Theme C: Incentive Mechanisms for Participation

Overall Goals

- Design algorithms (e.g. auctions, multi-round games), where the inputs are provided by correlated strategic agents (e.g. by a network of supply chains), and the result is truthful, max. welfare, etc.
- Design mechanisms for dynamic environments, where the actors are myopic and have asymmetric, limited information, and are subject to external forces (weather, pricing, trade)
- Incentivize players to strategically experiment and learn from data generated by players with incentives.

Challenges

- Policies and mechanisms are highly dynamic and must adapt to intrinsic and extrinsic factors to the network, such as changes in weather, prices, and trade policy.
- Small versus big players: big players (e.g. companies) are consortia of small players; their interests are correlated but not identical
- In general, incentive compatible mechanisms do not compose well.

Problem setup for dynamics on networks

Goal. Study dynamics on networks, where agents continually adapt to the observed state of the world.

Challenge.

- Players may take repeated myopic decisions that are in their best interest in the short term, but not the long term (e.g., each agent acts based on local information from its neighborhood).
- Some myopic actions – e.g., cutting production costs in places difficult to observe – can lead to suboptimal products or the market unraveling completely over time.

Approaches.

- Model production markets where players are connected, design and study classes of mechanisms that lead to growth of the market in the long run under different types of behavior of the players (e.g. regret minimization, multiplicative weight updates, proportional updates).

Our prior work on production markets

Setup

- Set of n producers, each makes good according to their production function
- The players bid on the goods at the market, exchange their goods for other goods they need, then use the bundles acquired for the next production.
- A simple dynamic is proportional response, where players update their bids proportionally to how useful the investments were in the past round.

Main result.

- This dynamic leads to growth of the market in the long term (whenever growth is possible) but also creates unbounded inequality, i.e. very rich and very poor players emerge over time.

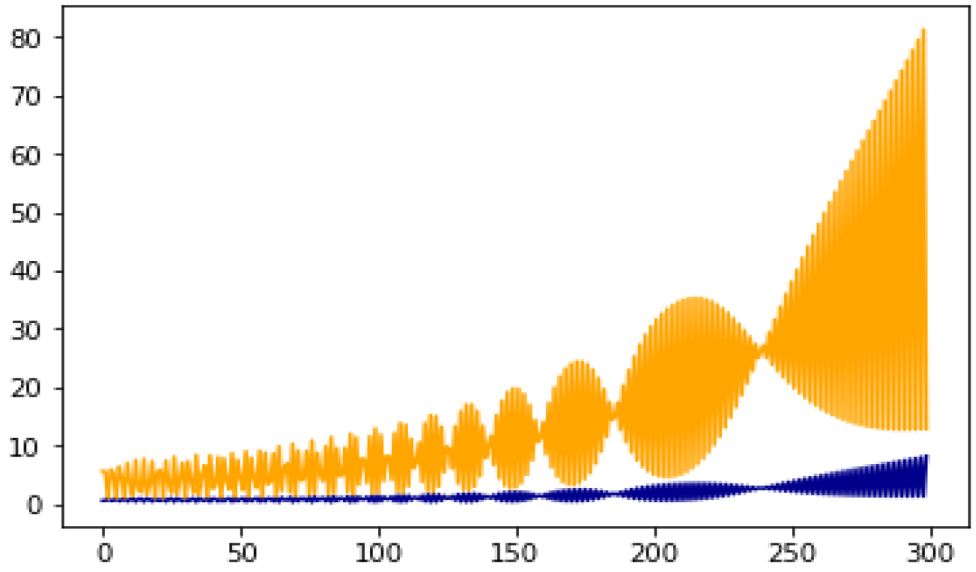

Oceanic markets with big and small players

Goal.

- Study market games with small versus big players: big players (e.g. companies) are consortia of small players; their interests are correlated but not identical.
- Related literature on theory of “oceanic games” (starting with Milnor and Shapley ‘77)

Example. Parametric cooperation

- Let cooperation parameter λ can model how aligned the interests of two small players are.
- Say Alice and Bob are dividing some land. Alice has value v_A for her piece and Bob has v_B for his, then Alice’s utility for this allocation is $u_A = v_A + \lambda\, v_{B.}$
- If Alice and Bob have partially aligned interests (e.g. form a community), this could be modeled by $\lambda=½$, so Alice’s utility is $u_A = v_A + ½\, v_B$.
- In auction allocations and rent division, these are unexplored dimensions of the problem. ($\lambda = 0$ is independent, $\lambda = 1$ is fully cooperative, $\lambda = -1$ is zero-sum)

Research directions

- Resource allocation (auctions, markets) and learning problems where players are not independent of each other (e.g. have correlated utilities, are big vs small).

Cooperation can have tremendously beneficial effects

Goal.

- Design mechanisms that incentivize players to experiment – too much competition between players can reduce risk tolerance.

Example and our prior work.

- Companies acquiring data and experimenting with different strategies for their products, too much competition makes players risk-adverse as failure is heavily penalized.
- Formalized as multiple players learning at the same time in the bandit model.
- When the decision is between a risky action and a predictable one, competing players (playing a zero sum game) explore less than a single player, while cooperating players explore more.

Research directions

- How should market mechanisms be set up so that players explore at the correct rate? What learning algorithms should the players use? (e.g. poly-time problems vs. NP-hard)
- What happens when people hide information from competitors?

Incentives to share rewards and move distributions require innovative mechanisms

Given a complex supply chain or distribution system, who should pay for a contamination scenario?

- Extensive local testing at each node may be too expensive.
- Idea: limited local testing and group testing
- Desiderata: low sample complexity, low interaction, the testing preserves the incentives.

Mechanisms that influence global distribution

- How can one design a mechanism M (e.g. a system of prices, rewards, fines), so that by repeatedly applying M the population moves from the initial distribution to the desired distribution?
- This problem requires modelling the feedback of participants to a mechanism, (e.g. networks of polya urn-like processes).

Auctions in the wild

This setting requires auctions and markets for selling products, but optimal auction theory cannot always be used to allocate resources:

- Sellers don’t have entire distributions of valuations as is assumed in optimal auction theory, but can use information from past sales (e.g. prices at which corn was sold). Mechanisms must be designed from samples.
- The market is inherently dynamic -- e.g., a participant may be unable to deliver the target production due to extrinsic factors (weather). This can require designing contracts between buyers and sellers. How should auctions be designed with contracts?
- We want polynomial time, economic efficiency, revenue and welfare guarantees.
- Optimal mechanisms can be very complex and hard to use – which will inhibit adoption. Addressing this will require us to understand tradeoffs between efficiency and “simplicity” (or complexity – computational, communication, posted prices versus auctions).

Measuring properties in networks of strategic agents from samples.

Application-based example.

- Based on resources allocated to them, farmers adjust their production or level of effort made. One would like to learn their incentives and valuations.
- Put simply, what can be inferred from a reluctant supply chain based on the samples we observe.

CS Formalism.

- Infer valuations of the players from their responses (actions) to a mechanism. We will start by assuming players are Bayesian or learning from data and minimizing regret over time.
- When the agents form a network, infer properties of a network from samples, such as the edge weights (this is typically studied when the nodes are not strategic).