-coordination game. This talk further explores the relation between higher order learning dynamics and Nash equilibrium. Our notion of "uncoupled" is that a player's learning dynamics depend on its own actions and its own evolving payoff vector, much like various no regret learning algorithms. Such dynamics do not depend explicitly on the utility functions of other agents as well as their actions. In the context of finite games, we first show that for any uncoupled higher order learning dynamics, there exists a game for which the dynamics do not converge to the unique mixed strategy Nash equilibrium. Simultaneously, for any game with an isolated mixed strategy Nash equilibrium, there exist uncoupled higher order learning dynamics that can lead to Nash equilibrium for that game as well as nearby games. Informally, any uncoupled higher order learning dynamics can be destabilized by the right game, and any mixed strategy Nash equilibrium can be stabilized (locally) by the right uncoupled higher order dynamics. We further show that stabilizing mixed strategy equilibria can come at the cost of the learning dynamics themselves possessing an inherent internal instability. This is joint work with Sarah Toonsi.

24 James Siderius

Massachusetts Institute of Technology, USA
A Model of Online Misinformation

Abstract

We present a network model of online content sharing where agents sequentially observe an article and decide whether to share it with others. This content may or may not contain misinformation. Agents gain utility from positive social media interactions but do not want to be called out for propagating misinformation. We characterize the (Bayesian-Nash) equilibria of this social media game and show sharing exhibits strategic complementarity. Our first main result establishes that the impact of homophily on content virality is non-monotone: homophily reduces the broader circulation of an article, but it creates echo chambers that impose less discipline on the sharing of low-reliability content. This insight underpins our second main result, which demonstrates that social media platforms interested in maximizing engagement tend to design their algorithms to create more homophilic communication patterns ("filter bubbles"). We show that platform incentives to boost the diffusion of an article are particularly pronounced for low-reliability content likely to contain misinformation and when there is greater polarization and more divisive content. Finally, we discuss various regulatory solutions to such platform-manufactured misinformation.

25 Fernando Vega-Redondo

Bocconi University, Italy
Riot Networks and the Tullock Paradox: An application to the Egyptian Arab Spring

Abstract

We study a dynamic model of collective action – for concreteness, we speak of a riot – in which agents interact through, and learn from, a co-evolving social network. We consider two different scenarios on how agents form their expectations when changing their behaviour. In one of them, conceived as a "benchmark", they are assumed to be completely informed of the prevailing state (action profile and network). Instead, in the alternative scenario, agents are assumed to shape their expectations about the state from a combination of local observation and social learning (modelled à la DeGroot). In both cases we provide a complete characterization of the long-run behaviour of the system. While the first assumption of complete information is common, the second one is arguably more realistic. Furthermore, we show that only the

latter assumption yields the following twin conclusion: a significant long-run probability of successful collective action and a meaningful time scale of convergence to this state of affairs. This, we argue, suggests a plausible route to understanding what otherwise seems a puzzle, i.e. how do very large populations attain ("coordinate on") collective action. Finally, we illustrate the empirical potential of the model by showing that it can be efficiently estimated for the so-called Egyptian Arab Spring using large-scale cross sectional data on agents' choices and their network derived from Twitter.

26 Xavier Venel

Luiss University, Italy

Competition between lobbies in a Degroot framework

Abstract

We analyze a class of stochastic games where two lobbies compete by influencing the opinions in a society. We assume that the opinions evolve according to De Groot opinion formation and that the decisions of the lobbies change the structure of the network representing the society. Based on two recent papers, I will present several results on the regularity of discounted Nash equilibrium payoffs when players become patient and on how the players should play in some particular cases.

27 Hoi-To Wai

The Chinese University of Hong Kong, China

On Multi-agent Performative Prediction Game over Multiplex Networks

Abstract

This talk presents recent results on a multi-agent performative prediction (Multi-PP) game described by a multiplex, i.e., multi-layer, network. We concentrate on a setting motivated by distributed personalized learning where agents cooperate on a network (called agent network) via a graph regularization design; while during training, the data samples drawn can be influenced by the prediction models of neighbouring agents on another network (called population network). We formulate the Multi-PP game as a multiplex network game and study the interplay between the agent and population networks.

We present the following findings on the Multi-PP game. First, we analyze sufficient conditions for the existence of performative stable equilibrium (PSE) and Nash equilibrium (NE), which are fixed points of distributed repeated minimization procedures. Second, we analyze the changes on the equilibria with perturbed data distributions, and derive the closed-form solutions where the network topologies are explicit. Our results reveal effects of the networks' structure such as asymmetric links, and it yields an interpretation via the Bonacich node centrality, to list a few. We also discuss a special case that illustrates the benefit of full cooperation among agents, if the latter is possible to attain.