Abstracts

Asymptotically Optimal Competitive Ratio for Online Allocation of Reusable Resources

Vineet Goyal[1], Garud Iyengar[1], and Rajan Udwani[2(✉)]

[1] Columbia University, New York, NY 10027, USA
{vgoyal,garud}@ieor.columbia.edu
[2] UC Berkeley, Berkeley, CA 94720, USA
rudwani@berkeley.edu

Abstract. We consider the problem of online allocation (matching, budgeted allocations, and assortments) of reusable resources where an adversarial sequence of resource requests is revealed over time and allocated resources are used/rented for a stochastic duration, drawn independently from known resource usage distributions. This problem is a fundamental generalization of well studied models in online resource allocation and assortment optimization. Previously, it was known that the greedy algorithm that simply makes the best decision for each arriving request is 0.5 competitive against clairvoyant benchmark that knows the entire sequence of requests in advance. We give a new algorithm that is $(1-1/e)$ competitive for arbitrary usage distributions and large resource capacities. This is the best possible guarantee for the problem.

Designing the optimal online policy for allocating reusable resources requires a reevaluation of the key trade off between conserving resources for future requests and being greedy. Resources that are currently in use may return "soon" but the time of return and types of future requests are both uncertain. At the heart of our algorithms is a new quantity that factors in the potential of reusability for each resource by (computationally) creating an asymmetry between identical units of the resource. We establish a performance guarantee for our algorithms by constructing a feasible solution to a novel LP *free* system of constraints. More generally, these ideas lead to a principled approach for integrating stochastic and combinatorial elements (such as reusability, customer choice, and budgeted allocations) in online resource allocation problems.

The full version of this paper is available at https://arxiv.org/pdf/2002.02430.pdf.

Keywords: Online resource allocation · Reusable resources · LP free analysis · Optimal competitive ratio

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 543, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Dynamic Bipartite Matching Market with Arrivals and Departures

Naonori Kakimura[1(✉)] and Donghao Zhu[2(✉)]

[1] Keio University, Yokohama, Japan
kakimura@math.keio.ac.jp
[2] Technical University of Munich, Munich, Germany
donghao.zhu@in.tum.de

Abstract. In this paper, we study a matching market model on a bipartite network where agents on each side arrive and depart stochastically by a Poisson process. For such a dynamic model, we design a mechanism that decides not only which agents to match, but also when to match them, to minimize the expected number of unmatched agents. The main contribution of this paper is to achieve theoretical bounds on the performance of local mechanisms with different timing properties. We show that an algorithm that waits to thicken the market, called the *Patient* algorithm, is exponentially better than the *Greedy* algorithm, i.e., an algorithm that matches agents greedily. This means that waiting has substantial benefits on maximizing a matching over a bipartite network. We remark that the Patient algorithm requires the planner to identify agents who are about to leave the market, and, under the requirement, the Patient algorithm is shown to be an optimal algorithm. We also show that, without the requirement, the Greedy algorithm is almost optimal. In addition, we consider the *1-sided algorithms* where only an agent on one side can attempt to match. This models a practical matching market such as a freight exchange market and a labor market where only agents on one side can make a decision. For this setting, we prove that the Greedy and Patient algorithms admit the same performance, that is, waiting to thicken the market is not valuable. This conclusion is in contrast to the case where agents on both sides can make a decision and the non-bipartite case by [Akbarpour et al., *Journal of Political Economy*, 2020].

Keywords: Bipartite matching · Markov chain · Online algorithm

The full version is available at https://arxiv.org/abs/2110.10824.

N. Kakimura—Supported by JSPS KAKENHI Grant Numbers JP17K00028, JP18H05291, 20H05795, and 21H03397.

D. Zhu—Supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 277991500/GRK2201.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 544, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Static Pricing for Multi-unit Prophet Inequalities (Extended Abstract)

Shuchi Chawla[1], Nikhil R. Devanur[2], and Thodoris Lykouris[3](✉)

[1] The University of Texas at Austin, Austin, USA
shuchi@cs.utexas.edu
[2] Amazon Research, Sunnyvale, USA
[3] Massachusetts Institute of Technology, Cambridge, USA
lykouris@mit.edu

The *prophet inequality* problem constitutes one of the cornerstones of online decision-making. A designer knows a set of n distributions from which random variables are sequentially realized in an arbitrary order. Once a random variable is realized, the designer decides whether to accept it or not; *at most one* realized random variable can be accepted. The objective is to maximize the value of the variable accepted, and the performance of the algorithm is evaluated against the *ex-post* maximum realized. In a beautiful result, Samuel-Cahn showed that a simple static threshold policy achieves the optimal competitive ratio for this problem. Samuel-Cahn's algorithm determines a threshold p such that the probability that there exists a realization exceeding the threshold is exactly $\frac{1}{2}$, and then accepts the first random variable that exceeds the threshold. This algorithm achieves a competitive ratio of $\frac{1}{2}$ against the ex-post optimum; no online algorithm, even one with adaptive thresholds, can obtain better performance.

Over the last few years, many extensions of the basic prophet inequality to more general feasibility constraints have been studied, and tight bounds on the competitive ratio have been established. However, one simple natural extension has largely remained open: where the designer is allowed to accept $k > 1$ random variables for some small value of k. This is called the *multi-unit* prophet inequality. When k is relatively large, then it is known that static threshold policies can achieve a competitive ratio of $1 - \Theta\left(\sqrt{\frac{\log(k)}{k}}\right)$ which goes to 1 as $k \to \infty$, and this ratio is asymptotically tight. However, (for example,) for $k = 2$ or 3, prior to our work, the best known competitive ratio of static thresholds remained $\frac{1}{2}$. Our work addresses this gap by answering the following question: *Can a static threshold policy achieve a better competitive ratio than* $\frac{1}{2}$ *for small* $k = 2, 3, \ldots$*?*

We develop an algorithm for finding a static threshold policy for the multi-unit prophet inequality that is sensitive to the supply k. Our algorithm is simple and practical. For any fixed price p, it estimates two statistics: (1) the fraction of items expected to be sold at that price, $\mu_k(p)$, and (2) the probability that not all units will sell out before all the customers have been served, $\delta_k(p)$. We then pick the static price p at which these two quantities are equal: $\mu_k(p) = \delta_k(p)$; this generalizes Samuel-Cahn's algorithm and proof via a min-max approach and shows that the worst-case competitive ratio is attained for Poisson distributions.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, pp. 545–546, 2022.
https://doi.org/10.1007/978-3-030-94676-0

The competitive ratio of our policy for $k = 2, \cdots, 5$ is $0.585, 0.630, 0.660$, and 0.682 respectively, and scales as $1 - \Omega(\sqrt{\log k/k})$ for large k.

The full version can be found here: https://arxiv.org/abs/2007.07990.

Fairness Maximization Among Offline Agents in Online-Matching Markets

Will Ma[1], Pan Xu[2(✉)], and Yifan Xu[3]

[1] Columbia University, New York, NY 10027, USA
[2] New Jersey Institute of Technology, Newark, NJ 07102, USA
pxu@njit.edu
[3] Key Lab of CNII, MOE, Southeast University, Nanjing, China
xyf@seu.edu.cn

Abstract. Matching markets involve heterogeneous agents (typically from two parties) who are paired for mutual benefit. During the last decade, matching markets have emerged and grown rapidly through the medium of the Internet. They have evolved into a new format, called *Online Matching Markets* (OMMs), with examples ranging from crowdsourcing to online recommendations to ridesharing. There are two features distinguishing OMMs from traditional matching markets. One is the dynamic arrival of one side of the market: we refer to these as *online agents* while the rest are *offline agents*. Examples of online and offline agents include keywords (online) and sponsors (offline) in Google Advertising; workers (online) and tasks (offline) in Amazon Mechanical Turk (AMT); riders (online) and drivers (offline when restricted to a short time window) in ridesharing. The second distinguishing feature of OMMs is the real-time decision-making element.

However, studies have shown that the algorithms making decisions in these OMMs leave disparities in the match rates of offline agents. For example, tasks in neighborhoods of low socioeconomic status rarely get matched to gig workers, and drivers of certain races/genders get discriminated against in matchmaking. In this paper, we propose online matching algorithms which optimize for either individual or group-level fairness among offline agents in OMMs. We present two linear-programming (LP) based sampling algorithms, which achieve online competitive ratios at least 0.725 for individual fairness maximization (IFM) and 0.719 for group fairness maximization (GFM), respectively. There are two key ideas helping us break the barrier of $1 - 1/\mathsf{e}$. One is *boosting*, which is to adaptively re-distribute all sampling probabilities only among those offline available neighbors for every arriving online agent. The other is *attenuation*, which aims to balance the matching probabilities among offline agents with different mass allocated by the benchmark LP. We conduct extensive numerical experiments and results show that our boosted version of sampling algorithms are not only conceptually easy to implement but also highly effective in practical instances of fairness-maximization-related models.

Here is the arXiv link to the full version: https://arxiv.org/abs/2109.08934.

Keywords: Fairness maximization · Online-Matching Markets

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

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 547, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Funding Public Projects: A Case for the Nash Product Rule

Florian Brandl[1], Felix Brandt[2], Matthias Greger[2], Dominik Peters[3], Christian Stricker[2], and Warut Suksompong[4](✉)

[1] Hausdorff Center for Mathematics, Universität Bonn, Bonn, Germany
florian.brandl@uni-bonn.de

[2] Institut für Informatik, Technische Universität München, Munich, Germany
{brandt,matthias.greger,christian.stricker}@tum.de

[3] Department of Computer Science, University of Toronto, Toronto, Canada
dominik@cs.toronto.edu

[4] School of Computing, National University of Singapore, Singapore, Singapore
warut@comp.nus.edu.sg

Abstract. We study a mechanism design problem where a community of agents wishes to fund public projects via voluntary monetary contributions by the community members. This serves as a model for public expenditure without an exogenously available budget, such as participatory budgeting or voluntary tax programs, as well as donor coordination when interpreting charities as public projects and donations as contributions. Our aim is to identify a mutually beneficial distribution of the individual contributions. In the preference aggregation problem that we study, agents with linear utility functions over projects report the amount of their contributions, and the mechanism determines a socially optimal distribution of the money. We identify a specific mechanism—the Nash product rule—which picks the distribution that maximizes the product of the agents' utilities weighted by their contributions. This rule arises naturally from a simple, dynamic procedure. The Nash product rule is Pareto efficient, and we prove that it satisfies attractive incentive properties: it spends each agent's contribution only on projects the agent finds acceptable, and agents are strongly incentivized to participate. We also derive impossibility theorems that show that strengthened versions of these two axioms are incompatible with Pareto efficiency.

Keywords: Public goods provision · Collective decision making · Participation incentives

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 5969/1-1, BR 2312/11-1, and BR 2312/12-1, by the European Research Council (ERC) under grant number 639945, and by an NUS Start-up Grant. The full version of the paper is available at http://arxiv.org/pdf/2005.07997.pdf.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 548, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Screening with Limited Information: The Minimax Theorem and a Geometric Approach

Zhi Chen[1], Zhenyu Hu[2], and Ruiqin Wang[2](✉)

[1] City University of Hong Kong, Kowloon Tong, Kowloon, Hong Kong SAR
zhi.chen@cityu.edu.hk

[2] National University of Singapore, 21 Lower Kent Ridge Road, 119077 Singapore, Singapore
bizhuz@nus.edu.sg, ruiqin_wang@u.nus.edu

Abstract. A seller seeks a selling mechanism to maximize the worst-case revenue obtained from a buyer whose valuation distribution lies in a certain ambiguity set. For a generic convex ambiguity set, we show via the minimax theorem that strong duality holds between the problem of finding the optimal robust mechanism and a minimax pricing problem where the adversary first chooses a worst-case distribution and then the seller decides the best posted price mechanism. This observation connects prior literature that separately studies the primal (robust mechanism design) and problems related to the dual (*e.g.*, robust pricing, buyer-optimal pricing and personalized pricing). We provide a geometric approach to analytically solving the minimax pricing problem (and the robust pricing problem) for several important ambiguity sets such as the ones with mean and various dispersion measures, and with the Wasserstein metric. The solutions are then used to construct the optimal robust mechanism and to compare with the solutions to the robust pricing problem.

Keywords: Robust mechanism design · Moment condition · Mean-preserving contraction · Wasserstein metric

The full paper can be found at http://ssrn.com/abstract=3940212.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 549, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Optimal DSIC Auctions for Correlated Private Values: Ex-Post Vs. Ex-Interim IR

Ido Feldman[1] and Ron Lavi[1,2](✉)

[1] Technion – Israel Institute of Technology, Haifa, Israel
[2] University of Bath, Bath, UK

Abstract. We study Dominant-Strategy Incentive-Compatible (DSIC) revenue-maximizing auctions ("optimal" auctions) for a single-item and correlated private values. We give tight bounds on the ratio of the revenue of the optimal Ex-Post Individually Rational (EPIR) auction and the revenue of the optimal Ex-Interim Individually Rational (EIIR) auction. This bound is expressed as a non-decreasing function of the expected social welfare of the underlying distribution. In particular, we show a class of distributions on which this ratio cannot be lower bounded by any positive number. Thus, the restriction to EPIR auctions, which has been the de-facto standard in the computer science literature on auctions with correlated values, may significantly reduce the revenue that can be possibly extracted, as the revenue extracted by an EPIR auction might be an arbitrarily small fraction of the revenue extracted by an EIIR auction.

Keywords: Optimal auctions · Correlated private values · Full surplus extraction · The look-ahead auction

This research was partially supported by the ISF-NSFC joint research program (grant No. 2560/17). We thank Ronny Lempel for stimulating discussions.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 550, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Throttling Equilibria in Auction Markets

Xi Chen, Christian Kroer, and Rachitesh Kumar(✉)

Columbia University, New York, NY 10027, USA
{xc2198,christian.kroer,rk3068}@columbia.edu

Abstract. Throttling is a popular method of budget management for online ad auctions in which the platform modulates the participation probability of an advertiser in order to smoothly spend her budget across many auctions. In this work, we investigate the setting in which all of the advertisers simultaneously employ throttling to manage their budgets, and we do so for both first-price and second-price auctions. We analyze the structural and computational properties of the resulting equilibria. For first-price auctions, we show that a unique equilibrium always exists, is well-behaved and can be computed efficiently via tâtonnement-style decentralized dynamics. In contrast, for second-price auctions, we prove that even though an equilibrium always exists, the problem of finding an equilibrium is PPAD-complete, there can be multiple equilibria, and it is NP-hard to find the revenue maximizing one. Finally, we compare the equilibrium outcomes of throttling to those of multiplicative pacing, which is the other most popular and well-studied method of budget management.

The full paper can be found at https://arxiv.org/abs/2107.10923.

Keywords: Auctions · Budget constraints · Computational advertising · Computational complexity · PPAD

X. Chen—Supported by NSF grants IIS-1838154 and CCF-1703925.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 551, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Generalized Nash Equilibrium Problems with Mixed-Integer Variables

Tobias Harks and Julian Schwarz(✉)

Augsburg University, Universitätsstraße 2, 86159 Augsburg, Germany
julian.schwarz@math.uni-augsburg.de

Abstract. We study generalized Nash equilibrium problems (GNEPs) with non-convex strategy spaces and non-convex cost functions. This general class of games includes the important case of games with mixed-integer variables for which only a few results are known in the literature. We present a new approach to characterize equilibria via a convexification technique using the Nikaido-Isoda function. To any given instance I of the GNEP, we derive a convexified instance I^{conv} and show that every feasible strategy profile for I is an equilibrium if and only if it is an equilibrium for I^{conv} and the convexified cost functions coincide with the initial ones. Based on this general result we identify important classes of GNEPs which allow us to reformulate the equilibrium problem via standard optimization problems.

1. We first define *quasi-linear* GNEPs in which for fixed strategies of the opponent players, the cost function of every player is linear and the convex hull of the respective strategy space is polyhedral. For this game class we reformulate the equilibrium problem for I^{conv} as a standard (non-linear) optimization problem.
2. We then study GNEPs with *joint constraint sets*. We introduce the new class of *projective-closed GNEPs* for which we show that I^{conv} falls into the class of jointly convex GNEPs. As an important application, we show that general GNEPs with shared binary sets $\{0, 1\}^k$ are projective-closed.
3. We demonstrate the applicability of our results by presenting a numerical study regarding the computation of equilibria for a class of quasi-linear and projective-closed GNEPs. It turns out that our characterization of a projective-closed GNEP via a jointly convex GNEP leads to an efficiently solvable reformulation of the original non-convex GNEP.

Keywords: Generalized Nash equilibrium problem · Mixed-integer variables · Nonconvex and discrete strategy space

---

This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - HA 8041/4-1.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 552, 2022.
https://doi.org/10.1007/978-3-030-94676-0

In Which Matching Markets Do Costly Compatibility Inspections Lead to a Deadlock?

Nicole Immorlica[1], Yash Kanoria[2], and Jiaqi Lu[3(✉)]

[1] Microsoft Research, Redmond, USA
[2] Columbia Business School, New York, USA
[3] The Chinese University of Hong Kong, Shenzhen, China
ujiaqi@cuhk.edu.cn

With the aim of understanding congestion in matching markets, we study a matching market with N women and $M = \alpha N$ men who want to match with each other. An agent pair must perform a costly inspection to verify compatibility prior to matching with each other, and we assume they are willing to perform the inspection only if it is "mutually desirable", i.e., they mutually rank each other as their favorite potential partner who remains under consideration. The inspection and matching process progresses iteratively in the market as matches form (in the case of successful inspections) and incompatibilities are revealed. We ask which large random markets suffer from an information deadlock, i.e., in which markets will a constant fraction of agents get stuck waiting for a mutually desirable inspection to become available. We prove, by building on the machinery of message passing and density evolution from statistical physics, that the existence of an information deadlock is governed by the men-to-women ratio α, the average degree of women (or men) and the probability that an inspection is successful. We find a phase transition between the information deadlock regime and the deadlock-free regime (where a vanishingly small fraction of agents are stuck waiting) and study the dependence of deadlock and its size on market primitives. We find, e.g., that well connected markets suffer from deadlocks, and holding the degree of women fixed there is a deadlock for α below a certain threshold.

A complete version is available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3697165.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 553, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Contest Design with Threshold Objectives

Edith Elkind, Abheek Ghosh(✉), and Paul W. Goldberg

Department of Computer Science, University of Oxford, Oxford, UK
{edith.elkind,abheek.ghosh,paul.goldberg}@cs.ox.ac.uk

Abstract. We study contests where the designer's objective is an extension of the widely studied objective of maximizing the total output: The designer gets zero marginal utility from a player's output if the output of the player is very low or very high. We consider two variants of this setting, which correspond to two objective functions: *binary threshold*, where a player's contribution to the designer's utility is 1 if her output is above a certain threshold, and 0 otherwise; and *linear threshold*, where a player's contribution is linear in her output if the output is between a lower and an upper threshold, and becomes constant below the lower and above the upper threshold. For both of these objectives, we study (1) *rank-order allocation* contests, which assign prizes based on players' rankings only, and (2) general contests, which may use the numerical values of the players' outputs to assign prizes. We characterize the contests that maximize the designer's objective and indicate techniques to efficiently compute them. We also prove that for the linear threshold objective, a contest that distributes the prize equally among a fixed number of top-ranked players offers a factor-2 approximation to the optimal rank-order allocation contest.

Keywords: Contest theory · Mechanism design · All-pay auctions

Full version is available at https://arxiv.org/abs/2109.03179.

A. Ghosh—Supported by Clarendon Fund and SKP (Pathak) Scholarship.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 554, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Confounding Equilibria for Platforms with Private Information on Promotion Value

Yonatan Gur[1], Gregory Macnamara[2], Ilan Morgenstern[1(✉)], and Daniela Saban[1]

[1] Stanford University Graduate School of Business, Stanford, CA 94305, USA
{ygur,ilanmor,dsaban}@stanford.edu

[2] Facebook Core Data Science, 1 Hacker Way, Menlo Park, CA 94025, USA

Keywords: Information design · Bayesian learning · Online marketplaces

Extended Abstract

Online marketplaces allow consumers to evaluate, compare, and purchase products while providing a channel for third-party sellers to reach a broad consumer base and increase demand for their products. As platforms seek to maintain a large consumer base, many platforms prioritize increasing consumer surplus by offering competitively priced products. At the same time, it is common practice in such marketplaces to let sellers determine their own price, but such flexibility may result in higher prices that reduce consumer surplus.

In this paper, we consider a platform facilitating trade between sellers and buyers with the objective of maximizing consumer surplus. Even though in many such marketplaces prices are set by revenue-maximizing sellers, platforms can influence prices through (i) price-dependent promotion policies that can increase demand for a product by featuring it in a prominent position on the webpage and (ii) the information revealed to sellers about the value of being promoted. Identifying effective joint information design and promotion policies is a challenging dynamic problem as sellers can sequentially learn the promotion value from sales observations and update prices accordingly. We introduce the notion of *confounding* promotion policies, which are designed to prevent a Bayesian seller from learning the promotion value (at the expense of the short-run loss of diverting consumers from the best product offering). Leveraging these policies, we characterize the maximum long-run average consumer surplus that is achievable through joint information design and promotion policies when the seller sets prices myopically. We then establish that these strategies are supported in a Bayesian Nash equilibrium, by showing that the seller's best response to the platform's optimal policy is to price myopically at every history. Moreover, the equilibrium we identify is platform-optimal within the class of horizon-maximin equilibria, in which strategies are not predicated on precise knowledge of the horizon length, and are designed to maximize payoff over the worst-case horizon. Our analysis allows one to identify practical long-run average optimal platform policies for a broad range of demand models.

M. Feldman et al. (Eds.): WINE 2021, LNCS 13112, p. 555, 2022.
https://doi.org/10.1007/978-3-030-94676-0

Author Index