The Price is (Probably) Right: Learning Market Equilibria from Samples

Omer Lev, Neel Patel, Vignesh Viswanathan and Yair Zick

Abstract

Equilibrium computation in markets usually considers settings where player valuation functions are known. We consider the setting where player valuations are *unknown*; using a PAC learning-theoretic framework, we analyze some classes of common valuation functions, and provide algorithms which output direct PAC equilibrium allocations, not estimates based on attempting to learn valuation functions. Since there exist trivial PAC market outcomes with an unbounded worst-case efficiency loss, we lower-bound the efficiency of our algorithms. While the efficiency loss under general distributions is rather high, we show that in some cases (e.g., unit-demand valuations), it is possible to find a PAC market equilibrium with significantly better utility.

1 Introduction

Do markets admit equilibrium allocations? This question has been extensively studied for many years [Varian, 1974]; more recently, the econ/CS community devoted significant effort to understanding when one can *efficiently compute* market equilibria. Much of this literature assumes that one has full access to player valuations over bundles of goods, an unrealistic assumption in many instances: combinatorial valuations are often difficult to elicit (especially for large markets), precluding any possibility of running full-information market algorithms. Machine learning techniques offer a compromise – assuming access to a *partial* dataset, we can learn player valuations, and use the learned valuations as a proxy. However, this approach raises several issues too: market valuations are often complex, and require a large number of samples to learn without overfitting. Moreover, even if we assume that player valuations have a simple structure, it is not immediately obvious that an *exact* equilibrium for the approximate valuations acts as an *approximate* equilibrium for the exact valuations; as we shall show, this may not be the case.

Our work explores a relatively new paradigm: instead of learning valuations, we focus on directly learning *market equilibria* from data. We build upon the framework of Jha and Zick [2020], and adopt the *PAC solution learning* framework. Jha and Zick [2020] show that in order to ensure that a market outcome (i.e. an allocation of items to players, as well as item prices) is likely to be a market equilibrium, it suffices to show that it is consistent with the data. That is, the prices and item allocation they induce are such that no player has a sample in the data they can afford and would rather have over their allocation. Finding a consistent market outcome is trivial: setting the price of all the goods to infinity would ensure consistency. However, this outcome would be very inefficient. Our goal is thus to learn *approximately efficient* PAC market equilibria.

Our Contribution

We study Fisher markets with indivisible goods under different classes of valuation functions, and propose algorithms which output an efficient PAC market equilibrium. That is, each player receives, with high probability, their most preferred affordable bundle of goods. We examine two main classes of valuation functions: *unit-demand* (Section 4) and *additive* (Section 5) valuations. For each class, we provide a tight, distribution-independent,

1

efficiency bound. We also show that, under more favorable distributions, we can achieve far better efficiency guarantees for unit-demand valuations. In addition to this, we examine markets with *single minded* (Section D) and *submodular* (Section E) valuations in the appendix.

1.1 Related Work

There is a rich body of classical literature on market equilibria with indivisible goods [Eisenberg, 1961, Varian, 1974, Kelso and Crawford, 1982, Paes-Leme and Wong, 2017, Gul and Stacchetti, 1999], exchange economies [Bikhchandani and Mamer, 1997] and Fisher markets [Babaioff et al., 2017, Borodin et al., 2016, Segal-Halevi, 2017, Brânzei et al., 2016, Brânzei et al., 2015]. In recent years there is a significant renewed interest in computing market outcomes, such as fair allocation [Kurokawa et al., 2016, Farhadi et al., 2019], optimal pricing [Guruswami et al., 2005, Shen et al., 2019], approximate equilibria [Budish, 2011, Barman and Krishnamurthy, 2019] and markets with divisible goods [Devanur et al., 2008]. However, the above do not address learning approximately efficient market solutions from data.

There exists a fast-growing body of literature on learning game-theoretic solutions from data, in cooperative games [Balkanski et al., 2017, Sliwinski and Zick, 2017, Balcan et al., 2015, Igarashi et al., 2019, Jha and Zick, 2020], auctions [Devanur et al., 2016, Brero et al., 2018, Cole and Roughgarden, 2014, Morgenstern and Roughgarden, 2015] and optimization [Balkanski and Singer, 2017a,b, Rosenfeld et al., 2018]. Some of this literature propose methods to learn market outcomes as well: Murray et al. [2020] and Kroer et al. [2019] examine the simpler case with divisible goods and additive valuations, Shen et al. [2019] examine markets with a single item and Viqueira and Greenwald [2020] propose a method to learn market outcomes indirectly from noisy valuations. However, to the best of our knowledge, there exists no prior work that attempts to learn market outcomes in combinatorial markets from samples.

2 Model and Preliminaries

We study the *Fisher market* model; there is a set of *players*, $N = \{1, 2, \ldots, n\}$ and a set of *goods*, $G = \{g_1, g_2, \ldots, g_k\}$. Each player i has a *budget* $b_i \in \mathbb{R}^+$ and a *valuation function* $v_i : 2^G \to \mathbb{R}_+ \cup \{0\}$ which assigns a value $v_i(S)$ for each bundle of goods $S \subseteq G$. We assume that no two players have the same budget, and that $b_1 > \cdots > b_n$. This is a standard assumption, and is not a significant loss of generality: it is mostly done to induce some priority order among players, and ensure that equilibria exist. When budgets are equal, one can introduce small perturbations (this is the method used by Budish [2011]). An *allocation* in such a market is a tuple $(\mathcal{A}, \vec{p})$, where $\mathcal{A} = \{A_1, A_2, \ldots, A_n\}$ is the allocation vector and $\vec{p} = \{p_1, p_2, \ldots, p_k\}$ is the price vector. Some A_is may be empty i.e. if $A_i = \emptyset$, then the player receives nothing. We define the *affordable set* $D_i(\vec{p}, b_i)$ as the set of affordable bundles for player i given a price vector $\vec{p}$:

$$D_i(\vec{p}, b_i) = \{S \subseteq G : \sum_{g_j \in S} p_j \leq b_i\}$$

An allocation is a *Walrasian* equilibrium (or simply an equilibrium) if all players are allocated the best possible set of goods they can afford i.e. $A_i \in D_i(\vec{p}, b_i)$ for all $i \in N$, and for all $i \in N$:

$$v_i(A_i) \in \arg\max\{v_i(S) : S \in D_i(\vec{p}, b_i)\}$$

Jha and Zick [2020] define a learning-theoretic equilibrium notion based on the *probably approximately correct* (PAC) framework [Anthony and Bartlett, 1999] called *PAC Equilibria*. An allocation $(\mathcal{A}, \vec{p})$ is a PAC Equilibrium if it is unlikely that a sample from a distribution $\mathcal{D}$ (over bundles in G), under the same prices, is both better than the current allocation and affordable for any player i. It is often easier to discuss learning-theoretic notions in terms of their *expected loss*; here, the loss is a function of player valuations v and budgets $\vec{b}$, a bundle $S \subseteq G$, and the proposed outcome $(\mathcal{A}, \vec{p})$:

$$L_{v,\vec{b}}(S, \mathcal{A}, \vec{p}) = \begin{cases} 1 & \text{if } \exists i \in N, v_i(A_i) < v_i(S) \\ & \wedge S \in D_i(\vec{p}, b_i) \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

We omit the v and $\vec{b}$ subscripts when they are clear from context. An allocation is an ϵ-PAC equilibrium with respect to $\mathcal{D}$ if its expected loss, denoted $L_{\mathcal{D}}(\mathcal{A}, \vec{p})$, is lower than ϵ

$$L_{\mathcal{D}}(\mathcal{A}, \vec{p}) \triangleq \mathbb{E}_{S \sim \mathcal{D}}[L_{v,\vec{b}}(S, A, \vec{p})] < \epsilon \tag{2}$$

ϵ-PAC equilibria are somewhat similar to ϵ-PAC approximations [Anthony and Bartlett, 1999]: given a function $u : 2^G \to \mathbb{R}$, we say that $\bar{u}$ is an ϵ-PAC approximation of u w.r.t. $\mathcal{D}$ if $\Pr_{S \sim \mathcal{D}}[u(S) \neq \bar{u}(S)] < \epsilon$.

We follow a standard model of learning from samples: we are given players' budgets $b_1, b_2, \ldots, b_n$, m input samples $S_1, S_2, \ldots, S_m$ drawn i.i.d. from a distribution $\mathcal{D}$, and player valuations over the samples: $v_i(S_j)$ for all $i \in N$ and $j \in [m]$. Our goal is to find algorithms, whose input is a set of i.i.d. sampled bundles and valuations over them, that output a PAC Equilibrium (as per Equation (2)) with probability $\geq 1 - \delta$ (over the randomization of sampling m i.i.d. samples from $\mathcal{D}$). In other words, if $(\mathcal{A}, \vec{p})$ is the output of some learning algorithm, then the PAC guarantee is

$$\Pr_{S_1, \ldots, S_m \overset{\text{iid}}{\sim} \mathcal{D}}[L_{\mathcal{D}}(\mathcal{A}, \vec{p}) < \epsilon] \geq 1 - \delta$$

The number of samples needed, m, should be polynomial in the number of players, the number of goods, and in $\frac{1}{\epsilon}, \log \frac{1}{\delta}$. As mentioned in Section 1, PAC equilibria are not guaranteed to be efficient; in what follows we explore market *stability* and allocative *efficiency*.

The *efficiency ratio* of an allocation $\mathcal{A}$ is the ratio of the total welfare (or utility) of $\mathcal{A}$ to that of the optimal equilibrium allocation i.e.

$$ER_v(\mathcal{A}) = \frac{\sum_{i=1}^n v_i(A_i)}{\sum_{i=1}^n v_i(A_i^*)} \tag{3}$$

where $\mathcal{A}^*$ is a welfare-maximizing equilibrium allocation; Unlike simpler settings (e.g. rent division [Gal et al., 2017]), market outcomes need not maximize social welfare in Fisher markets with indivisible goods.

Due to space constraints, some of the proofs have been omitted or replaced by proof sketches. The full proofs can be found in the appendix.

3 Computing PAC Equilibria

We first discuss some sufficient conditions for finding a PAC Equilibrium from samples, starting with a simple observation: if we are able to approximate player valuation functions v using an underestimate $\bar{v}$, then any *exact* equilibrium for $\bar{v}$ is a PAC equilibrium for v.

Proposition 3.1. *Let $v_1, \ldots, v_n : 2^G \to \mathbb{R}$ be a player valuation profile; let $(\bar{v}_i)_{i \in N}$ be $\frac{\epsilon}{n}$-PAC approximations of $(v_i)_{i \in N}$ w.r.t. $\mathcal{D}$, such that for all $i \in N$ and all $S \subseteq G$, $\bar{v}_i(S) \le v_i(S)$. If $(\mathcal{A}, \vec{p})$ is a market equilibrium under $\bar{v}$, then $(\mathcal{A}, \vec{p})$ is an ϵ-PAC equilibrium for v w.r.t. $\mathcal{D}$.*

Jha and Zick [2020] prove that a PAC equilibrium can be directly learned using only $\mathcal{O}(k)$ samples if one can efficiently compute a *consistent solution*, that is, a market outcome that has zero loss on the samples. More precisely, we say that a mechanism $\mathcal{M}$ outputs a *consistent solution* if for any given set of samples $\mathcal{S} = \{S_1, \ldots, S_m\}$, $\mathcal{M}$ outputs $(\mathcal{A}, \vec{p})$ such that the *empirical loss* $\hat{L}(\mathcal{A}, \vec{p})$ is 0:

$$\hat{L}(\mathcal{A}, \vec{p}) \triangleq \frac{1}{m} \sum_{j=1}^{m} L(S_j, \mathcal{A}, \vec{p}) = 0$$

Theorem 3.2 (Jha and Zick [2020]). *Suppose that an algorithm $\mathcal{M}$ takes as input a set of m samples of goods $\mathcal{S}$ drawn i.i.d. from an unknown distribution $\mathcal{D}$, and outputs a consistent equilibrium allocation. If $m \in \mathcal{O}\left(\frac{1}{\epsilon}\left(k \log \frac{1}{\epsilon} + \log \frac{1}{\delta}\right)\right)$ then the allocation output by $\mathcal{M}$ is an ϵ-PAC Equilibrium w.p. $\ge 1 - \delta$.*

Proposition 3.1 and Theorem 3.2 provide two paths to computing PAC market equilibria: either compute an equilibrium for a PAC underestimate, or directly learn market outcomes from samples. As we mention above, our objective is finding market outcomes with provable social welfare guarantees, with respect to the *true* player valuation profile.

4 Unit Demand Markets

We begin our exploration with a fundamental class of market valuations: unit-demand markets. In a unit demand market, the value of each bundle $S \subseteq G$ is the value of the most valuable good in S, i.e. for all $i \in N$, $v_i(S) = \max_{g \in S} v_i(\{g\})$. We make the standard assumption that players have distinct values for goods, i.e. that $v_i(\{g\}) \ne v_i(\{g'\})$ if $g \ne g'$; this is mostly done to break ties (see Budish [2011]). Unit demand markets correspond to room/housing allocation scenarios where each tenant can only stay in a single room/buy a single home [Alkan et al., 1991, Aragones, 1995, Gal et al., 2017], or to gaming "loot boxes", in which players care mainly about the most valuable item.

The standard data-driven approach is to PAC learn the valuation functions, and output an equilibrium allocation for the learned valuations. We refer to this method as *indirect learning*, and to outcomes computed in this manner as indirectly learned outcomes. For unit demand markets this can be done quite easily, by estimating the value of each item as the value of the least valuable sample that contains it (creating a PAC approximation for the valuations), and then allocating the items first to the player with the largest budget, who gets their most valued item; then the player with the second largest budget, who gets their most valued item which is still available, and so on .

Such an algorithm has a simple guarantee of efficiency:

Proposition 4.1. *If $(\mathcal{A}, \vec{p})$ is the output of an algorithm calculating PAC approximations of unit demand valuations and then allocating goods in decreasing order of player budgets, then $ER_v(\mathcal{A}) \ge \frac{1}{\sigma}$ where $\sigma = \max_{i \in N} \frac{\max_{g \in G} v_i(\{g\})}{\min_{g \in G} v_i(\{g\})}$, the maximal ratio between a players valuation for two different items.*

The main drawback with such an algorithm is that it does not output a PAC Equilibrium. Consider the example below:

Example 4.1. Consider a setting where $N = \{1, 2\}$ and $G = \{g_1, g_2, g_3\}$. Player budgets are $b_1 = 2, b_2 = 1$. Player valuations satisfy

$$v_1(\{g_1, g_2\}) = 5; v_1(\{g_3\}) = 3$$
$$v_2(\{g_1, g_2\}) = 4; v_2(\{g_3\}) = 2$$

We observe a distribution $\mathcal{D}$ which samples uniformly at random two sets: $\{g_1, g_2\}$ and $\{g_3\}$. We can thus reasonably assume that we observe both bundles with high probability after a small number of i.i.d. samples. Approximating preferences would yield:

$$\bar{v}_1(\{g_1\}) = \bar{v}_1(\{g_2\}) = 5; \bar{v}_1(\{g_3\}) = 3$$
$$\bar{v}_2(\{g_1\}) = \bar{v}_2(\{g_2\}) = 4; \bar{v}_2(\{g_3\}) = 2$$

A valuation-approximating algorithm allocates one item from g_1, g_2 to player 1 and the other to player 2, and allocates g_3 to player 2. We set the price of g_1 to 2 and the price of g_2 to 1. Assume w.l.o.g. that g_1 is assigned to player 1; it is possible that $v_1(\{g_1\}) = 0$ and $v_1(\{g_2\}) = 5$, in which case player 1 demands g_3. In that case, the probability of observing a sample (namely $\{g_3\}$) which player 1 demands is $\frac{1}{2}$, not an arbitrarily low $\epsilon > 0$, so this approach does not yield an ϵ-PAC equilibrium.

The bad behavior in Example 4.1 is not due to some intrinsic failure of the valuation-approximating algorithm; it is impossible to learn a consistent underestimate of a unit demand valuation. Consider again the setting in Example 4.1: it is impossible to determine whether $v_1(\{g_1\}) = 5$ or $v_1(\{g_2\}) = 5$; indeed, the only viable underestimate sets both items' values to 0. However, doing so yields $\bar{v}_1(\{g_1, g_2\}) = 0 < v_1(\{g_1, g_2\})$, an inconsistency. To conclude, the indirect approach does not yield a PAC Equilibrium. Let us turn our attention to directly learning PAC market outcomes from samples. We refer to this method as *direct solution learning*, and any outcome computed from this method as a directly learned equilibrium. Algorithm 1 directly learns a PAC equilibrium in the unit-demand setting.

Algorithm 1 iterates over all players in decreasing order of budget, and allocates the smallest bundle of goods from all available goods with the highest possible value. We use two properties of unit demand valuations, formalized in the following lemma.

Lemma 4.2. *Given two bundles of goods $S, T \subseteq G$ and some player $i \in N$ with unit demand valuations, if no two goods have the same value for i then*

1. *If $v_i(S) = v_i(T) = c$ then $v_i(S \cap T) = c$ as well.*

2. *If $v_i(S) > v_i(T)$ then $v_i(S) = v_i(S \setminus T)$*

Using Lemma 4.2, we identify the smallest most valued bundle B_i^1 for player i, and allocate it to the player if it contains no previously allocated items; otherwise, we remove all such samples from $\mathcal{S}$, since we know such items are already priced out of their budget by previous players, and we cannot use them to get information on the next most valued set of goods for this player. We continue to identify the next most valued bundle of minimal size for player i. We repeat this process until we identify the smallest subset of most valued items among unallocated items. If we allocate a bundle to i after t steps, then player i receives a bundle containing their t-th most valuable good – denoted B_i^t; we then price the items in B_i^t such that their total price is b_i. Note that all samples that contain B_i^t have a price of $\geq b_i$, which guarantees that no player $i' > i$ can afford them.

We repeat this procedure for all players. At the end of the **_for loop_** (Algorithm 1, line 3), we allocate any leftover goods to player n for free, and assign any good which is not present in the sample set to player 1 at a price of 0.

We first show that Algorithm 1 outputs a consistent outcome.

Algorithm 1: Directly Learning Equilibria for Unit Demand Valuations

---

Input: A set of samples $\mathcal{S}$, player valuations and budgets $b_1 > \cdots > b_n$

1: $Alloc \leftarrow \emptyset$
2: Allocate unobserved goods to player 1 at price 0
3: **for** $i \leftarrow 1$ **to** n **do**
4:     $\mathcal{S}_i^1 \leftarrow \mathcal{S}$; $c \leftarrow False$; $t \leftarrow 1$
5:     **while** $c = False$ **do**
6:        $C_i^t \leftarrow$ some set in $\arg\max_{T \in \mathcal{S}_i^t} v_i(T)$
7:        $\mathcal{L}_i^t \leftarrow \{T \in \mathcal{S}_i^t | v_i(T) = v_i(C_i^t)\}$
8:        $B_i^t \leftarrow \bigcap_{T \in \mathcal{L}_i^t} T$
9:        $B_i^t = B_i^t \setminus \bigcup_{T \in \mathcal{S} | v_i(T) < v_i(C_i^t)} T$
10:        **if** $B_i^t \cap Alloc \neq \emptyset$ **then**
11:           $t \leftarrow t + 1$; $\mathcal{S}_i^t \leftarrow \mathcal{S}_i^{t-1} \setminus \mathcal{L}_i^t$
12:        **end**
13:        **else**
14:           $c \leftarrow True$; $Alloc \leftarrow Alloc \cup B_i^t$
15:           $A_i \leftarrow B_i^t$ and price of each $g \in B_i^t$ is $\frac{b_i}{|B_i|}$
16:        **end**
17:     **end**
18: **end**
19: Allocate the leftover goods to player n at price 0

Theorem 4.3. *Algorithm 1 outputs a consistent market outcome.*

Proof. Let the output of Algorithm 1 be $(\mathcal{A}, \vec{p})$. Let us assume there is a sample $S \in \mathcal{S}$ such that for some player i, $v_i(S) > v_i(A_i)$; we need to show that $S \notin D_i(\vec{p}, b_i)$. Consider the items not available to player i when it is their turn to select a bundle, referred to as *Alloc* in Algorithm 1. If $S \cap Alloc \neq \emptyset$, then S must contain some previously allocated bundle $A_{i'}$, where $b_{i'} > b_i$; thus the price of S is greater than b_i, and S is not demanded by i. If S can be allocated to player i and is one of the most valued bundles at the time, player i selects their bundle (i.e. $S \in \mathcal{L}_i^t$), then $B_i^t \subseteq S$; in particular, $v_i(S) = v_i(B_i^t)$. Otherwise, $v_i(B_i^t) > v_i(S)$ therefore $v_i(A_i) \geq v_i(S)$ and i would not demand S. $\square$

While Algorithm 1 outputs a consistent outcome, it offers an efficiency guarantee of $\frac{1}{\min\{n,k\}}$, under the minor assumption that player valuations are normalised with respect to their budget.

Proposition 4.4. *If for all $i \in N$, $\max_{g \in G} v_i(\{g\}) = b_i$, Algorithm 1 outputs an allocation $(\mathcal{A}, \vec{p})$ with $ER_v(\mathcal{A}) \geq \frac{1}{\min\{n,k\}}$.*

Proposition 4.4 offers a rather weak bound: the same efficiency ratio can be achieved by allocating all goods to the player with the highest budget. However, the bound is tight, and is an outcome of "bad" distributions. We show that there exists sample sets for which no allocation can guarantee an efficiency greater than $\frac{1}{\min\{n,k\}}$.

Theorem 4.5. *Let $\mathcal{S}, v(\mathcal{S})$ be a set of samples along with its valuations; let $\mathcal{V}$ be the set of unit demand valuation profiles consistent with the set of samples and are budget normalised i.e. $\max_{g \in G} v_i(\{g\}) = b_i$ for all the players $i \in N$, and $\mathcal{B} \subset \mathbb{R}_+^n$ be the set of all feasible*

budgets i.e. the set of all budgets in $\mathbb{R}^n_+$ such that $b_1 > b_2 > \cdots > b_n$. Then, we have

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{S \subseteq 2^G, \vec{b} \in \mathcal{B}} ER_v(\mathcal{A}) \leq \frac{1}{\min\{n, k\} - \delta}$$

for any $\delta \in (0, n)$ where $\mathcal{A}$ is a consistent allocation with respect to the samples.

Proof Sketch. Suppose the only sample you have is the set of goods G. Then the only consistent allocation which can guarantee a non-zero efficiency is one which allocates the entire set of goods to player 1. Any allocation which partitions the set of goods and allocates it to multiple players cannot offer any efficiency guarantees. This leaves us with allocating all goods to player 1. There exist allocations which provide a total utility of $\geq (\min\{n, k\} - \delta) \times b_1$; however, allocating the set of goods to player 1 guarantees a utility of b_1, which yields the upper bound $\frac{1}{\min\{n,k\}-\delta}$. $\qquad\square$

While Algorithm 1 offers no reasonable welfare guarantees for general distributions, its performance guarantees improve significantly under certain distributional assumptions. Specifically, this holds true if $\mathcal{D}$ is a product distribution with a bounded probability of sampling each good. Recall that $\mathcal{D}$ is a product distribution over G if there exist values $p_1, \ldots, p_k \in [0, 1]$ such that for every $S \subseteq G$, $\Pr_{\mathcal{D}}[S] = \prod_{g_j \in S} p_j$. Product distributions offer more amenable welfare guarantees for two reasons: first, by definition, the presence of a particular good in the sample is independent of the presence of any other good (offering us a better chance of observing players' valuations for individual items); second, goods are sampled with non-zero probability (thus we observe all goods in some bundle with high probability). Theorem 4.8 shows that Algorithm 1 outputs a PAC equilibrium with an efficiency ratio of 1 with exponentially high probability, when samples are drawn i.i.d. from a product distribution; the proof requires that player preference orders over items are sufficiently distinct. Before we prove Theorem 4.8, we present two technical results – Lemma 4.6 and Lemma 4.7 – which we use to prove Theorem 4.8.

Lemma 4.6. *In unit demand markets with unequal budgets and strict preferences over items, any equilibrium allocation assigns player i the best possible available good, i.e. $\{g_i^*\}$ equals $\arg\max_{g \in \mathcal{G}_i} v_i(g)$ ($\mathcal{G}_1 = \arg\max_{g \in G} v_1(g)$ and for $i > 1$, $\mathcal{G}_i = G \setminus \{\bigcup_{l=1}^{i-1} \mathcal{G}_l\}$). Moreover, all equilibria have the same social welfare $\sum_i v_i(g_i^*)$.*

In Lemma 4.6, we show that the social welfare for any equilibrium for unit demand players is unique and each player i gets the good g_i^*. Therefore to show that the efficiency of Algorithm 1 is 1 with high probability, it is sufficient to show that Algorithm 1 assigns g_i^* for all i with high probability.

We now present Lemma 4.7, in which we prove that for any player i, if $\mathcal{S}_i^t$ at $t-$th iteration of the ***while loop*** in Algorithm 1 contains more than k^2 samples then the corresponding B_i^t contains only the best available good for player i in $\bigcup_{S \in \mathcal{S}_i^t} S$, with high probability.

Lemma 4.7. *Suppose that $\mathcal{D}$ is a product distribution such that for all $g \in G$, $1 - \sqrt{2e^{-1/k} - 1} < \Pr_{S \in \mathcal{D}}(g \in S) < \frac{1}{2} + \frac{\sqrt{2e^{-1/k} - 1}}{2}$. If $|\mathcal{S}_i^t| \geq k^2$ (at the $t-$th iteration of the **while loop** in Algorithm 1 for player i), the corresponding B_i^t equals $\{\hat{g}_i\}$ to player i with at least $1 - e^{-\frac{k}{2}}$ probability, where*

$$\hat{g}_i \in \arg\max\{v_i(\{g\}) : g \in \bigcup_{S \in \mathcal{S}_i^t} S\}$$

We are now ready to prove Theorem 4.8. We show that when we assume agent preferences sufficiently differ – no two agents have exactly the same favorite $\mathcal{O}(\log(\max(n, k)))$ goods – Algorithm 1 is optimal with high probability..

Theorem 4.8. *Suppose that $\mathcal{D}$ is a product distribution, such that $\Pr_{S \sim \mathcal{D}}[g \in S] \in [\alpha, \beta]$. Assume that for every agent i, $|\{g \in G : v_i(g) > v_i(g_i^*)\}| < \frac{\max\{\log n, \log k\}}{\log(\frac{1}{1-\beta})}$[1].*

If $k > 3$, $1 - \sqrt{2e^{-1/k} - 1} \leq \alpha$ and $\beta \leq \frac{1}{2} + \frac{\sqrt{2e^{-1/k} - 1}}{2}$, the output of Algorithm 1, $(\mathcal{A}, \vec{p})$, satisfies

$$\Pr[ER_v(\mathcal{A}) = 1] \geq 1 - \frac{2n \max\{\log n, \log k\}}{\log\left(\frac{1}{(1-\beta)}\right)} e^{-\frac{\max\{k, n\}}{4}}$$

Proof Sketch. When we have $\max\{k^4, n^2 k^2\}$ samples, if the condition on the valuation functions is satisfied, then for every player i, there is some t for which $|S_i^t| \geq k^2$ and $g_i^* \in \arg\max\{v_i(\{g\}) : g \in \bigcup_{S \in \mathcal{S}_i^t} S\}$. Therefore, using induction and Lemma 4.7, every player gets allocated g_i^* w.h.p., resulting in an efficiency ratio of 1. $\square$

As β decreases (provided $\beta > 1 - \sqrt{2e^{-1/k} - 1}$), the condition in Theorem 4.8 on the difference between players' preferences becomes less stringent. Moreover, if either n or k is large, the exponential term in the probability guarantee dominates, and Algorithm 1 is highly likely to output an efficient outcome. However, if β is smaller, the efficiency guarantee is less likely to hold. Note that when $\beta = 1$, i.e., there is a good g that appears in all samples, the performance of Algorithm 1 depends on which player gets g. If the most preferred good for all players is g, Algorithm 1 allocates g to player 1 and will not be able to continue: it is impossible to identify the second preferred good (and beyond). Therefore, Algorithm 1 has an efficiency $\geq \frac{1}{\rho n}$ since we can only guarantee that the highest budget player will receive their optimal equilibrium allocation.

We can generalize the efficiency bound in Theorem 4.8 for any preference order over the items for all players. We observe that with at least $\max(k^4, n^2 k^2)$ samples, the first $\sim \max(\log k, \log n)$ players will be assigned g_i^* with high probability. We show the efficiency guarantee for algorithm 1 for any preference order in Proposition 4.9, and its connection to the disparity in valuation functions between agents.

Proposition 4.9. *If $\mathcal{D}$ is a product distribution such that for all $g_j \in G$, $1 - \sqrt{2e^{-1/k} - 1} < \Pr_{S \in \mathcal{D}}(g \in S) < \frac{1}{2} + \frac{\sqrt{2e^{-1/k} - 1}}{2}$ and $k > 3$. Then, with exponentially high probability, Algorithm 1 allocates goods with an efficiency ratio $ER_v(\mathcal{A}) \geq \frac{\log n}{\rho n \log\left(\frac{1}{1-\beta}\right)}$ where $\rho = \max_{g \in G} \frac{\max_{i \in N} v_i(g)}{\min_{i \in N} v_i(g)}$ and $\beta = \max_{g_j \in G} \Pr_{S \in \mathcal{D}}(g \in S)$.*

Furthermore, in Corollary 4.10, we show the efficiency bound when each good is sampled i.i.d. w.p. $\frac{1}{2}$.

Corollary 4.10. *If the distribution $\mathcal{D}$ is uniform over the set 2^G and $k > 3$, with exponentially high probability. Algorithm 1 allocates goods with an efficiency $ER_v(\mathcal{A}) > \frac{\log n}{\rho n}$ where*

$$\rho = \max_{g \in G} \frac{\max_{i \in N} v_i(g)}{\min_{i \in N} v_i(g)}$$

using a polynomial number of samples.

5 Additive Markets

In additive markets, each player has additive valuations. The valuation of a bundle is equal to the sum of the valuations of every good in that bundle: $v_i(S) = \sum_{g \in S} v_i(\{g\})$. While

---

[1]g_i^* is defined as in Lemma 4.6: $\{g_i^*\} = \arg\max_{g \in \mathcal{G}_i} v_i(g)$ ($\mathcal{G}_1 = \arg\max_{g \in G} v_1(g)$) and for $i > 1$, $\mathcal{G}_i = G \setminus \{\bigcup_{l=1}^{i-1} \mathcal{G}_l\}$.

additive valuations are PAC-Learnable, we cannot use Proposition 3.1 to learn a PAC-Equilibrium since in a lot of cases, we cannot learn an underestimate of the valuations. This can be seen using Example 4.1.

Although additive Fisher markets with indivisible goods have recently received a lot of attention, there are still many open questions regarding the efficient computation of a market clearing equilibrium. Babaioff et al. [2017] examine the specific case where there are only two players and Brânzei et al. [2015] show that it is computationally intractable to decide if a market has a competitive equilibrium when budgets are equal. This dearth of positive algorithmic results means that even if we could accurately learn the valuation of each good (which is not guaranteed and depends on the samples), we may not be able to compute an equilibrium in polynomial time. In this paper, we take a different approach and attempt to learn an equilibrium directly (using Theorem 3.2); however, our outcome is not necessarily market clearing.

Our approach is described in Algorithm 2. The algorithm has three steps. First, we pre-process the samples to ensure that there are no proper subsets in the samples. This is done to ensure that no sample which is a superset of another sample is allocated. We can remove the supersets and replace them by the set difference between the superset and the subset: we can derive the value of this bundle under additive valuations, as executed in the function PreProcess.

The second step allocates samples to players. To each player, the algorithm allocates the favourite sample among all the unallocated samples. Here, a sample is unallocated if no good in the sample has been allocated. It then prices each good equally such that the total price is equal to the budget of the player.

The last step ensures consistency, it checks each of the original samples to see if a player prefers it over their own sample and can afford it. If there exists such a player, the algorithm proceeds to set the price of one of the goods in the sample to infinity to ensure that no player can afford it. This good is chosen as follows: if the sample has an unallocated good, then the unallocated good is chosen. If the sample does not have an unallocated good, the algorithm takes away a good from the sample which belonged to the player with the least budget and then sets its price to infinity. We refer to the act of setting the price of a good to infinity as burning a good.

It is easy to see because of the third step that the algorithm is always consistent. It also worth noting that as long as we can underestimate the valuation in Line 25 in Algorithm 2, we will always end up with a consistent outcome. This means that this algorithm could be modified for any class of valuations to output a consistent outcome. In Section E, we show how this can be done for submodular valuations.

We now prove two efficiency bounds for our algorithm. These bounds hold only for additive valuations. To start with, we show that when the valuations are budget normalised, then the efficiency is inversely related to the number of goods. Before that, we show that no good in player 1's initially allocated sample gets taken away in Lemma 5.1.

Lemma 5.1. *In Algorithm 2, no good in player 1's initially allocated sample gets taken away.*

Theorem 5.2. *When $\forall i \in N, \max_{g \in G} v_i(\{g\}) = b_i$, then Algorithm 2 outputs an allocation with $ER_v(\mathcal{A}) \geq \frac{1}{k}$*

Proof. Algorithm 2 always ensures the first player has a bundle with valuation at least b_1. If the first player's favourite good is not present in any sample, he receives at price 0 resulting in a valuation of at least b_1.

If the first player's favourite good is present in the samples, then there exists a sample (with the first player's favourite good in it) which is valued at at least b_1 by the first player.

Algorithm 2: Consistent Allocation For Additive Markets

    Input: A set of samples $\mathcal{S}$, player valuations for these samples $v(\mathcal{S})$ and budgets
            $b_1 > b_2 > \cdots > b_n$

1 $\mathcal{S}', \tilde{v}(\mathcal{S}') = \texttt{PreProess}(\mathcal{S}, v(\mathcal{S}))$

2   **for** $i \leftarrow 1$ *to* n **do**

3 $B_i = $ some set in $\arg\max_{T \in \mathcal{S}'} \tilde{v}_i(T)$

4   Allocate B_i to player i i.e. $A_i = B_i, \tilde{v}_i(A_i) = \tilde{v}_i(B_i)$

5   Set $p_g = \frac{b_i}{|B_i|}$ $\forall g \in B_i$

6 $\mathcal{S}' = \mathcal{S}' \setminus \bigcup_{S \in \mathcal{S}' : S \cap B_i \neq \emptyset} S$

7   **end**

8   **while** $\exists i \in N, S \in \mathcal{S}$ *s.t* $\tilde{v}_i(A_i) < v_i(S) \wedge \sum_{g \in S} p_g \leq b_i$ **do**

9   **if** $\exists g \in S$ *s.t.* $g \notin \bigcup_{i \in N} A_i$ **then**

10   Set $p_g = \infty$

11   **end**

12   **else**

13 $j = \arg\min_{i \in N : S \cap A_i \neq \emptyset} b_i$

14 $g = $ any good in $A_j \cap S$

15 $p_g = \infty$

16 $A_j = A_j \setminus \{g\}$

17 $\tilde{v}_j(A_j) = 0$

18   **end**

19   **end**

20   Allocate all leftover goods to player 1 at price 0

21   **Function** $\texttt{PreProess}(\mathcal{S}, v(\mathcal{S}))$**:**

22 $\mathcal{S}' = \mathcal{S}$

23 $\tilde{v}(\mathcal{S}') = v(\mathcal{S})$

24   **while** $\exists S', S'' \in \mathcal{S}'$ *s.t.* $S' \subsetneq S''$ **do**

25 $\tilde{v}_i(S'') = \tilde{v}_i(S'') - \tilde{v}_i(S')$ $\forall i \in N$

26 $S'' = S'' \setminus S'$

27   **end**

28   **return** $\mathcal{S}', \tilde{v}(\mathcal{S}')$

Since the first player is allocated his favourite sample, he is given a bundle whose value is at least b_1. By Lemma 5.1, none of these goods are taken away and their final utility is at least b_1.

Since the largest amount of value a good can give a player is b_1. The total utility of any allocation is upper bounded by kb_1. This gives us the following bound:

$$ER_v(\mathcal{A}) = \frac{\sum_{i \in N} v_i(A_i)}{\sum_{i \in N} v_i(A_i^*)} \geq \frac{v_1(A_1)}{\sum_{i \in N} v_i(A_i^*)} \geq \frac{b_1}{kb_1} = \frac{1}{k}$$

$\square$

We now show that this bound is tight for general distributions.

Theorem 5.3. *Let $\mathcal{S}, v(\mathcal{S})$ be a set of samples along with its valuations, $\mathcal{V}$ be the set of additive valuation function profiles which are consistent with the set of samples and are budget normalised i.e. $\max_{g \in G} v_i(\{g\}) = b_i$ for all the players $i \in N$ and $\mathcal{B} \subset \mathbb{R}_+^n$ be the set of all feasible budgets i.e. the set of all budgets in $\mathbb{R}_+^n$ such that $b_1 > b_2 > \cdots > b_n$. Then, we have*

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, \vec{b} \in \mathcal{B}} ER_v(\mathcal{A}) \leq \frac{1}{k - \delta}$$

for any $\delta \in (0, k)$ where $\mathcal{A}$ is a consistent allocation with respect to the samples.

Our next bound does not require the valuations to be normalised but imposes conditions on the samples and depends on the disparity in the valuations of goods.

In Proposition 5.4, we show that when samples are disjoint, the efficiency varies inversely with the disparity in valuations

Proposition 5.4. *When all the samples in $\mathcal{S}$ are pairwise disjoint, then Algorithm 2 outputs an allocation with $ER_v(\mathcal{A}) \geq \frac{1}{\rho}$ where $\rho = \max_{g \in G} \frac{\max_{i \in N} v_i(\{g\})}{\min_{i \in N} v_i(\{g\})}$*

6   Conclusions and Future Work

This work shows the benefit of directly learning equilibrium states, instead of learning utility functions, and calculating equilibria states from them. We deal with several valuation function families, and in all of them show algorithms to produce a PAC-approximation, with our results being tight, i.e., no better approximation can be guranteed.

We believe that this work is the tip of the iceberg in showing how PAC learning can help in reaching economic, game-theoretic results, directly from the data, without using the data to construct intermediate steps (such as learning utility functions). Plenty of problems are still open – from expanding results to a larger family of functions (XOS, gross substitutes), to further type of results (e.g., other desirable states beyond equilibria).

References

Ahmet Alkan, Gabrielle Demange, and David Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991.

M. Anthony and P. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.

Enriqueta Aragones. A derivation of the money rawlsian solution. *Social Choice and Welfare*, 12(3):267–276, 1995.

M. Babaioff, N. Nisan, and I. Talgam-Cohen. Competitive equilibria with indivisible goods and generic budgets. *arXiv preprint arXiv:1703.08150*, 2017.

Maria-Florina Balcan and Nicholas J.A. Harvey. Learning submodular functions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 793–802, 2011.

M.F. Balcan, A.D. Procaccia, and Y. Zick. Learning cooperative games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 475–481, 2015.

Eric Balkanski and Yaron Singer. Minimizing a submodular function from samples. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NIPS)*, pages 814–822, 2017a.

Eric Balkanski and Yaron Singer. The sample complexity of optimizing a convex function. In *Proceedings of the 30th Conference on Computational Learning Theory (COLT)*, pages 275–301, 2017b.

Eric Balkanski, Umar Syed, and Sergei Vassilvitskii. Statistical cost sharing. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 6221–6230, 2017.

Siddharth Barman and Sanath Krishnamurthy. On the proximity of markets with integral equilibria. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1748–1755, 2019.

S. Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413, 1997.

Allan Borodin, Omer Lev, and Tyrone Strangway. Budgetary effects on pricing equilibrium in online markets. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 95–103, Singapore, 2016.

Simina Brânzei, Hadi Hosseini, and Peter Bro Miltersen. Characterization and computation of equilibria for indivisible goods. In *Proceedings of the 8th International Symposium on Algorithmic Game Theory (SAGT)*, pages 244–255, 2015.

Simina Brânzei, Yuezhou Lv, and Ruta Mehta. To give or not to give: Fair division for single minded valuations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 123–129, 2016.

Gianluca Brero, Benjamin Lubin, and Sven Seuken. Combinatorial auctions via machine learning-based preference elicitation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 128–136, 2018.

Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economics*, 119(6):1061–1103, 2011.

Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 243–252, 2014.

Nikhil R Devanur, Christos H Papadimitriou, Amin Saberi, and Vijay V Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the Association for Computing Machinery*, 55(5):22, 2008.

Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 426–439, 2016.

Edmund Eisenberg. Aggregation of utility functions. *Management Science*, 7(4):337–350, 1961.

Alireza Farhadi, Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Sébastien Lahaie, David Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research*, 64(1):1–20, 2019. ISSN 1076-9757.

Ya'akov (Kobi) Gal, Moshe Mash, Ariel D. Procaccia, and Yair Zick. Which is the fairest (rent division) of them all? *Journal of the Association for Computing Machinery*, 64(6): 39:1–39:22, 2017.

F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.

Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, Frank McSherry, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173, 2005.

Ayumi Igarashi, Jakub Sliwinski, and Yair Zick. Forming probably stable communities with limited interactions. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 2053–2060, 2019.

Tushant Jha and Yair Zick. A learning framework for distribution-based game-theoretic solution concepts. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 355–377, 2020.

Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982. ISSN 00129682, 14680262.

Christian Kroer, Alexander Peysakhovich, Eric Sodomka, and Nicolas E Stier-Moses. Computing large market equilibria using abstractions. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 745–746, 2019.

David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 523–529, 2016.

Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 136–144, 2015.

Riley Murray, Christian Kroer, Alex Peysakhovich, and Parikshit Shah. Robust market equilibria with uncertain preferences. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2192–2199, 2020.

Renato Paes-Leme and Sam Chiu-Wai Wong. Computing walrasian equilibria: Fast algorithms and structural properties. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 632–651, 2017.

Nir Rosenfeld, Eric Balkanski, Amir Globerson, and Yaron Singer. Learning to optimize combinatorial functions. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 4374–4383, 2018.

Erel Segal-Halevi. Competitive equilibrium for almost all incomes. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1267–1275, 2017.

Weiran Shen, Sebastien Lahaie, and Renato Paes-Leme. Learning to clear the market. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 5710–5718, 2019.

J. Sliwinski and Y. Zick. Learning hedonic games. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2730–2736, 2017.

H. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

Enrique Areyan Viqueira and Amy Greenwald. Learning competitive equilibria in noisy combinatorial markets. In *Proceedings of the 2nd Games, Agents, and Incentives Workshop (GAIW@AAMAS 2020)*, 2020.

A Missing Proofs from Section 3

Proposition 3.1. *Let $v_1, \ldots, v_n : 2^G \to \mathbb{R}$ be a player valuation profile; let $(\bar{v}_i)_{i \in N}$ be $\frac{\epsilon}{n}$-PAC approximations of $(v_i)_{i \in N}$ w.r.t. $\mathcal{D}$, such that for all $i \in N$ and all $S \subseteq G$, $\bar{v}_i(S) \leq v_i(S)$. If $(\mathcal{A}, \vec{p})$ is a market equilibrium under $\bar{v}$, then $(\mathcal{A}, \vec{p})$ is an ϵ-PAC equilibrium for v w.r.t. $\mathcal{D}$.*

Proof. By the union bound,

$$\Pr_{S \sim \mathcal{D}}[\exists i \in N \mid v_i(S) \neq \bar{v}_i(S)] \leq \epsilon,$$

which implies:

$$\Pr_{S \sim \mathcal{D}}[\forall i, (v_i(S) > v_i(A_i)) \wedge (S \in D_i(\vec{p}, b_i))]$$
$$= \Pr_{S \sim \mathcal{D}}[\forall i, (v_i(S) > v_i(A_i)) \wedge (S \in D_i(\vec{p}, b_i)) \wedge \forall i, v_i(S) = \bar{v}_i(S)] +$$
$$\Pr_{S \sim \mathcal{D}}[\forall i, (v_i(S) > v_i(A_i)) \wedge (S \in D_i(\vec{p}, b_i)) \wedge \exists i, v_i(S) \neq \bar{v}_i(S)]$$
$$\leq \Pr_{S \sim \mathcal{D}}[\forall i, (\bar{v}_i(S) > v_i(A_i)) \wedge (S \in D_i(\vec{p}, b_i))] + \epsilon \tag{4}$$
$$\leq \Pr_{S \sim \mathcal{D}}[\forall i, (\bar{v}_i(S) > \bar{v}_i(A_i)) \wedge (S \in D_i(\vec{p}, b_i))] + \epsilon \leq \epsilon \tag{5}$$

The transition to (4) is due to the fact that $v_i(S) \geq \bar{v}_i(S)$; the transition to (5) is due to the fact that $(\mathcal{A}, \vec{p})$ is an exact market equilibrium for $\bar{v}$. $\square$

B Missing Proofs from Section 4

Proposition 4.1. *If $(\mathcal{A}, \vec{p})$ is the output of an algorithm calculating PAC approximations of unit demand valuations and then allocating goods in decreasing order of player budgets, then $ER_v(\mathcal{A}) \geq \frac{1}{\sigma}$ where $\sigma = \max_{i \in N} \frac{\max_{g \in G} v_i(\{g\})}{\min_{g \in G} v_i(\{g\})}$, the maximal ratio between a players valuation for two different items.*

Proof. If a good is valued at 0 by some player, then σ is undefined. So, we only consider the case where all goods have a non-zero valuation.

Let $\mathcal{A}^*$ be the socially optimal allocation. When $n \geq k$, in the optimal equilibrium allocation, the top k players budget-wise get a good each and the rest of the players get nothing. Indirect learning also allocates a good to each of the top k players budget wise. If g_i is the good allocated to player i in the algorithm, then the value to player i in the optimal equilibrium allocation is bounded by $\sigma v_i(g_i)$. This gives us the efficiency $ER_v(\mathcal{A})$,

$$\frac{\sum_{i=1}^n v_i(A_i)}{\sum_{i=1}^n v_i(A_i^*)} = \frac{\sum_{i=1}^k v_i(A_i)}{\sum_{i=1}^k v_i(A_i^*)} \geq \frac{\sum_{i=1}^k v_i(A_i)}{\sum_{i=1}^k \sigma v_i(A_i)} = \frac{1}{\sigma}$$

Similarly, when $n \leq k$, every player gets one good and some goods may be left unallocated. Indirect learning also allocates a good to each player. If g_i is the good allocated to player $i < n$ in the algorithm, then the value to player i in the optimal equilibrium allocation is bounded by $\sigma v_i(g_i)$. This gives us

$$ER_v(\mathcal{A}) = \frac{\sum_{i=1}^n v_i(A_i)}{\sum_{i=1}^n v_i(A_i^*)} \geq \frac{\sum_{i=1}^n v_i(A_i)}{\sum_{i=1}^n \sigma v_i(A_i)} = \frac{1}{\sigma}$$

$\square$

Lemma 4.2. *Given two bundles of goods $S, T \subseteq G$ and some player $i \in N$ with unit demand valuations, if no two goods have the same value for i then*

1. *If $v_i(S) = v_i(T) = c$ then $v_i(S \cap T) = c$ as well.*

2. *If $v_i(S) > v_i(T)$ then $v_i(S) = v_i(S \setminus T)$*

Proof. Since player i has a unit-demand valuation,

$$v_i(S) = \max\{v_i(S \cap T), v_i(S \setminus T)\} \tag{6}$$
$$v_i(T) = \max\{v_i(S \cap T), v_i(T \setminus S)\}; \tag{7}$$

since all items have different values, it must be the case that $v_i(S) > v_i(S \cap T)$ or $v_i(S) > v_i(S \setminus T)$. Suppose that $v_i(S) = v_i(T) = c$ and $c > v_i(S \cap T)$; then it must be the case that $v_i(S) = v_i(S \setminus T) = c$ and $v_i(T) = v_i(T \setminus S) = c$. However, this implies that there are two disjoint goods: $g \in S \setminus T$ and $g' \in T \setminus S$ that are equally valued by i, a contradiction. We have thus proven (6).

Similarly, if $v_i(S) > v_i(T)$ and $v_i(S) = v_i(S \cap T)$, we get that

$$v_i(S) > v_i(T) \geq v_i(S \cap T) = v_i(S),$$

a contradiction. Therefore $v_i(S) = v_i(S \setminus T)$ which proves (7). $\qquad\square$

Proposition 4.4. *If for all $i \in N$, $\max_{g \in G} v_i(\{g\}) = b_i$, Algorithm 1 outputs an allocation $(\mathcal{A}, \vec{p})$ with $ER_v(\mathcal{A}) \geq \frac{1}{\min\{n,k\}}$.*

Proof. Irrespective of the samples and the distribution, Algorithm 1 ensures that the player with the highest budget (Player 1) gets their best possible allocation i.e. $v_1(A_1) = b_1$. When $n \leq k$, Given the normalisation w.r.t. player budgets, the utility of the optimal equilibrium allocation has to be less than the sum of all the budgets i.e. $\sum_{i \in N} v_i(A_i^*) \leq \sum_{i \in N} b_i \leq n b_1$. When $n > k$, only k players can get a good which means the upper bound on the utility of the optimal equilibrium allocation will be $\leq k b_1$. From this the upper bound on the utility of the optimal equilibrium allocations will be $\leq \min\{k, n\} b_1$ Therefore $ER_v(\mathcal{A})$ is

$$\frac{\sum_{i \in N} v_i(A_i)}{\sum_{i \in N} v_i(A_i^*)} \geq \frac{b_1}{\sum_{i \in N} v_i(A_i^*)} \geq \frac{b_1}{\min\{n, k\} b_1} = \frac{1}{\min\{n, k\}} \tag{8}$$

$\qquad\square$

Theorem 4.5. *Let $\mathcal{S}, v(\mathcal{S})$ be a set of samples along with its valuations; let $\mathcal{V}$ be the set of unit demand valuation profiles consistent with the set of samples and are budget normalised i.e. $\max_{g \in G} v_i(\{g\}) = b_i$ for all the players $i \in N$, and $\mathcal{B} \subset \mathbb{R}_+^n$ be the set of all feasible budgets i.e. the set of all budgets in $\mathbb{R}_+^n$ such that $b_1 > b_2 > \cdots > b_n$. Then, we have*

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, \vec{b} \in \mathcal{B}} ER_v(\mathcal{A}) \leq \frac{1}{\min\{n, k\} - \delta}$$

for any $\delta \in (0, n)$ where $\mathcal{A}$ is a consistent allocation with respect to the samples.

Proof. Consider a market with n players and k goods. Define a set of unit demand valuation function profiles $\mathcal{V}'$ as follows: each player has one good for which $v_i(g) = b_i$ and every other good has value 0 for this player. We refer to the good with non-zero valuation as the favourite good of player i. Also, let no two players in the top $\min\{n, k\}$ players budget wise have the same favourite good. This set of valuations satisfies our budget normalisation condition.

Define the budget vector $\{b_1, b_2, \ldots, b_n\}$ as follows: for every player $b_i = b_1 - \delta_i$ where $\delta_1 = 0$, $0 < \delta_2 < \delta_3 < \cdots < \delta_n$ and $\sum_{i \in N} \delta_i = \delta b_1$. Let us call this vector of budgets $\vec{b'}$.

Now, suppose the only sample we have is the set of goods G ($\mathcal{S} = \{G\}$) and $v_i(G) = b_i$ for all $i \in N$. This satisfies our budget normalisation condition and is consistent with all the valuation function profiles in $\mathcal{V'}$.

Note that for any valuation profile in $v \in \mathcal{V'}$, the best equilibrium allocation is where the top $\min\{n, k\}$ players get their favourite good. This allocation gives us a total value of

$$\sum_{i \in N} v_i(A_i^*) = \sum_{i=1}^{\min\{n,k\}} b_i = \min\{n, k\}b_1 - \sum_{i=1}^{\min\{n,k\}} \delta_i$$
$$\geq \min\{n, k\}b_1 - \delta b_1 \tag{9}$$

Suppose that allocation $\mathcal{A}$ allocates all the goods to one player. The maximum total utility that $\mathcal{A}$ can guarantee is b_1 and this arises when the entire bundle is allocated to player 1. Allocating the entire bundle to any other player will give us a strictly lower utility since all other players have a lower budget. This allocation gives us an efficiency

$$\min_{v \in \mathcal{V}} ER_v(\mathcal{A}) \leq \frac{b_1}{\min\{n, k\}b_1 - \delta b_1} = \frac{1}{\min\{n, k\} - \delta}$$

since the maximum utility that the optimal equilibrium allocation can get among all the valuation function profiles consistent with $\mathcal{S}$ is lower bounded by Equation (9).

If this is not the case and $\mathcal{A}$ allocates goods to more than one player, then we show that the maximum utility that $\mathcal{A}$ can guarantee is 0. Let $\mathcal{A}$ allocate non-empty bundles to players in $\{i_1, i_2, \ldots, i_{n'}\}$. Therefore, the bundles $\{A_{i_1}, A_{i_2}, \ldots, A_{i'_n}\}$ are non-empty. There exists a valuation function in $\mathcal{V'}$ such that the favourite good of i_1 is in A_{i_2}, the favourite good of i_2 is in A_{i_3} and so on till finally, the favourite good of $i_{n'}$ is in A_{i_1}. All the players in $\{i_1, i_2, \ldots, i_{n'}\}$ have different favourite goods here implying that all the players in $\{i_1, i_2, \ldots, i_{n'}\}$ which are in the top $\min\{n, k\}$ players have different favourite goods. For those players in the top $\min\{n, k\}$ who are not allocated any goods, we can set their favourite good such that no two players in the top $\min\{n, k\}$ have the same favourite good. This valuation profile is in $\mathcal{V'}$ and is consistent with $\mathcal{S}$. The optimal equilibrium utility in this case is non-zero trivially and therefore the efficiency guaranteed by this allocation is 0.

This means, given the set of samples and the set of budgets as defined above, we cannot guarantee an efficiency greater than $\frac{1}{\min\{n,k\}-\delta}$. This means that

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, b \in \mathcal{B}} ER_v(\mathcal{A}) \leq$$
$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} = \{G\}, \vec{b} = \vec{b'}} ER_v(\mathcal{A}) \leq \frac{1}{\min\{n, k\} - \delta}$$

This concludes the proof. $\qquad \square$

Lemma 4.6. *In unit demand markets with unequal budgets and strict preferences over items, any equilibrium allocation assigns player i the best possible available good, i.e. $\{g_i^*\}$ equals $\arg\max_{g \in \mathcal{G}_i} v_i(g)$ ($\mathcal{G}_1 = \arg\max_{g \in G} v_1(g)$ and for $i > 1$, $\mathcal{G}_i = G \setminus \{\bigcup_{l=1}^{i-1} \mathcal{G}_l\}$). Moreover, all equilibria have the same social welfare $\sum_i v_i(g_i^*)$.*

Proof. We prove this result by induction on i. Since player 1's budget, b_1, is the highest, in any equilibrium allocation they should be allocated item $g_1^* = \arg\max_{g' \in \mathcal{G}_1} v_1(g')$. Note that since players have strict preferences over items, g_1^* is unique. If player 1 is not allocated g_1^*, then either it is unallocated (and has a price of 0), in which case player 1 demands it,

contradicting that it is an equilibrium. Otherwise, it is allocated to another player j whose budget is $b_j < b_i$, in which case the price of g_1^* is less than b_1, and player 1 demands it. Assume that the claim holds for players $1, \ldots, i-1$, and consider player i. By the induction hypothesis, $g_i^* \notin \mathcal{G}_l$ for any $l < i$. If $g_i^* \neq \arg\max_{g \in \mathcal{G}_i} v_i(g)$ then $\arg\max_{g \in \mathcal{G}_i} v_i(g)$ is allocated to some other player $b_{i'}$ where $i' > i$ which is impossible because then player i can afford the good $\arg\max_{g \in \mathcal{G}_i} v_i(g)$. The above argument shows that any market equilibrium should assign g_i^* to player i and g_i^* is the best possible good which can be assigned to player i, which implies that for any equilibrium allocation $\mathcal{A}$, $v_i(A_i) = v_i(g_i^*)$. Therefore the social welfare for any equilibrium is $\sum_{i \in N} v_i(g_i^*)$ $\qquad\square$

Lemma 4.7. *Suppose that $\mathcal{D}$ is a product distribution such that for all $g \in G$, $1 - \sqrt{2e^{-1/k} - 1} < \Pr_{S \in \mathcal{D}}(g \in S) < \frac{1}{2} + \frac{\sqrt{2e^{-1/k}-1}}{2}$. If $|\mathcal{S}_i^t| \geq k^2$ (at the t-th iteration of the **while loop** in Algorithm 1 for player i), the corresponding B_i^t equals $\{\hat{g}_i\}$ to player i with at least $1 - e^{-\frac{k}{2}}$ probability, where*

$$\hat{g}_i \in \arg\max\{v_i(\{g\}) : g \in \bigcup_{S \in \mathcal{S}_i^t} S\}$$

Proof. Let $\hat{g}_i$ be the most valued good for player i in $\bigcup_{S \in \mathcal{S}_i^t} S$; we set $\alpha = \min_{g \in G} \Pr_{S \sim \mathcal{D}}(g \in S)$ and $\beta = \max_{g \in G} \Pr_{S \sim \mathcal{D}}(g \in S)$.

At the t-th iteration of the **while loop** in Algorithm 1 for player i; $B_i^t \neq \{\hat{g}_i\}$ if and only if:

(a) $\hat{g}_i$ is not present in any sample in $\mathcal{S}_i^t$, or (b) There exists another good $g' \neq \hat{g}_i$ which appears in all samples which contain $\hat{g}_i$, and does not appear in samples which do not contain $\hat{g}_i$ in $\mathcal{S}_i^t$. When this happens, $\{\hat{g}_i, g'\} \subseteq B_i^t$. Both these events occur with an exponentially low probability. The probability of not sampling $\hat{g}_i$ in $\geq k^2$ samples in $\mathcal{S}_i^t$ is

$$\leq (1-\alpha)^{k^2} \leq (\sqrt{2e^{-1/k} - 1})^{k^2} \leq (2e^{-1/k} - 1)^{k^2/2}$$
$$= 2^{k^2/2} e^{-k/2} \left(1 - \frac{1}{2e^{-1/k}}\right)^{k^2/2}$$
$$\leq 2^{k^2/2} e^{-k/2} \left(\frac{1}{2}\right)^{k^2/2} \leq e^{-k/2} \quad (\text{when } k \geq 3)$$

The probability that a good $g'(\neq \hat{g}_i)$ is present or absent together with $\hat{g}_i$ in a sample $S \in \mathcal{S}_i^t$ is

$$\Pr_{S \sim \mathcal{D}}(g' \in S) \Pr_{S \sim \mathcal{D}}(\hat{g}_i \in S)+$$
$$\left(1 - \Pr_{S \sim \mathcal{D}}(g' \in S)\right)\left(1 - \Pr_{S \sim \mathcal{D}}(\hat{g}_i \in S)\right) \tag{10}$$

The upper bound on (10) for any product distribution is $\gamma^2 + (1-\gamma)^2$ where $\gamma = \min(\alpha, 1-\beta)$. When given $\geq k^2$ samples in $\mathcal{S}_i^t$, the probability that a good $g'(\neq g^*)$ is present or absent together with $\hat{g}_i$ in all samples is

$$\leq (\gamma^2 + (1-\gamma)^2)^{k^2} = e^{-k^2 \log\left(\frac{1}{\gamma^2 + (1-\gamma)^2}\right)} \tag{11}$$

Equation (11) is $< e^{-k}$ only when $\gamma \in \left(\frac{1}{2} \pm \frac{\sqrt{2e^{-1/k}-1}}{2}\right)$ in particular

$$\min(\alpha, 1-\beta) > \frac{1}{2} - \frac{\sqrt{2e^{-1/k} - 1}}{2}$$

This implies that

$$\frac{1}{2} - \frac{\sqrt{2\mathrm{e}^{-1/k} - 1}}{2} < \Pr_{S\sim\mathcal{D}}(g \in S) < \frac{1}{2} + \frac{\sqrt{2\mathrm{e}^{-1/k} - 1}}{2}$$

Since this is true for all goods, using the union bound, the probability that $B_i^t = \{\hat{g}_i\}$ is $\leq (k-1) \cdot \mathrm{e}^{-k} \leq \mathrm{e}^{-k/2}$. □

Theorem 4.8. *Suppose that $\mathcal{D}$ is a product distribution, such that $\Pr_{S\sim\mathcal{D}}[g \in S] \in [\alpha, \beta]$. Assume that for every agent i, $|\{g \in G : v_i(g) > v_i(g_i^*)\}| < \frac{\max\{\log n, \log k\}}{\log(\frac{1}{1-\beta})}$ [2].*

If $k > 3$, $1 - \sqrt{2\mathrm{e}^{-1/k} - 1} \leq \alpha$ and $\beta \leq \frac{1}{2} + \frac{\sqrt{2\mathrm{e}^{-1/k}-1}}{2}$, the output of Algorithm 1, $(\mathcal{A}, \vec{p})$, satisfies

$$\Pr[ER_v(\mathcal{A}) = 1] \geq 1 - \frac{2n \max\{\log n, \log k\}}{\log\left(\frac{1}{(1-\beta)}\right)} \mathrm{e}^{-\frac{\max\{k,n\}}{4}}$$

Proof. Let $\alpha = \min_{g\in G} \Pr_{S\sim\mathcal{D}}(g \in S)$ and $\beta = \max_{g\in G} \Pr_{S\sim\mathcal{D}}(g \in S)$. For simplicity, let $\Psi = \max\{\log n, \log k\}/\log(\frac{1}{1-\beta})$

We claim that with $\max(k^4, n^2 k^2)$ samples, for all $i \in N$,

$$\Pr[\forall i' \leq i \; A_{i'} = \{g_{i'}^*\}] \geq 1 - 2i \cdot \Psi \mathrm{e}^{-\frac{\max\{k,n\}}{4}}$$

We prove our claim by induction on i. For player 1 Lemma 4.7 shows that we require $\geq k^2$ samples in $\mathcal{S}$ for algorithm 1 to allocate $\{g_1^*\}$ to player 1 with at least $1 - 1 \cdot \mathrm{e}^{-\frac{k}{4}}$ probability. By induction hypothesis, we assume that our claim is true for first $i-1$ players. In other words, when we compute A_i for player i, the set of already allocated items, *Alloc*, consists, with high probability, of $\{g_1^*, \ldots, g_{i-1}^*\}$. For player i, let g_i^* be the t−th preferred good. By Lemma 4.6, we know that $g_i^* \in \arg\max_{g\in G\backslash Alloc} v_i(g)$, and hence that player i's $t-1$ most favorite goods are in *Alloc*.

Let $g_i^{t'}$ be the t'-th most preferred good for player i. Let $B_i^{t'}$s correspond to $B_i^{t'}$ at the t' iteration of the ***while loop*** of Algorithm 1 (line 5). Now the probability that Algorithm 1 assigns good $\{g_i^*\}$ to player i is at least

$$\Pr[A_i = \{g_i^*\}] \geq \Pr\left[\bigcap_{t'=1}^t \{\text{The set } B_i^{t'} = \{g_i^{t'}\}\}\right] \tag{12}$$

$$= \prod_{t'=1}^t \Pr\left[B_i^{t'} = \{g_i^{t'}\} \mid \forall l < t' \; B_i^l = \{g_i^l\}\right]$$

The event in Equation (12) occurs when Algorithm 1 outputs $A_i = \{g_i^*\}$. As the first $(t-1)$ favorite goods of player i are in *Alloc*, Algorithm 1 allocates $\{g_i^t\}$ to player i which is the same as $\{g_i^*\}$.

Consider the case when $k > n$. For player i, if $B_i^l = \{g_i^l\}$ for $l < t'$, then by Lemma 4.7, if $|\mathcal{S}_i^{t'}| \geq k^2$ then with high probability $B_i^{t'} = \{g_i^{t'}\}$. The probability of sampling from a sample set of size k^2 at least one sample which does not contain $g_i^1, g_i^2, \ldots, g_i^{t'-1}$ in $\mathcal{S}_i^{t'}$ (i.e., that will remain in $\mathcal{S}_i^{t'}$) is at least $(1 - (1-\beta)^{t'-1})^{k^2}$. Since we assume $t' \leq t \leq \max\{\log n, \log k\}/\log(\frac{1}{1-\beta})$, the probability of sampling at least one sample that will remain in $\mathcal{S}_i^{t'}$ from a sample set of size at least k^2 is:

$$\leq (1 - (1-\beta)^{t'})^{k^2} \leq \left(1 - \frac{1}{k}\right)^{k^2} < \mathrm{e}^{-k}$$

---

[2]g_i^* is defined as in Lemma 4.6: $\{g_i^*\} = \arg\max_{g\in\mathcal{G}_i} v_i(g)$ ($\mathcal{G}_1 = \arg\max_{g\in G} v_1(g)$ and for $i > 1$, $\mathcal{G}_i = G \backslash \{\bigcup_{l=1}^{i-1} \mathcal{G}_l\}$).

Therefore, with $|\mathcal{S}| \geq k^4$ samples (which can be viewed as k^2 different sets of samples, each of size k^2), using the union bound, the probability that there are less than k^2 samples in $\mathcal{S}_i^{t'}$ is $\leq k^2 \mathrm{e}^{-k} \leq \mathrm{e}^{-k/4}$ (for $k > 3$). Which implies that for each $t' \leq t$;

$$\Pr\left[B_i^{t'} = \{g_i^{t'}\} \mid \forall l < t' \; B_i^l = \{g_i^l\}\right]$$
$$\geq (1 - \mathrm{e}^{-k/4})(1 - \mathrm{e}^{-k/2}) \geq 1 - 2\mathrm{e}^{-k/4}$$

similarly when $n > k$, with $n^2 k^2$ many samples, $\Pr\left[B_i^{t'} = \{g_i^{t'}\} \mid \forall l < t' \; B_i^l = \{g_i^l\}\right] \geq 1 - 2\mathrm{e}^{-n/4}$

Therefore the probability that Algorithm 1 assigns good $\{g_i^*\}$ to player i is

$$\geq \prod_{t'=1}^{t-1} \Pr\left[B_i^{t'} = \{g_i^{t'}\} \mid \forall l < t' \; B_i^l = \{g_i^l\}\right]$$
$$\geq (1 - 2\mathrm{e}^{-\max\{n,k\}/4})^t \geq 1 - 2\Psi\mathrm{e}^{-\frac{\max\{k,n\}}{4}}$$

Now by union bound for first $(i - 1)$ players and guarantees for player i we get,

$$\Pr[\forall i' \leq i \; A_{i'} = \{g_{i'}^*\}] \geq 1 - 2i\Psi\mathrm{e}^{-\frac{\max\{k,n\}}{4}}$$

Setting $i = n$ concludes the proof. $\qquad\square$

Proposition 4.9. *If $\mathcal{D}$ is a product distribution such that for all $g_j \in G$, $1 - \sqrt{2\mathrm{e}^{-1/k} - 1} < \Pr_{S \in \mathcal{D}}(g \in S) < \frac{1}{2} + \frac{\sqrt{2\mathrm{e}^{-1/k} - 1}}{2}$ and $k > 3$. Then, with exponentially high probability, Algorithm 1 allocates goods with an efficiency ratio $ER_v(\mathcal{A}) \geq \frac{\log n}{\rho n \log\left(\frac{1}{1-\beta}\right)}$ where $\rho = \max_{g \in G} \frac{\max_{i \in N} v_i(g)}{\min_{i \in N} v_i(g)}$ and $\beta = \max_{g_j \in G} \Pr_{S \in \mathcal{D}}(g \in S)$.*

Proof. Assume that player i gets good g_i^* in the optimal equilibrium allocation; let $\alpha = \min_{g \in G} \Pr_{S \sim \mathcal{D}}(g \in S)$ and let β be $\max_{g \in G} \Pr_{S \sim \mathcal{D}}(g \in S)$. Using a similar argument to that in Theorem 4.8, the probability that Algorithm 1 assigns g_i^* to player i for $i \leq \max(\log n, \log k)/\log(\frac{1}{1-\beta})$ is at least

$$1 - \frac{\max\{\log n, \log k\}^2}{\log^2\left(\frac{1}{(1-\beta)}\right)} \mathrm{e}^{-\frac{\max\{k,n\}}{4}}$$

The second part of the proof uses the above result to show efficiency bounds. The efficiency ratio $ER_v(\mathcal{A})$ is

$$\frac{\sum_{i=1}^n v_i(A_i)}{\sum_{i=1}^n v_i(A_i^*)} \geq$$
$$\frac{\sum_{i=1}^{\log n/\log(\frac{1}{1-\beta})} v_i(A_i^*)}{\sum_{i=1}^{\log n/\log(\frac{1}{1-\beta})} v_i(A_i^*) + \sum_{i=\log n/\log(\frac{1}{1-\beta})+1}^n v_i(A_i^*)}$$

Let us assume that the minimum utility achieved by any player among the first $\log n/\log(\frac{1}{1-\beta})$ players is c. This makes the $ER_v(\mathcal{A})$,

$$\geq \frac{(\log n/\log(\frac{1}{1-\beta}))c}{(\log n/\log(\frac{1}{1-\beta}))c + \sum_{i=\log n/\log(\frac{1}{1-\beta})+1}^n v_i(A_i^*)}$$

For the remaining players, the optimal utility is bounded by ρc since anything higher would violate the equilibrium condition. This is because if any remaining player (say i^*) receives a

20

bundle with value $> \rho c$, then some player (say i) with a higher budget who currently has a value of c for their allocated bundle will have a value $> c$ for the bundle allocated to player i^*. This violates the equilibrium condition since player i can afford A_{i^*} and strictly prefers A_{i^*} to their allocation. This implies that the efficiency is

$$\geq \frac{(\log n / \log(\frac{1}{1-\beta}))c}{(\log n / \log(\frac{1}{1-\beta}))c + (n - \log n / \log(\frac{1}{1-\beta}))\rho c}$$

$$= \frac{(\log n / \log(\frac{1}{1-\beta}))}{\rho n + (1-\rho)\log n / \log(\frac{1}{1-\beta})} > \frac{\log n}{\rho n \log(\frac{1}{1-\beta})}$$

$\square$

Corollary 4.10. *If the distribution $\mathcal{D}$ is uniform over the set 2^G and $k > 3$, with exponentially high probability. Algorithm 1 allocates goods with an efficiency $ER_v(\mathcal{A}) > \frac{\log n}{\rho n}$ where*

$$\rho = \max_{g \in G} \frac{\max_{i \in N} v_i(g)}{\min_{i \in N} v_i(g)}$$

using a polynomial number of samples.

Proof. This extends directly from Proposition 4.9. When $k \geq \frac{1}{\log 1.6}$, the lower bound and upper bound constraints on the probability of sampling each good in Proposition 4.9 improve such that the product distribution where $\Pr_{S \in \mathcal{D}}(g \in S) = \frac{1}{2} \quad \forall g \in G$ satisfies the constraints. This means we can directly apply the results of Proposition 4.9: the uniform distribution is a product distribution with $\alpha = \beta = \frac{1}{2}$ (which satisfy the constraints specified in Proposition 4.9). Thus, the efficiency ratio of the uniform distribution is at least $\frac{\log n}{\rho n \log 2} > \frac{\log n}{\rho n}$. $\square$

C Missing Proofs from Section 5

Lemma 5.1. *In Algorithm 2, no good in player 1's initially allocated sample gets taken away.*

Proof. The bundle that player 1 is allocated is either a sample or a subset of a sample. Let's call this parent sample S. The price of S is at least b_1 which is unaffordable to all other players and therefore, no other player can demand it. Any sample intersecting with this sample (say S') may be affordable to other players and the algorithm may burn a good from this sample. However, S' will have a good $g \notin A_1$ because the PreProcess step ensures that no samples are allocated which are proper supersets of other samples. This good either remains unallocated or is allocated to a player with lower budget. Either way, it gets burnt first to ensure consistency leaving the goods in A_1 untouched. $\square$

Theorem 5.3. *Let $\mathcal{S}, v(\mathcal{S})$ be a set of samples along with its valuations, $\mathcal{V}$ be the set of additive valuation function profiles which are consistent with the set of samples and are budget normalised i.e. $\max_{g \in G} v_i(\{g\}) = b_i$ for all the players $i \in N$ and $\mathcal{B} \subset \mathbb{R}_+^n$ be the set of all feasible budgets i.e. the set of all budgets in $\mathbb{R}_+^n$ such that $b_1 > b_2 > \cdots > b_n$. Then, we have*

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, \vec{b} \in \mathcal{B}} ER_v(\mathcal{A}) \leq \frac{1}{k - \delta}$$

for any $\delta \in (0, k)$ where $\mathcal{A}$ is a consistent allocation with respect to the samples.

Proof. Consider a market with n players and k goods. We divide this proof into two parts.

When $n \geq k$: Define a set of additive valuation function profiles $\mathcal{V}'$ as follows: each player has one good for which $v_i(g) = b_i$ and every other good has value 0 for this player. We refer to the good with non-zero valuation as the favourite good of player i. Also, let no two players in the top k players budget wise have the same favourite good. This set of valuations satisfies our budget normalisation condition.

Define the budget vector $\{b_1, b_2, \ldots, b_n\}$ as follows: for every player $b_i = b_1 - \delta_i$ where $\delta_1 = 0$, $0 < \delta_2 < \delta_3 < \cdots < \delta_n$ and $\delta_n = \frac{\delta b_1}{k}$. Let us call this vector of budgets $\vec{b'}$.

Now, suppose the only sample we have is the set of goods G ($\mathcal{S} = \{G\}$) and $v_i(G) = b_i$ for all $i \in N$. This is consistent with all the valuation function profiles in $\mathcal{V}'$.

Note that for any valuation profile $v \in \mathcal{V}'$, the best equilibrium allocation is where the top k players get their favourite good. This allocation gives us a total value of

$$
\begin{aligned}
\sum_{i \in N} v_i(A_i^*) &= \sum_{i=1}^{k} b_i \\
&= kb_1 - \sum_{i=1}^{k} \delta_i \\
&\geq kb_1 - \sum_{i=1}^{k} \delta_n \\
&= kb_1 - \delta b_1
\end{aligned}
\tag{13}
$$

Suppose that allocation $\mathcal{A}$ allocates all the goods to one player. The maximum total utility that $\mathcal{A}$ can guarantee is b_1 and this arises when the entire bundle is allocated to player 1. Allocating the entire bundle to any other player will give us a strictly lower utility since all other players have a lower budget. This allocation gives us an efficiency

$$
\min_{v \in \mathcal{V}} ER_v(\mathcal{A}) \leq \frac{b_1}{kb_1 - \delta b_1} = \frac{1}{k - \delta}
$$

since the maximum utility that the optimal equilibrium allocation can get among all the valuation function profiles consistent with $\mathcal{S}$ is lower bounded by Equation 13.

If this is not the case and $\mathcal{A}$ allocates goods to more than one player, then we show that the maximum utility that $\mathcal{A}$ can guarantee is 0. Let $\mathcal{A}$ allocate non-empty bundles to players in $\{i_1, i_2, \ldots, i_{n'}\}$. Therefore, the bundles $\{A_{i_1}, A_{i_2}, \ldots, A_{i_{n'}}\}$ are non-empty. There exists a valuation function in $\mathcal{V}'$ such that the favourite good of i_1 is in A_{i_2}, the favourite good of i_2 is in A_{i_3} and so on till finally, the favourite good of $i_{n'}$ is in A_{i_1}. All the players in $\{i_1, i_2, \ldots, i_{n'}\}$ have different favourite goods here implying that all the players in $\{i_1, i_2, \ldots, i_{n'}\}$ which are in the top k players have different favourite goods. For those players in the top k who are not allocated any goods, we can set their favourite good such that no two players in the top k have the same favourite good. This valuation profile is in $\mathcal{V}'$ and is consistent with $\mathcal{S}$. The optimal equilibrium utility in this case is non-zero trivially and therefore the efficiency guaranteed by this allocation is 0.

This means, given the set of samples and the set of budgets as defined above, we cannot guarantee an efficiency greater than $\frac{1}{k-\delta}$. This means that

$$
\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, b \in \mathcal{B}} ER_v(\mathcal{A}) \leq \min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} = \{G\}, \vec{b} = \vec{b'}} ER_v(\mathcal{A}) \leq \frac{1}{k - \delta}
$$

When $n < k$: Let G' be a subset of G such that $|G'| = n$. Define a set of additive valuation function profiles $\mathcal{V}'$ as follows: each player has one good in G' for which $v_i(g) = b_i$ and every

other good in G' has value 0 for this player. We refer to this good with non-zero valuation as the favourite good of player i. Also, let no two players have the same favourite good. Each of the goods in $G \setminus G'$ is valued by exactly one player in N at a value equal to their budget. Note that it is not necessary for all the goods in $G \setminus G'$ to be valued by the same player. All the valuations in $\mathcal{V}$ satisfy the budget normalisation property.

Define the budget vector $\{b_1, b_2, \ldots, b_n\}$ as follows: for every player $b_i = b_1 - \delta_i$ where $\delta_1 = 0$, $0 < \delta_2 < \delta_3 < \cdots < \delta_n$ and $\delta_n = \frac{\delta b_1}{k}$. Let us call this vector of budgets $\vec{b'}$.

Now, suppose the only sample we have is the set G' ($\mathcal{S} = \{G'\}$) and $v_i(G') = b_i$ for all $i \in N$. This is consistent with all the valuation function profiles in $\mathcal{V}'$.

Note that for any valuation profile $v \in \mathcal{V}'$, the best equilibrium allocation is where all the players get their favourite good and the goods in $G \setminus G'$ are given to the only player who values them at a non-zero value.

This allocation gives us a total value of

$$\sum_{i \in N} v_i(A_i^*) \geq \sum_{i=1}^{k} b_n$$
$$= kb_1 - k\delta_n$$
$$= kb_1 - \delta b_1 \tag{14}$$

Before we prove the highest utility a consistent allocation can guarantee, we first show that no allocation can guarantee any utility from any good in the set $G \setminus G'$ even when the valuation function profile is in $\mathcal{V}'$. If the allocation allocates all the goods in $G \setminus G'$ to one player (say i), there exists a valuation function profile with the same set of favourite goods where all the goods in $G \setminus G'$ is valued by some player $j \neq i$. If this is not the case and the allocation allocates the good in $G \setminus G'$ to multiple players (say $\{j_1, j_2, \ldots, j_{n'}\}$), then there exists a valuation function profile in $\mathcal{V}$ with the same favourite goods such that all the goods given to j_2 are valued by j_1, all the goods given to j_3 are valued by j_2 and so on till finally, all the goods given to j_1 are valued by $j_{n'}$. Either way, there exists a valuation function for which no good in $G' \setminus G$ provides any value. Therefore, we only need to look at the utility guaranteed by goods in G'.

Now, suppose that allocation $\mathcal{A}$ allocates all the goods in G' to one player. The maximum total utility that $\mathcal{A}$ can guarantee is b_1 and this arises when the entire bundle is allocated to player 1. This is because G' guarantees a utility of b_1 and $G \setminus G'$ cannot guarantee a non-zero utility. Allocating the entire bundle to any other player will give us a strictly lower utility since all other players have a lower budget. This allocation gives us an efficiency

$$\min_{v \in \mathcal{V}} ER_v(\mathcal{A}) \leq \frac{b_1}{kb_1 - \delta b_1} = \frac{1}{k - \delta}$$

since the maximum utility that the optimal equilibrium allocation can get among all the valuation function profiles consistent with $\mathcal{S}$ is lower bounded by Equation 14.

If this is not the case and $\mathcal{A}$ allocates goods in G' to more than one player, then we show that the maximum utility that $\mathcal{A}$ can guarantee is 0. Let $\mathcal{A}$ allocate non-empty subsets of G' to players in $\{i_1, i_2, \ldots, i_{n'}\}$. Therefore, the bundles $\{A_{i_1}, A_{i_2}, \ldots, A_{i'_n}\}$ are non-empty. There exists a valuation function in $\mathcal{V}'$ such that the favourite good of i_1 is in A_{i_2}, the favourite good of i_2 is in A_{i_3} and so on till finally, the favourite good of $i_{n'}$ is in A_{i_1}. All the players in $\{i_1, i_2, \ldots, i_{n'}\}$ have different favourite goods. For those players who are not allocated any goods, we can set their favourite good such that no two players have the same favourite good. Furthermore, we can choose a valuation function in $\mathcal{V}$ with these favourite goods such that no utility is guaranteed by the goods in $G \setminus G'$. The optimal equilibrium utility in this case is non-zero trivially and therefore the efficiency guaranteed by this allocation is 0.

This means, given the set of samples and the set of budgets as defined above, we cannot guarantee an efficiency greater than $\frac{1}{k-\delta}$. This means that

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, b \in \mathcal{B}} ER_v(\mathcal{A}) \leq$$

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S}=\{G'\}, \vec{b}=\vec{b'}} ER_v(\mathcal{A}) \leq \frac{1}{k-\delta}$$

This concludes the proof. $\qquad\square$

Proposition 5.4. *When all the samples in $\mathcal{S}$ are pairwise disjoint, then Algorithm 2 outputs an allocation with $ER_v(\mathcal{A}) \geq \frac{1}{\rho}$ where $\rho = \max_{g \in G} \frac{\max_{i \in N} v_i(\{g\})}{\min_{i \in N} v_i(\{g\})}$*

Proof. When all the samples are pairwise disjoint, all the goods are allocated and none of them are burnt. This is because, each player gets their favourite sample that has not been allocated yet. All samples that have been allocated to players with higher budgets are unaffordable to this player. Therefore, if v_g is the amount of utility gained by the player who is allocated good g in $\mathcal{A}$ and v_g^* refers to the same for allocation A^*, then

$$ER_v(\mathcal{A}) = \frac{\sum_{g \in G} v_g}{\sum_{g \in G} v_g^*} \geq \frac{\sum_{g \in G} v_g}{\sum_{g \in G} \rho v_g} = \frac{1}{\rho}$$

$\qquad\square$

D   Single Minded Markets

In single minded markets, each player has a particular bundle of goods, $D_i \subseteq G$ they desire; every bundle that does not contain D_i has no value i.e.

$$v_i(S) = \begin{cases} 1 & D_i \subseteq S \\ 0 & \text{otherwise.} \end{cases}$$

We show that a PAC underestimate for single-minded valuations can be efficiently learned, and an equilibrium for single-minded valuations can be efficiently computed. Therefore, using Proposition 3.1, a PAC Equilibrium is computable in polynomial time.

Proposition D.1. *The class of single minded valuation functions can be efficiently PAC learned, such that the learned valuation function weakly underestimates players' true valuations.*

Proof. From a given set of samples $\mathcal{S}$, set $\bar{D}_i = \bigcap_{S \in \mathcal{S}: v_i(S) > 0} S$. If, for a player $i \in N$, no sample has $v_i(S) > 0$, then set $\bar{D}_i = G$. This learned valuation is consistent and weakly lower than the actual valuations since $D_i \subseteq \bar{D}_i$ (i.e., a sample containing a set of items that is in D_i but not all of $\bar{D}_i$ will be given a value 0 instead of 1).

The total number of possible valuation functions, i.e., size of the hypothesis class, is 2^k (the number of possible choices for D_i). Thus, in order to PAC-learn D_i, we need a number of samples polynomial in $\frac{1}{\epsilon}$, $\log \frac{1}{\delta}$ and $\log |\mathcal{H}| \in \mathcal{O}(k)$ (a classic learning result for finite hypothesis classes, see Anthony and Bartlett, 1999). $\qquad\square$

Brânzei et al. [2016] present an algorithm to compute equilibria under equal budgets. We extend this algorithm to settings where all budgets are different.

Theorem D.2. *Algorithm 3 outputs a market equilibrium for single minded players with all different budgets.*

Algorithm 3: Competitive Equilibrium for Single Minded Valuations and Different Budgets

---

Notation: b_i^* is the remaining budget for player i; prices are represented by $\vec{p}$.

1: $\vec{p} = \vec{0}$; $\vec{b}^* = \{b_1, b_2, \ldots, b_n\}$
2: $\mathcal{B} = \{D_1, D_2, \ldots, D_n\}$
3: **for** *each* $g_j \in G$ **do**
4:     **if** g_j *is only demanded by one player* **then**
5:        Allocate g_j to that player at $p_j = 0$
6:     **end**
7:     **else if** g_j *is demanded by multiple players* **then**
8:        $p_j \leftarrow \texttt{SetPrice}(g_j, \vec{b}^*, \mathcal{B})$
9:        Allocate g_j to the player that can afford it at price p_j
10:       $\texttt{UpdateDemand}\,(\mathcal{B}, \vec{p}, \vec{b}^*)$
11:     **end**
12: **end**
13: Allocate all unallocated goods to player n at price 0
14: **Function** $\texttt{SetPrice}$ $(g_j, \vec{b}^*, \mathcal{B})$:
15:     $s = \arg\max_{i \in N \wedge g_j \in D_i} b_i^*$
16:     $t = \arg\max_{i \in N \setminus s \wedge g_j \in D_i} b_i^*$
17:     $p_j = b_t^* + \frac{b_s^* - b_t^*}{n^2}$
18:     $b_s^* = b_s^* - p_j$
19:     **while** $\exists i \neq s : b_i^* = b_s^*$ **do**
20:       $b_s^* = b_s^* - \frac{b_s^* - b_t^*}{n^2}$, $p_j = p_j + \frac{b_s^* - b_t^*}{n^2}$
21:     **end**
22:     **return** p_j
23: **Function** $\texttt{UpdateDemand}$ $(\mathcal{B}, \vec{p}, \vec{b}^*)$:
24:     **for** $i \in N$ **do**
25:       **if** $(D_i \neq \emptyset) \wedge (p(D_i) > b_i^*)$ **then**
26:         $D_i = \emptyset$
27:       **end**
28:     **end**

Proof. Algorithm 3 iteratively allocates goods while keeping track of players' remaining budgets. If a good is demanded by multiple players, it is priced such that only one player can afford it, and allocated to that player. The $\texttt{SetPrice}$ function ensures that no two players have the same remaining budget, by slightly increasing the price; this ensures that there are no ties when selecting the next player to allocate a good to.

All players either get their desired set or a subset of their desired set if it is unaffordable. Thus the resulting allocation is an equilibrium: players who do not receive their desired set are not able to afford it. □

The key difference between our approach and that of Brânzei et al. [2016] is how over-demanded goods are priced. Brânzei et al. [2016] assign the good to the player with the smallest desired set at a price equal to their budget. In our case, player budgets differ and therefore, ties cannot be broken by desired set size; rather, we instead break ties by remaining budgets.

Computing an equilibrium with total welfare at least K has been shown to be NP-hard by Brânzei et al. [2015] when players have equal budgets. In Theorem D.3, we show this for

our setting as well.

Theorem D.3. *It is NP-Complete to decide if a single minded market has an equilibrium with total welfare at least K*

Proof. We use a reduction from the NP-Complete problem SET PACKING:

> *Given a collection $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ of finite sets, all of which are a subset of a universal set $U = \{e_1, e_2, \ldots e_m\}$, and a positive integer $K \leq n$, does $\mathcal{C}$ contain at least K mutually disjoint sets? .*

Given a collection $\mathcal{C}$, a universal set U and an integer K, construct a market with n players $N = \{1, 2, \ldots, n\}$, $m + n$ goods $G = \{g_1, g_2, \ldots, g_{m+n}\}$. Let each player have an arbitrary non-zero budget b_i and desired set $D_i = C_i \cup \{g_{m+i}\}$. We show that there exists an equilibrium with total welfare at least K if and only if $\mathcal{C}$ has a disjoint collection of K sets.

If an equilibrium with total welfare at least K exists, then there are at least K players who receive their desired set. This means that the C_i's for all the players who receive their desired must be disjoint; otherwise, the equilibrium allocation would not be feasible. Therefore, there are at least K sets in $\mathcal{C}$ which are mutually disjoint.

If there are at least K sets which are mutually disjoint, we can construct an equilibrium as follows:

Assuming w.l.o.g. the sets $C_1, C_2, \ldots, C_K$ are mutually disjoint, for each $i \in \{1, 2, \ldots, K\}$, assign each good in the bundle C_i a non-zero price such that the total price is equal to b_i and price good g_{m+i} at zero. For each $i \in \{K + 1, K + 2, \ldots, n\}$, assign each good g_{m+i} a price equal to b_i. For all the goods whose prices have not been defined so far, set them to zero. Now, allocate the first K players their desired set and for every player $i \in \{K + 1, K + 2, \ldots, n\}$, allocate the good g_{m+i}. Lastly, assign all the remaining goods to player n. This allocation has a total welfare at least K since K players get their desired set. It is also easy to verify that the above allocation is an equilibrium since any player $i \in \{K + 1, K + 2, \ldots, n\}$ (who do not get their desired set) cannot afford their desired set. This concludes our proof. $\qquad\square$

Theorem D.5 shows that despite this, it is possible to compute a PAC equilibrium with an efficiency ratio $\geq \frac{1}{\min\{n,k\}}$. We now turn to establishing the efficeincy bounds of the algorithm.

Lemma D.4. *Algorithm 3 assigns at least one player its desired set.*

Proof. Let us first define a few terms which will help us with the proof. At any point in the algorithm, a player is *in the running* if they can afford their desired set, and is *eliminated* otherwise. In Algorithm 3, all players start out in the running and get eliminated as the prices increase. Once a player gets eliminated, they will stay that way till the end of the algorithm since prices of goods never decrease and therefore will never afford their bundle again; indeed, Algorithm 3 sets players' demands to $\emptyset$ once they are eliminated.

We prove inductively that before and after any good is allocated, at least one player is still in the running. We assume w.l.o.g. that goods are considered in the order $g_1, \ldots, g_k$.

For the first good g_1, all players are in the running before the allocation. There are three cases which determine how the good is allocated: if at most one player demands g_1 then prices remain 0, and all players are still in the running; if multiple players demand g_1, then one player is allocated g_1 and the remaining players are eliminated (lines 8-10); however, the player who received g_1 remains in the running. In both cases, at least one player remains in the running after the good is allocated. Now, let us assume this is true for goods $g_1, \ldots, g_{i-1}$. For g_i, there exists at least one player who is in the running before g_i is allocated by the inductive hypothesis. If no more than one player demands g_i then

the price of g_i is 0, and no player is eliminated; otherwise, all players who demand g_i get eliminated, except for the player who receives g_i, who is still in the running. Thus, there is at least one player in the running when we reach g_k. This player receives their desired set; otherwise, the allocation is not an equilibrium which contradicts Theorem D.2. □

Theorem D.5. *Let $(\mathcal{A}, \vec{p})$ be the output of Algorithm 3 on valuations learned as in Proposition D.1; then $ER_v(\mathcal{A}) \geq \frac{1}{\min\{n,k\}}$.*

Proof. From Lemma D.4, we get that at least one player will receive his desired set. This desired set is the learned desired set which is a superset of the actual desired set (see Proposition D.1). Therefore, the player who receives his learned desired set also receives his actual desired set. This means that the total welfare obtained is at least 1. The maximum welfare any allocation can obtain is $\min\{n,k\}$ since the total number of players getting their desired set is upper bounded by k and n. Thus, the efficiency of the computed PAC Equilibrium is $\geq \frac{1}{\min\{n,k\}}$ □

Similar to unit demand markets, we show that our result in Theorem D.5 is tight and no algorithm can guarantee a better efficiency.

Theorem D.6. *Let $\mathcal{S}, v(\mathcal{S})$ be a set of samples along with its valuations, $\mathcal{V}$ be the set of single minded valuation function profiles which are consistent with the set of samples and $\mathcal{B} \subset \mathbb{R}_+^n$ be the set of all feasible budgets i.e. the set of all budgets in $\mathbb{R}_+^n$ such that $b_1 > b_2 > \cdots > b_n$. Then, we have*

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, \vec{b} \in \mathcal{B}} ER_v(\mathcal{A}) \leq \frac{1}{\min\{n,k\}}$$

where $\mathcal{A}$ is a consistent allocation with respect to the samples.

Proof. Consider a market with n players and k goods. Define a set of single minded valuation function profiles $\mathcal{V}'$ as follows: the desired set of each player consists of only one good. This good is referred to as the desired good. Furthermore, let no two players in the top $\min\{n,k\}$ players budget wise have the same desired good.

Define the budget vector $\{b_1, b_2, \ldots, b_n\}$ as any budget vector such that $b_1 > b_2 > \cdots > b_n$. Let us call this vector of budgets $\vec{b}'$.

Now, suppose the only sample we have is the set of goods G ($\mathcal{S} = \{G\}$) and $v_i(G) = 1$ for all $i \in N$. This sample set is consistent with all the valuation function profiles in $\mathcal{V}'$.

Note that for any valuation profile $v \in \mathcal{V}'$, the best equilibrium allocation is where the top $\min\{n,k\}$ players get their desired good. This allocation gives us a total value of $\min\{n,k\}$.

Suppose that allocation $\mathcal{A}$ allocates all the goods to one player. The maximum total utility that $\mathcal{A}$ can guarantee is 1. This allocation gives us an efficiency

$$\min_{v \in \mathcal{V}} ER_v(\mathcal{A}) \leq \frac{1}{\min\{n,k\}}$$

since the maximum utility that the optimal equilibrium allocation can get among all the valuation function profiles consistent with $\mathcal{S}$ is lower bounded by $\min\{n,k\}$.

If this is not the case and $\mathcal{A}$ allocates goods to more than one player, then we show that the maximum utility that $\mathcal{A}$ can guarantee is 0. Let $\mathcal{A}$ allocate non-empty bundles to players in $\{i_1, i_2, \ldots, i_{n'}\}$. Therefore, the bundles $\{A_{i_1}, A_{i_2}, \ldots, A_{i'_n}\}$ are non-empty. There exists a valuation function in $\mathcal{V}'$ such that the desired good of i_1 is in A_{i_2}, the desired good of i_2 is in A_{i_3} and so on till finally, the desired good of $i_{n'}$ is in A_{i_1}. All the players in $\{i_1, i_2, \ldots, i_{n'}\}$ have different desired goods here implying that all the players in $\{i_1, i_2, \ldots, i_{n'}\}$ which are

in the top $\min\{n, k\}$ players have different desired goods. For those players in the top $\min\{n, k\}$ who are not allocated any goods, we can set their desired good such that no two players in the top $\min\{n, k\}$ have the same desired good. This valuation profile is in $\mathcal{V}'$ and is consistent with $\mathcal{S}$. The optimal equilibrium utility in this case is non-zero trivially and therefore the efficiency guaranteed by this allocation is 0.

This means, given the set of samples and the set of budgets as defined above, we cannot guarantee an efficiency greater than $\frac{1}{\min\{n,k\}}$. This means that

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} \subseteq 2^G, b \in \mathcal{B}} ER_v(\mathcal{A}) \leq$$

$$\min_{v \in \mathcal{V}} \max_{\mathcal{A}} \min_{\mathcal{S} = \{G\}, \vec{b} = \vec{b'}} ER_v(\mathcal{A}) \leq \frac{1}{\min\{n, k\}}$$

This concludes the proof. $\qquad\square$

E   Submodular Markets

In submodular markets, each player has monotone submodular valuations i.e. each player's valuation function $v_i : 2^G \mapsto \mathbb{R}^+ \cup \{0\}$ satisfies the following three conditions:

(a) $v_i(\emptyset) = 0$

(b) For any two $S, T \subseteq G$ such that $S \subseteq T$, $v_i(S) \leq v_i(T)$.

(c) For any two $S, T \subseteq G$,

$$v_i(S) + v_i(T) \geq v_i(S \cup T) + v_i(S \cap T) \tag{15}$$

The class of monotone submodular valuations contains the class of additive valuations, as well as many others. This increase in complexity comes with an even greater dearth of positive algorithmic results. In addition to this, monotone submodular valuations cannot be efficiently PAC learned [Balcan and Harvey, 2011]. So, we cannot use Proposition 3.1 to learn a PAC Equilibrium.

We, instead, use a direct learning approach similar to that of additive markets but modify our algorithm slightly due to two reasons. First, the pre-process step that worked for additive valuations will not work for submodular valuations since we cannot accurately determine the value of the bundle that results when you remove a subset from a set. However, we can underestimate it using equation (15) as follows: given two sets $A, B \subseteq G$ such that $A \subseteq B$, then by substituting $S = B \setminus A$ and $T = A$ in equation (15) we get

$$v_i(B \setminus A) \geq v_i(B) - v_i(A)$$

Therefore, $v_i(B) - v_i(A)$ gives us an underestimate of $v_i(B \setminus A)$. Furthermore, the inequality does not change if we replace $v_i(B)$ with an underestimate of $v_i(B)$.

Second, because we have to underestimate valuations, our efficiency guarantee may not hold. In order to prevent this, we modify our algorithm so that it can use extra information about the valuations. This is done using an additional input parameter c_i for all $i \in N$ which specifies an underestimate of the value of the highest valued good i.e., for all $i \in N$: $c_i \leq \max_{g \in G} v_i(\{g\})$. Note that when there is no available information about the value of c_i, we can set $c_i = 0$.

The algorithm has been described in Algorithm 4. The algorithm has the same three steps as that of Algorithm 2 but the first two steps are modified to work for submodular valuations.

The PreProcess step removes any supersets from the set $\mathcal{S}$ and replaces them with the set difference between the superset and the subset. It also computes the set of goods which could have a value $\geq c_i$ and stores it in the set F_i. Note that F_i is never empty and has a value of at least c_i to player i. The following lemma proves it.

Lemma E.1. *In the set $\mathcal{F}$ output by the PreProcess function of Algorithm 4, $F_i \neq \emptyset$ and $v_i(F_i) \geq c_i \quad \forall i \in N$.*

Proof. There exists at least one good g such that $v_i(\{g\}) \geq c_i$ by definition. Any sample with g will have value at least c_i by the monotone property.

Refer to the definition of F_i in Line 39 in Algorithm 4. If the good g is not present in any of the samples, then this good is included in F_i. If this good is present in any of the samples then this good will be present in $\bigcup_{S \in \mathcal{S}: v_i(S) \geq c_i} S$ and will not be present in $\bigcup_{S' \in \mathcal{S}: v_i(S') < c_i} S'$. Therefore, the good will be included in F_i.

Since all goods with $v_i(\{g\}) \geq c_i$ will be present in F_i and there is at least one good such that $v_i(\{g\}) \geq c_i$, the theorem follows immediately. $\square$

We then use this in the second step to give a player a bundle of value at least c_i when no other sample guarantees a value of at least c_i. Of course, this is not applicable when an element of F_i has been allocated to some other player.

The third step remains the same and ensures consistency since $\tilde{v}_i$ is an underestimate of v_i. So, if for any $S \in \mathcal{S}$, $v_i(S) > v_i(A_i)$, then, $v_i(S) > \tilde{v}_i(A_i)$.

We now show that when valuations are budget normalised, then the algorithm has an efficiency of at least $\frac{1}{k}$. But before we do that, we show that even in this algorithm, none of player 1's goods get taken away.

Lemma E.2. *In Algorithm 4, none of player 1's goods get taken away.*

Proof. If the first player is not allocated F_i, then he is allocated a subset of a sample or a sample. Let's call the parent sample S. No player can afford this sample since it has a price which is at least b_1. However, some players may be able to afford and prefer a sample (say S') which intersects with the allocated bundle. However, since the pre-process step ensures that no supersets are allocated, there will be at least one good in S' which is not allocated to the first player. This good either remains unallocated or is allocated to a player with lower budget. Either way, this good is burnt first to ensure consistency leaving the first player's allocated bundle intact.

If the first player is allocated F_i, this means that no set in $\mathcal{S}'$ can guarantee a value of at least c_i.

In such a scenario, any sample $S \in \mathcal{S}$ which contains a good $g \in F_i$ also contains a good $g' \notin F_i$. Assume for contradiction that this is not the case. Then there exists at least one sample which is a subset of F_i. Let S be minimal such that $S \subseteq F_i$ and $S \in \mathcal{S}$. This means (from the way we define F_i(Line 39)), $v_i(S) \geq c_i$.

S must have a subset $S' \in \mathcal{S}$ such that $v_i(S') < c_i$. This is because if it does not have a subset, then S will be in $\mathcal{S}'$ and $v_i(S) \geq c_i$ resulting in a contradiction. Furthermore, since S is minimal by our assumption, $v_i(S') < c_i$.

From the way we define F_i, since $v_i(S') < c_i$, we get $S' \cap F_i = \emptyset$. Since $S' \subsetneq S$, there are certain elements in S which are not present in F_i which is a contradiction.

Now, since any sample $S \in \mathcal{S}$ which contains a good $g \in F_i$ also contains a good $g' \notin F_i$, the good g' remains unallocated or belongs to a player with lower budget. Either way, it gets burnt first to ensure consistency leaving the goods in F_1 intact. $\square$

This brings us to our final proof. When we have budget normalised valuations, then Algorithm 4 gives us an allocation with efficiency at least $\frac{1}{k}$

Algorithm 4: Submodular Markets Consistent Allocation

Input: A set of samples $\mathcal{S}$, player valuations for these samples $v(\mathcal{S})$, budgets
$\quad\quad b_1 > b_2 > \cdots > b_n$ and $c_i \le \max_{g \in G} v_i(\{g\}) \forall i \in N$

1 $\mathcal{S}', \tilde{v}(\mathcal{S}'), \mathcal{F} = \texttt{PreProess}(\mathcal{S}, v(\mathcal{S}))$

2 for $i \leftarrow 1$ *to* n **do**

3 $\quad$ **if** $\tilde{v}_i(S') < c_i \ \forall S' \in \mathcal{S}' \wedge F_i \cap \bigcup_{j=1}^{i-1} A_j = \emptyset$ **then**

4 $\quad\quad$ Allocate F_i to player i i.e. $A_i = F_i, \tilde{v}_i(A_i) = c_i$

5 $\quad$ **end**

6 $\quad$ **else**

7 $\quad\quad$ $B_i = $ some set in $\arg\max_{T \in \mathcal{S}'} \tilde{v}_i(T)$

8 $\quad\quad$ Allocate B_i to player i i.e. $A_i = B_i, \tilde{v}_i(A_i) = \tilde{v}_i(B_i)$

9 $\quad$ **end**

10 $\quad$ Set $p_g = \frac{b_i}{|A_i|} \ \forall g \in A_i$

11 $\quad$ $\mathcal{S}' = \mathcal{S}' \setminus \bigcup_{S \in \mathcal{S}': S \cap A_i \ne \emptyset} S$

12 end

13 while $\exists i \in N, S \in \mathcal{S} \ s.t \ \tilde{v}_i(A_i) < v_i(S) \wedge \sum_{g \in S} p_g \le b_i$ **do**

14 $\quad$ **if** $\exists g \in S \ s.t. \ g \notin \bigcup_{i \in N} A_i$ **then**

15 $\quad\quad$ Set $p_g = \infty$

16 $\quad$ **end**

17 $\quad$ **else**

18 $\quad\quad$ $j = \arg\min_{i \in N: S \cap A_i \ne \emptyset} b_i$

19 $\quad\quad$ $g = $ any good in $A_j \cap S$

20 $\quad\quad$ $p_g = \infty$

21 $\quad\quad$ $A_j = A_j \setminus \{g\}$

22 $\quad\quad$ $\tilde{v}_j(A_j) = 0$

23 $\quad$ **end**

24 end

25 Allocate all leftover goods to player 1 at price 0

26 Function $\texttt{PreProess}(\mathcal{S}, v(\mathcal{S}))$:

27 $\quad$ $\mathcal{S}' = \mathcal{S}$

28 $\quad$ $\tilde{v}(\mathcal{S}') = v(\mathcal{S})$

29 $\quad$ **for** $S' \in \mathcal{S}'$ **do**

30 $\quad\quad$ **while** $\exists S \in \mathcal{S} \ s.t. \ S \subsetneq S'$ **do**

31 $\quad\quad\quad$ $\tilde{v}_i(S') = \tilde{v}_i(S') - v_i(S) \quad \forall i \in N$

32 $\quad\quad\quad$ $S' = S' \setminus S$

33 $\quad\quad$ **end**

34 $\quad$ **end**

35 $\quad$ **if** $\exists S', S'' \in \mathcal{S}' \ s.t. \ S' \subsetneq S''$ **then**

36 $\quad\quad$ Remove S' from $\mathcal{S}'$

37 $\quad$ **end**

38 $\quad$ $\mathcal{F} = \{F_1, F_2, \ldots, F_n\}$

39 $\quad$ $F_i = \left(\bigcup_{S \in \mathcal{S}: v_i(S) \ge c_i} S \setminus \bigcup_{S' \in \mathcal{S}: v_i(S') < c_i} S' \right) \cup \left(G \setminus \bigcup_{S \in \mathcal{S}} S \right) \quad \forall i \in N$

40 $\quad$ **return** $\mathcal{S}', \tilde{v}(\mathcal{S}'), \mathcal{F}$

Theorem E.3. *When $\max_{g \in G} v_i(\{g\}) = b_i$, then Algorithm 4 outputs an allocation with efficiency $ER_v(\mathcal{A}) \geq \frac{1}{k}$*

Proof. When valuations are budget normalised, we can set $c_i = b_i$ for every player. If there exists a sample in $\mathcal{S}'$ with utility at least b_1, then the first player will get allocated a sample with utility at least b_1. If not, then the first player will be allocated F_1 which has value at least b_1 (Lemma D.4). Using Lemma E.2, none of these goods are taken away from the first player and so his final utility will be at least b_1.

Since the maximum utility a good can give a player is upper bounded by b_1, the utility of the optimal equilibrium is upper bounded by kb_1. This gives us the following efficiency bound:

$$ER_v(\mathcal{A}) = \frac{\sum_{i \in N} v_i(A_i)}{\sum_{i \in N} v_i(A_i^*)} \geq \frac{v_1(A_1)}{\sum_{i \in N} v_i(A_i^*)} \geq \frac{b_1}{kb_1} = \frac{1}{k}$$

□

Since additive valuations are a subset of monotone submodular valuations, Theorem 5.3 applies in this case as well. This means the bound in Theorem E.3 is tight.

Omer Lev
Ben-Gurion University of the Negev
Beersheba, Israel
Email: `omerlev@bgu.ac.il`

Neel Patel
University of Southern California
Los Angeles, CA, USA
Email: `neelbpat@usc.edu`

Vignesh Viswanathan
University of Massachusetts, Amherst
Amherst, MA, USA
Email: `vviswanathan@umass.edu`

Yair Zick
University of Massachusetts, Amherst
Amherst, MA, USA
Email: `yzick@umass.edu`