

Learning Competitive Equilibria in Noisy Combinatorial Markets

Extended Abstract

Enrique Areyan Viqueira
Brown University
Providence, RI
eareyan@brown.edu

Cyrus Cousins
Brown University
Providence, RI
cyrus_cousins@brown.edu

Amy Greenwald
Brown University
Providence, RI
amy_greenwald@brown.edu

ABSTRACT

We present a methodology to robustly estimate the competitive equilibria (CE) of combinatorial markets under the assumption that buyers do not know their precise valuations for bundles of goods, but instead can only provide noisy estimates. We first show tight lower- and upper-bounds on the buyers' utility loss, and hence the set of CE, given a uniform approximation of one market by another. We then present two probably-approximately-correct algorithms for learning CE with finite-sample guarantees. The first is a baseline and the second leverages a connection between the first welfare theorem of economics and uniform approximations to adaptively prune value queries when it is determined that they are provably not part of a CE. Extensive experimentation shows that pruning achieves better estimates than the baseline with far fewer samples.

KEYWORDS

Competitive Equilibria Learning, Noisy Combinatorial Markets, PAC Algorithms for Combinatorial Markets

ACM Reference Format:

Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. 2021. Learning Competitive Equilibria in Noisy Combinatorial Markets: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*, Online, May 3–7, 2021, IFAAMAS, 3 pages.

1 OVERVIEW

Combinatorial Markets (CMs) are a class of markets in which buyers are interested in acquiring bundles of goods for which their values can be arbitrary. Real-world examples of CMs include: spectrum auctions [5] allocation of landing and take-off slots at airports [2]; internet ad placement [6]; and procurement of bus routes [4]. An outcome of a CM is an assignment of bundles to buyers together with prices for the goods. A competitive equilibrium (CE) is an outcome of particular interest in CMs and other well-studied economic models [3, 13]. In a CE, buyers are utility-maximizing (i.e., they maximize their utilities among all feasible allocations at the posted prices) and the seller maximizes its revenue (again, over all allocations at the posted prices).

One of the defining features of CMs is that they afford buyers the flexibility to express complex preferences, which in turn has the potential to increase market efficiency. However, the extensive

expressivity of these markets presents challenges for both the market maker and the buyers. With an exponential number of bundles in general, it is infeasible for a buyer to evaluate them all. We thus present a model of noisy buyer valuations: e.g., buyers might use approximate or heuristic methods to obtain value estimates [7]. In turn, the market maker chooses an outcome in the face of uncertainty about the buyers' valuations. We call the objects of study in this work *noisy combinatorial markets* (NCM) to emphasize that buyers do not have direct access to their values for bundles, but instead can only noisily estimate them.¹

In this work,² we formulate a mathematical model of NCMs. Our goal is then to design learning algorithms with rigorous finite-sample guarantees that approximate the competitive equilibria of NCMs. First, we present tight lower- and upper-bounds on the set of CE, given uniform approximations of buyers' valuations. We then present two learning algorithms. The first one—Elicitation Algorithm; EA—serves as a baseline. It uses Hoeffding's inequality [9] to produce said uniform approximations. Our second algorithm—Elicitation Algorithm with Pruning; EAP—leverages the first welfare theorem of economics to adaptively prune value queries when it determines that they are provably not part of a CE.

An interesting tradeoff arises between computational and sample efficiency. To prune a value query and retain rigorous guarantees on the quality of the learned CE, we must solve a welfare-maximizing problem whose complexity grows with the size of the market. Consequently, at each iteration of EAP, for each value query, we face a choice. Either solve said welfare-maximizing problem and potentially prune the value query (thereby saving on future samples), or defer attempts to prune the value query, until more is known about the market. To combat this problem, we show that an upper bound on the optimal welfare's value (rather than the precise value) suffices to obtain rigorous guarantees on the learned CE's quality. Such upper bounds can be found easily, by solving a relaxation of the welfare-maximization problem. Reminiscent of designing admissible heuristics in classical search problems, this methodology applies to any combinatorial market, but at the same time allows for the application of domain-dependent knowledge to compute these upper bounds, when available. Empirically, we show that solving a computationally cheap relaxation of the welfare-maximization problem instead of the exact welfare-maximization problem yields substantial sample and computational savings in a large market.

¹Our setup generalizes the standard model of combinatorial markets by allowing buyers' values to be drawn from (possibly unknown) probability distributions. In the standard model, these distributions would be known and degenerate.

²An earlier version of this paper was presented at the 2nd Games, Agents, and Incentives Workshop [1]. An extended version of this paper can be found on arXiv [12].

2 THEORETICAL CONTRIBUTIONS

In this extended abstract, we introduce the minimum mathematical notation needed to convey our work’s main ideas. For a detailed mathematical treatment, please see the arXiv version [12].

A *combinatorial market* M is defined by a set of goods $G = [m]$ and a set of buyers $N = [n]$. A *bundle* of goods is a set of goods $S \subseteq G$. Each buyer i is characterized by their preferences over bundles, represented as a valuation function $v_i : 2^G \rightarrow \mathbb{R}_+$, where $v_i(S) \in \mathbb{R}_+$ is buyer i ’s value for bundle S .

Given a market M , an *allocation* $\mathcal{S} = (S_1, \dots, S_n)$ denotes an assignment of goods to buyers, where $S_i \subseteq G$ is the bundle assigned to buyer i . A *pricing profile* $\mathcal{P} = (P_1, \dots, P_n)$ is a vector of n pricing functions, one function $P_i : 2^G \rightarrow \mathbb{R}_+$ per buyer, each mapping bundles to prices, $P_i(S) \in \mathbb{R}_+$. A pair $(\mathcal{S}, \mathcal{P})$ is an outcome. We study approximations of CE, a fundamental outcome class.

Definition 2.1 (Approximate Competitive Equilibria). Let $\varepsilon > 0$. An outcome $(\mathcal{S}, \mathcal{P})$ is a ε -competitive equilibrium (ε -CE) if:

- (RM) for any allocation $\mathcal{S}' : \sum_{i \in N} P_i(S_i) \geq \sum_{i \in N} P_i(S'_i)$
- (ε -UM) $\forall i \in N, T \subseteq G : v_i(S_i) - P_i(S_i) + \varepsilon \geq v_i(T) - P_i(T)$

Revenue-maximization (RM) ensures the seller maximizes its revenue over all possible allocations under the outcome’s pricing. ε -Utility maximization (ε -UM) ensures that buyers maximize (almost, up to additive error ε) their utility (value minus prices) for their allocation under the outcome’s pricing. For $\alpha \geq 0$, we denote by $\mathcal{CE}_\alpha(M)$ the set of all α -approximate CE of M .

M' is called an ε -uniform approximation of M if $\|M - M'\|_\infty = \max_{(i,S) \in N \times G} |v_i(S) - v'_i(S)| \leq \varepsilon$. Our main theoretical contribution is to show that an ε -uniform approximation M' of a market M preserves the set of CE, up to an additive error:

THEOREM 2.2 (COMPETITIVE EQUILIBRIUM APPROXIMATION).

Let $\varepsilon > 0$. If M and M' are markets such that $\|M - M'\|_\infty \leq \varepsilon$, then $\mathcal{CE}(M) \subseteq \mathcal{CE}_{2\varepsilon}(M') \subseteq \mathcal{CE}_{4\varepsilon}(M)$.

In Theorem 2.2, consider M as an unobservable market of interest and M' as an approximation of M learned by constructing confidence intervals around buyers’ values. Then, Theorem 2.2 establishes that the unobservable set of CE of M are well preserved by the observable set of CE of M' . Moreover, the approximation error can be as low as desired provided ε is a user-controlled parameter (e.g., using Hoeffding’s inequality [9]).

Next, we summarize our model of noisy combinatorial markets, our elicitation and pruning algorithm, and experimental results.

3 ELICITATION & PRUNING ALGORITHM

In a noisy combinatorial market, we have no access to a buyer i ’s valuation function v_i . Instead, we pose the existence of an abstract set of conditions $\mathcal{X}$, whose elements $x \in \mathcal{X}$ define a conditional valuation function $v_i(S, x)$. Set $\mathcal{X}$ models those unobservable factors that influence i ’s value for S . Then, given a distribution $\mathcal{D}$ over $\mathcal{X}$, we define the expected combinatorial market as the market where i ’s values are given by $v_i(S, \mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}}[v_i(S, x)]$. Our goal is to efficiently learn $v_i(S, \mathcal{D})$. Intuitively, we deem a learning procedure to be efficient if samples are concentrated only on those pairs (i, S) that a CE could comprise.

Algorithm 1 is an informal version of our elicitation with pruning (EAP). The algorithm takes an oracle $\mathcal{O}$ that models the elicitation

Algorithm 1 Elicitation Algo. with Pruning (EAP) (Informal)

Input: Oracle $\mathcal{O}$ that takes $\delta, t, \{(i, S)\} \subseteq N \times 2^G$, and outputs for each (i, S) , a $1 - \delta$ -confidence interval around $v_i(S)$ from t samples. A sampling schedule t , a failure probability schedule δ , a pruning budget schedule π , and a target approximate error ε .

Output: Value estimates $\hat{v}_i(S)$, for all (i, S) , failure probability $\hat{\delta}$, and CE error $\hat{\varepsilon}$.

```

1:  $\mathcal{I} \leftarrow N \times 2^G$       {Initially, all buyers-bundle pairs are active}
2: for  $k \in 1, \dots, |t|$  do
3:    $(\{\hat{v}_i\}_{(i,S) \in \mathcal{I}}, \hat{\varepsilon}) \leftarrow \mathcal{O}(\mathcal{I}, t_k, \delta_k)$  {Call oracle on active pairs.}
4:   if  $\hat{\varepsilon} \leq \varepsilon$  or  $k = |t|$  or  $\mathcal{I} = \emptyset$ , then
     return  $(\{\hat{v}_i\}_{i \in N}, \sum_{l=1}^k \delta_l, \hat{\varepsilon})$ 
5:    $\mathcal{I}_{\text{PRUNE}} \leftarrow \emptyset$ 
6:    $\mathcal{I}_{\text{CANDIDATES}} \leftarrow$  a subset of  $\mathcal{I}$  of size at most  $\pi_k$ 
7:   for  $(i, S) \in \mathcal{I}_{\text{CANDIDATES}}$  do
8:     if  $(i, S)$  is prunable, then  $\mathcal{I}_{\text{PRUNE}} \leftarrow \mathcal{I}_{\text{PRUNE}} \cup (i, S)$ 
9:   end for
10:   $\mathcal{I} \leftarrow \mathcal{I} \setminus \mathcal{I}_{\text{PRUNE}}$ 
11: end for
```

of buyers’ values via noisy value queries (i.e., buyer–bundle pairs), where the market maker eliciting i ’s value for S receives not full information but only confidence intervals around $v_i(S, \mathcal{D})$. The algorithm progressively elicits buyers’ valuations via repeated calls to $\mathcal{O}$. Additionally, between calls to $\mathcal{O}$, EAP searches for value queries that provably no CE can comprise. All such queries then cease to be part of the index set with which $\mathcal{O}$ is called in future iterations. A technical contribution was to develop a connection between ε -uniform approximations and the first welfare theorem of economics that allows us to prune value queries and provably preserve the guarantees of Theorem 2.2.

4 EXPERIMENTS

After establishing the correctness of EAP, we evaluated its empirical performance using both synthetic unit-demand valuations and two spectrum auction value models: the Global Synergy Value Model (GSVM) [8] and the Local Synergy Value Model (LSVM) [11]. Unit-demand valuations are a class of valuations central to the literature on economics and computation [10] where buyers have no complements in their valuation of goods, while GSVM and LSVM model situations in which buyers’ valuations encode complements, a considerably more challenging situation.

In all three models, we measure the average quality of learned CE via our algorithms, compared to the CE of the corresponding certain market (i.e., here, “certain” means lacking uncertainty), as a function of the number of samples. We find that EAP consistently yields better error guarantees than a baseline using far fewer samples, because it successfully prunes buyers’ valuations (i.e., it ceases querying for buyers’ values on bundles of goods that a CE provably does not comprise), even without any *a priori* knowledge of the market’s combinatorial structure.

ACKNOWLEDGMENTS

This work was supported in part by NSF Award CMMI-1761546 and in part by DARPA grant FA8750.

REFERENCES

- [1] Enrique Areyan Viqueira and Amy Greenwald. 2020. Learning Competitive Equilibria in Noisy Combinatorial Markets. In *Proceedings of the 2nd Games, Agents, and Incentives Workshop (GAIW@ AAMAS 2020)*.
- [2] MO Ball, G Donohue, and K Hoffman. 2006. Auctions for the safe, efficient and equitable allocation of airspace system resources. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions*.
- [3] Sushil Bikhchandani and John W Mamer. 1997. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory* 74, 2 (1997), 385–413.
- [4] Estelle Cantillon and Martin Pesendorfer. 2006. Auctioning bus routes: The London experience. . (2006).
- [5] Peter Cramton et al. 2002. Spectrum auctions. *Handbook of telecommunications economics* 1 (2002), 605–639.
- [6] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.
- [7] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. 1999. Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999. 2 Volumes, 1450 pages*, Thomas Dean (Ed.). Morgan Kaufmann, 548–553. <http://ijcai.org/Proceedings/99-1/Papers/079.pdf>
- [8] Jacob K Goeree and Charles A Holt. 2010. Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior* 70, 1 (2010), 146–169.
- [9] Wassily Hoeffding. 1994. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*. Springer, 409–426.
- [10] Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.
- [11] Tobias Scheffel, Georg Ziegler, and Martin Bichler. 2012. On the impact of package selection in combinatorial auctions: an experimental study in the context of spectrum auction design. *Experimental Economics* 15, 4 (2012), 667–692.
- [12] Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. 2021. Learning Competitive Equilibria in Noisy Combinatorial Markets. arXiv:2101.09551 [cs.GT]
- [13] Léon Walras. 2003. *Elements of Pure Economics: Or the Theory of Social Wealth*. Routledge. <https://books.google.com/books?id=hwjRD3z0Qy4C>