Abstracts book

Submission index: [17], [22], [25], [27], [28], [29], [30], [31], [34], [35], [39], [40], [42], [44], [52], [57], [66], [72], [78], [80], [90], [92], [98], [103], [105], [107], [124], [125], [131], [136], [151], [152], [155], [156], [158], [159], [162], [169], [183], [188], [189], [190], [193], [195], [199], [203], [207], [219], [220], [224], [229], [232], [237], [239], [243], [257], [266], [270], [278], [279], [280], [282], [285], [288], [308], [309], [313], [314], [315], [319], [326], [340], [341], [343], [349], [351], [355], [362], [363], [370], [373], [385], [397], [400], [408], [411], [412], [422], [423], [428], [430], [431], [444], [450], [456], [463], [483], [486], [491]

[17] *One Dollar Each Eliminates Envy*

Johannes Brustle (McGill University), Jack Dippel (McGill University), Vishnu Narayan (McGill University), Mashbat Suzuki (McGill University) and Adrian Vetta (McGill University).

Abstract

We study the fair division of a collection of m indivisible goods amongst a set of n agents. Whilst envy-free allocations typically do not exist in the indivisible goods setting, envy-freeness can be achieved if some amount of a divisible good (money) is introduced. Specifically, Halpern and Shah (SAGT 2019, pp.374-389) showed that, given additive valuation functions where the marginal value of each item is at most one dollar for each agent, there always exists an envy-free allocation requiring a subsidy of at most $(n-1) m$ dollars. The authors also conjectured that a subsidy of $n-1$ dollars is sufficient for additive valuations. We prove this conjecture. In fact, a subsidy of at most one dollar per agent is sufficient to guarantee the existence of an envy-free allocation. Further, we prove that for general monotonic valuation functions an envy-free allocation always exists with a subsidy of at most $2(n-1)$ dollars per agent. In particular, the total subsidy required for monotonic valuations is independent of the number of items.

---

[22] *Two-sided random matching markets: ex-ante equivalence of the deferred acceptance procedures*

Simon Mauras (IRIF).

Abstract

Stable matching in a community consisting of N men and N women is a classical combinatorial problem that has been the subject of intense theoretical and empirical study since its introduction in 1962 in a seminal paper by Gale and Shapley.

When the input preference profile is generated from a distribution, we study the output distribution of two stable matchings procedures: women-proposing-deferred-acceptance and men-proposing-deferred-acceptance. We show that the two procedures are ex-ante equivalent: that is, under certain conditions on the input distribution, their output distribution are identical.

In terms of technical contributions, we generalize (to the non-uniform case) an integral formula, due to Knuth and Pittel, which gives the probability that a fixed matching is stable. Using an inclusion-exclusion principle on the set of rotations, we give a new formula which gives the probability that a fixed matching is the women/men-optimal stable matching. We show that those two probabilities are equal with an integration by substitution.

---

[25] *Efficient Matchmaking in Assignment Games with Application to Online Platforms*

Peng Shi (University of Southern California).

Abstract

An assignment game is a model of two-sided markets with transfer payments, and can be used to study the matching between heterogeneous customers and service providers with endogenous determination of prices. We study how a market intermediary can best facilitate matches in an assignment game so as to achieve a good outcome while minimizing the number of bits of information agents need to communicate. Here, ``good outcomes'' are formalized as ε-

stable, which means that no pair of agents can deviate and both benefit by more than ε. We show that an ε-stable outcome can be found using only O(log n) bits of communication per agent whenever the horizontal component of one side's preferences is predictable or dense. ("Dense" means that those agents have high preferences for many potential partners.) But matchmaking requires Ω(sqrt{n}) bits per agent when the horizontal preferences of both sides are unpredictable and sparse, meaning that both sides have high preferences only for a vanishing proportion of the market. We propose near optimal sequential protocols for both dense and sparse preferences, as well as a near-optimal simultaneous protocol for thick markets with dense preferences on both sides. The protocols yield prescriptive insights for online platforms in the home services industry.

---

[27] *Strategic Formation and Reliability of Supply Chain Networks*

Victor Amelkin (University of Pennsylvania) and Rakesh Vohra (University of Pennsylvania).

Abstract

Supply chains are the backbone of the global economy. Disruptions to them can be costly. Centrally managed supply chains invest in ensuring their resilience. Decentralized supply chains, however, must rely upon the self-interest of their individual components to maintain the resilience of the entire chain.

We examine the incentives that independent self-interested agents have in forming a resilient supply chain network in the face of production disruptions and competition. In our model, competing suppliers are subject to yield uncertainty (they deliver less than ordered) and congestion (lead time uncertainty or, "soft" supply caps). Competing retailers must decide which suppliers to link to based on both price and reliability.

In the presence of yield uncertainty only, the resulting equilibrium supply chain networks are sparse. Retailers concentrate their links on a single supplier, counter to the idea that they should mitigate yield uncertainty by diversifying their supply base. This happens because retailers benefit from supply variance. It suggests that competition will amplify output uncertainty. When congestion is included as well, the resulting networks are denser and resemble bipartite expander graphs that have been proposed in the supply chain literature, thereby, providing the first example of endogenous formation of resilient supply chain networks, without resilience being explicitly encoded in payoffs. Finally, we show that a supplier's investments in improved yield can make it worse off. This happens because high production output saturates the market, which, in turn lowers prices and profits for participants.

---

[28] *Quick or cheap? Breaking points in dynamic markets*

Panayotis Mertikopoulos (CNRS), Heinrich Nax (ETH Zurich) and Bary Pradelski (CNRS).

Abstract

We examine two-sided markets where players arrive stochastically over time and are drawn from a continuum of types. The cost of matching a client and provider varies, so a social planner is faced with two contending objectives: a) to reduce players' waiting time before getting matched; and b) to form efficient pairs in order to reduce matching costs. We show that such markets are characterized by a quick or cheap dilemma: Under a large class of distributional assumptions, there is no `free lunch', i.e., there exists no clearing schedule that is simultaneously optimal along both objectives. We further identify a unique breaking point signifying a stark reduction in matching cost contrasted by an increase in waiting time. Generalizing this model, we identify two regimes: one, where no free lunch exists; the other, where a window of opportunity opens to achieve a free lunch. Remarkably, greedy scheduling is never optimal in this setting.

---

[29] *Prophet Inequalities with Linear Correlations and Augmentations*

Nicole Immorlica (Microsoft Research), Sahil Singla (Princeton University and the Institute for Advanced Study) and Bo Waggoner (University of Colorado).

Abstract

In a classical online decision problem, a decision-maker who is trying to maximize her value inspects a sequence of arriving items to learn their values (drawn from known distributions), and decides when to stop the process by taking the current item. The goal is to prove a ``prophet inequality'': that she can do approximately as well as a prophet with foreknowledge of all the values. In this work, we investigate this problem when the values are allowed to be correlated. Since non-trivial guarantees are impossible for arbitrary correlations, we consider a natural ``linear'' correlation structure introduced by Bateni et al.~\cite{BDHS-ESA15} as a generalization of the common-base value model of Chawla et al.~\cite{CMS-EC10Journal15}.

A key challenge is that threshold-based algorithms, which are commonly used for prophet inequalities, no longer guarantee good performance for linear correlations. We relate this roadblock to another ``augmentations'' challenge that might be of independent interest: many existing prophet inequality algorithms are not robust to slight increase in the values of the arriving items. We leverage this intuition to prove bounds (matching up to constant factors) that decay gracefully with the amount of correlation of the arriving items. We extend these results to the case of selecting multiple items by designing a new $(1+o(1))$ approximation ratio algorithm that is robust to augmentations.

---

[30] *Menu-size Complexity and Revenue Continuity of Buy-many Mechanisms*

Shuchi Chawla (University of Wisconsin-Madison), Yifeng Teng (University of Wisconsin-Madison) and Christos Tzamos (University of Wisconsin-Madison).

Abstract

We study the multi-item mechanism design problem where a monopolist sells n heterogeneous items to a single buyer. We focus on buy-many mechanisms, a natural class of mechanisms frequently used in practice. The buy-many property allows the buyer to interact with the mechanism multiple times instead of once as in the more commonly studied buy-one mechanisms. This imposes additional incentive constraints and thus restricts the space of mechanisms that the seller can use.

In this paper, we explore the qualitative differences between buy-one and buy-many mechanisms focusing on two important properties: revenue continuity and menu-size complexity.

Our first main result is that when the value function of the buyer is perturbed multiplicatively by a factor of $1\pm\e$, the optimal revenue obtained by buy-many mechanisms only changes by a factor of $1 \pm \textrm{poly}(n,\e)$. In contrast, for buy-one mechanisms, the revenue of the resulting optimal mechanism after such a perturbation can change arbitrarily.

Our second main result shows that under any distribution of arbitrary valuations, finite menu size suffices to achieve a $(1-\e)$-approximation to the optimal buy-many mechanism. We give tight upper and lower bounds on the number of menu entries as a function of n and $\e$. On the other hand, such a result fails to hold for buy-one mechanisms as even for two items and a buyer with either unit-demand or additive valuations, the menu-size complexity of approximately optimal mechanisms is unbounded.

---

[31] *Optimal Mechanism Design for Single-Minded Agents*

Nikhil Devanur (Amazon), Kira Goldner (Columbia University), Raghuvansh Saxena (Princeton University), Ariel Schvartzman (Princeton University) and S. Matthew Weinberg (Princeton University).

Abstract

We consider optimal (revenue maximizing) mechanism design in the interdimensional setting, where one dimension is the ‘value’ of the buyer, and the other is a ‘type’ that captures some auxiliary information. A prototypical example of this is the FedEx Problem, for which Fiat et al. [2016] characterize the optimal mechanism for a single agent. Another example of this is when the type encodes the buyer’s budget [DW17]. The question we address is how far can such characterizations go? In particular, we consider the setting of single-minded agents. A seller has heterogenous items. A buyer has a valuation v for a specific subset of items S, and obtains value v if and only if he gets all the items in S (and potentially some others too). We show the following results. 1. Deterministic mechanisms (i.e. posted prices) are optimal for distributions that satisfy the “declining marginal revenue” (DMR) property. In this case we give an explicit construction of the optimal mechanism. 2. Without the DMR assumption, the result depends on the structure of the minimal directed acyclic graph (DAG) representing the partial order among types. When the DAG has out- degree at most 1, we characterize the optimal mechanism `a la FedEx; this can be thought of as a generalization of the FedEx characterization since FedEx corresponds to a DAG that is a line. 3. Surprisingly, without the DMR assumption and when the DAG has at least one node with an out- degree of at least 2, then we show that there is no hope of such a characterization. The minimal such example happens on a DAG with 3 types. We show that in this case the menu complexity is unbounded in that for any M, there exist distributions over (v,S) pairs such that the menu complexity of the optimal mechanism is at least M. 4. For the case of 3 types, we also show that for all distributions there exists an optimal mechanism of finite menu complexity. This is in contrast to the case where you have 2 heterogenous items with additive utilities for which the menu complexity could be uncountably infinite [MV07; DDT15]. In addition, we prove that optimal mechanisms for Multi-Unit Pricing (without a DMR assumption) can have unbounded menu complexity as well, and we further propose an extension where the menu complexity of optimal mechanisms can be countably infinite, but not uncountably infinite. Taken together, these results establish that optimal mechanisms in interdimensional settings are both surprisingly richer than single-dimensional settings, yet also vastly more structured than multi-dimensional settings.

[34] *Designing Informative Rating Systems: Evidence from an Online Labor Market*

Nikhil Garg (Stanford University) and Ramesh Johari (Stanford University).

Abstract

Platforms critically rely on rating systems to learn the quality of market participants. In practice, however, these ratings are often highly inflated, and therefore not very informative. In this paper, we first investigate whether the platform can obtain less inflated ratings by altering the {\em meaning} and {\em relative importance} of the {\em levels} in the rating system. Second, we seek a principled approach to make these choices in the design of the rating system.

First, we analyze the results of a randomized controlled trial on an online labor market in which an additional question was added to the feedback form. Between treatment conditions, we vary the question phrasing and answer choices; in particular, the treatment conditions include several {\em positive-skewed verbal rating scales} with descriptive phrases or adjectives providing specific interpretation for each rating level. The online labor market test reveals that current inflationary norms can in fact be countered by re-anchoring the meaning of the levels of the rating system. In particular, the positive-skewed verbal rating scales yield rating distributions that significantly reduce rating inflation and are much more informative about seller quality.

Second, we develop a model-based framework to compare and select among rating system designs, and apply this framework to the data obtained from the online labor market test. Our simulations demonstrate that our model-based framework for scale design and optimization can identify the most informative rating system and substantially improve the quality of information obtained over baseline designs.

Overall, our study illustrates that rating systems that are informative in practice {\em can} be designed, and demonstrates {\em how} to design them in a principled manner.

---

[35] *Incentive-Compatible Selection Mechanisms for Forests*

Oren Dean (Technion - Israel Institute of Technology), Moshe Tennenholtz (Technion - Israel Institute of Technology) and Yakov Babichenko (Technion - Israel Institute of Technology).

Abstract

Given a directed forest-graph, a probabilistic selection mechanism is a probability distribution over the vertex set. A selection mechanism is incentive-compatible (IC), if the probability assigned to a vertex does not change when we alter its outgoing edge (or even remove it). The quality of a selection mechanism is the worst-case ratio between the expected progeny under the mechanism's distribution and the maximal progeny in the forest. In this paper we prove an upper bound of 4/5 and a lower bound of 1/ln16~0.36 for the quality of any IC selection mechanism. The lower bound is achieved by two novel mechanisms and is a significant improvement to the results of Babichenko et al. (WWW 2018). The first, simpler mechanism, has the nice feature of generating distributions which are fair (i.e., monotone and proportional). The downside of this mechanism is that it is not exact (i.e., the probabilities might sum-up to less than 1). Our second, more involved mechanism, is exact but not fair. We also prove an impossibility for an IC mechanism that is both exact and fair and has a positive quality.

---

[39] *Voluntary Disclosure and Personalized Pricing*

S. Nageeb Ali (Pennsylvania State University), Greg Lewis (Microsoft Research) and Shoshana Vasserman (Stanford University).

Abstract

A concern central to the economics of privacy is that firms may use consumer data to price discriminate. A common response is that consumers should have control over their data and the ability to choose how firms access it. Since firms draw inferences based on both the data seen as well as the consumer's disclosure choices, the strategic implications of this proposal are unclear. We investigate whether such measures improve consumer welfare in monopolistic and competitive environments. We find that consumer control can guarantee gains for every consumer type relative to both perfect price discrimination and no personalized pricing. This result is driven by two ideas. First, consumers can use disclosure to amplify competition between firms. Second, consumers can share information that induces a seller---even a monopolist---to make price concessions. Furthermore, whether consumer control improves consumer surplus depends on both the technology of disclosure and the competitiveness of the marketplace. In a competitive market, simple disclosure technologies such as ``track / do-not-track'' suffice for guaranteeing gains in consumer welfare. However, in a

monopolistic market, welfare gains require richer forms of disclosure technology whereby consumers can decide how much information they would like to convey.

---

[40] *Matching Algorithms for Blood Donation*

Duncan Mcelfresh (University of Maryland), Sergey Pupyrev (Facebook), Christian Kroer (Columbia University), Karthik Sankararaman (Facebook), John Dickerson (University of Maryland), Eric Sodomka (Facebook), Zack Chauvin (Facebook) and Neil Dexter (Facebook).

Abstract

Managing perishable inventory, such as blood stock awaiting use by patients in need, has been a topic of research for decades. Yet, most research focuses on the effective use of this scarce resource across the blood supply chain, and assumes the supply of blood itself can be impacted only via coarse policy levers. In this paper, empowered by the recently-deployed Facebook Blood Donation tool, we choose instead to model the first stage of the full blood supply chain---that is, the supply of blood itself---as a matching market. Here, potential blood donors are matched to donation centers by way of a central recommendation engine; that engine can at some cost prompt (e.g., via push notification) individual donors to donate to a preferred center or centers. Potential donors may have other constraints (e.g., maximum allowable frequency of prompts) or preferences (e.g., geographic, social) that should be taken into account. We develop policies for matching potential blood donors to donation centers under these constraints and preferences, and under simple models of demand shocks to the system. We test these policies in both computational simulations and real-world experiments on the Facebook Blood Donation tool. In both the simulated and real experiments, these matching-based notification policies significantly increase the expected number of blood donations.

---

[42] *Correlation Neglect in Student-to-School Matching*

Ran Shorrer (Penn State), Chloe Tergiman (Penn State) and Alex Rees-Jones (Cornell).

Abstract

A growing body of evidence suggests that decision-makers fail to account for correlation in signals that they receive. We study the relevance of this mistake in students' interactions with school-choice matching mechanisms. In a lab experiment presenting simple and incentivized school-choice scenarios, we find that subjects tend to follow optimal application strategies when schools' admissions decisions are determined independently. However, when schools rely on a common priority---inducing correlation in admissions---decision making suffers: application strategies become substantially more aggressive and fail to include attractive ``safety'' options. We document that this pattern holds even within-subject, with significant fractions of participants applying to different programs when correlation is varied but all payoff-relevant elements are held constant. We provide a battery of tests suggesting that this phenomenon is at least partially driven by correlation neglect, and we discuss implications that arise for the design and deployment of student-to-school matching mechanisms.

---

[44] *One for all, all for one--von Neumann, Wald, Rawls, and Pareto*

Mehmet Ismail (King's College London).

Abstract

Applications of the maximin criterion extend beyond economics to statistics, computer science, politics, and operations research. However, the maximin criterion---be it von Neumann's, Wald's, or Rawls'---draws fierce criticism due to its extremely pessimistic stance. I propose a novel concept, dubbed the optimin criterion, which is based on (Pareto) optimizing the worst-case payoffs of tacit agreements. The optimin criterion generalizes and unifies results in various fields: It not only coincides with (i) Wald's statistical decision-making criterion when Nature is antagonistic, (ii) the core in cooperative games when the core is nonempty, though it exists even if the core is empty, but also generalizes (iii) Nash equilibrium in n-person constant-sum games, (iv) stable matchings in matching models, and (v) competitive equilibrium in the Arrow-Debreu economy. Moreover, every Nash equilibrium satisfies the optimin criterion in an auxiliary game.

---

[52] *Ordered Search with Boundedly Rational Consumers*

Mauro Papi (University of Aberdeen).

Abstract

The literature on ordered search has assumed consumers to search optimally. In contrast, I investigate a price-competition model in which consumers are boundedly rational and firms use persuasive advertising to influence the consumers' aspiration price, i.e., the price regarded as 'satisfactory'. I consider various variants of the model capturing different consumers' second-best strategies if they do not find any satisfactory price. I derive a number of results including (i) predictions about the correlation between firm prominence and important market indicators, such as profits and conversion rates, that can help explain empirical evidence, (ii) extensions of the basic model, such as an analysis of the relationship between the consumer satisfaction rate and firms profits, and (iii) policy implications of the model.

---

[57] ***Minimum Earnings Regulation and the Stability of Marketplaces***

Arash Asadpour (Lyft Marketplace Labs), Ilan Lobel (New York University) and Garrett van Ryzin (Lyft Marketplace Labs).

Abstract

We build a model to study the implications of utilization-based minimum earning regulations of the kind recently enacted by New York City for its ride-hailing providers. We identify the precise conditions under which a utilization-based minimum earnings rule causes marketplace instability, where stability is defined as the ability of platforms to keep wages bounded while maintaining the current flexible (free-entry) work model. We also calibrate our model using publicly available data, showing the limited power of the law to increase earnings within an open marketplace. We argue that affected ride-hailing companies might respond to the law by reducing driver flexibility.

---

[66] ***The Equilibrium Existence Duality: Equilibrium with Indivisibilities & Income Effects***

Elizabeth Baldwin (University of Oxford), Omer Edhan (University of Manchester), Ravi Jagadeesan (Harvard University), Paul Klemperer (University of Oxford) and Alexander Teytelboym (University of Oxford).

Abstract

We show that, with indivisible goods, the existence of competitive equilibrium fundamentally depends on agents' substitution effects, not their income effects. Our Equilibrium Existence Duality allows us to transport results on the existence of equilibrium from transferable-utility economies to settings with income effects. One consequence is that net substitutability—which is a strictly weaker condition than gross substitutability—is sufficient for the existence of equilibrium. We also extend the "demand types'' classification of valuations (Baldwin and Klemperer, 2019), and state a "Unimodularity Theorem with Income Effects'' that gives conditions on the patterns of substitution effects that guarantee the existence of competitive equilibrium.

---

[72] ***Balancing Agent Retention and Waiting Time in Service Platforms***

Marcelo Olivares (Universidad de Chile), Andres Musalem (Universidad de Chile) and Daniel Yung (Universidad de Chile).

Abstract

In many service industries the speed of service and support by experienced employees are two major drivers of service quality. When demand for a service is variable and the staffing requirements cannot be adjusted in real-time, choosing capacity levels requires making a trade-off between service speed and operating costs. Online service platforms have crowdsourcing of a large pool of employees with flexible working hours that are compensated through piece-rates. While this business model can operate at low levels of utilization without increasing operating costs, a different trade-off emerges: the service platform must control employee turnover, which may increase when employees are working at low levels of utilization. Hence, to make staffing decisions and manage workload, it is necessary to empirically measure the trade-off between customer conversion and employee retention. In this context, we study an online service platform that operates with a pool of flexible agents working remotely to sell auto insurance. We develop an econometric approach to model customer behavior that captures two key features of outbound calls: customer time-sensitivity and employee heterogeneity. We find a strong impact of waiting time on customer behavior: conversion rates drop by 33% when the time to make the first outbound call increases from 5 to 30 minutes. In addition, we use a survival model to measure how agent retention is affected by the assigned workload and find that a 10% increase in workload translates into a 25% percentage decrease in weekly agent attrition. These empirical models of customer and agent behavior are

combined to illustrate how to balance customer conversion and employee retention, showing that both are relevant to plan staffing and allocate workload in the context of an on-demand service platform.

---

[78] *Proportionality and the Limits of Welfarism*

Dominik Peters (Carnegie Mellon University) and Piotr Skowron (TU Berlin).

Abstract

We study two influential voting rules proposed in the 1890s by Phragmen and Thiele, which elect a committee of k candidates which proportionally represents the voters. Voters provide their preferences by approving an arbitrary number of candidates. Previous work has proposed proportionality axioms satisfied by Thiele's rule (now known as Proportional Approval Voting, PAV) but not by Phragmen's rule. By proposing two new proportionality axioms (laminar proportionality and priceability) satisfied by Phragmen but not Thiele, we show that the two rules achieve two distinct forms of proportional representation. Phragmen's rule ensures that all voters have a similar amount of influence on the committee, and Thiele's rule ensures a fair utility distribution.

Thiele's rule is a welfarist voting rule (one that maximizes a function of voter utilities). We show that no welfarist rule can satisfy our new axioms, and we prove that no such rule can satisfy the core. Conversely, some welfarist fairness properties cannot be guaranteed by Phragmen-type rules. This formalizes the difference between the two types of proportionality. We then introduce an attractive committee rule which satisfies a property intermediate between the core and extended justified representation (EJR). It satisfies laminar proportionality, priceability, and is computable in polynomial time. We show that our new rule provides a logarithmic approximation to the core. On the other hand, PAV provides a factor-2 approximation to the core, and this factor is optimal for rules that are fair in the sense of the Pigou--Dalton principle.

---

[80] *Machine Learning Instrument Variables for Causal Inference*

Amandeep Singh (The Wharton School, University of Pennsylvania), Kartik Hosanagar (The Wharton School, University of Pennsylvania) and Amit Gandhi (The Wharton School, University of Pennsylvania).

Abstract

Instrumental variables (IVs) are a commonly used technique for causal inference from observational data. In practice, the variation induced by IVs can be limited, which yields imprecise or biased estimates of causal effects and renders the approach ineffective for policy decisions. We confront this challenge by formulating the problem of constructing instrumental variables from candidate exogenous data as a machine learning problem. We propose a novel algorithm, called MLIV (machine-learned instrumental variables), which allows learning of instruments and causal inference to be simultaneously performed from sample data. Further, $O(\sqrt{n})$ consistency and asymptotic normality of our estimators hold under standard regularity conditions. Simulations and application to real-world data demonstrate that the algorithm is highly effective and significantly improves the performance of causal inference from observational data.

---

[90] *Fairness-Efficiency Tradeoffs in Dynamic Fair Division*

David Zeng (Two Sigma) and Alexandros Psomas (Google).

Abstract

We investigate the tradeoffs between fairness and efficiency when allocating indivisible items over time. Suppose T items arrive over time and must be allocated upon arrival, immediately and irrevocably, to one of n agents. Agent i assigns a value $v_{it} \in [0,1]$ to the t-th item to arrive and has an additive valuation function. If the values are chosen by an adaptive adversary, that gets to see the (random) allocations of items 1 through $t-1$ before choosing v_{it}, it is known that the algorithm that minimizes maximum pairwise envy simply allocates each item uniformly at random; the maximum pairwise envy is then sublinear in T, namely $\tilde{O}(\sqrt{T/n})$. If the values are independently and identically drawn from an adversarially chosen distribution D, it is also known that, under some mild conditions on D, allocating to the agent with the highest value --- a Pareto efficient allocation --- is envy-free with high probability.

In this paper we study fairness-efficiency tradeoffs in this setting and provide matching upper and lower bounds under a spectrum of progressively stronger adversaries. On one hand we show that, even against a non-adaptive adversary, there is no algorithm with sublinear maximum pairwise envy that Pareto dominates the simple algorithm that allocates each item uniformly at random. On the other hand, under a slightly weaker adversary regime where item values are drawn

from a known distribution and are independent with respect to time, i.e. v_{it} is independent of $v_{i\hat{t}}$ but possibly correlated with $v_{\hat{i}t}$, optimal (in isolation) efficiency is compatible with optimal (in isolation) fairness. That is, we give an algorithm that is Pareto efficient ex-post and is simultaneously optimal with respect to fairness: for each pair of agents i and j, either i envies j by at most one item (a prominent fairness notion), or i does not envy j with high probability. En route, we prove a structural (and constructive) result about allocations of divisible items that might be of independent interest: there always exists a Pareto efficient divisible allocation where each agent i either strictly prefers her own bundle to the bundle of agent j, or, if she is indifferent, then i and j have identical allocations and the same value (up to multiplicative factors) for all the items that are allocated to them.

---

[92] *Optimal Persuasion via Bi-Pooling*

Itai Arieli (Technion), Yakov Babichenko (Technion), Rann Smorodinsky (Technion) and Takuro Yamashita (University of Toulouse Capitole).

Abstract

The canonical Bayesian persuasion setting studies a model where an informed agent, the Sender, can partially share his information with an uninformed agent, the Receiver. The Receiver's utility is a function of the state of nature and the Receiver's action while the Sender's is only a function of the Receiver's action. The classical results characterize the Sender's optimal information disclosure policy whenever the state space is finite. In this paper we study the same setting where the state space is an interval on the real line. We introduce the class of bi-pooling policies and the induced distribution over posteriors which we refer to as bi-pooling distributions. We show that this class of distributions characterizes the set of optimal distributions in the aforementioned setting. Every persuasion problem admits an optimal bi-pooling distribution as a solution. Conversely, for every bi-pooling distribution there exists a persuasion problem in which the given distribution is the unique optimal one. We leverage this result to study the structure of the price function (see [Dworzak Martini 2019]) in this setting and to identify optimal information disclosure policies.

---

[98] *Does Quality Improve with Customer Voice? Evidence from the Hotel Industry*

Uttara Ananthakrishnan (UW), Davide Proserpio (University of Southern California) and Siddartha Sharma (CMU).

Abstract

In this paper, we empirically study whether firms improve their quality based on reviews left by their customers in a dynamic quality environment. We do so by analyzing the US hotel industry using data from two major online review platforms: TripAdvisor and Expedia. Using management response as a proxy for whether hotels read and listen to consumer reviews, and a difference-in-differences strategy, we demonstrate that hotels make improvements in quality by paying attention to the reviews they receive. Moreover, we show that these improvements are primarily made by low-rated hotels that have more room for improvement, and by chain hotels likely because of their lower operational marginal cost. To pin down the underlying mechanism, we analyze the text of reviews and responses using novel tools for natural language processing, and show that: (i) hotels that listen to customers improve on issues that are frequently mentioned in their reviews, and (ii) hotels that use canned responses are less likely to see improvements in quality. Overall, our results suggest that online user-generated reviews form a feedback mechanism through which consumers make themselves heard by businesses and contribute to changes in the quality of those businesses.

---

[103] *Combinatorial Ski Rental and Online Bipartite Matching*

Hanrui Zhang (Duke University) and Vincent Conitzer (Duke University).

Abstract

We consider a combinatorial variant of the classical ski rental problem --- which we call combinatorial ski rental --- where multiple resources are available to purchase and to rent, and are demanded online. Moreover, the costs of purchasing and renting are potentially combinatorial. The dual problem of combinatorial ski rental, which we call combinatorial online bipartite matching, generalizes the classical online bipartite matching problem into a form where constraints, induced by both offline and online vertices, can be combinatorial. We give a 2-competitive (resp. $e / (e - 1)$-competitive) deterministic (resp. randomized) algorithm for combinatorial ski rental, and an $e / (e - 1)$-competitive algorithm for combinatorial online bipartite matching. All these ratios are optimal given simple lower bounds inherited from the respective well-studied special cases. We also prove information-theoretic impossibility of constant-factor algorithms when any part of our assumptions is considerably relaxed.

[105] *Characterization of group-strategyproof mechanisms for facility location in strictly convex space*

Pingzhong Tang (Tsinghua University), Dingli Yu (Princeton University) and Shengyu Zhao (Institute for Interdisciplinary Information Sciences, Tsinghua University).

Abstract

We characterize the class of group-strategyproof mechanisms for the single facility location game in any unconstrained strictly convex space. A mechanism is group-strategyproof, if no group of agents can misreport so that all its members are strictly better off. A strictly convex space is a normed vector space where $\|x+y\|<2$ holds for any pair of different unit vectors $x \neq y$, e.g., any L_p space with $p\in (1,\infty)$.

We show that any deterministic, unanimous, group-strategyproof mechanism must be dictatorial, and that any randomized, unanimous, translation-invariant, group-strategyproof mechanism must be 2-dictatorial. Here a randomized mechanism is 2-dictatorial if the lottery output of the mechanism must be distributed on the line segment between two dictators' inputs. A mechanism is translation-invariant if the output of the mechanism follows the same translation of the input.

Our characterization directly implies that any (randomized) translation-invariant approximation algorithm satisfying the group-strategyproofness property has a lower bound of 2-approximation for maximum cost (whenever $n \geq 3$), and $n/2 - 1$ for social cost. We also find an algorithm that 2-approximates the maximum cost and $n/2$-approximates the social cost, proving the bounds to be (almost) tight.

---

[107] *Budget-Constrained Incentive Compatibility for Stationary Mechanisms*

Santiago Balseiro (Columbia University), Anthony Kim (Columbia University), Mohammad Mahdian (Google Research) and Vahab Mirrokni (Google Research).

Abstract

Motivated by online advertising applications, we study incentive properties of stationary mechanisms that satisfy budget constraints in expectation at a stationary equilibrium. We consider a general repeated auction setting where a seller sells identical items to buyers with budget constraints and the buyers' value distributions can be arbitrarily correlated. We introduce the novel notion of budget-constrained incentive compatibility (BCIC) under which each buyer chooses an optimal strategy among budget-feasible strategies. Armed with the notion of BCIC, we characterize Bayesian optimal mechanisms that satisfy the budget constraints in expectation with respect to the profit, utility and welfare objectives in the restricted setting where the buyers' value distributions are independent. Furthermore, in the general setting where the buyers' value distributions are correlated, we analyze the incentive properties of several budget management mechanisms that the seller can implement to control the buyers' expenditures on their behalf. These mechanisms have been previously studied in the literature and include: throttling, thresholding, bid shading, reserve pricing and two versions of multiplicative boosting. Our paper provides the first systematic study on the incentive properties of different budget management mechanisms and shows that some popular mechanisms are not incentive compatible even when restricting attention to budget-feasible deviations.

---

[124] *Fair Prediction with Endogenous Behavior*

Christopher Jung (University of Pennsylvania), Sampath Kannan (University of Pennsylvania), Changhwa Lee (University of Pennsylvania), Mallesh Pai (Rice University), Aaron Roth (University of Pennsylvania) and Rakesh Vohra (University of Pennsylvania).

Abstract

There is increasing regulatory interest in whether machine learning algorithms deployed in consequential domains (e.g. in criminal justice) treat different demographic groups "fairly." However, there are several proposed notions of fairness, typically mutually incompatible. Using criminal justice as an example, we study a model in which society chooses an incarceration rule. Agents of different demographic groups differ in their outside options (e.g. opportunity for legal employment) and decide whether to commit crimes. We show that equalizing type I and type II errors across groups is consistent with the goal of minimizing the overall crime rate; other popular notions of fairness are not.

---

[125] *Data Linkages and Incentives*

Annie Liang (University of Pennsylvania) and Erik Madsen (New York University).

Abstract

Many firms, such as banks and insurers, condition their level of service on a consumer's perceived "quality," for instance their creditworthiness. Increasingly, firms have access to consumer segmentations derived from auxiliary data on behavior, and can link outcomes across individuals in a segment for prediction. How does this practice affect consumer incentives to exert (socially-valuable) effort, e.g. to repay loans? We show that the impact of a linkage on behavior depends crucially on whether the linkage reflects quality (via correlations in types) or a shared circumstance (via common shocks to observed outcomes).

---

[131] *An Economic Analysis of Difficulty Adjustment Algorithms in Proof-of-Work Blockchain Systems*

Shunya Noda (University of British Columbia), Kyohei Okumura (University of Tokyo) and Yoshinori Hashimoto (BUIDL, Ltd.).

Abstract

The design of the difficulty adjustment algorithm (DAA) of the Bitcoin system is vulnerable as it dismisses miners' strategic responses to policy changes. We develop an economic model of the Proof-of-Work based blockchain system. Our model allows miners to pause operation when the expected reward is below the shutdown point. Hence, the supply of aggregate hash power can be elastic in the cryptocurrency price and the difficulty target of the mining puzzle. We prove that, when the hash supply is elastic, the Bitcoin DAA fails to adjust the block arrival rate to the targeted level. In contrast, the DAA of another blockchain system, Bitcoin Cash, is shown to be stable even when the cryptocurrency price is volatile and the supply of hash power is highly elastic. We also provide empirical evidence and simulation results supporting the model's prediction. Our results indicate that the current Bitcoin system might collapse if a sharp price reduction lowers the reward for mining denominated in fiat money. However, this crisis can be prevented through the upgrading of DAA.

---

[136] *Driver Surge Pricing*

Nikhil Garg (Stanford University) and Hamid Nazerzadeh (University of Southern California).

Abstract

Ride-hailing marketplaces like Uber and Lyft use dynamic pricing, often called surge, to balance the supply of available drivers with the demand for rides. We study pricing mechanisms for such marketplaces from the perspective of drivers, presenting the theoretical foundation that has informed the design of Uber's new additive driver surge mechanism.

We present a dynamic stochastic model to capture the impact of surge pricing on driver earnings and their strategies to maximize such earnings. In this setting, some time periods (surge) are more valuable than others (non-surge), and so trips of different time lengths vary in the opportunity cost they impose on drivers.

First, we show that multiplicative surge, historically the standard on ride-hailing platforms, is not incentive compatible in a dynamic setting. We then propose a structured, incentive-compatible pricing mechanism. This closed-form mechanism has a simple form and is well-approximated by Uber's new additive surge mechanism. Finally, through both numerical analysis and real data from a ride-hailing marketplace, we show that additive surge is more incentive compatible in practice than is multiplicative surge.

---

[151] *The Effects of Influencer Advertising Disclosure Regulations: Evidence From Instagram*

Daniel Ershov (Toulouse School of Economics, Universite Capitole 1 Toulouse) and Matthew Mitchell (Rotman School of Management, University of Toronto).

Abstract

We collect data from fifty top Instagram influencers in Germany and Spain from 2014 to 2019. Germany experienced changes in disclosure regulation for social media sponsorship during the sample period. Using a difference-in-difference approach, we study the impact of the the rules on the content of posts and the nature of interaction of followers with the posts. On the content side, we measure directly whether the posts include suggested disclosure terms and show variable but substantial adoption of disclosure. We use both an approach based on a fixed list of words that are associated with

sponsorship (i.e. links, mentions of brands, use of words like sale and coupon) as well as natural language processing to assess the likelihood that a post was sponsored. We show that likely sponsored content use may have increased after changes in disclosure and that followers may have been negatively affected. On the other hand, there is evidence that consumers' reaction to sponsored posts, measured by likes, may be quite different under stricter disclosure rules, suggesting that the rules could have a substantial impact on information transmission.

---

[152] *The Secretary Recommendation Problem*

Niklas Hahn (Goethe University Frankfurt/Main), Martin Hoefer (Goethe University Frankfurt/Main) and Rann Smorodinsky (Technion - Israel Institute of Technology).

Abstract

In this paper we revisit the basic variant of the classical secretary problem. We propose a new approach in which we separate between an agent that evaluates the secretary performance and one that has to make the hiring decision. The evaluating agent (the sender) signals the quality of the candidate to the hiring agent (the receiver) who must make a decision. Whenever the two agents' interests are not fully aligned, this induces an information transmission (signaling) challenge for the sender. We study the sender's optimization problem subject to persuasiveness constraints of the receiver for several variants of the problem.

Our results quantify the loss in performance for the sender due to online arrival. We provide optimal and near-optimal persuasive mechanisms that recover at least a constant fraction of a natural utility benchmark for the sender. The separation of evaluation and decision making can have a substantial impact on the approximation results. While in some scenarios, techniques and results closely mirror the conditions in the standard secretary problem, we also explore conditions that lead to very different characteristics.

---

[155] *Sequential Fundraising and Social Insurance*

Amir Ban (Weizmann Institute of Science) and Moran Koren (Stanford University).

Abstract

Seed fundraising for ventures often takes place by sequentially approaching potential contributors, whose decisions are observed by other contributors. The fundraising succeeds when a target number of investments is reached. When a single investment suffices, this setting resembles the classic information cascades model. but when more than one investment is needed, the solution is radically different, and exhibits surprising complexities. We analyze a setting where contributors' levels of information are i.i.d. draws from a known distribution, and find strategies in equilibrium for all. We show that participants rely on social insurance, i.e., invest despite having unfavorable private information, relying on future player strategies to protect them from loss. Delegation is an extreme form of social insurance, where a contributor will unconditionally invest, effectively delegating the decision to future players. In a typical fundraising, early contributors will invest unconditionally, stopping when the target is ``close enough'', thus de facto delegating the business of determining fundraising success or failure to the last contributors.

---

[156] *Convex Optimization for Bundle Size Pricing Problem*

Xiaobo Li (National University of Singapore), Hailong Sun (National University of Singapore) and Chung Piaw Teo (National University of Singapore).

Abstract

We study the bundle size pricing (BSP) problem where a monopolist sells bundles of products to customers, and the price of each bundle depends only on the size (number of items) of the bundle. Although this pricing mechanism is attractive in practice, finding optimal bundle prices is difficult since it involves characterizing distributions of the maximum partial sums of order statistics. In this paper, we propose to solve the BSP problem under a class of choice model using only the first and second moments of customer valuations. We show that the BSP problem under this model is convex and can be efficiently solved using off-the-shelf solvers. Our approach is flexible in optimizing any given bundle sizes, and numerical results show that it performs very well compared with state-of-the-art heuristics.

---

[158] *The Edgeworth Conjecture with Small Coalitions and Approximate Equilibria in Large Economies*

Siddharth Barman (Indian Institute of Science) and Federico Echenique (Caltech).

Abstract

We revisit the connection between bargaining and equilibrium in exchange economies, and study its algorithmic implications. We consider bargaining outcomes to be allocations that cannot be blocked (i.e., profitably re-traded) by coalitions of small size, and show that these allocations must be approximate Walrasian equilibria. Our results imply that deciding whether an allocation is approximately Walrasian can be done in polynomial time, even in economies for which finding an equilibrium is known to be computationally hard.

---

[159] *On the Effect of Positive Discrimination on Multistage Selection Problems in the Presence of Implicit Variance*

Vitalii Emelianov (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG), Nicolas Gast (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG), Krishna Gummadi (MPI-SWS) and Patrick Loiseau (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG).

Abstract

Positive discrimination (i.e., intentionally favoring underprivileged groups in decision making) is often used in selection problems such as hiring or college admission to reduce inequalities based on demographic attributes such as gender, ethnicity or sexual orientation. Yet, its effect is debated: it is often seen as potentially harmful to the decision making, introducing a trade-off between fairness and quality of the selected candidates. In recent work, however, Kleinberg and Raghavan [Proc. of ITCS '18] showed that, in the presence of implicit bias, positive discrimination (specifically the so-called Rooney rule) can increase the utility of a selection process.

We argue that even if bias was corrected, a fundamental difference in the quality estimate of candidates from different groups will remain: their variance will inevitably differ, be it only for simple statistical reasons (one has more data, hence more confidence in the estimation, for the majority group). We term this phenomenon implicit variance and, in this paper, we ask the question: can positive discrimination be beneficial to the utility of a selection process in the presence of implicit variance (even without implicit bias)? To answer this question, we propose a simple model in which candidates have a true latent quality that is drawn from a normal distribution. To make the selection, an evaluator receives an unbiased estimate of the quality of each candidate, with normal noise, but whose variance depends on the candidate's group. Then, we show that positive discrimination (specifically the so-called four-fifths rule) does improve the average quality of selected candidates for every selection ratio. We extend our model to a two-stage selection process where the true quality is observed at the second stage and analyze how our results are changed in that case. We finally discuss multiple extensions of our results, in particular to different distributions of the true latent quality.

---

[162] *Product Ranking on Online Platforms*

Mahsa Derakhshan (University of Maryland), Negin Golrezaei (MIT), Vahideh Manshadi (Yale University) and Vahab Mirrokni (Google).

Abstract

On online platforms, consumers face an abundance of options that are displayed in the form of a position ranking. Only products placed in the first few positions are readily accessible to the consumer, and she needs to exert effort to access more options. For such platforms, we develop a two-stage sequential search model where in the first stage, the consumer sequentially screens positions to learn the preference weight of the products placed in them and forms a consideration set. In the second stage, she learns the additional idiosyncratic utility that she can derive from each product and chooses the highest-utility product within her consideration set. For this model, we first characterize the optimal sequential search policy of a welfare-maximizing consumer. We then study how platforms with different objectives should rank products. We focus on two objectives: (i) maximizing the platform's market share and (ii) maximizing the consumer's welfare. Somewhat surprisingly, we show that ranking products in decreasing order of their preference weights does not necessarily maximize market share or consumer welfare. Such a ranking may shorten the consumer's consideration set due to the externality effect of high-positioned products on low-positioned ones, leading to insufficient screening. We then show that both problems—maximizing market share and maximizing consumer welfare—are NP-complete. We develop novel near-optimal polynomial-time ranking algorithms for each objective. Further, we show that even though ranking products in decreasing order of their preference weights is suboptimal, such a ranking enjoys strong performance guarantees for both objectives.

[169] *On Optimal Ordering in the Optimal Stopping Problem*

Shipra Agrawal (Columbia University), Jay Sethuraman (Columbia University) and Xingyu Zhang (Columbia University).

Abstract

In the classical optimal stopping problem, a player is given a sequence of random variables X_1,..., X_n with known distributions. After observing the realization of X_i, the player can either accept the observed reward from X_i and stop, or reject the observed reward from X_i and continue to observe the next variable X_{i+1} in the sequence. Under any fixed ordering of the random variables, an optimal stopping policy, one that maximizes the player's expected reward, is given by the solution of a simple dynamic program. In this paper, we investigate a relatively less studied question of selecting the order in which the random variables should be observed so as to maximize the expected reward at the stopping time. Perhaps surprisingly, we demonstrate that this ordering problem is NP-hard even in a very restricted case where each random variable X_i has a distribution with 3-point support of form {0, m_i, 1}, and provide an FPTAS. We also provide a simple O(n^2) algorithm for finding an optimal ordering in the case of 2-point distributions. Further, we demonstrate the benefits of order selection, by proving a novel prophet inequality for 2-point distributions that shows the optimal ordering can achieve an expected reward within a factor of 1.25 of the expected hindsight maximum; this is an improvement over the corresponding factor of 2 for the worst-case ordering.

---

[183] *Unpaired Kidney Exchange: Overcoming Double Coincidence of Wants without Money*

Akbarpour Mohammad (Stanford GSB), Julien Combe (CREST - Ecole polytechnique), He Yinghua (Rice University), Hiller Victor (LEMMA - University Paris 2), Shimer Robert (University of Chicago) and Olivier Tercieux (PSE).

Abstract

We propose a new matching algorithm - Unpaired kidney exchange - to tackle the problem of double coincidence of wants without using money. The fundamental idea is that "memory" can serve as a medium of exchange. In a dynamic matching model with heterogeneous agents, we prove that average waiting time under the Unpaired algorithm is close-to optimal, and substantially less than the standard pairwise and chain exchange algorithms. We evaluate this algorithm using a rich dataset of the kidney patients in France. Counterfactual simulations show that the Unpaired algorithm can match nearly 57% of the patients, with an average waiting time of 424 days (state-of-the-art algorithms match about 31% with an average waiting time of 675 days or more). The optimal algorithm performs only slightly better: it matches 58% of the patients and leads to an average waiting time of 410 days. The Unpaired algorithm confronts two incentive-related practical challenges. We address those challenges via a practical version of the Unpaired algorithm that employs kidneys from the deceased donors waiting list. The practical version can match nearly 87% of patient-donor pairs, while reducing the average waiting time to about 141 days.

---

[188] *An Experiment on Network Density and Sequential Learning*

Krishna Dasaratha (Harvard University) and Kevin He (Caltech and University of Pennsylvania).

Abstract

We conduct a sequential social-learning experiment where subjects take turns guessing a hidden state based on private signals and the guesses of a subset of their predecessors. A network determines the observable predecessors, and we compare subjects' accuracy on sparse and dense networks. Accuracy gains from social learning are twice as large on sparse networks compared to dense networks. Models of naive inference where agents ignore correlation between observations predict this comparative static in network density, while the finding is difficult to reconcile with rational-learning models.

---

[189] *Incentivizing Exploration with Selective Data Disclosure*

Nicole Immorlica (Microsoft), Jieming Mao (Google), Aleksandrs Slivkins (Microsoft) and Zhiwei Steven Wu (University of Minnesota).

Abstract

We study the design of rating systems that incentivize (more) efficient social learning among self-interested agents. Agents arrive sequentially and are presented with a set of possible actions, each of which yields a positive reward with an

unknown probability. A disclosure policy sends messages about the rewards of previously-chosen actions to arriving agents. These messages can alter agents' incentives towards exploration, taking potentially sub-optimal actions for the sake of learning more about their rewards. Prior work achieves much progress with disclosure policies that merely recommend an action to each user, but relies heavily on standard, yet very strong rationality assumptions.

We study a particular class of disclosure policies that use messages, called unbiased subhistories, consisting of the actions and rewards from a subsequence of past agents. Each subsequence is chosen ahead of time, according to a predetermined partial order on the rounds. We posit a flexible model of frequentist agent response, which we argue is plausible for this class of "order-based" disclosure policies. We measure the success of a policy by its regret, i.e., the difference, over all rounds, between the expected reward of the best action and the reward induced by the policy. A disclosure policy that reveals full history in each round risks inducing herding behavior among the agents, and typically has regret linear in the time horizon T. Our main result is an order-based disclosure policy that obtains regret $\tilde{O}(\sqrt{T})$. This regret is known to be optimal in the worst case over reward distributions, even absent incentives. We also exhibit simpler order-based policies with higher, but still sublinear, regret. These policies can be interpreted as dividing a sublinear number of agents into constant-sized focus groups, whose histories are then revealed to future agents.

---

[190] *Surrogate Scoring Rules*

Yiling Chen (Harvard University), Yang Liu (University of California, Santa Cruz) and Juntao Wang (Harvard University).

Abstract

Strictly proper scoring rules (SPSR) are incentive compatible for eliciting information about random variables from strategic agents when the principal can reward agents after the realization of the random variables. They also quantify the quality of elicited information, with more accurate predictions receiving higher score in expectation. In this paper, we extend such scoring rules to settings where a principal elicits private probabilistic beliefs but only has access to agents' reports. We name our solution \emph{Surrogate Scoring Rules} (SSR). SSR build on a bias correction step and an error rate estimation procedure for a reference answer defined using agents' reports. We show that, with one bit of information about the prior distribution of the random variables, SSR in a multi-task setting recover SPSR in expectation, as if having access to the ground truth. Therefore, a salient feature of SSR is that they quantify the quality of information despite the lack of ground truth, just as SPSR do for the \emph{with} ground truth setting. As a by-product, SSR induce \emph{dominant truthfulness} in reporting.

Our work complements the proper scoring rule literature via extending existing SPSR to operate when there is no clean ground truth verification. Because of the non-existence of verification, our setting falls into the classical information elicitation without verification (IEWV) domain, which has focused on eliciting discrete signals. Therefore our work also contributes to the peer prediction literature via providing a scoring rule that elicits continuous probabilistic beliefs, an approach that rewards accuracy instead of correlation, and a mechanism that achieves truthfulness in \emph{dominant strategy} in a multi-task setting. Our method is verified both theoretically and empirically using data collected from real human forecasters.

---

[193] *The Value of Observability in Dynamic Pricing*

Jose Correa (Universidad de Chile), Dana Maria Pizarro (Universidad de Chile) and Gustavo Vulcano (Universidad Torcuato Di Tella).

Abstract

Research on dynamic pricing has been growing during the last four decades due to its use in practice by a variety of companies as well as the several model variants that can be considered. In this work, we consider the particular pricing problem where a firm wants to sell one item to a single buyer along an infinite time horizon in order to maximize her expected revenue. The firm commits to a price function over an infinite horizon. The buyer has a private value for the item and purchases at the time when his utility is maximized. In our model, the buyer is more impatient than the seller and we study how important is to observe the buyer time arrival in terms of the seller's expected revenue. When the seller can observe the arrival of the buyer, she can make the price function contingent on the buyer's arrival time. On the contrary, when the seller cannot observe the arrival, her price function is fixed at time zero for the whole horizon. The {value of observability} is defined as the worst case ratio between the expected revenue of the seller when she observes the buyer's arrival and that when the seller does not observe the time when the buyer arrives. Our main result is to prove that in a very general setting, the value of observability is at most 4.911. To obtain this result we fully characterize the observable setting and use this solution to construct a random and periodic price function for the unobservable case.

[195] *Adversarial Perturbations of Opinion Dynamics in Networks*

Jason Gaitonde (Cornell University), Jon Kleinberg (Cornell University) and Eva Tardos (Cornell University).

Abstract

In this paper we study the connections between network structure, opinion dynamics, and an adversary's power to artificially induce disagreements. We approach these questions by extending models of opinion formation in the mathematical social sciences to represent scenarios, familiar from recent events, in which external actors have sought to destabilize communities through sophisticated information warfare tactics via fake news and bots. In many instances, the intrinsic goals of these efforts are not necessarily to shift the overall sentiment of the network towards a particular policy, but rather to induce discord. These perturbations will diffuse via opinion dynamics on the underlying network, through mechanisms that have been analyzed and abstracted through work in computer science and the social sciences.

Here we investigate the properties of such attacks, considering optimal strategies both for the adversary seeking to create disagreement and for the entities tasked with defending the network from attack. By employing spectral techniques, we show that for different formulations of these types of objectives, different regimes of the spectral structure of the network will limit the adversary's capacity to sow discord; in fact, somewhat surprisingly, the entire spectrum can be relevant, rather than just the extreme eigenvectors. Via the strong connections between spectral and structural properties of graphs, we are able to qualitatively describe which networks are most vulnerable or resilient against these perturbations. We then consider the algorithmic task of a network defender to mitigate these sorts of adversarial attacks by insulating nodes heterogeneously; we show that, by considering the geometry of this problem, this optimization task can be efficiently solved via convex programming. Finally, we generalize these results to allow for two network structures, where the opinion dynamics process and the measurement of disagreement become uncoupled; for instance, this may arise when opinion dynamics are controlled an online community via social media, while disagreement is measured along "real-world" connections. We characterize conditions on the relationship between these two graphs that will determine how much power the adversary gains when this occurs.

---

[199] *Stability and Learning in Strategic Queuing Systems*

Jason Gaitonde (Cornell University) and Eva Tardos (Cornell University).

Abstract

Bounding the price of anarchy, which quantifies the damage to social welfare due to selfish behavior of the participants, has been an important area of research. In this paper, we study this phenomenon in the context of a game modeling queuing systems: routers compete for servers, where packets that do not get service will be resent at future rounds, resulting in a system where the number of packets at each round depends on the success of the routers in the previous rounds. We model this as an (infinitely) repeated game, where the system holds a state (number of packets held by each queue) that arises from the results of the previous round. We assume that routers satisfy the no-regret condition, e.g. they use learning strategies to identify the server where their packets get the best service.

Classical work on repeated games makes the strong assumption that the subsequent rounds of the repeated games are independent (beyond the influence on learning from past history). The carryover effect caused by packets remaining in this system makes learning in our context result in a highly dependent random process. We analyze this random process and find that if the capacity of the servers is high enough to allow a centralized and knowledgeable scheduler to get all packets served even with double the packet arrival rate, and queues use no-regret learning algorithms, then the expected number of packets in the queues will remain bounded throughout time, assuming older packets have priority. This paper is the first to study the effect of selfish learning in a queuing system, where the learners compete for resources, but rounds are not all independent: the number of packets to be routed at each round depends on the success of the routers in the previous rounds.

---

[203] *Vertically Disintegrated Platforms*

Christoph Aymanns (Quantco / Systemic Risk Centre LSE), Tarik Roukny (KU Leuven) and Mathias Dewatripont (ULB).

Abstract

Understanding the implications of digital platform strategies on user welfare has become a crucial issue for policy makers and economists alike. Recently, distributed ledger technologies such as the blockchain have introduced new forms of platform governance in the form of vertically disintegrated platforms (VDP). In this paper, we develop a framework that specifically focuses on such platforms and study the implications of vertical disintegration on user welfare. A VDP

mediates between users and processors that enable interactions between users. The VDP controls the price structure but the price level is set in a competitive equilibrium between processors; proceeds from operating the platform service are split between processors and the VDP. We find that the welfare ordering between traditional (integrated) and disintegrated platforms crucially depends on the platform cost structure and the regulatory conditions in the market for processors. When the cost of integrating processors depends on the transaction volume, the VDP can produce higher welfare. In contrast, when this cost is fixed, an unregulated VDP is less desirable than the integrated platform. However, regulatory limits on the VDP power over processors can make user welfare under a VDP dominate. Within our framework, we analyze recent applications in the payment industry, including central bank digital currencies, and show how different design choices affect user welfare.

---

[207] *Dynamic Privacy Choices*

Shota Ichihashi (Bank of Canada).

Abstract

I study a dynamic model of consumer privacy and platform data collection. Consumers choose how actively to use a platform in each period. The activity generates information about consumers, which benefits the platform but hurts consumers. In the long-run, consumers lose privacy and receive a low payoff, but choose a high activity level. To induce this outcome, the platform adopts a privacy policy that gradually expands the scope of data collection. Platform competition is less likely when consumers are receiving lower payoffs from an incumbent. Regulating data collection can lower consumer and total welfare in the long-run.

---

[219] *Escaping Saddle Points in Constant Dimensional Spaces: an Agent-based Modeling Perspective*

Grant Schoenebeck (University of Michigan) and Fang-Yi Yu (University of Michigan).

Abstract

We study a large family of stochastic processes captured by the fact that they update a limited amount in each step, e.g., agent-based models where one agent at a time updates their state where the step size is not too large. A key question is how this family of stochastic processes are approximated by their mean-field approximations. Prior work shows that the stochastic processes escape repelling fixed points and saddle points in polynomial time.

We provide a tight analysis: for any non-attracting fixed point in any stochastic process in this family, we show that for a sufficiently small constant $\epsilon >0$, the process will be ϵ-far away from the fixed point in $O(n\log n)$ time with high probability. We also show that it takes time $\Omega(n\log n)$ to escape such a fixed point with constant probability. This shows our result is optimal up to a multiplicative constant.

We leverage the above result to show that with high probability, these stochastic processes are arbitrarily close to an attracting fixed point in $O(n\log n)$ time.

We show the power of our results by applying them to several settings: evolutionary game theory, opinion formation dynamics, and stochastic gradient descent in a finite-dimensional space.

---

[220] *Credible, Truthful, and Bounded-Round Mechanisms via Cryptographic Commitments*

Matheus Venturyne Xavier Ferreira (Princeton University) and S. Matthew Weinberg (Princeton University).

Abstract

We consider the sale of a single item to multiple buyers by a revenue-maximizing seller. Recent work of Akbarpour and Li formalizes credibility as an auction desideratum, and prove that the only optimal, credible, strategyproof auction is the ascending price auction.

In contrast, when buyer valuations are MHR, we show that the mild additional assumption of a cryptographically secure commitment scheme suffices for a simple two-round auction which is optimal, credible, and strategyproof.

We extend our analysis to the case when buyer valuations are α-strongly regular for any $\alpha > 0$, up to arbitrary ε in credibility. Interestingly, we also prove that this construction cannot be extended to regular distributions, nor can the ε be removed with multiple bidders.

[224] *On the (in-)approximability of Bayesian Mechanism Design for a Combinatorial Buyer*

Natalie Collina (Google) and S. Matthew Weinberg (Princeton University).

Abstract

We consider a revenue-maximizing single seller with m items for sale to a single buyer whose value v for the items is drawn from a known distribution D of support k. A series of works by Cai et al. establishes that when each v in the support of D is additive or unit-demand (or c-demand) that the revenue-optimal auction can be found in poly(m,k) time.

We show that going barely beyond this, even to matroid-based valuations (a proper subset of Gross Substitutes), results in strong hardness of approximation. Specifically, even on instances with m items and $k \leq m$ valuations in the support of D, it is not possible to achieve a $1/m^{1-\varepsilon}$-approximation for any $\varepsilon>0$ to the revenue-optimal mechanism for matroid-based valuations in (randomized) poly-time unless NP $\subseteq$ RP (note that a 1/k-approximation is trivial).

Cai et al.'s main technical contribution is a black-box reduction from revenue maximization for valuations in class V to optimizing the difference between two values in class V. Our main technical contribution is a black-box reduction in the other direction, establishing that their reduction is essentially tight.

---

[229] *Prizes on Crowdsourcing Platforms: An Equilibrium Analysis of Competing Contests*

Konstantinos Stouras (University College Dublin), Sanjiv Erat (University of California, San Diego) and Kenneth C. Lichtendahl Jr. (Google and University of Virginia).

Abstract

On a typical crowdsourcing platform solvers can self-select which (if any) of the concurrently running contests to participate in. Thus, firms who offer prizes and organize contests on these platforms are competing among themselves (for solver participation and effort). We formalize and model this competition among contests and examine the equilibrium outcomes. Our analysis reveals that, in general, there is a unique dominant strategy for each firm to offer multiple identical prizes. Moreover, when the quality of submitted solutions is sufficiently noise-driven (as opposed to effort-driven), we find that a single winner-take-all reward is the unique equilibrium allocation. Our analytical framework integrates and extends prior results of the monopolistic contest.

---

[232] *Blind Dynamic Control of a Shared Transportation System via Mirror Backpressure*

Pengyu Qian (Columbia Business School) and Yash Kanoria (Columbia Business School).

Abstract

We study the problem of maximizing payoff generated over a period of time in a general class of closed queueing networks with finite, fixed number of supply units which circulate in the system. Demand arrives stochastically, and serving a demand unit (customer) causes a supply unit to relocate from the "origin" to the "destination" of the customer. We consider general controls including entry control, pricing, and assignment. Motivating applications include shared transportation platforms and scrip systems.

Inspired by the mirror descent algorithm for optimization and the backpressure policy for network control, we introduce a novel family of Mirror Backpressure (MBP) control policies. The MBP policies are simple and practical, and crucially do not need any statistical knowledge of the demand (customer) arrival rates.

Under mild conditions, we show that MBP policies lose at most O(K/T+1/K) payoff per customer relative to the optimal policy that knows the demand arrival rates, where K is the number of supply units and T is the total number of customers over the time horizon. The key technical challenge we overcome is that the number of supply units at any node can never be negative. Simulation results in a realistic ridehailing environment support our theoretical findings.

---

[237] *Incentive Auction Design Alternatives: A Simulation Study*

Neil Newman (The University of British Columbia), Kevin Leyton-Brown (The University of British Columbia), Paul Milgrom (Stanford University) and Ilya Segal (Stanford University).

Abstract

Over 13 months in 2016—17 the US Federal Communications Commission (FCC) conducted an “incentive auction” to repurpose radio spectrum from broadcast television to wireless internet. This paper revisits from a computational perspective the descending clock “reverse” auction used to procure broadcast rights. We investigate the quantitative significance of various aspects of the design by running extensive simulations, leveraging a reverse auction simulator and realistic models of bidder values.

---

[239] *Differentially Private Call Auctions and Market Impact*

Emily Diana (University of Pennsylvania), Hadi Elzayn (University of Pennsylvania), Michael Kearns (University of Pennsylvania), Aaron Roth (University of Pennsylvania), Saeed Sharifi-Malvajerdi (University of Pennsylvania) and Juba Ziani (University of Pennsylvania).

Abstract

We propose and analyze differentially private (DP) mechanisms for call auctions as an alternative to the complex and ad-hoc privacy efforts that are common in modern electronic markets. We prove that the number of shares cleared in the DP mechanisms compares favorably to the non-private optimal and provide a matching lower bound. We analyze the incentive properties of our mechanisms and their behavior under natural no-regret learning dynamics by market participants. We include simulation results and connections to the finance literature on market impact.

---

[243] *Pandora's Box Problem with Order Constraints*

Shant Boodaghians (University of Illinois at Urbana-Champaign), Federico Fusco (Sapienza University of Rome), Philip Lazos (Sapienza University of Rome) and Stefano Leonardi (Sapienza University of Rome).

Abstract

The Pandora's Box Problem, originally formalized by Weitzman in 1979, models selection from set of random, alternative options, when evaluation is costly. This includes, for example, the problem of hiring a skilled worker, where only one hire can be made, but the evaluation of each candidate is an expensive procedure.

Weitzman showed that the Pandora's Box Problem admits an elegant, simple solution, where the options are considered in decreasing order of reservation value, i.e. the value that reduces to zero the expected marginal gain for opening the box. We study for the first time this problem when order --- or precedence --- constraints are imposed between the boxes. We show that, despite the difficulty of defining reservation values for the boxes which take into account both in-depth and in-breath exploration of the various options, greedy optimal strategies exist and can be efficiently computed for tree-like order constraints.

We also prove that finding approximately optimal adaptive search strategies is NP-hard when certain matroid constraints are used to further restrict the set of boxes which may be opened, or when the order constraints are given as reachability constraints on a DAG. We complement the above result by giving approximate adaptive search strategies based on a connection between optimal adaptive strategies and non-adaptive strategies with bounded adaptivity gap for a carefully relaxed version of the problem.

---

[257] *The Multiplayer Colonel Blotto Game*

Enric Boix-Adserà (MIT), Benjamin Edelman (Harvard University) and Siddhartha Jayanti (MIT).

Abstract

We initiate the study of the natural multiplayer generalization of the classic continuous Colonel Blotto game. The two-player Blotto game, introduced by Borel as a model of resource competition across n simultaneous fronts, has been studied extensively for a century and seen numerous applications throughout the social sciences. Our work defines the multiplayer Colonel Blotto game and derives Nash equilibria for various settings of k (number of players) and n. We also introduce a ``Boolean'' version of Blotto that becomes interesting in the multiplayer setting. The main technical difficulty of our work, as in the two-player theoretical literature, is the challenge of coupling various marginal distributions into a joint distribution satisfying a strict sum constraint. In contrast to previous works in the continuous setting, we derive our couplings algorithmically in the form of efficient sampling algorithms.

---

[266] *Loot Box Pricing and Design*

Ningyuan Chen (University of Toronto), Adam Elmachtoub (Columbia University), Michael Hamilton (University of Pittsburgh) and Xiao Lei (Columbia University).

Abstract

In the online video game industry, a significant portion of the revenue is generated from microtransactions, where a small amount of real-world currency is exchanged for virtual items to be used in the game. One popular way to conduct microtransactions is via a loot box, which is a random bundle of virtual items whose contents are not revealed until after purchase. In this work, we consider how to optimally price and design loot boxes from the perspective of a revenue-maximizing video game company. Our paper provides the first formal treatment of loot boxes, with the aim to provide customers, companies, and regulatory bodies with insights into this popular selling strategy.

We consider two types of loot boxes: a traditional one where customers can receive (unwanted) duplicates, and a unique one where customers are guaranteed to never receive duplicates. We show that as the number of virtual items grows large, the unique box strategy is asymptotically optimal, while the traditional box strategy only garners 1/e of the optimal revenue. On the other hand, unique box strategies leaves almost zero customer surplus, while traditional box strategies leaves positive surplus. Further, when designing traditional and unique loot boxes, we show it is asymptotically optimal to allocate the items uniformly, even when the item valuation distributions are highly heterogeneous. We also show that when the seller purposely misrepresents the allocation probabilities, then their revenue may increase significantly and thus strict regulation is needed. Finally, we show that if the seller allows customers to salvage unwanted items then loot box designs allow sellers to flexibly trade-off between revenue and customer surplus, and can be guaranteed to dominate the revenue of separate selling.

---

[270] *Information Design for Congested Social Services: Optimal Need-Based Persuasion*

Jerry Anunrojwong (Agoda, MIT, Chulalongkorn University), Krishnamurthy Iyer (University of Minnesota) and Vahideh Manshadi (Yale School of Management).

Abstract

We study the effectiveness of information design in reducing congestion in social services catering users with varied levels of need. In the absence of price discrimination and centralized admission, the provider relies on sharing information about wait times to improve welfare. We consider a stylized model with heterogeneous users who differ in their private outside options: {\em low-need} users have an acceptable outside option to the social service, whereas {\em high-need} users have no viable outside option. Upon arrival, a user decides to wait for the service by joining an unobservable first-come-first-serve queue, or leave and seek her outside option. To engineer improved social outcomes, the service provider seeks to reduce congestion, by persuading more low-need users to avail their outside option, and thus better serving high-need users. We characterize the Pareto-optimal signaling mechanisms and compare their welfare against several benchmarks. We show that if either type is the overwhelming majority of the population, information design does not provide improvement over sharing full information or no information. On the other hand, when the population is a mixture of the two types, information design not only Pareto dominates full-information and no-information mechanisms, in some regimes, it achieves the same welfare as the ``first-best'', i.e., the Pareto-optimal centralized admission policy with knowledge of users' types.

---

[278] *Feasible Joint Posterior Beliefs*

Itai Arieli (Technion), Yakov Babichenko (Technion), Fedor Sandomirskiy (Technion) and Omer Tamuz (Caltech).

Abstract

We study the set of possible joint posterior belief distributions of a group of agents who share a common prior regarding a binary state and who observe some information structure. Our main result is that, for the two agent case, a quantitative version of Aumann's Agreement Theorem provides a necessary and sufficient condition for feasibility. We use our characterization to construct joint belief distributions in which agents are informed regarding the state, and yet receive no information regarding the other's posterior. We also study the related Bayesian persuasion problems with a single sender and multiple receivers, and, relatedly, explore the extreme points of the set of feasible distributions.

---

[279] *To Infinity and Beyond: Scaling Economic Theories via Logical Compactness*

Yannai A. Gonczarowski (Microsoft Research), Scott Duke Kominers (Harvard University) and Ran Shorrer (Penn State University).

Abstract

Many economic-theoretic models incorporate finiteness assumptions that, while introduced for simplicity, play a real role in the analysis. Such assumptions introduce a conceptual problem, as results that rely on finiteness are often implicitly nonrobust, for example they may rely on edge effects or artificial boundary conditions. Here, we present a unified method that allows us to remove finiteness assumptions, such as those on datasets, market sizes, and time horizons. We then apply our approach to a variety of revealed preference, matching, and exchange settings.

The key to our approach is Logical Compactness, a core result from Propositional Logic. Building on Compactness, in a revealed-preference setting, we reprove Reny's infinite-data version of Afriat's theorem and (newly) prove an infinite-data version of McFadden and Richter's characterization of rationalizable stochastic datasets. In a matching setting, we reprove existence results implied by Fleiner's analysis, and prove both the strategy-proofness of the man-optimal stable mechanism in infinite markets, and an infinite-market version of Nguyen and Vohra's existence result for near-feasible stable matchings with couples. In a trading-network setting, we prove that Hatfield et al.'s result on existence of Walrasian equilibrium extends to infinite markets. Finally, in a dynamic matching setting, we prove that Pereyra's existence result extends to a doubly-infinite time horizon.

---

[280] *Experimental Design in Two-Sided Platforms: An Analysis of Bias*

Ramesh Johari (Stanford University), Hannah Li (Stanford University) and Gabriel Weintraub (Stanford University).

Abstract

We develop an analytical framework to study experimental design in two-sided platforms. In the settings we consider, customers rent listings; rented listings are occupied for some amount of time, then become available. Platforms typically use two common designs to study interventions in such settings: customer-side randomization (CR), and listing-side randomization (LR), along with associated estimators. We develop a stochastic model and associated mean field limit to capture dynamics in such systems, and use our model to investigate performance of these estimators. Good experimental design depends on market balance: we show that in highly demand-constrained markets, CR is unbiased, while LR is biased; conversely, in highly supply-constrained markets, LR is unbiased, while CR is biased. We also study a design based on two-sided randomization (TSR) where both customers and listings are randomized to treatment and control, and show that appropriate choices of such designs can yield low bias in intermediate regimes of market balance.

---

[282] *The Complexity of Interactively Learning a Stable Matching by Trial and Error*

Ehsan Emamjomeh-Zadeh (University of Southern California), Yannai A. Gonczarowski (Microsoft Research) and David Kempe (University of Southern California).

Abstract

In a stable matching setting, we consider a query model that allows for an interactive learning algorithm to make precisely one type of query: proposing a matching, the response to which is either that the proposed matching is stable, or a blocking pair (chosen adversarially) indicating that this matching is unstable. For one-to-one matching markets, our main result is an essentially tight upper bound of O(n^2 log n) on the deterministic query complexity of interactively learning a stable matching in this coarse query model, along with an efficient randomized algorithm that achieves this query complexity with high probability. For many-to-many matching markets in which participants have responsive preferences, we first give an interactive learning algorithm whose query complexity and running time are polynomial in the size of the market if the maximum quota of each agent is bounded; our main result for many-to-many markets is that the deterministic query complexity can be made polynomial (more specifically, O(n^3 log n)) in the size of the market even for arbitrary (e.g., linear in the market size) quotas.

---

[285] *A Learning Framework for Distribution-Based Game-Theoretic Solution Concepts*

Tushant Jha (IIIT Hyderabad) and Yair Zick (National University of Singapore).

Abstract

The past few years have seen several works exploring learning economic solutions from data; these include optimal auction design, function optimization, stable payoffs in cooperative games and more.

In this work, we provide a unified learning-theoretic methodology for modeling such problems, and establish tools for determining whether a given solution concept can be efficiently learned from data.

Our learning theoretic framework generalizes a notion of function space dimension --- the graph dimension --- adapting it to the solution concept learning domain.

We identify sufficient conditions for efficient solution learnability, and show that results in existing works can be immediately derived using our methodology.

Finally, we apply our methods in other economic domains, yielding learning variants of competitive equilibria and Condorcet winners.

---

[288] *Optimal Communication-Distortion Tradeoff in Voting*

Debmalya Mandal (Columbia University), Nisarg Shah (University of Toronto) and David Woodruff (Carnegie Mellon University).

Abstract

In recent work, \citet{MPSW19} study a novel framework for the winner selection problem in voting, in which a voting rule is seen as a combination of an elicitation rule and an aggregation rule. The elicitation rule asks voters to respond to a query based on their preferences over a set of alternatives, and the aggregation rule aggregates voter responses to return a winning alternative. They study the tradeoff between the communication complexity of a voting rule, which measures the number of bits of information each voter must send in response to its query, and its distortion, which measures the quality of the winning alternative in terms of utilitarian social welfare. They prove upper and lower bounds on the communication complexity required to achieve a desired level of distortion, but their bounds are not tight. Importantly, they also leave open the question whether the best randomized rule can significantly outperform the best deterministic rule.

We settle this question in the affirmative. For a winner selection rule to achieve distortion d with m alternatives, we show that the communication complexity required is $\Theta~{\frac{m}{d}}$ when using deterministic elicitation, and $\Theta~{\frac{m}{d^3}}$ when using randomized elicitation; both bounds are tight up to logarithmic factors. Our upper bound leverages recent advances in streaming algorithms. To establish our lower bound, we derive a new lower bound on a multi-party communication complexity problem.

We then study the k-selection problem in voting, where the goal is to select a set of k alternatives. For a k-selection rule that achieves distortion d with m alternatives, we show that the best communication complexity is $\Theta~{\frac{m}{kd}}$ when the rule uses deterministic elicitation and $\Theta~{\frac{m}{kd^3}}$ when the rule uses randomized elicitation. Our optimal bounds yield the non-trivial implication that the k-selection problem becomes strictly easier as k increases.

---

[308] *EFX Exists for Three Agents*

Bhaskar Ray Chaudhury (MPI-INF), Jugal Garg (University of Illinois at Urbana-Champaign) and Kurt Mehlhorn (MPI-INF).

Abstract

We study the problem of allocating a set of indivisible items among agents with additive valuations in a fair manner. Envy-freeness up to any item (EFX) is arguably the most compelling fairness concept for this problem. However, despite significant efforts by many researchers for several years, its existence has not been settled beyond the simple case of two agents. In this paper, we break this barrier by showing that an EFX allocation always exists for three agents! Our proof is algorithmic and quite involved. Furthermore, we also falsify a conjecture of Caragiannis et al. by showing an instance with three agents for which there is a partial EFX allocation (some items are not allocated) with higher Nash welfare than that of any complete EFX allocation.

---

[309] *Fair Cake Division Under Monotone Likelihood Ratios*

Siddharth Barman (Indian Institute of Science) and Nidhi Rathi (Indian Institute of Science).

Abstract

This work develops algorithmic results for the classic cake-cutting problem in which a divisible, heterogeneous resource (modeled as a cake) needs to be partitioned among agents with distinct preferences. We focus on a standard formulation of cake cutting wherein each agent must receive a contiguous piece of the cake. While multiple hardness results exist in this setup for finding fair/efficient cake divisions, we show that, if the value densities of the agents satisfy the monotone

likelihood ratio property (MLRP), then strong algorithmic results hold for various notions of fairness and economic efficiency.

Addressing cake-cutting instances with MLRP, first we develop an algorithm that finds cake divisions (with connected pieces) that are envy-free, up to an arbitrary precision. The time complexity of our algorithm is polynomial in the number of agents and the bit complexity of an underlying Lipschitz constant. We obtain similar positive results for maximizing social (utilitarian) and egalitarian welfare. In addition, we show that, under MLRP, the problem of maximizing Nash social welfare admits a fully polynomial-time approximation scheme (FPTAS).

Many distribution families are also known to bear MLRP. In particular, this property holds if all the value densities belong to any one of the following families: Gaussian (with the same variance), linear, binomial, Poisson, and exponential distributions. Furthermore, it is known that linear translations of any log-concave function satisfy MLRP. Therefore, our results also hold when the value densities of the agents are linear translations of the following (log-concave) distributions: Laplace, gamma, beta, Subbotin, chi-square, Dirichlet, and logistic. Hence, through MLRP, the current work obtains novel cake-cutting algorithms for multiple distribution families.

---

[313] *Online Stochastic Max-Weight Matching: prophet inequality for vertex and edge arrival models*

Tomer Ezra (Tel Aviv University), Michal Feldman (Tel Aviv University), Nick Gravin (ITCS, Shanghai University of Finance and Economics) and Zhihao Gavin Tang (ITCS, Shanghai University of Finance and Economics).

Abstract

We provide prophet inequality algorithms for online weighted matching in general (non-bipartite) graphs, under two well-studied arrival models, namely edge arrival and vertex arrival. The weight of each edge is drawn independently from an a-priori known probability distribution. Under edge arrival, the weight of each edge is revealed upon arrival, and the algorithm decides whether to include it in the matching or not. Under vertex arrival, the weights of all edges from the newly arriving vertex to all previously arrived vertices are revealed, and the algorithm decides which of these edges, if any, to include in the matching. To study these settings, we introduce a novel unified framework of em batched prophet inequalities that captures online settings where elements arrive in batches; in particular it captures matching under the two aforementioned arrival models. Our algorithms rely on the construction of suitable online contention resolution scheme (OCRS). We first extend the framework of OCRS to batched-OCRS, we then establish a reduction from batched prophet inequality to batched OCRS, and finally we construct batched OCRSs with selectable ratios of 0.337 and 0.5 for edge and vertex arrival models, respectively. Both results improve the state of the art for the corresponding settings. For the vertex arrival, our result is tight. Interestingly, a pricing-based prophet inequality with comparable competitive ratios is unknown.

---

[314] *On Sustainable Equilibria*

Srihari Govindan (Department of Economics, University of Rochester), Rida Laraki (CNRS (Lamsade-PSL) and University of Liverpool (Computer Science Department)) and Lucas Pahl (Department of Economics, University of New South Wales).

Abstract

Following the ideas laid out in Myerson (1996), Hofbauer (2000) defined an equilibrium of a game as sustainable if it can be made the unique equilibrium of a game obtained by deleting a subset of the strategies that are inferior replies to it, and then adding others. Hofbauer also formalized Myerson's conjecture about the relationship between the sustainability of an equilibrium and its index: for generic games, an equilibrium is sustainable iff its index is +1. von Schemde and von Stengel (2008) proved this conjecture for bimatrix games. This paper shows that the conjecture is true for all finite games.

---

[315] *Graphons and Contagion in Large Networks*

Francesca Parise (MIT), Selman Erol (Tepper School of Business, Carnegie Mellon University) and Alexander Teytelboym (University of Oxford).

Abstract

We consider a threshold contagion process over networks sampled from a graphon, which we interpret as a stochastic network formation model. We investigate whether the contagion outcome in the sampled networks can be predicted by

only exploiting information about the graphon. To this end, we formally define a threshold contagion process on a graphon. Our main results show that contagion on large but finite sampled networks is well approximated by contagion on a graphon. We then give several examples of tractable analysis of contagion on graphons for different choices of agents' thresholds (both deterministic or drawn uniformly at random). Finally, we give sufficient conditions on the locations of vanishingly small seed sets that generate non-trivial contagion.

---

[319] *A General Framework for Endowment Effects in Combinatorial Markets*

Tomer Ezra (Tel Aviv University), Michal Feldman (Tel Aviv University) and Ophir Friedler (Tel Aviv University).

Abstract

The endowment effect, coined by Nobel Laureate Richard Thaler, posits that people tend to inflate the value of items they own. This bias has been traditionally studied mainly using experimental methodology. Recently, Babaioff et al. proposed a specific formulation of the endowment effect in combinatorial markets, and showed that the existence of Walrasian equilibrium with respect to the endowed valuations extends from gross substitutes to submodular valuations, but provably fails to extend to XOS valuations.

We propose to harness the endowment effect further. To this end, we introduce a principle-based framework that captures a wide range of different formulations of the endowment effect (including that of Babaioff et al.). We equip our framework with a partial order over the different formulations, which (partially) ranks them from weak to strong, and provide algorithms for computing endowment equilibria with high welfare for sufficiently strong endowment effects, as well as non-existence results for weaker ones.

Our main results are the following:

1) For markets with XOS valuations, we provide an algorithm that, for any sufficiently strong endowment effect, given an arbitrary initial allocation S, returns an endowment equilibrium with at least as much welfare as in S. In particular, the socially optimal allocation can be supported in an endowment equilibrium; moreover, every such endowment equilibrium gives at least half of the optimal social welfare. Evidently, the negative result of Babaioff et al. for XOS markets is an artifact of their specific formulation.

2) For markets with arbitrary valuations, we show that bundling leads to a sweeping positive result. In particular, if items can be prepacked into indivisible bundles, we provide an algorithm that, for a wide range of endowment effects, given an arbitrary initial allocation S, computes an endowment equilibrium with at least as much welfare as in S. The algorithm runs in poly time with poly many value (resp., demand) queries for submodular (resp., general) valuations. This result is essentially a black-box reduction from the computation of an approximately-optimal endowment equilibrium with bundling to the algorithmic problem of welfare approximation.

---

[326] *Finite State Machines Play Extensive-Form Games*

Jakub Cerny (Nanyang Technological University), Branislav Bosansky (Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University in Prague) and Bo An (Nanyang Technological University).

Abstract

Finite state machines are a well-known representation of strategies in (in)finitely repeated or stochastic games. Actions of players correspond to states in the machine and the transition between machine-states are caused by observations in the game. For extensive-form games (EFGs), machines can act as a formal grounding for abstraction methods used for solving large EFGs and as a domain-independent approach for generating sufficiently compact abstractions. We show that using machines of a restricted size in EFGs can both (i) reduce the theoretical complexity of computing some solution concepts, including Strong Stackelberg Equilibrium (SSE), (ii) as well as bring new practical algorithms that compute near-optimal equilibria considering only a fraction of strategy space. Our contributions include (1) formal definition and theoretical characterization of machine strategies in EFGs, (2) formal definitions and complexity analysis for solution concepts and their computation when restricted to small classes of machines, (3) new algorithms for computing SSE in general-sum games and Nash Equilibrium in zero-sum games that both directly use the concept of machines. Experimental results on two different domains show that the algorithms compute near-optimal strategies and achieve significantly better scalability compared to previous state-of-the-art algorithms.

---

[340] *Escaping Cannibalization? Correlation-Robust Pricing for a Unit-Demand Buyer*

Moshe Babaioff (Microsoft Research), Michal Feldman (Tel Aviv University), Yannai A. Gonczarowski (Microsoft Research), Brendan Lucier (Microsoft Research) and Inbal Talgam-Cohen (Technion—Israel Institute of Technology).

Abstract

A single seller wishes to sell multiple items to a single unit-demand buyer. We consider a robust version of this revenue maximization pricing problem, where the seller knows the buyer's marginal distributions of values for each item, but not the joint distribution, and wishes to maximize worse-case revenue over possible correlation structures. We devise a computationally-efficient (polynomial in the support size of the marginals) algorithm that computes the worst-case joint distribution for any setting of item prices. And yet, in sharp contrast to the additive buyer case (Carroll, 2017), we show that it is NP-hard to approximate the optimal choice of prices to within any factor better than n^(1/2-ε). For the special case of marginal distributions that satisfy the monotone-hazard-rate property, we show how to guarantee a constant fraction of the optimal worst-case revenue by posting item prices. This pricing equates revenue across all possible correlations and can be computed efficiently.

---

[341] *Online Policies for Efficient Volunteer Crowdsourcing*

Vahideh Manshadi (Yale University) and Scott Rodilitz (Yale University).

Abstract

Nonprofit crowdsourcing platforms such as food recovery organizations rely on volunteers to perform time-sensitive tasks. Thus, their success crucially depends on efficient volunteer utilization and engagement. To encourage volunteers to complete a task, platforms use nudging mechanisms to notify a subset of volunteers with the hope that at least one of them responds positively. However, since excessive notifications may reduce volunteer engagement, the platform faces a trade-off between notifying more volunteers for the current task and saving them for future ones. Motivated by these applications, we introduce the online volunteer notification problem, a generalization of online stochastic bipartite matching where tasks arrive following a known time-varying distribution over task types. Upon arrival of a task, the platform notifies a subset of volunteers with the objective of minimizing the number of missed tasks. To capture each volunteer's adverse reaction to excessive notifications, we assume that a notification triggers a random period of inactivity, during which she will ignore all notifications. However, if a volunteer is active and notified, she will perform the task with a given pair-specific match probability that captures her preference for the task. We develop two online randomized policies that achieve constant-factor guarantees which are close to the upper-bounds we establish for the performance of any online policy. Our policies as well as hardness results are parameterized by the minimum discrete hazard rate of the inter-activity time distribution. The design of our policies relies on two modifications of an ex ante feasible solution: (1) properly scaling down the notification probability prescribed by the ex-ante solution, and (2) sparsifying that solution. Further, in collaboration with Food Rescue U.S., a volunteer-based food recovery platform, we demonstrate the effectiveness of our policies by testing them on the platform's data from various locations across the U.S.

---

[343] *A Complete Characterization of Infinitely Repeated Two-Player Games having Computable Strategies with no Computable Best Response under Limit-of-Means Payoff*

Jakub Dargaj (Department of Computer Science, University of Copenhagen) and Jakob Grue Simonsen (Department of Computer Science, University of Copenhagen).

Abstract

It is well-known that for infinitely repeated games, there are computable strategies that have best responses, but no computable best responses. These results were originally proved for either specific games (e.g., Prisoner's dilemma), or for classes of games satisfying certain conditions not known to be both necessary and sufficient.

We present a complete characterization in the form of simple necessary and sufficient conditions for the existence of a computable strategy without a computable best response under limit-of-means payoff. We further refine the characterization by requiring the strategy profiles to be Nash equilibria or subgame-perfect equilibria, and we show how the characterizations entail that it is efficiently decidable whether an infinitely repeated game has a computable strategy without a computable best response.

---

[349] *Learning through the Grapevine: the Impact of Message Mutation, Transmission Failure, and Deliberate Bias*

Matthew Jackson (Stanford University), Suraj Malladi (Stanford University) and David McAdams (Duke Fuqua School of Business).

Abstract

We examine how well agents learn when information is noisily relayed from person to person. Our results suggest how communication platforms can improve learning without observing messages or knowing which ones are false. If noise only takes the form of random mutations and transmission failures, then there is a sharp threshold such that a receiver learns fully if she has access to more chains than the threshold and nothing otherwise. Moreover, simple information-processing rules can perform as well as fully Bayesian learning. However, if some (even small) number of agents deliberately distort message content, learning may be impossible with any number of chains as distance from primary sources grows. Platform designers can improve learning by limiting how often any agent can forward a given message. This policy does not require the ability to fact-check, respects privacy, increases the fraction of true to false messages, and is implemented by some real-world communication platforms for this purpose.

---

[351] *Pricing for the Stars: Dynamic Pricing in the Presence of Rating Systems*

Andre Stenzel (University of Mannheim and MaCCI), Christoph Wolf (Bocconi University) and Peter Schmidt (QuantCo).

Abstract

We study dynamic pricing in the presence of product ratings. A monopolist sells a good of unknown quality to short-lived heterogeneous consumers who observe aggregate ratings reflecting past reviews. Long-run outcomes depend on the sensitivity of the rating system to incoming reviews and the degree to which reviews internalize the purchase price. When internalization is high, low prices induce good reviews. For low internalization, good reviews obtain with high prices via selection on consumer tastes. Sensitivity benefits the seller due to easier ratings management, but may harm consumers by exacerbating upward pricing pressure when internalization is low.

---

[355] *Judged in Hindsight: Regulatory Incentives in Approving Innovations*

Suraj Malladi (Stanford University).

Abstract

I study how limited information and ex-post evaluation by third parties with the benefit of hindsight affect how regulators approve innovations. In the face of ambiguity over innovation characteristics, such regulators opt for limited or delayed product approval, even when they are not waiting for new information to arrive. When evidence is costly for firms to generate but can be selectively reported, the regulator delegates information acquisition to the firm with the objective of minimizing max-regret. This model can explain patterns of correlation between firm costs and benefits of approval, why regulators drag their feet on approval even in the face of strong positive evidence, and support for regulatory sandboxes even when they do not hasten learning.

---

[362] *Dynamic Stochastic Matching Under Limited Time*

Ali Aouad (London Business School) and Omer Saritac (London Business School).

Abstract

Motivated by centralized matching markets, we study an online stochastic matching problem on edge-weighted graphs, where the agents' arrivals and abandonments are stochastic and heterogeneous. The problem is formulated as a continuous-time Markov decision process (MDP) under the average-cost criterion. While the MDP is computationally intractable, we design simple matching algorithms that achieve constant-factor approximations in cost-minimization and reward-maximization settings. Specifically, we devise a 3-approximation algorithm for cost minimization on graphs satisfying a metric-like property. We develop a (e-1)/(2e)-approximation algorithm for reward maximization on arbitrary bipartite graphs. Our matching algorithms possess a simple greedily-like structure. In extensive experiments, we simulate the matching operations of a car-pooling platform using real-world taxi demand data. The newly-developed algorithms have the potential to improve cost efficiency by 10% in certain realistic market conditions against the widely used batching algorithms.

---

[363] *A Formal Separation Between Strategic and Nonstrategic Behavior*

James Wright (University of Alberta) and Kevin Leyton-Brown (The University of British Columbia).

Abstract

It is common in multiagent systems to make a distinction between ``strategic'' behavior and other forms of intentional but ``nonstrategic'' behavior: typically, that strategic agents model other agents while nonstrategic agents do not. However, a crisp boundary between these concepts has proven elusive. This problem is pervasive throughout the game theoretic literature on bounded rationality and particularly critical in parts of the behavioral game theory literature that make an explicit distinction between the behavior of ``nonstrategic'' level-$0$ agents and ``strategic'' higher-level agents (e.g., the level-k and cognitive hierarchy models). Overall, work discussing bounded rationality rarely gives clear guidance on how the rationality of nonstrategic agents must be bounded, instead typically just singling out specific decision rules and informally asserting them to be nonstrategic (e.g., truthfully revealing private information; randomizing uniformly). In this work, we propose a new, formal characterization of nonstrategic behavior. Our main contribution is to show that it satisfies two properties: (1) it is general enough to capture all purportedly ``nonstrategic'' decision rules of which we are aware in the behavioral game theory literature; (2) behavior that obeys our characterization is distinct from strategic behavior in a precise sense.

---

[370] *An Improved Approximation Algorithm for Maximin Shares*

Jugal Garg (University of Illinois at Urbana-Champaign) and Setareh Taki (University of Illinois at Urbana-Champaign).

Abstract

We study the problem of fair allocation of m indivisible items among n agents with additive valuations using the popular notion of maximin share (MMS) as our measure of fairness. An MMS allocation provides each agent a bundle worth at least her maximin share. While it is known that such an allocation need not exist, a series of remarkable work provided 2/3 approximation algorithms in which each agent receives a bundle worth at least 2/3 times her maximin share. More recently, Ghodsi et al. [EC'18] showed the existence of 3/4-MMS allocation and a PTAS to find a 3/4-\epsilon MMS allocation. Most of the previous work utilize intricate algorithms and require agents' approximate MMS values, which are computationally expensive to obtain.

In this paper, we develop a new approach that gives a simple algorithm for showing the existence of a 3/4-MMS allocation. Furthermore, our approach is powerful enough to be easily extended in two directions: First, we get a strongly polynomial time algorithm to find a 3/4-MMS allocation, where we do not need to approximate the MMS values at all. Second, we show that there always exists a 3/4+1/(2n)-MMS allocation, breaking the barrier of 3/4. This considerably improves the approximation guarantee for small n. We note that 3/4 was the best factor known for n>4.

---

[373] *Consensus-Halving: Does It Ever Get Easier?*

Aris Filos-Ratsikas (University of Liverpool), Alexandros Hollender (University of Oxford), Katerina Sotiraki (Massachusetts Institute of Technology) and Manolis Zampetakis (Massachusetts Institute of Technology).

Abstract

In the epsilon-Consensus-Halving problem, there are n agents with valuations over the interval [0,1], and the goal is to divide the interval into pieces and assign a label ``+'' or ``-'' to each piece, such that every agent values the total amount of ``+'' and the total amount of ``-'' almost equally. The problem was recently proven by Filos-Ratsikas and Goldberg (2019) to be the first ``natural'' complete problem for the computational class PPA, answering a decade-old open question.

In this paper, we examine the extent to which the problem becomes easy to solve, if one restricts the class of valuation functions. To this end, we provide the following contributions. First, we obtain a strengthening of the PPA-hardness result of Filos-Ratsikas and Goldberg (2019), to the case when agents have piecewise uniform valuations with only two blocks. We obtain this result via a new reduction, which is in fact conceptually much simpler than the corresponding one in [Filos-Ratsikas and Goldberg 2019]. Then, we consider the case of single-block (uniform) valuations and provide a parameterized algortihm for solving the algorithm for any epsilon, as well as a polynomial-time algorithm for epsilon=1/2; these are the first algorithmic results for the problem. Finally, an important application of our new techniques is the first hardness result for a generalization of Consensus-Halving, the Consensus-1/k-Division problem. In particular, we prove that epsilon-Consensus-1/3-Division is PPAD-hard.

---

[385] *Resource-Aware Protocols for Network Cost-Sharing Games*

Giorgos Christodoulou (University of Liverpool), Vasilis Gkatzelis (Drexel University), Alkmini Sgouritsa (University of Liverpool) and Mohamad Latifian (Sharif University of Technology).

Abstract

We study the extent to which decentralized cost-sharing protocols can achieve good price of anarchy (PoA) bounds in network cost-sharing games with n agents. We focus on the model of resource-aware protocols, where the designer has prior access to the network structure and can also increase the total cost of an edge (overcharging), and we study classes of games with concave or convex cost functions. We first consider concave cost functions and our main result is a cost-sharing protocol for symmetric games on directed acyclic graphs that achieves a PoA of $2+\epsilon$ for some arbitrary small positive ε, which improves to $1+\epsilon$ for games with at least two players. We also achieve a PoA of 1 for series-parallel graphs and show that no protocol can achieve a PoA better than $Omega(\sqrt{n})$ for multicast games. We then also consider convex cost functions and prove analogous results for series-parallel networks and multicast games, as well as a lower bound of $\Omega(n)$ for the PoA on directed acyclic graphswithout the use of overcharging.

---

[397] *Meddling Metrics: the Effects of Measuring and Constraining Partisan Gerrymandering on Voter Incentives*

Brian Brubach (University of Maryland, College Park), Aravind Srinivasan (University of Maryland, College Park) and Shawn Zhao (Montgomery Blair High School).

Abstract

Gerrymandering is the process of drawing electoral district maps in order to manipulate the outcomes of elections. Partisan gerrymandering occurs when political parties use this practice to gain an advantage. Increasingly, computers are involved in both drawing biased, partisan districts and in attempts to measure and regulate this practice. Several of the most high-profile proposals to measure partisan gerrymandering involve the use of past voting data. Prior work primarily studies the ability of these metrics to detect gerrymandering. However, it does not account for how legislation based on the metrics could affect voter behavior. We show that even in a two-party election, using past voting data can affect strategyproofness. We further focus on the proposal to ban “outlier maps”, maps which appear biased toward a particular party when compared to a random sampling of legal maps. We introduce a game which models the iterative sequence of voting and redrawing districts under the restriction that outlier maps are forbidden. Using this game, we illustrate strategies for a majority party to increase its seat count by voting strategically. This leads to a heuristic for gaming the system when outliers are banned. Additionally, we address some questions from the recent US Supreme Court case, Rucho v. Common Cause, that relate to our model. Finally, we apply a modified version of our heuristic to past North Carolina voter data and show that the strategies can be applied on real states.

---

[400] *Information Choice in Auctions*

Nina Bobkova (Rice University).

Abstract

Bidders are uncertain about their valuation for an object and choose about which component to learn. Their valuation consists of a common value component (which matters to all bidders) and a private value component (which is relevant only to individual bidders). Learning about a private component yields independent estimates, whereas learning about a common component leads to correlated information between bidders. I identify conditions for the second-price auction, such that bidders only learn about their private component: an independent private value framework and an efficient outcome arise endogenously. In a first-price auction, every robust equilibrium is inefficient under certain conditions.

---

[408] *Online Matching with Stochastic Rewards: Optimal Competitive Ratio via Path Based Formulation*

Rajan Udwani (Columbia University) and Vineet Goyal (Columbia University).

Abstract

The problem of online matching with stochastic rewards is a generalization of the online bipartite matching problem where each edge has a probability of success. When a match is made it succeeds with the probability of the corresponding edge. Introducing this model, Mehta and Panigrahi (FOCS 2012) focused on the special case of identical edge probabilities. Comparing against a deterministic offline LP, they showed that the Ranking algorithm of Karp et al. (STOC 1990) is 0.534 competitive and proposed a new online algorithm with an improved guarantee of 0.567 for vanishingly small probabilities. For the case of vanishingly small but heterogeneous probabilities Mehta et al.\ (SODA

2015), gave a 0.534 competitive algorithm against the same LP benchmark. For the more general vertex-weighted version of the problem, to the best of our knowledge, no results being $1/2$ were previously known even for identical probabilities.

We focus on the vertex-weighted version and give two improvements. First, we show that a natural generalization of the Perturbed-Greedy algorithm of Aggarwal et al.\ (SODA 2011), is $(1-1/e)$ competitive when probabilities decompose as a product of two factors, one corresponding to each vertex of the edge. This is the best achievable guarantee as it includes the case of identical probabilities and in particular, the classical online bipartite matching problem. Second, we give a deterministic 0.596 competitive algorithm for the previously well studied case of case of fully heterogeneous but vanishingly small edge probabilities. A key contribution of our approach is the use of novel path-based formulations and a generalization of the primal-dual scheme of Devanur et al.\ (SODA 2013). These allow us to compare against the natural benchmark of clairvoyant (offline) algorithms that know the sequence of arrivals and the edge probabilities in advance, but not the outcomes of potential matches. These ideas may be of independent interest in other online settings with post-allocation stochasticity.

---

[411] *The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify*

David Holtz (MIT Sloan School of Management), Benjamin Carterette (Spotify), Praveen Chandar (Spotify), Zahra Nazari (Spotify), Henriette Cramer (Spotify) and Sinan Aral (MIT Sloan School of Management).

Abstract

It remains unknown whether personalized recommendations increase or decrease the diversity of content people consume. In this paper, we present results from a large-scale, randomized field experiment on Spotify testing the effect of personalized recommendations on consumption diversity. In the experiment, both control and treatment users were given podcast recommendations, with the sole aim of increasing podcast consumption. However, the recommendations provided to treatment users were personalized based on their music listening history, whereas control users were recommended the most popular podcasts among users in their demographic group. Consistent with previous studies, we find that the treatment increased the average number of podcast streams per user by 28.90%. However, we also find the treatment decreased the average individual-level diversity of podcast streams by 11.51%, and increased the aggregate diversity of podcast streams by 5.96%, indicating that personalized recommendations have the potential to create patterns of consumption that are homogenous within users and diverse across users, a pattern reflecting Balkanization. Our results provide evidence of an "engagement-diversity trade-off" when recommendations are optimized solely to drive consumption: while personalized recommendations increase user engagement, they also affect the diversity of content that users consume. This shift in consumption diversity can affect user retention and lifetime value, and also impact the optimal strategy for content producers. We also observe evidence that our experimental treatment affected streams originating from sections of Spotify's app that were not directly affected by the experiment, suggesting that exposure to personalized recommendations can also affect the content that users consume organically. We believe these findings highlight the need for both academics and practitioners to continue investing in approaches to personalization that explicitly take into account the diversity of content recommended to users.

---

[412] *Continuous Credit Networks and Layer 2 Blockchains: Monotonicity and Sampling*

Geoffrey Ramseyer (Stanford University) and Ashish Goel (Stanford University).

Abstract

To improve transaction rates, many cryptocurrencies have implemented so-called ``Layer-2'' transaction protocols, where payments are routed across networks of private payment channels. However, for a given transaction, not every network state provides a feasible route to perform the payment; in this case, the transaction must be put on the public ledger. The payment channel network thus multiplies the transaction rate of the overall system; the less frequently it fails, the higher the multiplier.

We build on earlier work on credit networks and show that this network liquidity problem is connected to the combinatorics of graphical matroids. Earlier work could only analyze the (unnatural) scenario where transactions had discrete sizes. In this work, we give an analytical framework that removes this assumption. This enables meaningful liquidity analysis for real-world parameter regimes.

Superficially, it might seem like the continuous case would be harder to examine. However, removing this assumption lets us make progress in two important directions. First, we give a partial answer to the ``monotonicity conjecture'' that previous work left open. This conjecture asks that the network's performance not degrade as capacity on any edge increases. And second, we construct here a network state sampling procedure with much faster asymptotic performance than off-the-shelf Markov chains ($O(|E|\beta(|E|))$, where $\beta(x)$ is the complexity of solving a linear program on x constraints.)

We obtain our results by mapping the underlying graphs to convex bodies and then showing that the liquidity and sampling problems reduce to bounding and computing the volumes of these bodies. The transformation relies crucially on the combinatorial properties of the underlying graphic matroid, as do the proofs of monotonicity and fast sampling.

---

[422] *How Does Market Thickness impact Inequity? Evidence from Ridesharing*

Soheil Ghili (Yale University) and Vineet Kumar (Yale University).

Abstract

This paper develops a strategy with simple implementation and limited data requirements to identify spatial distortion of supply from demand resulting in inequity in access to supply across regions in transportation markets. We apply our method to ride-level, multi-platform data from New York City (NYC) and show that for smaller rideshare platforms, supply tends to be disproportionately concentrated in more densely populated areas.We estimate the minimum required platform size to avoid geographical supply distortions, which informs the current policy debate in NYC around whether ridesharing platforms should be downsized. We find the minimum required size to be approximately 3.5M rides/month for NYC, implying that downsizing Lyft or Via–but not Uber–can increase geographical inequity.

---

[423] *Simple, Credible, and Approximately-Optimal Auctions*

Costis Daskalakis (MIT), Maxwell Fishelson (MIT), Brendan Lucier (Microsoft), Santhoshini Velusamy (Harvard University) and Vasilis Syrgkanis (Microsoft).

Abstract

We identify the first credible mechanism for multi-item additive auctions that achieves a constant factor of the optimal revenue. This is one instance of a more general framework for designing two-part tariff auctions, adapting the duality framework of Cai et al. (2016). Given a (not necessarily incentive compatible) auction format A satisfying certain technical conditions, our framework augments the auction with a personalized entry fee for each bidder, which must be paid before the auction can be accessed. These entry fees depend only on the prior distribution of bidder types, and in particular are independent of realized bids. Our framework can be used with many common auction formats, such as simultaneous first-price, simultaneous second-price, and simultaneous all-pay auctions. If all-pay auctions are used, we prove that the resulting mechanism is credible in the sense that the auctioneer cannot benefit by deviating from the stated mechanism after observing agent bids. If second-price auctions are used, we obtain a truthful O(1)-approximate mechanism with fixed entry fees that are amenable to tuning via online learning techniques.

---

[428] *Allocation with Correlated Information: Too good to be true*

Deniz Kattwinkel (Bonn University).

Abstract

A principal can allocate an indivisible good to an agent. The agent privately learns the value of the good while the principal privately learns the cost. Value and cost are correlated. The agent wants to have the good in any case. The principal wants to allocate whenever the value exceeds the cost. She cannot use monetary transfers to screen the agent.

I study how the principal utilizes her information in the optimal mechanism: when the correlation is negative, she bases her decision only on the costs, and when the correlation is positive, she screens the agent. To this end, she forgoes her best allocation opportunities: when the agent reports high valuations but her own costs are low. Under positive correlation, these realizations are unlikely; the principal will find them too good to be true. In contrast to standard results, this optimal mechanism may not allocate to a higher value agent with higher probability. I discuss applications to intra-firm allocations, task-delegation, and industry self-regulation.

---

[430] *Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation*

Rupert Freeman (Microsoft), Nisarg Shah (University of Toronto) and Rohit Vaish (Rensselaer Polytechnic Institute).

Abstract

We study the problem of allocating indivisible goods among agents. When randomization is allowed, it is possible to achieve compelling notions of fairness such as envy-freeness, which states that no agent should prefer any other agent's

allocation to her own. When allocations must be deterministic, achieving exact fairness is impossible but approximate notions such as envy-freeness up to one good can be guaranteed. Our goal in this work is to achieve both, by constructing a randomized allocation that is exactly fair ex-ante and approximately fair ex-post. We show that ex-ante envy-freeness can be achieved in combination with ex-post envy-freeness up to one good. If we additionally require economic efficiency, we obtain an impossibility result. However, we show that economic efficiency and ex-ante envy-freeness can be simultaneously achieved if we slightly relax our ex-post fairness guarantee. On our way, we characterize the well-known Maximum Nash Welfare allocation rule in terms of a recently introduced fairness guarantee that applies to groups of agents, not just individuals.

---

[431] *Portfolio Compression in Financial Networks: Incentives and Systemic Risk*

Steffen Schuldenzucker (Goethe University Frankfurt) and Sven Seuken (University of Zurich).

Abstract

We study portfolio compression, a post-trade mechanism that eliminates cycles in a financial network. We study the incentives for banks to engage in compression and its systemic effects in terms of all banks' equities. We show that, contrary to conventional wisdom, compression may be socially and individually detrimental and incentives may be misaligned with social welfare. We show that these effects depend on the parameters of the financial system and the compression in a complex and non-monotonic way. We then present sufficient conditions under which compression is incentivized for participating banks or a Pareto improvement for all banks. Our results contribute to a better understanding of the implications of recent regulatory policy.

---

[444] *Queue Lengths as Constantly Adapting Prices: Allocative Efficiency Under Random Dynamics*

Itai Ashlagi (Stanford University), Jacob Leshno (University of Chicago), Pengyu Qian (Columbia University) and Amin Saberi (Stanford University).

Abstract

This paper studies a dynamic economy, in which unit-demand agents and items of finitely many different types arrive randomly over an infinite horizon. Items are allocated to agents through decentralised queues, with a separate queue for each type. Agents have private values over items, and their utility is quasi-linear in waiting costs. Upon arrival, they have to make an irrevocable decision on whether to join one of the queues or leave.

In these markets, the waiting times play the role of prices in guiding the allocation and rationing items. However, they are inherently stochastic, as the randomness in the arrival of items and agents to constant fluctuations in waiting times. We quantify the allocative efficiency loss resulting from this fluctuation as a function of the marginal waiting cost of having an additional agent in the queue. To do this, we link the adjustment of queue lengths, which act as prices, to stochastic gradient descent.

---

[450] *Small Steps to Accuracy: Incremental Belief Updaters Are Better Forecasters*

Pavel Atanasov (Pytho LLC), Jens Witkowski (Frankfurt School of Finance & Management), Lyle Ungar (University of Pennsylvania), Barbara Mellers (University of Pennsylvania) and Philip Tetlock (University of Pennsylvania).

Abstract

We study the belief updating patterns of real-world forecasters and relate those patterns to forecaster accuracy. We distinguish three aspects of belief updating: frequency of updates, magnitude of updates, and each forecaster's confirmation propensity (i.e., a forecaster's tendency to restate her preceding forecast). Drawing on data from a four-year forecasting tournament that elicited over 400,000 probabilistic predictions on almost 500 geopolitical questions, we find that the most accurate forecasters make frequent, small updates, while low-skill forecasters are prone to make infrequent, large revisions or to confirm their initial judgments. Relating these findings to behavioral and psychometric measures, we find that high-frequency updaters tend to demonstrate deeper subject-matter knowledge and more open-mindedness, access more information, and improve their accuracy over time. Small-increment updaters tend to score higher in fluid intelligence and obtain their advantage from superior accuracy in their initial forecasts. Hence, frequent, small forecast revisions are a reliable signal of skill.

[456] *Multi-Item Mechanisms without Item-Independence: Learnability via Robustness*

Johannes Brustle (McGill University), Yang Cai (Yale University) and Constantinos Daskalakis (Massachusetts Institute of Technology).

Abstract

We study the sample complexity of learning revenue-optimal multi-item auctions. We obtain the first set of positive results that go beyond the standard but unrealistic setting of item-independence. In particular, we consider settings where bidders' valuations are drawn from correlated distributions that can be captured by Markov Random Fields or Bayesian Networks -- two of the most prominent graphical models. We establish parametrized sample complexity bounds for learning an up-to-ε optimal mechanism in both models, which scale polynomially in the size of the model, i.e.~the number of items and bidders, and only exponential in the natural complexity measure of the model, namely either the largest in-degree (for Bayesian Networks) or the size of the largest hyper-edge (for Markov Random Fields).

We obtain our learnability results through a novel and modular framework that involves first proving a robustness theorem. We show that, given only ``approximate distributions'' for bidder valuations, we can learn a mechanism whose revenue is nearly optimal simultaneously for all ``true distributions'' that are close to the ones we were given in Prokhorov distance. Thus, to learn a good mechanism, it suffices to learn approximate distributions. When item values are independent, learning in Prokhorov distance is immediate, hence our framework directly implies the main result of Gonczarowski and Weinberg~\cite{GonczarowskiW18}. When item values are sampled from more general graphical models, we combine our robustness theorem with novel sample complexity results for learning Markov Random Fields or Bayesian Networks in Prokhorov distance, which may be of independent interest. Finally, in the single-item case, our robustness result can be strengthened to hold under an even weaker distribution distance, the L\'evy distance.

---

[463] *Variance-Weighted Estimators to Improve Sensitivity in Online Experiments*

Kevin Liou (Facebook) and Sean Taylor (Lyft).

Abstract

Running studies with high sensitivity is the gold standard in online controlled experimentation. As companies increasingly rely on experiments to make product decisions, being able to precisely measure a change in a metric is important. Over the past several years, various methods to increase sensitivity in experiments have been proposed, including methods using pre-experimental data, machine learning approaches, and implementation of more advanced sampling strategies. However, prior work has not explored efficiency improvements from considering heterogeneity in the variance of individual experimental units. In this paper, we propose a lower variance average treatment effect estimator that relies on modeling the individual variance of units in an experiment using historical data. Using these heterogeneous variance estimates, we show that one can increase the precision of the treatment effect, and prove that the coefficient of variation of the noise of the sample population is a sufficient statistic for determining the magnitude of possible variance reduction. We provide empirical results from case studies at Facebook demonstrating the effectiveness of this approach, where the average experiment achieved a 17% decrease in variance with minimal impact on bias.

---

[483] *More Revenue from Two Samples via Factor Revealing SDPs*

Constantinos Daskalakis (Massachusetts Institute of Technology) and Manolis Zampetakis (Massachusetts Institute of Technology).

Abstract

We consider the classical problem of selling a single item to a single bidder whose value for the item is drawn from a regular distribution F, in a ``data-poor'' regime where F is not known to the seller, and very few samples from F are available. Prior work (Dhangwatnotai et al. '10) has shown that one sample from F can be used to attain a 1/2-factor approximation to the optimal revenue, but it has been challenging to improve this guarantee when more samples from F are provided, even when two samples from F are provided. In this case, the best approximation known to date is 0.509, achieved by the Empirical Revenue Maximizing (ERM) mechanism (Babaioff et al. '18). We improve this guarantee to 0.558, and provide a lower bound of 0.65. Our results are based on a general framework, based on factor-revealing Semidefinite Programming relaxations aiming to capture as tight as possible a superset of regular distributions and product measures of these distributions. The framework is general and can be applied in more abstract settings to evaluate the performance of a policy chosen using independent samples from a distribution and applied on a fresh sample from that same distribution.

[486] *Optimizing Offer Sets in Sub-Linear Time*

Vivek Farias (MIT), Andrew Li (Carnegie Mellon University) and Deeksha Sinha (MIT).

Abstract

Personalization and recommendations are now accepted as core competencies in just about every online setting, ranging from media platforms to e-commerce to social networks. While the challenge of estimating user preferences has garnered significant attention, the operational problem of using such preferences to construct personalized offer sets to users is largely still open, particularly in modern settings where a massive number of items and a millisecond response time requirement mean that even enumerating all of the items is impossible. Faced with such settings, existing techniques are either (a) entirely heuristic with no principled justification, or (b) theoretically sound, but simply too slow to work.

Thus motivated, we propose an algorithm for personalized offer set optimization that runs in time sub-linear in the number of items while enjoying a uniform performance guarantee. Our algorithm works for an extremely general class of problems and models of user choice that includes the mixture MNL model as a special case. We achieve a sub-linear runtime by leveraging the dimensionality reduction from learning an accurate latent factor model, along with existing sub-linear time approximate near neighbor algorithms. Our algorithm can be entirely data-driven, relying on samples of the user, where a 'sample' refers to the user interaction data typically collected by firms. We evaluate our approach on a massive content discovery dataset from Outbrain that includes millions of advertisements. Results show that our implementation indeed runs on the order of milliseconds, with increased performance relative to existing fast heuristics.

---

[491] *Biased Programmers? Or Biased Data? A Field Experiment in Operationalizing AI Ethics*

Bo Cowgill (Columbia), Fabrizio Dell'Acqua (Columbia University), Augustin Chaintreau (Columbia), Nakul Verma (Columbia) and Samuel Deng (Columbia).

Abstract

Why do biased beliefs arise, and what interventions can prevent them? We quantify the effects of these using a field experiment using machine learning to predict human capital. We randomly assign $\approx$400 AI engineers to predict standardized test scores of OECD residents under different experimental conditions. We then assess the resulting predictive algorithms using out-of-sample test scores and through randomized, audit manipulations of algorithmic inputs. We find that biased beliefs are mostly caused by biased training data. However, simple reminders about bias are almost half as effective as fully de-biasing training data. We find mixed results on technical education and programmer demographics. Programmers who understand technical guidance successfully reduce bias. However, many do not follow the advice, resulting in algorithms that are \emph{worse} than programmers given a simple reminder. Predictions by female and minority AI programmers do not exhibit less bias or discrimination. However prediction errors are correlated within engineering demographics, creating bias reductions from cross-demographic averaging. We also find no effects of our incentive treatments and no evidence that programmers' implicit associations between gender and math (measured through an implicit association test) are correlated with bias in code.