Abstracts

Strategyproof and Proportionally Fair Facility Location

Haris Aziz[1], Alexander Lam[1(✉)], Barton E. Lee[2], and Toby Walsh[1,3]

[1] UNSW Sydney, Sydney, Australia
{haris.aziz,alexander.lam1,t.walsh}@unsw.edu.au
[2] ETH Zürich, Zürich, Switzerland
bartonlee@ethz.ch
[3] Data61 CSIRO, Eveleigh, Australia

Abstract. We focus on a simple, one-dimensional collective decision problem (often referred to as the facility location problem) and explore issues of strategyproofness and proportional fairness. We present several characterization results for mechanisms that satisfy strategyproofness and varying levels of proportional fairness. We also characterize one of the mechanisms as the unique equilibrium outcome for any mechanism that satisfies natural fairness and monotonicity properties. Finally, we identify strategyproof and proportionally fair mechanisms that provide the best welfare-optimal approximation among all mechanisms that satisfy the corresponding fairness axiom.

ArXiV link: https://arxiv.org/pdf/2111.01566.pdf

Keywords: Facility location · Mechanism design · Social choice · Fairness · Strategyproofness

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 351, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Robust Misspecified Models and Paradigm Shifts

Cuimin Ba(✉)

University of Pennsylvania, Philadelphia, PA 19104, USA
cuiminba@sas.upenn.edu

Abstract. This paper studies which misspecified models are likely to persist when the decision maker compares her model with competing models. I present a framework where the agent learns about an action-dependent data-generating process and makes decisions repeatedly. Aware of potential model misspecification, she uses a Bayes factor criterion to switch between models according to how well they fit the data. The main result provides a characterization of persistent models based on the model-induced equilibrium, properties of the learning process such as priors and the switching threshold, and the set of competing models that may arise. I show that misspecified models can be robust against a wide range of competing models—including the true data-generating process, despite the agent having an infinite amount of data. Moreover, simple misspecified models with entrenched priors may have even better robustness properties than correctly specified models. I use these results to provide learning foundations for the persistence of systemic biases in two canonical applications: first, in a natural class of effort-choice problems, overconfidence in one's ability is more robust than underconfidence; second, an oversimplified binary view in politics trumps a correct view and leads to polarization when individuals consume media without fully recognizing the media bias.

Keywords: Robust misspecified models · Learning with misspecified models · Self-confirming equilibrium · Berk-Nash equilibrium

The full paper is available at https://cuiminba.github.io/Papers/Robust_Models.pdf.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 352, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Information Design in Allocation with Costly Verification

Yi-Chun Chen[1], Gaoji Hu[2(✉)], and Xiangqian Yang[3]

[1] Department of Economics and Risk Management Institute, National University of Singapore, Singapore 117570, Singapore
ecsycc@nus.edu.sg

[2] School of Economics, Shanghai University of Finance and Economics, Shanghai 200433, China
hugaoji@sufe.edu.cn

[3] Department of Economics, Hunan University, Changsha 410082, China
yangxiangqian@hnu.edu.cn

Abstract. This paper studies information design in the context of allocation with costly verification à la [1]. Particularly, a principal who values an object allocates it to one or more agents. Agents learn private information (signals) from an information designer about the allocation payoff to the principal. Monetary transfer is not available but the principal can costly verify agents' private signals. The information designer can influence the agents' signal distributions, based upon which the principal maximizes the allocation *surplus*. An agent's utility is simply the *probability* of obtaining the good.

With a single agent, we characterize (i) the agent-optimal information, (ii) the principal-worst information, and (iii) the principal-optimal information. For concrete examples, making the signal distribution the least informative is principal-worst and the most informative being principal-optimal. An agent-optimal information pools information above a cutoff signal and fully reveals information below the cutoff. Even though the objectives of the principal and the agent are not directly comparable, any agent-optimal information is principal-worst, but not the converse.

With multiple agents, agent-optimal information maximizes the total probability of agents' obtaining the good. Compared with the prior distribution, under some agent-optimal information, all agents can be better off; while under some other agent-optimal information, some agents get worse off. Moreover, agent-optimal informations may deliver different payoffs to the principal, which implies that an agent-optimal information need *not* be principal-worst.

The principal's payoff under the principal-worst information provides an upper bound for the payoff that can be achieved by a "robust" mechanism which does not depend on details of the agent's type distribution. We find a robust mechanism that does achieve such an upper bound payoff, which is therefore an optimal robust mechanism. Moreover, allowing for correlated distributions does not affect the result.

https://ssrn.com/abstract=4245445.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, pp. 353–354, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Reference

1. Ben-Porath, E., Dekel, E., Lipman, B.L.: Optimal allocation with costly verification. Am. Econ. Rev. **104**(12), 3779–3813 (2014)

Optimal Private Payoff Manipulation Against Commitment in Extensive-form Games

Yurong Chen[1(✉)], Xiaotie Deng[1(✉)], and Yuhao Li[2(✉)]

[1] Center on Frontiers of Computing Studies, School of Computer Science, Peking University, Beijing 100871, China
{chenyurong,xiaotie}@pku.edu.cn
[2] Columbia University, New York, NY 10027, USA
yuhaoli@cs.columbia.edu

Abstract. To take advantage of strategy commitment, a useful tactic of playing games, a leader must learn enough information about the follower's payoff function. However, this leaves the follower a chance to provide fake information and influence the final game outcome. Through a carefully contrived payoff function misreported to the learning leader, the follower may induce an outcome that benefits him more, compared to the ones when he truthfully behaves.

We study the follower's optimal manipulation via such strategic behaviors in extensive-form games. Followers' different attitudes are taken into account. An optimistic follower maximizes his true utility among all game outcomes that can be induced by some payoff function. A pessimistic follower only considers misreporting payoff functions that induce a unique game outcome. For all the settings considered in this paper, we characterize all the possible game outcomes that can be induced successfully. We show that it is polynomial-time tractable for the follower to find the optimal way of misreporting his private payoff information. Our work completely resolves this follower's optimal manipulation problem on an extensive-form game tree.

Full version of the paper can be found at https://arxiv.org/abs/2206.13119.

Keywords: Stackelberg equilibrium · Strategic behavior · Private information manipulation · Extensive-form games

This work is supported by Science and Technology Innovation 2030 -"The Next Generation of Artificial Intelligence" Major Project No. (2018AAA0100901).
Y. Li—Supported by NSF grants CCF-1563155, CCF-1703925, IIS-1838154, CCF-2106429 and CCF-2107187.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 355, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Optimal Feature-Based Market Segmentation and Pricing

Titing Cui[1(✉)] and Michael L. Hamilton[2]

[1] Katz Graduate School of Business, University of Pittsburgh, Pittsburgh, PA 15260, USA
tic54@pitt.edu

[2] Katz Graduate School of Business, University of Pittsburgh, Pittsburgh, PA 15260, USA
mhamilton@katz.pitt.edu

Abstract. In this work, we study semi-personalized pricing strategies where a seller uses features about their customers to segment the market, and customers are offered segment-specific prices. In general, finding jointly optimal market segmentation and pricing policies is computationally intractable, with practitioners often resorting to heuristic segment-then-price strategies. In response, we study how to optimize and analyze feature-based market segmentation and pricing under the assumption that the seller has a trained (noisy) regression model mapping features to valuations. First, we establish novel hardness and approximation results in the case when model noise is independent. Second, in the common cases when the noise in the model is log-concave, we show the joint segmentation and pricing problem can be efficiently solved, and characterize a number of attractive structural properties of the optimal feature-based market segmentation and pricing. Finally, we conduct a case study using home mortgage data, and show that compared to heuristic approaches, our optimal feature-based market segmentation and pricing policies can achieve nearly all of the available revenue with only a few segments.

The full paper can be found at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4151103.

Keywords: Market segmentation · Personalized pricing · Third degree price discrimination · Regression

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 356, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Competition Among Parallel Contests

Xiaotie Deng[1], Ningyuan Li[1], Weian Li[1], and Qi Qi[2(✉)]

[1] Center on Frontiers of Computing Studies, School of Computer Science, Peking University, Beijing, China
{xiaotie,liningyuan,weian_li}@pku.edu.cn
[2] Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China
qi.qi@ruc.edu.cn

Abstract. Contest depicts a scene in which many players compete for several designed prizes, capturing many realistic game-theoretical settings involving competition, and is an important part of mechanism design theory, which has attracted the attention of many researchers from the past to the present. So far, most of the research literature in contest theory has focused on the setting of a single contest and aimed to design the rewarding policy to achieve some specific goals. With the emergence of crowdsourcing competitions, contests are becoming increasingly popular. More and more contests are run in parallel nowadays.

In this paper, we investigate the model of multiple contests held in parallel, where each contestant selects one contest to join and each contest designer decides the prize structure to compete for the participation of contestants. We first analyze the strategic behaviors of contestants and completely characterize the symmetric Bayesian Nash equilibrium. As for the strategies of contest designers, when other designers' strategies are known, we show that computing the best response is NP-hard and propose a fully polynomial time approximation scheme (FPTAS) to output the ϵ-approximation best response. When other designers' strategies are unknown, we perform a worst-case analysis of one designer's strategy. An upper bound on the worst-case utility of any strategy is derived and taken as a benchmark. We propose a method to construct a strategy whose utility can guarantee a constant ratio of the benchmark.

The full version is available at https://arxiv.org/abs/2210.06866.

Keywords: Competition · Parallel contests · Equilibrium behavior · Best response · Safety level

This research was partially supported by the National Natural Science Foundation of China (NSFC) (No.62172012), and the Research Funds of Renmin University of China (22XNKJ07) and No.297522503709, Beijing Outstanding Young Scientist Program No.BJJWZYJH012019100020098, and Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the "Double-First Class" Initiative, Renmin University of China.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 357, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Revenue Management with Product Retirement and Customer Selection

Adam N. Elmachtoub[1], Vineet Goyal[1], Roger Lederman[2], and Harsh Sheth[1(✉)]

[1] Department of Industrial Engineering and Operations Research and Data Science Institute, Columbia University, New York, USA
{ae2516,vg2277,hts2112}@columbia.edu
[2] Amazon, Seattle, USA
rllederm@amazon.com

Abstract. We consider a multi-period, multi-product revenue management problem where in each period the seller has a fixed inventory of multiple substitutable products to sell over a fixed time horizon. In each time period, the seller chooses which subset of products to retire and also selects a customer to visit. When a product is retired, it becomes unavailable to all future customers. When a customer is selected, all available products – non-retired products with positive remaining inventory – are offered for the customer to choose from. The objective of the seller is to dynamically retire products and select customers in order to maximize the total expected revenue over a fixed time horizon. Such product retirement decisions are essential when the seller is not able to personalize the set of products offered to each customer.

When customers choose according to the same multinomial logit model, we show that a deterministic product retirement policy is asymptotically optimal as the inventories grow large. For multiple customer types, we give an asymptotically optimal policy for product retirement and customer selection when the upper bound linear program has an optimal solution with specific structure. We show that such solution can always be found when there are only two products. In the general case with multiple customer types and products, we design a linear programming-based policy that guarantees a constant fraction of the optimal dynamic retirement-selection policy. Finally, we show that our policies perform well in numerical experiments calibrated with real data, compared to natural benchmarks.

Full paper: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4033922.

Work of this author was conducted while employed by IBM Research.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 358, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Order Selection Problems in Hiring Pipelines

Boris Epstein(✉) and Will Ma

Graduate School of Business, Columbia University, New York, NY 10027, USA
{bepstein25,wm2428}@gsb.columbia.edu

Abstract. Motivated by hiring pipelines, we study two order selection problems in which applicants for a finite set of positions must be interviewed or made offers sequentially. There is a finite time budget for interviewing or making offers, and a stochastic realization after each decision, leading to computationally-challenging problems.

In the first problem we study sequential interviewing. In this setting, a firm must interview candidates to fill out k job positions. There is a pool of n candidates, and the values obtained by the firm from hiring each one of them are non-negative random variables sampled from known, independent distributions. The firm can interview candidates to learn the realization of their values. Up to T interviews can be sequentially carried out, after which k out of the T interviewed candidates are chosen for hire. We show that a computationally tractable, non-adaptive policy that must make offers immediately after interviewing is approximately optimal, assuming offerees always accept their offers.

In the second problem, there are again k positions but we assume that the n applicants have already been interviewed. They accept offers independently according to known probabilities. Moreover, offers can be sent in parallel, under the constraint that at each of the T time periods the amount of offers sent does not exceed the amount of positions remaining. We develop a computationally tractable policy that makes offers for the different positions in parallel, which is approximately optimal even relative to a policy that can make kT offers sequentially.

Our two results both generalize and improve the guarantees in the work of Purohit et al. [1] on hiring algorithms, from 1/2 and 1/4 to approximation factors that are at least $1 - 1/e \approx 63.2\%$. Our algorithms work by solving LP relaxations of the corresponding problems, and then rounding the optimal solutions to decide which candidates to interview and send offers to.

Keywords: Hiring · Order selection · Stochastic probing · Adaptivity gap

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, pp. 359–360, 2022.
https://doi.org/10.1007/978-3-031-22832-2

A full version of this paper can be found in https://arxiv.org/abs/2210.04059.

Reference

1. Purohit, M., Gollapudi, S., Raghavan, M.: Hiring under uncertainty. In: International Conference on Machine Learning, pp. 5181–5189. PMLR (2019)

Strategyproofness in Kidney Exchange with Cancellations

Itai Feigenbaum[1,2(✉)]

[1] Lehman College, City University of New York, Bronx, New York 10468, USA
itai.feigenbaum@lehman.cuny.edu
[2] The Graduate Center, City University of New York, New York, NY 10016, USA

Patients requiring kidney transplant may have proxy donors: people who want to donate a kidney to the patient, but can't due to medical incompatibility. p is called a proxy patient of d if d is a proxy donor of p. Some patients, called overloaded, have multiple proxy donors. The pool of a patient is the set of her proxy donors. Patients can swap proxy donors, so that each swapping patient ends up with a compatible donor. A matching is a set of planned transplants resulting from swaps, altruistic donations from donors without proxy patients, and donations to the waiting list of patients without proxy donors. $C_{\max}$ is the maximum cycle size allowed, and $\Lambda_{\max}$ is the maximum number of donations allowed from a pool.

In practice, many planned transplants get canceled. Cancellation of a transplant from donor d to patient p can be direct—due to reasons involving d and p, or indirect—due to cancellation of a transplant to d's proxy patient. We assume that donors in directly-but-not-indirectly canceled transplants are redirected to donate to the waiting list. We want matchings to maximize the objective of the expected number of actually executed transplants. Exact maximization introduces perverse incentives for overloaded patients, who can increase the probability they receive a kidney by hiding some of their proxy donors.

We design the SuperGreedy Algorithm, which incentivizes patients to fully reveal their pools: each patient maximizes the probability of receiving a kidney by full revelation. We then assume a uniformly constant direct cancellation probability $1-\alpha$ for all transplants. When $\alpha < \frac{1}{\Lambda_{\max}}$, we prove a bound of $\frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha}\max\{(1+2\alpha-2\alpha^2),(1-\alpha)(1+C_{\max}\alpha)\}$ on SuperGreedy's approximation ratio. Next, we implement SuperGreedy via integer programming, and simulate it on realistic data. Our results show much better performance than our theoretical bound. Specifically, we get an upper bound on the average approximation ratio of 1.142 when $\Lambda_{\max}=1$ and $\alpha=0.3$; as $\Lambda_{\max}$ increases and α decreases, the bound decreases, down to 1.016 when $\Lambda_{\max}=4$ and $\alpha=0.1$. A full version is available at http://www.itaifeigenbaum.com/WINEStrategicCancellations.pdf.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 361, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Tractable Fragments of the Maximum Nash Welfare Problem

Jugal Garg[1], Edin Husić[2], Aniket Murhekar[1(✉)], and László Végh[2]

[1] University of Illinois, Urbana-Champaign, USA
{jugal,aniket2}@illinois.edu
[2] London School of Economics and Political Science, London, UK
{e.husic,l.vegh}@lse.ac.uk

We study the problem of maximizing Nash welfare (MNW) while allocating indivisible goods to asymmetric agents. The Nash welfare of an allocation is the weighted geometric mean of agents' utilities, and the allocation with maximum Nash welfare is known to satisfy several desirable fairness and efficiency properties. However, computing such an MNW allocation is NP-hard, even for two agents with identical, additive valuations. Hence, we aim to identify tractable classes that either admit a PTAS, an FPTAS, or an exact polynomial-time algorithm. To this end, we design a PTAS for finding an MNW allocation for the case of asymmetric agents with identical, additive valuations, thus generalizing a similar result for symmetric agents [2]. We also extend our PTAS to compute a nearly Nash-optimal allocation which also satisfies the best fairness guarantee offered by the optimal MNW allocation (a weighted relaxation of envy-freeness); showing we do not need to compromise fairness for tractability. Our techniques can also be adapted to give (i) a PTAS for the problem of computing the optimal p-mean welfare, and (ii) a polynomial time algorithm for computing an MNW allocation for identical agents with k-ary valuations when k is a constant, where every agent has at most k different values for the goods. Next, we consider the special case where every agent finds at most two goods valuable, and show that this class admits an efficient algorithm, even for general monotone valuations. In contrast, we note that when agents can value three or more goods, maximizing Nash welfare is NP-hard, even when agents are symmetric and have additive valuations, showing our algorithmic result is essentially tight. Finally, we show that for constantly many asymmetric agents with additive valuations, the MNW problem admits an FPTAS. The full version of the paper is available at [1].

Work supported by the NSF Grant CCF-1942321 and ERC Starting Grant ScaleOpt.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, pp. 362–363, 2022.
https://doi.org/10.1007/978-3-031-22832-2

References

1. Garg, J., Husic, E., Murhekar, A., Végh, L.A.: Tractable fragments of the maximum Nash welfare problem. CoRR abs/2112.10199 (2021). arxiv.org/abs/2112.10199
2. Nguyen, T.T., Rothe, J.: Minimizing envy and maximizing average Nash social welfare in the allocation of indivisible goods. Discrete Appl. Math. **179**(C), 54–68 (2015)

Project Selection with Partially Verifiable Information

Sumit Goel and Wade Hann-Caruthers(✉)

California Institute of Technology, Pasadena, CA 91125, USA
{sgoel,whanncar}@caltech.edu

Abstract. We consider a principal agent project selection problem with asymmetric information. There are N projects and the principal must select exactly one of them. Each project provides some profit to the principal and some payoff to the agent and these profits and payoffs are the agent's private information. If the principal could use transfers, it could essentially sell the firm to the agent and extract the entire surplus. If transfers are not feasible and the agent is unconstrained in its reporting, the principal can do no better than to choose the ex-ante optimal project. However, the agent's ability to manipulate may be constrained due to environmental factors, or because it may be required to furnish evidence in support of its claims. Motivated by such considerations, we consider the problem under a natural partial verifiability constraint of no-overselling wherein the agent cannot report a project to be more profitable than it actually is.

To study this problem, we first characterize the set of implementable mechanisms. As we show, every implementable mechanism can be decomposed into two functions. The first maps each vector of reported profits $\boldsymbol{\pi}$ to a subset of projects, and the second maps each subset of projects T and vector of reported agent payoffs $\boldsymbol{\alpha}$ to the agent-preferred project in T. The first function can be understood as determining the set of projects the principal makes available for the agent to choose from. We show that such a function corresponds to an implementable mechanism if and only if it is increasing: if every project's reported profit under $\boldsymbol{\pi}'$ is (weakly) higher than under $\boldsymbol{\pi}$, then every project made available when $\boldsymbol{\pi}$ is reported must also be made available when $\boldsymbol{\pi}'$ is reported.

Using this characterization, we study the principal's problem of finding an optimal mechanism for two different objectives: maximizing expected profit and maximizing the probability of choosing the most profitable project. For both objectives, we find that in the case of two projects, the optimal mechanism takes the form of a simple cutoff mechanism. The simple structure of the optimal mechanism also allows us to find evidence in support of the well-known ally-principle which says

The full version of the paper is available at https://arxiv.org/abs/2007.00907.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, pp. 364–365, 2022.
https://doi.org/10.1007/978-3-031-22832-2

that the principal delegates more authority to an agent who shares their preferences. In particular, we find the optimal cutoff for the case where principal agent payoffs are distributed bivariate normal and show that it decreases as the payoffs become more correlated.

Keywords: Mechanism design · Partial verifiability · Cutoff mechanisms

Revenue Management Under a Price Alert Mechanism

Bo Jiang[1], Zizhuo Wang[2], and Nanxi Zhang[1,3(✉)]

[1] Research Institute of Interdisciplinary Science, Shanghai University of Finance and Economics, Shanghai, China
jiang.bo@mail.shufe.edu.cn

[2] School of Data Science, The Chinese University of Hong Kong, Shenzhen, Shenzhen 518172, China
wangzizhuo@cuhk.edu.cn

[3] Sauder School of Business, University of British Columbia, Vancouver, BC, Canada
nanxi.zhang@sauder.ubc.ca

Abstract. Many online platforms adopt a price alert mechanism to facilitate customers tracking the price changes. This mechanism allows customers to register their valuation to the system when they find the price is larger than the valuation on their arrival period. Once the price drops below the customers' registered price, a message will be sent to notify them. In this paper, we study the optimal pricing problem under this mechanism. First, when the customer's waiting time is one period, we show that it is optimal for the seller to use a threshold to decide whether to accept or reject a registered price, and the price trajectory under the optimal policy has a stochastic cyclic decreasing structure. When the customer's valuation is a uniform distribution, the analytical form of the optimal policy is further obtained. When the customer's patience level is two periods, we obtain the structure of the optimal policy by showing the asymmetric role each registered price plays in the optimal policy. Then we consider the case when the customer can stay in the system for an infinite number of periods. We derive an asymptotic optimal policy for this case. We find that adopting the price alert mechanism always increases social welfare; however, it may hurt the customer surplus when the seller has a large discount factor. Finally, we consider the case when the customers can strategically react to the price alert mechanism by timing their purchases and reporting false valuations. Using a Stackelberg's game model, we obtain the seller's optimal threshold type of policy. We

Link to the full paper: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4154861.

The first author's research is supported by the National Natural Science Foundation of China (Grants 72171141, 72150001 and 11831002), and Program for Innovative Research Team of Shanghai University of Finance and Economics. The second author's research is partly supported by the National Science Foundation of China (NSFC) Grant 72150002.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, pp. 366–367, 2022.
https://doi.org/10.1007/978-3-031-22832-2

show that the price alert mechanism can still be helpful to the seller, although the advantage diminishes when customers are very strategic.

Keywords: Price alert mechanism · Threshold property · Stochastic cyclic decreasing price · Strategic customer

Bicriteria Nash Flows over Time

Tim Oosterwijk[1(✉)], Daniel Schmand[2], and Marc Schröder[3]

[1] Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081, HV Amsterdam, The Netherlands
t.oosterwijk@vu.nl
[2] University of Bremen, Bibliothekstraße 5, 28359 Bremen, Germany
schmand@uni-bremen.de
[3] Maastricht University, Tongersestraat 53, 6211, LM Maastricht, The Netherlands
m.schroder@maastrichtuniversity.nl

Traffic congestion imposes a huge economic loss to the economy. As such, there has been a huge effort to understand congestion using theoretical models. The dynamic model that gained most attention for modelling traffic is the deterministic fluid queuing model, already introduced by Vickrey [4]. A common drawback of most of the models is the simplified assumption that road network users only aim for minimizing their arrival time [1,2]. However, in traffic networks in particular, users are not always that single-minded. In this paper we extend the state-of-the-art game theoretic traffic models with a multi-criteria objective function. We assume that users try to minimize costs subject to arriving at the sink before a given deadline. Here, costs could be thought of as an intrinsic preference a user has regarding the different route choices and queuing dynamics only play a role for the arrival time of a user.

We determine the existence and the structure of Nash flows over time and fully characterize the price of anarchy for this model, which measures the ratio of the quality of the Nash flow and the optimal flow. We evaluate the quality both with respect to the throughput for a given deadline and the makespan for a given amount of flow. We prove the following three results. (i) In series-parallel graphs, both prices of anarchy are unbounded. (ii) In parallel path graphs the throughput-PoA is at most 2, or at most $\mathrm{e}/(\mathrm{e}-1)$ if all transit times are 0. Both bounds are tight. (iii) In parallel path graphs the makespan-PoA is at most $\mathrm{e}/(\mathrm{e}-1)$, independent of transit times, and this is tight. All our upper bounds are also valid for dynamic equilibria in the deterministic fluid queuing model.

The full version of the paper can be found in [3].

References

1. Cominetti, R., Correa, J., Larré, O.: Dynamic equilibria in fluid queueing networks. Oper. Res. **63**(1), 21–34 (2015)
2. Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. Theor. Comput. Syst. **49**(1), 71–97 (2011)
3. Oosterwijk, T., Schmand, D., Schörder, M.: Bicriteria nash flows over time. arXiv (2021). arxiv.org/abs/2111.08589
4. Vickrey, W.S.: Congestion theory and transport investment. Am. Econ. Rev. **59**(2), 251–260 (1969)

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 368, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Truthful Generalized Linear Models

Yuan Qiu[1(✉)], Jinyan Liu[2], and Di Wang[3]

[1] Georgia Institute of Technology, Atlanta, GA 30332, USA
yuan.qiu@gatech.edu
[2] Beijing Institute of Technology, Beijing 100081, China
jyliu@bit.edu.cn
[3] King Abdullah University of Science and Technology, Thuwal 23955, Saudi Arabia
di.wang@kaust.edu.sa

Abstract. In this paper we study estimating Generalized Linear Models (GLMs) in the case where the agents (individuals) are strategic or self-interested and they concern about their privacy when reporting data. Compared with the classical setting, here we aim to design mechanisms that can both incentivize most agents to truthfully report their data and preserve the privacy of individuals' reports, while their outputs should also close to the underlying parameter. In the first part of the paper, we consider the case where the covariates are sub-Gaussian and the responses are heavy-tailed where they only have the finite fourth moments. First, motivated by the stationary condition of the maximizer of the likelihood function, we derive a novel private and closed form estimator. Based on the estimator, we propose a mechanism which has the following properties via some appropriate design of the computation and payment scheme for several canonical models such as linear regression, logistic regression and Poisson regression: (1) the mechanism is $o(1)$-jointly differentially private (with probability at least $1 - o(1)$); (2) it is an $o(\frac{1}{n})$-approximate Bayes Nash equilibrium for a $(1 - o(1))$-fraction of agents to truthfully report their data, where n is the number of agents; (3) the output could achieve an error of $o(1)$ to the underlying parameter; (4) it is individually rational for a $(1 - o(1))$ fraction of agents in the mechanism; (5) the payment budget required from the analyst to run the mechanism is $o(1)$. In the second part, we consider the linear regression model under more general setting where both covariates and responses are heavy-tailed and only have finite fourth moments. By using an ℓ_4-norm shrinkage operator, we propose a private estimator and payment scheme that have similar properties as in the sub-Gaussian case.

Di Wang is supported in part by King Abdullah University of Science and Technology URF/1/4663-01-01, FCC/1/1976-49-01, and a funding from SDAIA-KAUST AI center. Jinyan Liu is partially supported by National Natural Science Foundation of China (NSFC Grant No.62102026).

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, pp. 369–370, 2022.
https://doi.org/10.1007/978-3-031-22832-2

The full version of the paper is available at https://arxiv.org/pdf/2209.07815.pdf.

Keywords: Generalized linear models · Bayesian game · Differential privacy · Sub-gaussian and heavy-tailed data · Truthful mechanism design

Algorithmic Challenges in Ensuring Fairness at the Time of Decision

Jad Salem[1(✉)], Swati Gupta[1], and Vijay Kamble[2]

[1] Georgia Institute of Technology, Atlanta, GA 30332, USA
{jsalem7,swatig}@gatech.edu
[2] University of Illinois Chicago, Chicago, IL 60607, USA
kamble@uic.edu

Abstract. Algorithmic decision-making in societal contexts such as retail pricing, loan administration, recommendations on online platforms, etc., often involves experimentation with decisions for the sake of learning, which results in perceptions of unfairness amongst people impacted by these decisions. It is hence necessary to embed appropriate notions of fairness in such decision-making processes. The goal of this paper is to highlight the rich interface between temporal notions of fairness and online decision-making through a novel meta-objective of ensuring *fairness at the time of decision*. Given some arbitrary comparative fairness notion for static decision-making (e.g., students should pay at most 90% of the general adult price), a corresponding online decision-making algorithm satisfies fairness at the time of decision if the said notion of fairness is satisfied for any entity receiving a decision *in comparison to all the past decisions*. We show that this basic requirement introduces new methodological challenges in online decision-making. We illustrate the novel approaches necessary to address these challenges in the context of stochastic convex optimization with bandit feedback under a comparative fairness constraint that imposes lower bounds on the decisions received by entities depending on the decisions received by everyone in the past. The paper showcases novel research opportunities in online decision-making stemming from temporal fairness concerns.

Keywords: Fairness · Online learning · Bandit convex optimization

This work was partially supported by NSF grant 2112533.
A complete version of this work can be found at https://arxiv.org/abs/2103.09287.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 371, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Eliminating Waste in Cadaveric Organ Allocation

Peng Shi(✉) and Junxiong Yin

USC Marshall School of Business, Los Angeles, CA 90089, USA
{pengshi,junxiong.yin}@marshall.usc.edu

Abstract. There is a shortage in the supply of cadaveric organs in most countries, but many successfully procured and medically tenable organs are currently being discarded. This wastage of cadaveric organs exacerbates the shortage in organ supply and the financial strains on healthcare systems. Many reforms have been or are currently being implemented to address the wastage problem. However, we show that waste will still be a problem as long as the allocation mechanism continues to prioritize patients by their waiting times, which incentivizes patients to reject organs of reasonable quality now to wait for better offers in the future. Such waiting is risky, as the patients' health conditions may deteriorate while they wait, and they may no longer be fit to receive transplants when the ideal offers come. Through analyzing a theoretical model, we show that the necessary and sufficient conditions to eliminating waste are to disincentivize waiting by allocating over-demanded organ types only to the patients who recently signed up for transplantation, and to give the patients who are not allocated their ideal organs an opportunity to take another offer. However, such a policy may be contentious as it no longer prioritizes patients by waiting times. Moreover, it may reduce the welfare of the patients who are most willing to wait. The benefits of eliminating waste should be weighed against these costs when making policy decisions.

Keywords: Market design · Matching markets · Organ allocation · Wait-list

Link to the full paper:
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4069084

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 372, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Matrix-Exact Covers of Minimum-Cost-Spanning-Tree Games

Zhibin Tan, Zhigang Cao, and Zhengxing Zou(✉)

School of Economics and Management, Beijing Jiaotong University, No. 3 Shangyuancun, Beijing 100044, China
{tanzhibin,zgcao,zhxzou}@bjtu.edu.cn

Abstract. The minimum-cost-spanning-tree (m.c.s.t.) game is a classical cooperative game model that has been extensively studied. Many important solutions for m.c.s.t. games depend on certain pruning operations that construct a new m.c.s.t. game by reducing the costs of certain edges such that the new game is simpler, and its core is a subset of the core of the original game. Examples include the classical irreducible graph, based on which the irreducible core (Bird, 1976), the Folk rule (Feltkamp et al., 1994; Bergantiños and Vidal-Puga, 2007), and the DK rule (Dutta and Kar, 2004) are defined, and the cycle-complete graph (Trudeau, 2012). However, these operations often make the relevant solutions use very little information of the original cost matrix. As criticized by Bogomolnaia and Moulin (2010), "this drastic pruning of the cost data throws away much information relevant to the fairness of the eventual cost sharing".

To answer this criticism, we address the problem of decreasing the connection costs as much as possible in an m.c.s.t. game such that its core does not change. We define the desired m.c.s.t. game as the matrix-exact cover of the original m.c.s.t. game, and show its existence and uniqueness by providing an explicit formula. Our results also imply that the set of all m.c.s.t. games with the same core possesses a meet-semilattice structure: if two m.c.s.t. games have the same core, and we construct a new m.c.s.t. game by defining the cost of each edge as the minimum between the corresponding connection costs of the two games, then the new game has the same core too.

Keywords: Cooperative games · Core · Exact games · Matrix-exact games · Meet-semilattice

A full version is available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4238837

Zhigang Cao was supported by the National Natural Science Foundation of China (Grants No. 71922003, 71871009, 72271016), and Beijing Natural Science Foundation (Grant No. Z220001). Zhengxing Zou was supported by the National Natural Science Foundation of China (Grant No. 72101015).

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 373, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Personalized Assortment Optimization Under Consumer Choice Models with Local Network Effects

Tong Xie[1] and Zizhuo Wang[2(✉)]

[1] Booth School of Business, University of Chicago, Chicago, IL 60637, USA
tong.xie@chicagobooth.edu

[2] School of Data Science, The Chinese University of Hong Kong, Shenzhen, Shenzhen 518172, China
wangzizhuo@cuhk.edu.cn

Abstract. In this paper, we introduce a consumer choice model in which each consumer's utility is affected by the purchase probabilities of his/her neighbors in a network. Such a consumer choice model is a general model to characterize consumer choice under network effect. We first characterize the choice probabilities under such a choice model. Then we consider the associated personalized assortment optimization problem. Particularly, the seller is allowed to offer a personalized assortment to each consumer, and the consumer chooses among the products according to the proposed choice model. We show that the problem is NP-hard even if the consumers form a star network. Despite of the complexity of the problem, we show that if the consumers form a star network, then the optimal assortment to the central consumer cannot be strictly larger than that without network effects; and the optimal assortment to each peripheral consumer must be a revenue-ordered assortment that is a subset of the optimal assortment without network effect. We also present a condition when revenue-ordered assortments can achieve a provable performance. Then in view of the fact that each node in a network can represent a group of consumers, we propose a novel idea in which the sellers are allowed to offer "randomized assortments" to each node in the network. We show that allowing for randomized assortments may further increase the revenue, and under certain conditions, the optimal assortment for the central consumer must be a combination of two adjacent revenue-ordered assortments and thus efficient algorithm can be developed. Finally, we extend the results to directed acyclic graphs (DAGs), showing that a mixture of adjacent revenue-ordered assortments is optimal under certain conditions.

Keywords: Revenue management · Consumer choice models · Network effects · Assortment optimization

Link to the full paper: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3788880.
The second author's research is partly supported by the National Science Foundation of China (NSFC) Grant 72150002.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 374, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Exploring the Tradeoff Between Competitive Ratio and Variance in Online-Matching Markets

Pan Xu(✉)

New Jersey Institute of Technology, Newark, NJ 07102, USA
pxu@njit.edu

Abstract. In this paper, we propose an online-matching-based model to study the assignment problems arising in a wide range of online-matching markets, including online recommendations, ride-hailing platforms, and crowdsourcing markets. It features that each assignment can request a random set of resources and yield a random utility, and the two (cost and utility) can be arbitrarily correlated with each other. We present two linear-programming-based parameterized policies to study the tradeoff between the competitive ratio (CR) on the total utilities and the variance on the total number of matches (unweighted version). The first one (SAMP) is to sample an edge according to the distribution extracted from the clairvoyant optimal, while the second (ATT) features a time-adaptive attenuation framework that leads to an improvement over the state-of-the-art competitive-ratio result. We also consider the problem under a large-budget assumption and show that SAMP achieves asymptotically optimal performance in terms of competitive ratio.

Here is the arXiv link to the full version: http://arxiv.org/abs/2209.07580.

Keywords: Online matching · Competitive ratio · Variance analysis

Pan Xu is partially supported by NSF CRII Award IIS-1948157.

K. A. Hansen et al. (Eds.): WINE 2022, LNCS 13778, p. 375, 2022.
https://doi.org/10.1007/978-3-031-22832-2

Author Index