Learning Competitive Equilibria in Noisy Combinatorial Markets

Enrique Areyan Viqueira
Brown University
eareyan@brown.edu

Amy Greenwald
Brown University
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 exactly know their valuations for bundle of goods but can instead only provide noisy samples of them. We first show tight lower- and upper-bounds on the set of CE, given a uniform approximation of one market by another. We then develop a learning framework for our setup, and present two probably-approximately-correct algorithms capable of learning CE with finite-sample guarantees. The first is a baseline that uses Hoeffding's inequality to produce a uniform approximation of buyers' valuations with high probability. The second leverages a novel connection between the first welfare theorem of economics and uniform approximations to adaptively prune value queries when it determines that they are provably not part of a CE. We experiment with our algorithms and show that the pruning algorithm is capable of achieving better estimates than the baseline with far fewer samples.

KEYWORDS

Noisy combinatorial markets; PAC algorithms; First welfare theorem of economics

1 INTRODUCTION

Combinatorial Markets (CMs) are a class of markets in which buyers are interested in acquiring bundles of goods. Real-world examples of these markets include: spectrum auctions, of which the 2014 Canadian 700 MHZ raised upwards of $5 billion [21]; allocation of landing and take-off slots at airports [3]; Internet ad placement [7]; and procurement of bus routes [5].

An outcome of a CM is an assignment of bundles to buyers together with prices for the goods, dictated by the market maker. A competitive equilibrium (CE) is an outcome of particular interest in CMs, and in other well-studied economic models [4, 20]. In a CE, buyers are utility-maximizing (i.e., they maximize their utilities at the posted prices) and the market clears, meaning excess supply is priced at zero.

One of the defining features of CMs is that they afford buyers the flexibility to express complex preferences over a wide variety of outcomes, 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 assume a model of noisy buyer valuations: e.g., buyers might use approximate or heuristic methods to obtain value estimates [8]. In turn, the market maker has to choose an outcome, despite any uncertainty about the buyers' valuations. We call these markets *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.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 2020, Auckland, New Zealand

https://doi.org/doi

In this work, 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.

Our first result is to show 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 [10] to produce said uniform approximations. Our second algorithm—Elicitation Algorithm with Pruning; EAP—leverages a novel connection between the first welfare theorem of economics and uniform approximations, which enables it to adaptively prune value queries when it determines that they are provably not part of a CE.

After establishing the correctness of our algorithms, we evaluate their empirical performance. Although our methodology is general enough to apply to any CM, it is well-known that gross substitutes (GS) is the largest class of valuation functions for which CE are guaranteed to exist [9, 11]. Since the goal of our methodology is to learn CE, we focus on GS valuations in our experiments. In particular, we test our algorithms using synthetic unit-demand valuations, a class of valuations central to the literature on economics and computation [14]. We measure the quality of learned CE in noisy unit-demand markets as compared to the CE of the underlying market. We find that EAP is capable of exploiting the combinatorial structure of unit-demand markets, to prune buyers' valuations for goods when said goods would never be allocated in a CE, even without any *a priori* knowledge of the market's structure. As a result, EAP often yields better error guarantees than EA using far fewer samples.

Related Work. The idea for this paper stemmed from the work on abstraction in Fisher markets by [12]. There, the authors tackle the problem of computing equilibria in large markets by creating an abstraction of the market, computing equilibria in the abstraction, and lifting those equilibria back to the original market. Likewise, we develop a pruning criterion which in effect builds an abstraction of any CM, in which we can then compute the CE, which are provably also approximate CE in the original market.

The mathematical formalism we adopt follows that of [17]. There, the authors propose a mathematical framework for empirical game-theoretic analysis [22], and algorithms that learn the Nash equilibria of so-called simulation-based games [18, 19]. In this paper, we extend this methodology to market equilibria, and provide analogous results in the case of CMs. Whereas the basic pruning criterion in games is straightforward—simply prune dominated strategies—the challenge in this work was to discover a pruning criterion that would likewise prune valuations that are provably not part of an equilibrium.

Another related line of research is concerned with learning valuation functions from data [1, 2, 13]. In contrast, our work emphasizes learning CE, rather than buyers' valuations. Indeed, our main conclusion is that CE often can be learned from just a subset of buyers' valuations. In this sense, our work complements other research on learning valuations, especially when the primary reason for doing so is to learn CE.

There is also a long line of work on preference elicitation in combinatorial auctions (e.g., [6]), where the goal is for an auctioneer to pose value queries in an intelligent order so as to minimize the computational burden on the bidders, while still clearing the auction.

Finally, our pruning criterion relies on a novel application of the first welfare theorem of economics. While prior work [15] has connected economic theory with algorithmic complexity, this work connects economic theory with learning theory.

2 MODEL

We write $\mathbb{X}_+$ to denote the set of positive values in a numerical set $\mathbb{X}$ including zero. Given an integer $k \in \mathbb{Z}$, we write $[k]$ to denote the first k integers, inclusive: i.e., $[k] = \{1, 2, \ldots, k\}$. Given a finite set of integers $Z \subset \mathbb{Z}$, we write 2^Z to denote the power set of Z.

A *combinatorial market* is defined by a set of goods and a set of buyers. We denote the set of goods by $G = [m]$, and the set of buyers by $N = [n]$. We index an arbitrary good by $j \in G$, and an arbitrary buyer by $i \in N$. 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 $v_i : 2^G \mapsto \mathbb{R}_+$, where $v_i(S) \in \mathbb{R}_+$ is buyer i's value for bundle S. We assume valuations are normalized so that $v_i(\emptyset) = 0$, for all $i \in N$. Using this notation, a combinatorial market—market, hereafter—is a tuple $M = (G, N, \{v_i\}_{i \in N})$.

Given a market M, an *allocation* $\mathcal{S} = (S_1, \ldots, S_n)$ denotes an assignment of goods to buyers, where $S_i \subseteq G$ is the bundle of goods assigned to buyer i. We consider only feasible allocations. An allocation $\mathcal{S}$ is *feasible* if $S_i \cap S_k = \emptyset$ for all $i, k \in N$ such that $i \neq k$. We denote the set of all feasible allocations of market M by $\mathcal{F}(M)$. The *welfare* of allocation $\mathcal{S}$ is defined as $w(\mathcal{S}) = \sum_{i \in N} v_i(S_i)$. A welfare-maximizing allocation $\mathcal{S}^*$ is a feasible allocation that yields maximum welfare among all feasible allocations, i.e., $\mathcal{S}^* \in \arg\max_{\mathcal{S} \in \mathcal{F}(M)} w(M)$. We denote by $w^*(M)$ the welfare of any welfare-maximizing allocation $\mathcal{S}^*$, i.e., $w^*(M) = w(\mathcal{S}^*) = \sum_{i \in N} v_i(S_i^*)$.

In this paper, we consider only linear and anonymous prices. Thus, a pricing $\mathbf{p} = (p_1, \ldots, p_m)$ is an assignment of prices to goods where $p_j \in \mathbb{R}_+$ is the price assigned to good j. The total price of bundle S is then given by $P(S) = \sum_{j \in S} p_j$. We refer to pair $(\mathcal{S}, \mathbf{p})$ as a *market outcome*—outcome, for short.

In this paper, we are interested in approximations of one market by another. We now define a mathematical framework in which to formalize such approximations. In what follows, whenever we decorate a market M, e.g., M', what we mean is that we decorate each of its components: i.e., $M' = (G', N', \{v'_i\}_{i \in N'})$.

It will be convenient to refer to a subset of buyer–bundle pairs. We use the notation $\mathcal{I} \subseteq N \times 2^G$ for this purpose.

Markets M and M' are *compatible* if $G = G'$ and $N = N'$. Whenever a market M is compatible with a market M', an outcome of M is also an outcome of M'. Given two compatible markets M and M', we measure the difference between them at $\mathcal{I}$ as $\|M - M'\|_{\mathcal{I}} = \max_{(i,S) \in \mathcal{I}} |v_i(S) - v'_i(S)|$. When $\mathcal{I} = N \times 2^G$, this difference is precisely the infinity norm. Given $\varepsilon > 0$, M and M' are called *ε-approximations* of one another if $\|M - M'\|_\infty \leq \varepsilon$.

Definition 2.1 (Competitive Equilibrium). Given a market M, an outcome $(\mathcal{S}, \mathbf{p})$ is a *competitive equilibrium* (CE) if:

(UM) $\forall i \in N, T \subseteq G : v_i(S_i) - P(S_i) \geq v_i(T) - P(T)$
(MC) If $j \notin \cup_{i \in N} S_i$ then $p_j = 0$

Definition 2.2 (Approximate Competitive Equilibria). Let $\varepsilon > 0$. An outcome $(\mathcal{S}, \mathbf{p})$ is a ε-competitive equilibrium (ε-CE) if it is a CE in which UM holds up to ε:

$$(\varepsilon\text{-UM}) \forall i \in N, T \subseteq G : v_i(S_i) - P(S_i) + \varepsilon \geq v_i(T) - P(T)$$

For $\alpha \geq 0$, we denote by $\mathcal{CE}_\alpha(M)$ the set of all α-approximate CE of M, i.e., $\mathcal{CE}_\alpha(M) = \{(\mathcal{S}, \mathbf{p}) : (\mathcal{S}, \mathbf{p}) \text{ is a } \alpha\text{-approximate CE of } M\}$. Note that $\mathcal{CE}_0(M)$ is the set of (exact) CE of market M, which we denote $\mathcal{CE}(M)$.

Theorem 2.3 (Competitive Equilibrium Approximation). *Let $\varepsilon > 0$. If M and M' are compatible markets such that $\|M - M'\|_\infty \leq \varepsilon$, then $\mathcal{CE}(M) \subseteq \mathcal{CE}_{2\varepsilon}(M') \subseteq \mathcal{CE}_{4\varepsilon}(M)$.*

Proof. We prove the following: $\mathcal{CE}_\alpha(M) \subseteq \mathcal{CE}_{\alpha+2\varepsilon}(M')$, for $\alpha \geq 0$. This result then implies $\mathcal{CE}(M) \subseteq \mathcal{CE}_{2\varepsilon}(M')$ when $\alpha = 0$; likewise, it (symmetrically) implies $\mathcal{CE}_{2\varepsilon}(M') \subseteq \mathcal{CE}_{4\varepsilon}(M)$ when $\alpha = 2\varepsilon$.

Let M and M' be compatible markets s.t. $\|M - M'\|_\infty \leq \varepsilon$. Suppose $(\mathcal{S}, \mathbf{p})$ is a α-competitive equilibrium of M. Our task is to show that $(\mathcal{S}, \mathbf{p})$, interpreted as an outcome of M', is a $(\alpha+2\varepsilon)$-competitive equilibrium of M'.

First, note that the MC condition is immediately satisfied, because $\mathcal{S}$ and $\mathbf{p}$ do not change when interpreting $(\mathcal{S}, \mathbf{p})$ as an outcome of M'. Thus, we need only show that the approximation holds for the UM condition:

$$\begin{aligned} v'_i(S_i) - P(S_i) &\geq v_i(S_i) - P(S_i) - \varepsilon, & \forall i, S_i \quad (1) \\ &\geq v_i(T) - P(T) - \alpha - \varepsilon, & \forall T \subseteq G \quad (2) \\ &\geq v'_i(T) - P(T) - \alpha - 2\varepsilon, & \forall T \subseteq G \quad (3) \end{aligned}$$

where (1) and (3) follow because $\|M - M'\|_\infty \leq \varepsilon$, and (2) follows because $(\mathcal{S}, \mathbf{p})$ is a α-approximate CE of M. □

3 LEARNING METHODOLOGY

We now present a formalism in which to model noisy combinatorial markets. Intuitively, a noisy market is one in which buyers' valuations over bundles are not known precisely; rather, only noisy samples are available.

Definition 3.1 (Conditional Combinatorial Markets). A *conditional combinatorial market* $M_{\mathcal{X}} = (\mathcal{X}, G, N, \{v_i\}_{i \in N})$ consists of a set of conditions $\mathcal{X}$, a set of goods G, a set of buyers N, and a set of conditional valuation functions $\{v_i\}_{i \in N}$, where $v_i : 2^G \times \mathcal{X} \mapsto \mathbb{R}_+$. Given a condition $x \in \mathcal{X}$, the value $v_i(S, x)$ is i's value for bundle $S \subseteq G$.

Definition 3.2 (Expected Combinatorial Market). Let $M_{\mathcal{X}} = (\mathcal{X}, G, N, \{v_i\}_{i \in N})$ be a conditional combinatorial market and let $\mathcal{D}$ be a distribution over $\mathcal{X}$. For all $i \in N$, define the expected valuation function $v_i : 2^G \mapsto \mathbb{R}_+$ by $v_i(S) = \mathbf{E}_{x \sim \mathcal{D}}[v_i(S, x)]$, and the corresponding expected combinatorial market as $M_{\mathcal{D}} = (G, N, \{v_i\}_{i \in N})$.

The goal of this work is to design algorithms that learn the approximate CE of expected combinatorial markets. We will learn their equilibria given access only to their empirical counterparts, which we define next.

Definition 3.3 (Empirical Combinatorial Market). Let $M_{\mathcal{X}} = (\mathcal{X}, G, N, \{v_i\}_{i \in N})$ be a conditional combinatorial market and let $\mathcal{D}$ be a distribution over $\mathcal{X}$. Denote by $\boldsymbol{x} = (x_1, \ldots, x_t) \sim \mathcal{D}$ a vector of t samples drawn from $\mathcal{X}$ according to distribution $\mathcal{D}$. For all $i \in N$, we define the empirical valuation function $\hat{v}_i : 2^G \mapsto \mathbb{R}_+$ by $\hat{v}_i(S) = \frac{1}{t}\sum_{l=1}^{t} v_i(S, x_l)$, and the corresponding empirical combinatorial market as $\hat{M}_{\boldsymbol{x}} = (G, N, \{\hat{v}_i\}_{i \in N})$.

Observation 1 (Learnability). *Let $M_{\mathcal{X}}$ be a conditional combinatorial market and let $\mathcal{D}$ be a distribution over $\mathcal{X}$. Let $M_{\mathcal{D}}$ and $\hat{M}_{\boldsymbol{x}}$ be the corresponding expected and empirical combinatorial markets. If, for some $\varepsilon, \delta > 0$, it holds that $\mathbb{P}\left(\|M_{\mathcal{D}} - \hat{M}_{\boldsymbol{x}}\| \leq \varepsilon\right) \geq 1 - \delta$, then the competitive equilibria of $M_{\mathcal{D}}$ are learnable: i.e, any competitive equilibrium of $M_{\mathcal{D}}$ is a 2ε-competitive equilibrium of $\hat{M}_{\boldsymbol{x}}$ with probability at least $1 - \delta$.*

Lemma 3.4 (Finite-Sample Bounds for Expected Combinatorial Markets via Hoeffding's Inequality). *Let $M_{\mathcal{X}}$ be a conditional combinatorial market, $\mathcal{D}$ a distribution over $\mathcal{X}$, and $\mathcal{I} \subseteq N \times 2^G$ an index set. Suppose that for all $x \in \mathcal{X}$ and $(i, S) \in \mathcal{I}$, it holds that $v_i(S, x) \in [0, c]$ where $c \in \mathbb{R}_+$. Then, with probability at least $1 - \delta$ over samples $\boldsymbol{x} = (x_1, \ldots, x_t) \sim \mathcal{D}$, it holds that $\|M_{\mathcal{D}} - \hat{M}_{\boldsymbol{x}}\|_{\mathcal{I}} \leq c\sqrt{\ln(2|\mathcal{I}|/\delta)/2t}$.*

Hoeffding's inequality is just one possible choice of concentration inequality. While it requires bounded noise, boundedness is not a limitation of our methodology. We could instead assume (unbounded) subgaussian or subexponential noise, and substitute the appropriate Chernoff bounds.

3.1 Baseline Algorithm

EA (Algorithm 1) is a preference elicitation algorithm for combinatorial markets. The algorithm places value queries, but is only assumed to elicit noisy values for bundles. The following guarantee follows immediately from Lemma 3.4.

Theorem 3.5 (Elicitation Algorithm Guarantees). *Let $M_{\mathcal{X}}$ be a conditional market, $\mathcal{D}$ be a distribution over $\mathcal{X}$, $\mathcal{I}$ an index set, $t \in \mathbb{N}_{>0}$ a number of samples, $\delta > 0$, and $c \in \mathbb{R}_+$. Suppose that for all $x \in \mathcal{X}$ and $(i, S) \in \mathcal{I}$, it holds that $v_i(S, x) \in [0, c]$. If EA outputs $(\{\hat{v}_i\}_{(i,S) \in \mathcal{I}}, \hat{\varepsilon})$ on input $(M_{\mathcal{X}}, \mathcal{D}, \mathcal{I}, t, \delta, c)$, then, with probability at least $1 - \delta$, it holds that $\|M_{\mathcal{D}} - \hat{M}_{\boldsymbol{x}}\|_{\mathcal{I}} \leq c\sqrt{\ln(2|\mathcal{I}|/\delta)/2t}$.*

Algorithm 1 Elicitation Algorithm (EA)

Input: $M_{\mathcal{X}}, \mathcal{D}, \mathcal{I}, t, \delta, c$
A conditional combinatorial market $M_{\mathcal{X}}$, a distribution $\mathcal{D}$ over $\mathcal{X}$, an index set $\mathcal{I}$, sample size t, failure probability δ, and valuation range c.
Output: Valuation estimates $\hat{v}_i(S)$, for all $(i, S) \in \mathcal{I}$, and an approximation error $\hat{\varepsilon}$.

$(x_1, \ldots, x_t) \sim \mathcal{D}$ {Draw t samples from $\mathcal{D}$}
for $(i, S) \in \mathcal{I}$ **do**
  $\hat{v}_i(S) \leftarrow \frac{1}{t}\sum_{l=1}^{t} v_i(S, x_l)$
end for
$\hat{\varepsilon} \leftarrow c\sqrt{\ln(2|\mathcal{I}|/\delta)/2t}$ {Compute error}
return $(\{\hat{v}_i\}_{(i,S) \in \mathcal{I}}, \hat{\varepsilon})$

Algorithm 2 Elicitation Algorithm with Pruning (EAP)

Input: $M_{\mathcal{X}}, \mathcal{D}, \boldsymbol{t}, \boldsymbol{\delta}, c, \varepsilon$
A conditional combinatorial market $M_{\mathcal{X}}$, a distribution $\mathcal{D}$ over $\mathcal{X}$, a sampling schedule $\boldsymbol{t}$, a failure probability schedule $\boldsymbol{\delta}$, a valuation range c, and a target approximation error ε.
Output: Valuation estimates $\hat{v}_i(S)$, for all (i, S), approximation errors $\hat{\varepsilon}_{i,S}$, failure probability $\hat{\delta}$, and CE error $\hat{\varepsilon}$.

$\mathcal{I} \leftarrow N \times 2^G$ {Initialize index set}
$(\hat{v}_i(S), \hat{\varepsilon}_{i,S}) \leftarrow (0, c/2), \forall (i, S) \in \mathcal{I}$ {Initialize outputs}
for $k \in 1, \ldots, |\boldsymbol{t}|$ **do**
  $(\{\hat{v}_i\}_{(i,S) \in \mathcal{I}}, \hat{\varepsilon}) \leftarrow \text{EA}(M_{\mathcal{X}}, \mathcal{D}, \mathcal{I}, t_k, \delta_k, c)$
  $\hat{\varepsilon}_{i,S} \leftarrow \hat{\varepsilon}, \forall (i, S) \in \mathcal{I}$ {Update error rates}
  if $\hat{\varepsilon} \leq \varepsilon$ or $k = |\boldsymbol{t}|$ or $\mathcal{I} = \emptyset$ **then**
    return $(\{\hat{v}_i\}_{i \in N}, \{\hat{\varepsilon}_{i,S}\}_{(i,S) \in N \times 2^G}, \sum_{l=1}^{k} \boldsymbol{\delta}_l, \hat{\varepsilon})$
  end if
  Let $\hat{M}$ be the market with valuations $\{\hat{v}_i\}_{(i,S) \in \mathcal{I}}$
  $\mathcal{I}_{\text{PRUNE}} = \emptyset$ {Initialize set of indices to prune}
  for $(i, S) \in \mathcal{I}$ **do**
    Let $\hat{M}_{-(i,S)}$ be the (i, S)-submarket of $\hat{M}$
    if $\hat{v}_i(S) + w^*(\hat{M}_{-(i,S)}) + 2\hat{\varepsilon}n < w^*(\hat{M})$ **then**
      $\mathcal{I}_{\text{PRUNE}} \leftarrow \mathcal{I}_{\text{PRUNE}} \cup (i, S)$
    end if
  end for
  $\mathcal{I} \leftarrow \mathcal{I} \setminus \mathcal{I}_{\text{PRUNE}}$
end for

3.2 Pruning Algorithm

EA elicits buyers' valuations for all bundles, but in certain situations, some buyer valuations are not relevant for computing a CE—although bounds on all of them are necessary to guarantee strong bounds on the set of CE (Theorem 2.3). For example, in a first-price auction for one good, it is enough to accurately learn the highest bid, but is not necessary to accurately learn all other bids, if it is known that they are lower than the highest. Since our goal is to learn CE, we present EAP (Algorithm 2), an algorithm that does not sample uniformly, but instead adaptively decides which value queries to prune so that, with provable guarantees, EAP's estimated market satisfies the conditions of Theorem 2.3.

EAP takes as input a sampling schedule $\boldsymbol{t}$, which is a sequence of strictly increasing integers, and a failure probability schedule

$\boldsymbol{\delta}$, where $\delta_k \in (0,1)$ and $\sum_k \delta_k \in (0,1)$. The algorithm progressively elicits buyers' valuations via repeated calls to EA. However, between calls to EA, EAP looks for value queries that are provably not part of a CE. All such queries (i.e., buyer–bundle pairs) then cease to be part of the index set with which EA is called in future iterations.

In what follows, we prove several intermediate results, which enable us to prove the main result of this section, Theorem 3.10, which establishes EAP's correctness. Specifically, the market learned by EAP—with potentially different numbers of samples for different (i,S) pairs—is enough to provably recover any CE of the underlying market.

Lemma 3.6 (Optimal Welfare Approximations). *Let M and M' be compatible markets such that they ε-approximate one another. Then $|w^*(M) - w^*(M')| \leq \varepsilon n$.*

Proof. Let $\mathcal{S}^*$ be a welfare-maximizing allocation for M and $\mathcal{U}^*$ be a welfare-maximizing allocation for M'. Let $w^*(M)$ be the maximum achievable welfare in market M. Then,

$$w^*(M) = \sum_{i\in N} v_i(\mathcal{S}_i^*) \geq \sum_{i\in N} v_i(\mathcal{U}_i^*) \geq \sum_{i\in N} v_i'(\mathcal{U}_i^*) - \varepsilon n = w^*(M') - \varepsilon n$$

The first inequality follows from the optimality of $\mathcal{S}^*$ in M, and the second from the ε-approximation assumption. Likewise, $w^*(M') \geq w^*(M) - \varepsilon n$, so the result holds. □

The key to this work was the discovery of a pruning criterion that removes (i,S) pairs from consideration if they are provably not part of any CE. Our check relies on computing the welfare of the market without the pair: i.e., in submarkets.

Definition 3.7. Given a market M and buyer–bundle pair (i,S), the (i,S)-*submarket* of M, denoted by $M_{-(i,S)}$, is the market obtained by removing all goods in S and buyer i from market M. That is, $M_{-(i,S)} = (G \setminus S, N \setminus \{i\}, \{v_k\}_{k\in N\setminus\{i\}})$.

Lemma 3.8 (Pruning Criteria). *Let M and M' be compatible markets such that $\|M - M'\|_\infty \leq \varepsilon$, (i,S) an arbitrary buyer, bundle pair, and $M'_{-(i,S)}$ the (i,S)-submarket of M'. If the following pruning criterion holds, then S is not allocated to i in any welfare-maximizing allocation of M:*

$$v_i'(S) + w^*(M'_{-(i,S)}) + 2\varepsilon n < w^*(M') \quad . \tag{4}$$

Proof. Let $\mathcal{S}^*, \mathcal{U}^*$, and $\mathcal{U}^*_{-(i,S)}$ be welfare-maximizing allocations of markets M, M', and $M'_{-(i,S)}$, respectively. Then,

$$\begin{aligned} w^*(M) &\geq w^*(M') - \varepsilon n & (5)\\ &> v_i'(S) + w^*(M'_{-(i,S)}) + \varepsilon n & (6)\\ &\geq v_i(S) - \varepsilon + w^*(M_{-(i,S)}) - \varepsilon(n-1) + \varepsilon n & (7)\\ &= v_i(S) + w^*(M_{-(i,S)}) & (8)\end{aligned}$$

The first inequality follows from Lemma 3.6. The second follows from Equation (4). The third follows the assumption that $\|M - M'\|_\infty \leq \varepsilon$, and by Lemma 3.6 applied to submarket $M_{-(i,S)}$. Therefore, the allocation in which S is allocated to i cannot be welfare-maximizing in market M. □

To establish the correctness of EAP, we rely on the following generalization of the first welfare theorem of economics, which handles additive errors.

Theorem 3.9 (First Welfare Theorem [16]). *For $\varepsilon > 0$, let $(\mathcal{S}, \mathbf{p})$ be an ε-competitive equilibrium of M. Then, $\mathcal{S}$ is a welfare-maximizing allocation of M, up to additive error εn.*

Theorem 3.10 (Elicitation Algorithm with Pruning Guarantees). *Let $M_{\mathcal{X}}$ be a conditional market, let $\mathcal{D}$ be a distribution over $\mathcal{X}$, and let $c \in \mathbb{R}_+$. Suppose that for all $x \in \mathcal{X}$ and $(i,S) \in \mathcal{I}$, it holds that $v_i(S,x) \in [0,c]$, where $c \in \mathbb{R}$. Let $\boldsymbol{t}$ be a sequence of strictly increasing integers, and $\boldsymbol{\delta}$ a sequence of the same length as $\boldsymbol{t}$ such that $\delta_k \in (0,1)$ and $\sum_k \delta_k \in (0,1)$. If* EAP *outputs $(\{\hat{v}_i\}_{i\in N}, \{\hat{\varepsilon}_{i,S}\}_{(i,S)\in N\times 2^G}, 1 - \sum_k \delta_k, \hat{\varepsilon})$ on input $(M_{\mathcal{X}}, \mathcal{D}, \boldsymbol{t}, \boldsymbol{\delta}, c, \varepsilon)$, then the following holds with probability at least $1 - \sum_k \delta_k$:*

(1) $\|M_{\mathcal{D}} - \hat{M}\|_{\mathcal{I}} \leq \hat{\varepsilon}_{i,S}$
(2) $\mathcal{CE}(M_{\mathcal{D}}) \subseteq \mathcal{CE}_{2\hat{\varepsilon}}(\hat{M}) \subseteq \mathcal{CE}_{4\hat{\varepsilon}}(M_{\mathcal{D}})$

Here $\hat{M}$ is the empirical market obtained via EAP*, i.e., the market with valuation functions given by $\{\hat{v}_i\}_{i\in N}$.*

Proof. To show part 1, note that at each iteration k of EAP, Line 5 updates the error estimates for each (i,S) after a call to EA (Line 4 of EAP) with input failure probability δ_k. Theorem 3.5 implies that each call to EA returns estimated values that are within $\hat{\varepsilon}$ of their expected value with probability at least $1 - \delta_k$. By union bounding all calls to EA within EAP, part 1 then holds with probability at least $1 - \sum_k \delta_k$.

To show part 2, note that only pairs (i,S) for which Equation (4) holds are removed from index set $\mathcal{I}$ (Line 12 of EAP). By Lemma 3.8, no such pair can be part of any approximate welfare-maximizing allocation of the expected market, $M_{\mathcal{D}}$. By Theorem 3.9, no such pair can be a part of any CE. Consequently, $\hat{M}$ contains accurate enough estimates (up to ε) of all (i,S) pairs that may participate in any CE. Part 2 then follows from Theorem 2.3. □

4 EXPERIMENTS

In this section, we evaluate the empirical performance of our learning algorithms. To our knowledge, there have been no analogous attempts to learning CE; hence, we do not include any baseline algorithms from the literature in our experiments. Rather, we compare the performance of EAP, our pruning algorithm, to EA, investigating the quality of the CE learned by both, as well as their sample efficiencies.

We let $U[a,b]$ denote the continuous uniform over the continuous range $[a,b]$, and $U\{k,l\}$, the discrete uniform distribution over the set $\{k, k+1, \ldots, l\}$, for $k \leq l \in \mathbb{N}$.

4.1 Experimental Setup

Since one of the main goals of these experiments is to study the quality of learned CE, we focus on GS valuations. In particular, we choose unit-demand valuations. A buyer i is endowed with unit-demand valuations if, for all $S \subseteq G$, $v_i(S) = \max_{j\in S} v_i(\{j\})$. In a unit-demand market, all buyers have unit-demand valuations. A unit-demand market can be compactly represented by matrix $\mathbf{V}$, where entry $v_{ij} \in \mathbb{R}_+$ is i's value for j, i.e., $v_{ij} = v_i(\{j\})$. In what

Distribution	$\varepsilon = 0.05$ $\hat{\mathbf{p}}_{\text{MIN}}$	$\varepsilon = 0.05$ $\hat{\mathbf{p}}_{\text{MAX}}$	$\varepsilon = 0.2$ $\hat{\mathbf{p}}_{\text{MIN}}$	$\varepsilon = 0.2$ $\hat{\mathbf{p}}_{\text{MAX}}$
Uniform	0.0018	0.0020	0.0074	0.0082
Preferred-Good	0.0019	0.0023	0.0080	0.0094
Preferred-Good-Distinct	0.0000	0.0020	0.0000	0.0086
Preferred-Subset	0.0019	0.0022	0.0076	0.0090

Table 1: Average *UM-Loss* for $\varepsilon \in \{0.05, 0.2\}$.

follows, we denote by $\mathcal{V}$ a random variable over a unit-demand valuations, represented by a matrix $\mathbf{V}$.

Our goal is to robustly evaluate our algorithms, so we experiment with a variety of qualitatively different inputs. In particular, we construct four different distributions over unit-demand markets: Uniform, Preferred-Good, Preferred-Good-Distinct, and Preferred-Subset. All distributions are parameterized by n and m, the number of buyers and goods, respectively. A uniform unit-demand market $\mathcal{V} \sim$ Uniform is such that for all i, j, $v_{ij} \sim U[0, 10]$. In a preferred-good unit-demand market, where $\mathcal{V} \sim$ Preferred-Good, each buyer i has a preferred good j_i, with $j_i \sim U\{1, \ldots, m\}$ and $v_{ij_i} \sim U[0, 10]$. Conditioned on v_{ij_i}, i's value for good $k \neq j_i$ is given by $v_{ik} = v_{ij_i}/2^k$. Distribution Preferred-Good-Distinct is similar to Preferred-Good, except that no two buyers have the same preferred good. Note that the Preferred-Good-Distinct distribution is only well defined if $n \leq m$. Finally, in a preferred-subset unit-demand market, where $\mathcal{V} \sim$ Preferred-Subset, each buyer i is interested in a subset of goods $i_G \subseteq G$, where i_G is drawn uniformly at random from the set of all bundles. Then, the value i has for j is given by $v_{ij} \sim U[0, 10]$, if $j \in i_G$; and 0, otherwise.

4.2 Simulation of Noisy Valuation Elicitation

In what follows, we fix a realization of a unit-demand market $\mathbf{V}$, drawn from one of the aforementioned unit-demand distributions. We also fix a condition set $\mathcal{X} = [a, b]$, where $a < b$. We then define the conditional unit-demand market $M_{\mathcal{X}}$, where $v_i(S, x_{ij}) = \max_{j \in S}\{v_{ij}\} + x_{ij}$, for $x_{ij} \in \mathcal{X}$. Conditional market $M_{\mathcal{X}}$ together with distribution $\mathcal{D}$ on $\mathcal{X}$ is the model from which our algorithms elicit noisy valuations from buyers. For these experiments, all noise is drawn i.i.d.

We experiment with three noise models, low, medium, and high, by adding noise drawn from $U[-.5, .5]$, $U[-1, 1]$, and $U[-2, 2]$, respectively. We choose $n, m \in \{5, 10, 15, 20\}^2$, and we fix the failure probability at $\delta = 0.1$.

4.3 Empirical UM Loss of EA

In our first set of experiments, we investigate the empirical quality of the CE in the markets learned by EA. To measure the quality of a CE $(\hat{\mathcal{S}}, \hat{\mathbf{p}})$ computed for a market $\hat{M}$ in another market M, we first define the metric

$$\textit{UM-Loss}_{M,i}(\hat{\mathcal{S}}, \hat{\mathbf{p}}) = \max_{S \subseteq G}(v_i(S) - \hat{P}(S)) - (v_i(\hat{S}_i) - \hat{P}(\hat{S}_i)),$$

i.e., the difference between the maximum utility i could have attained at prices $\hat{\mathbf{p}}$ and the utility i attains at the outcome $(\hat{\mathcal{S}}, \hat{\mathbf{p}})$. Our metric of interest is then

$$\textit{UM-Loss}_M(\hat{\mathcal{S}}, \hat{\mathbf{p}}) = \max_{i \in N} \textit{UM-Loss}_{M,i}(\hat{\mathcal{S}}, \hat{\mathbf{p}}),$$

which is a worst-case measure of utility loss over all buyers in the market. Note that it is not useful to incorporate the MC condition into a loss metric, because it is always satisfied.

In our experiments, given an empirical estimate $\hat{M}_{\boldsymbol{x}}$ of M, and a CE $(\hat{\mathcal{S}}, \hat{\mathbf{p}})$ in $\hat{M}_{\boldsymbol{x}}$, we measure $\textit{UM-Loss}_M(\hat{\mathcal{S}}, \hat{\mathbf{p}})$, i.e., the loss in M at prices $\hat{\mathbf{p}}$ of CE $(\hat{\mathcal{S}}, \hat{\mathbf{p}})$. Theorem 2.3 implies that if $\hat{M}_{\boldsymbol{x}}$ is an ε-approximation of M, then $\textit{UM-Loss}_M(\hat{\mathcal{S}}, \hat{\mathbf{p}}) \leq 2\varepsilon$. Moreover, Theorem 3.5 yields the same guarantees, but with probability at least $1 - \delta$, provided the ε-approximation holds with probability at least $1 - \delta$.

As a learned CE is a CE of a learned market, we require a means of computing the CE of a market—specifically a unit-demand market $\mathbf{V}$. To do so, we first solve for the[1] welfare-maximizing allocation $\mathcal{S}^*_{\mathbf{V}}$ of $\mathbf{V}$, by solving for the maximum weight matching in the bipartite graph whose weight matrix is given by $\mathbf{V}$. Fixing $\mathcal{S}^*_{\mathbf{V}}$, we then solve for prices via linear programming. In general, there might be many prices that couple with $\mathcal{S}^*_{\mathbf{V}}$ to form a CE of $\mathbf{V}$. For simplicity, we solve for two pricings given $\mathcal{S}^*_{\mathbf{V}}$, the revenue-maximizing $\mathbf{p}_{\text{MAX}}$ and revenue-minimizing $\mathbf{p}_{\text{MIN}}$, where revenue is defined as the sum of the prices.

For each distribution, we draw 50 markets, and for each such market $\mathbf{V}$, we run EA four times, each time to achieve guarantee $\hat{\varepsilon} \in \{0.05, 0.1, 0.15, 0.2\}$. EA then outputs empirical estimate $\hat{\mathbf{V}}$ for each $\mathbf{V}$. We compute outcomes $(\mathcal{S}^*_{\hat{\mathbf{V}}}, \hat{\mathbf{p}}_{\text{MAX}})$ and $(\mathcal{S}^*_{\hat{\mathbf{V}}}, \hat{\mathbf{p}}_{\text{MIN}})$, and measure $\textit{UM-Loss}_{\mathbf{V}}(\mathcal{S}^*_{\hat{\mathbf{V}}}, \hat{\mathbf{p}}_{\text{MAX}})$ and $\textit{UM-Loss}_{\mathbf{V}}(\mathcal{S}^*_{\hat{\mathbf{V}}}, \hat{\mathbf{p}}_{\text{MIN}})$, for all possible combinations of the experimental parameters. We then average across all market draws, for both the minimum and the maximum pricings, for each possible configuration of the experimental parameters.

Table 1 summarizes a subset of these results. The error guarantees are consistently met across the board, indeed by one or two orders of magnitude, and they degrade as expected: i.e., with higher values of ε. We note that the quality of the learned CE is roughly the same for all distributions, except in the case of $\hat{\mathbf{p}}_{\text{MIN}}$ and Preferred-Good-Distinct, where learning is more accurate. For this distribution, it is enough to learn the preferred good of each buyer. Then, one possible CE is to allocate each buyer its preferred good and price all goods at zero which yields near no *UM-Loss*. Note that, in general, pricing all goods at zero is not a CE unless the market has some special structure, like that of markets drawn from Preferred-Good-Distinct.

4.4 Sample Efficiency of EAP

In these experiments, we evaluate the sample efficiency of EAP. We say that algorithm A has better *sample efficiency* than algorithm B if A requires fewer samples than B to achieve a desired accuracy ε.

Our experimental design is as follows. Fixing a unit-demand market, and the following values of $\varepsilon \in \{0.05, 0.1, 0.15, 0.2\}$, we compute the number of samples $t(\varepsilon)$ that would be required for EA to achieve accuracy ε. We then use the following doubling strategy

[1] Since values v_{ij} are drawn from continuous distributions, we assume that the set of markets for which there are multiple welfare-maximizing allocations is of negligible size.

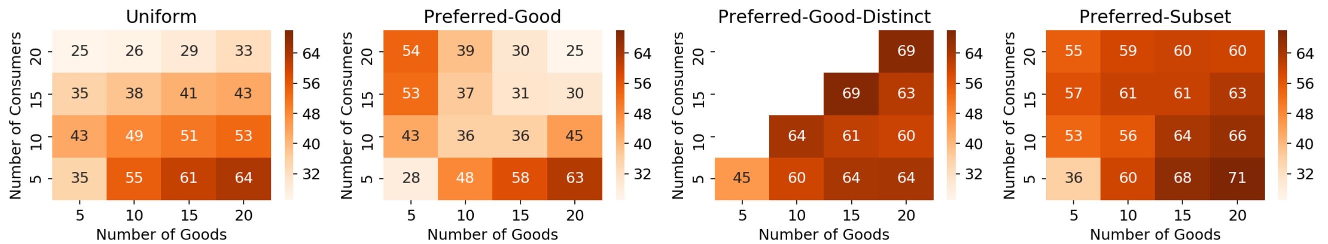

Figure 1: Average EAP **sample efficiency relative to** EA **for** $\hat{\varepsilon} = 0.05$**. Each** (i, j) **entry is annotated with the corresponding % saving.**

as a sampling schedule for EAP, $\boldsymbol{t}(t(\varepsilon)) = [t(\varepsilon)/4, t(\varepsilon)/2, t(\varepsilon), 2t(\varepsilon)]$, rounding to the nearest integer as necessary, and the following failure probability schedule $\boldsymbol{\delta} = [0.0125, 0.0125, 0.0125, 0.0125]$ so that the total failure probabilities sum to 0.1. Using these schedules, we run EAP with a desired accuracy of zero. We denote by $\varepsilon_{\text{EAP}}(\varepsilon)$ the approximation guarantee achieved by EAP upon termination.

For each distribution, we compute the average of the number of samples used by EAP across 50 independent market draws. We report samples used by EAP as a percentage of the number of samples used by EA to achieve the same guarantee, namely, $\varepsilon_{\text{EAP}}(\varepsilon)$, for each initial value of ε. Figure 1 depicts heat maps for all distributions, for $\varepsilon = 0.05$, where darker colors indicate more savings, and thus better EAP sample efficiency. A few trends arise, which we note are similar for other values of ε. For fixed number of buyers, EAP has better sample efficiency as the number of goods increases, because fewer goods can be allocated, which means that there are more candidate values to prune, resulting in more savings. The sample efficiency usually decreases as the number of buyers increases, which is to be expected, as the pruning criterion degrades with the number of buyers (Lemma 3.8). While savings exceed 30% across the board, we note that Uniform experiences the least savings—it has the least structure—and Preferred-Subset, the most. This shows that EAP is capable of exploiting the structure present in these distributions, despite not knowing anything about them *a priori*.

Finally, we note that sample efficiency quickly degrades for higher values of ε. In fact, for high enough values of ε (in our experiments, $\varepsilon = 0.2$), EAP might, on average, require more samples than EA to produce the same guarantee. Most of the savings are the result of pruning enough (i, j) pairs early enough: i.e., during the first few iterations of EAP. Four our experimental setup, when ε is large ($\varepsilon = 0.2$), then our sampling schedule does not allocate enough samples early on. While we report results for our particular choices of parameters, it is worth noting that when designing sampling schedules for EAP, one must allocate enough (but not too many) samples at the beginning of the schedule. Precisely how to determine this schedule is an empirical question, largely dependent on the particular application at hand.

5 CONCLUSION AND FUTURE DIRECTIONS

In this paper, we define noisy combinatorial markets as a model of combinatorial markets in which buyers cannot feasibly express their valuations with complete certainty, but can instead provide only noisy samples of them by, for example, using approximate methods, heuristics, or truncating the run-time of a complete algorithm. For this model, we tackle the problem of learning competitive equilibria solely from samples of buyers' valuations for bundles of goods. We first show tight lower- and upper-bounds on the set of CE, given a uniform approximation of one market by another. We then develop learning algorithms that, with high probability, learn said uniform approximations using only finitely many samples. Leveraging the first welfare theorem of economics, we define a pruning criterion under which an algorithm can provably stop learning about buyers' valuations for bundles, without affecting the quality of the set of learned competitive equilibria. We embed these conditions in an algorithm that we show experimentally is capable of learning equilibria with substantially fewer samples than the baseline, provided the underlying market has some exploitable structure. Crucially, the algorithm need not know anything about this structure *a priori*. Our algorithm is general enough to work in any combinatorial market. An interesting future direction is to speed up learning by developing alternative pruning criteria that exploit the structure of special classes of valuations.

REFERENCES

[1] Maria-Florina Balcan and Nicholas JA Harvey. Learning submodular functions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 793–802, 2011.

[2] Maria Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. Learning valuation functions. In *Conference on Learning Theory*, pages 4–1, 2012.

[3] Michael Ball, George Donohue, and Karla Hoffman. Auctions for the safe, efficient, and equitable allocation of airspace system resources. *Combinatorial auctions*, 1, 2006.

[4] Sushil Bikhchandani and John W Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory*, 74(2):385–413, 1997.

[5] Estelle Cantillon and Martin Pesendorfer. Auctioning bus routes: The london experience. ., 2006.

[6] Wolfram Conen and Tuomas Sandholm. Preference elicitation in combinatorial auctions. pages 168–169, 01 2002.

[7] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.

[8] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *IJCAI*, volume 99, pages 548–553, 1999.

[9] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic theory*, 87(1):95–124, 1999.

[10] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.

[11] Alexander S Kelso Jr and Vincent P Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483–1504, 1982.

[12] Christian Kroer, Alexander Peysakhovich, Eric Sodomka, and Nicolas E Stier-Moses. Computing large market equilibria using abstractions. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 745–746, 2019.

[13] Sebastien M Lahaie and David C Parkes. Applying learning algorithms to preference elicitation. In *Proceedings of the 5th ACM conference on Electronic commerce*,

pages 180–188, 2004.

[14] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.

[15] Tim Roughgarden and Inbal Talgam-Cohen. Why prices need algorithms. In *Proceedings of the sixteenth acm conference on economics and computation*, pages 19–36, 2015.

[16] Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.

[17] Enrique Areyan Viqueira, Cyrus Cousins, Eli Upfal, and Amy Greenwald. Learning equilibria of simulation-based games. *arXiv preprint arXiv:1905.13379*, 2019.

[18] Yevgeniy Vorobeychik and Michael P Wellman. Stochastic search methods for nash equilibrium approximation in simulation-based games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pages 1055–1062. International Foundation for Autonomous Agents and Multiagent Systems, 2008.

[19] Yevgeniy Vorobeychik. Probabilistic analysis of simulation-based games. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 20(3):1–25, 2010.

[20] Léon Walras. *Elements of Pure Economics: Or the Theory of Social Wealth*. Elements of Pure Economics, Or the Theory of Social Wealth. Routledge, 2003.

[21] Michael Weiss, Benjamin Lubin, and Sven Seuken. Sats: A universal spectrum auction test suite. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 51–59. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

[22] Michael P Wellman. Methods for empirical game-theoretic analysis. In *AAAI*, pages 1552–1556, 2006.