

Do Prices Coordinate Markets?

Justin Hsu* Jamie Morgenstern† Ryan Rogers‡ Aaron Roth§ Rakesh Vohra¶

June 20, 2016

Abstract

Walrasian equilibrium prices have a remarkable property: they allow each buyer to purchase a bundle of goods that she finds the most desirable, while guaranteeing that the induced allocation over all buyers will globally maximize social welfare. However, this clean story has two caveats: first, the prices may induce indifferences. In fact, the minimal equilibrium prices *necessarily* induce indifferences. Accordingly, buyers may need to coordinate with one another to arrive at a socially optimal outcome—the prices alone are *not* sufficient to coordinate the market; second, although natural procedures converge to Walrasian equilibrium prices on a fixed population, in practice buyers typically observe prices without participating in a price computation process. These prices cannot be perfect Walrasian equilibrium prices, but instead somehow reflect distributional information about the market.

To better understand the performance of Walrasian prices when facing these two problems, we give two results. First, we propose a mild genericity condition on valuations under which the minimal Walrasian equilibrium prices induce allocations which result in low over-demand, no matter how the buyers break ties. In fact, under genericity the over-demand of any good can be bounded by 1, which is the best possible at the minimal prices. We demonstrate our results for unit demand valuations and give an extension to matroid based valuations (MBV), conjectured to be equivalent to gross substitute valuations (GS).

Second, we use techniques from learning theory to argue that the over-demand and welfare induced by a price vector converge to their expectations uniformly over the class of all price vectors, with respective sample complexity linear and quadratic in the number of goods in the market. These results make no assumption on the form of the valuation functions.

These two results imply that under a mild genericity condition, the exact Walrasian equilibrium prices computed in a market are guaranteed to induce both low over-demand and high welfare when used in a new market where agents are sampled independently from the same distribution, whenever the number of agents is larger than the number of commodities in the market.

*Department of Computer and Information Science, University of Pennsylvania. Email: justhsu@cis.upenn.edu
Supported in part by NSF grants CCF-1101389, TWC-1513694 and CNS-1065060 and by a grant from the Simons Foundation (#360368)

†Warren Center Postdoctoral Fellow, Departments of Computer and Information Science and Economics, University of Pennsylvania. Email: jamiemmt@cs.cmu.edu

‡Department of Applied Mathematics and Computational Science, University of Pennsylvania. Email: ryrogers@sas.upenn.edu

§Department of Computer and Information Sciences, University of Pennsylvania. Email: aaroth@cis.upenn.edu
Supported in part by an NSF grant CCF-1101389 and an Alfred P. Sloan Research Fellowship.

¶Economics Department, University of Pennsylvania. Email: rvohra@seas.upenn.edu

Contents

1	Introduction	2
2	Model	4
3	Lower bound	6
4	Unit demand	7
4.1	Swap Graph	8
4.2	Generic Valuations	9
4.3	Over-Demand	9
4.4	Welfare	11
5	Towards gross substitutes	14
5.1	Swap Graph with GS Valuations	15
5.2	Matroid Based Valuations	19
5.3	Structural results for MBV	20
5.4	Generic MBV	23
5.5	Properties of swap graph with GMBV	24
5.6	Bounding the Buyer In-Degree	28
5.7	Bounding welfare	32
6	Welfare and over-demand generalization	33
6.1	Results for Arbitrary Valuations	34
6.2	Learning Theory Preliminaries	36
6.3	Learning from Pricings and Learning with Compression	37
6.4	Over-Demand Concentration	39
6.5	Welfare Concentration	40
6.6	Lower bounds for learning	43
	References	45
A	Valuations over bundles of copies versus bundles of goods	48
B	Non-minimal Walrasian prices and genericity	49
C	Minimal demand bundles form a matroid basis	49
D	Constructing genericity via perturbation	50
E	Proof of Lemma 4.12	53
F	Omitted Learning Proofs	54
F.1	Proof of Lemma 6.12	54
F.2	The existence of an encodable tie-breaking rule for GMBV	56
F.3	Shattering arguments	57
F.4	Omitted concentration proofs	57

1 Introduction

The power of Walrasian equilibrium is often summarized with a pithy slogan: prices coordinate markets. However, this is not exactly true—a Walrasian equilibrium specifies a price for each good *and* an assignment of goods to buyers. The assignment is just as important as the prices, since there can be multiple bundles of goods which maximize a buyer’s utility at given prices. If buyers select arbitrarily among these bundles, they may over-demand some goods.

One way to avoid this problem is to assume strictly concave valuation functions ensuring that each buyer always has a unique, utility maximizing bundle of goods. However, this idea does not apply in many natural economic settings; for instance, when goods are indivisible. For another possible solution, perhaps we could use Walrasian prices that eliminate the coordination problem. For example, with a single good and distinct valuations, a price *strictly* between the largest and second largest buyer valuations will eliminate indifferences. However, how do we arrive at such prices? Relying on a coordinator with full knowledge of the market to solve the coordination problem defeats the purpose of markets.

In this work, we consider a different approach: natural assumptions on buyers to simplify the coordination problem. We focus our investigation on *minimal* Walrasian equilibrium prices, which result from many natural market dynamics¹ [Kelso and Crawford, 1982]. We aim to allocate bundles of m types of indivisible goods g , each with some supply s_g , to n buyers who have matroid based Valuations² and quasi-linear preferences; the assignment (“unit demand”) model is an important special case. We detail our model in Section 2.

We begin our technical results in Section 3, by showing that indifferences at the minimal Walrasian prices can be a serious problem—goods can be in the demand correspondence of *every* buyer, possibly leading to all buyers demanding the same good. Clever tie-breaking does not help; for any tie-breaking rule, the induced over-demand can be as large as $\Omega(n)$. (We provide a simple example in Section 3). Even worse, we observe that the minimal Walrasian equilibrium prices always induce indifferences.

Hence, we cannot hope to rule out all over-demand at the minimal prices. But, we give a “genericity” condition on buyer valuations that is the next-best thing: the over-demand for each good g will be at most 1, independent of its supply s_g and the tie-breaking strategy used by buyers. Therefore, as the supply grows, worst-case over-demand becomes negligible. We warm up with the assignment model in Section 4, where our condition is simple to state: buyer valuations for goods should be linearly independent over the coefficients $\{-1, 0, 1\}$. To follow, we generalize our techniques to the matroid based valuations case in Section 5. The situation is significantly more complicated, but the core genericity definition and proof strategy are in the same spirit.

After we show that exact minimal Walrasian prices generically induce low over-demand, a natural question is whether this property holds when the same prices are used on “similar buyers”; we turn to this question in Section 6. More formally, imagine a sample N_1 of n buyers drawn from an unknown distribution Π of buyer valuations. The goods are priced using the minimal Walrasian equilibrium prices, computed from the valuations of buyers in N_1 . Now, keeping these prices fixed, we draw a fresh sample N_2 of n buyers from Π , who each choose some bundle from their demand

¹Minimal Walrasian prices are focal in other ways: e.g., in matching markets they correspond to VCG prices.

²“Matroid based valuations” are a structured subclass of *gross substitutes* valuations. In fact, Ostrovsky and Paes Leme [2015] conjecture that the class of matroid based valuations is *equal* to the class of gross substitutes valuations, the largest class of valuations for which Walrasian equilibrium prices are guaranteed to exist.

correspondence at the given prices (breaking ties arbitrarily). Will the over-demand and welfare on N_2 be close to the over-demand and welfare on N_1 ?

If the supply s_g of a good g is small, the difference in over-demand between N_1 and N_2 may be large when compared to s_g . However, we show that if $s_g \geq \tilde{O}(m/\epsilon^2)$, the demand for any good g on sample N_2 will be within a $1 \pm \epsilon$ factor of the supply s_g of good g . Note that the supply requirement is independent of the market size n . Similarly, if the optimal welfare for N_1, N_2 is large, the induced welfare of the two markets will be within a $1 \pm \epsilon$ of one another (and within a $1 - \epsilon$ factor of the optimal welfare for N_2).

Furthermore, we are able to prove these bounds *without* any assumption on the structure of the valuation functions. This lack of structure makes it difficult to argue directly about notions of combinatorial dimension like VC dimension, and so we take a different approach which may be of independent interest. Using a recent compression argument of [Daniely and Shalev-Shwartz \[2014\]](#), we show that assuming fixed but unknown prices, the class of functions *predicting* a buyer’s demanded bundle at those prices is learnable using $\tilde{O}(m/\epsilon^2)$ many samples. Because this is a multi-class learning problem, learning does not imply uniform convergence. However, the binary problem of predicting demand for a *particular* good is a 1-dimensional projection of the bundle prediction problem, and hence is also learnable with the same number of samples. By a classical result of [Ehrenfeucht et al. \[1989\]](#), learning and uniform convergence have the same sample complexity in binary prediction problems. So, we can bound the VC dimension, and thus the sample complexity for uniform convergence for demand. Moreover, our bound is tight—even for unit demand buyers, the VC dimension of the class of demand predictors is $\Omega(m)$.

Welfare, unlike demand, corresponds to a real-valued prediction problem, and so the sample complexity needed for uniform convergence cannot be bounded by bounding the sample complexity of learning. Instead, we directly bound the pseudo-dimension of the class of welfare predictors by $\tilde{O}(m^2)$, again without making any assumptions about the form of the valuation functions. We show that if the optimal welfare is $\tilde{\Omega}(m^4\sqrt{n}/\epsilon^2)$, the welfare induced by the Walrasian prices p for N_1 when applied to N_2 is within a $1 - \epsilon$ factor of optimal.

Related work We follow a long line of work on understanding how markets behave under limited coordination. If buyers’ valuations are strictly concave and items are divisible, [Arrow and Debreu \[1954\]](#) show that there exist item prices p such that each agent has a unique utility-maximizing bundle at p , and that when each agent selects her unique utility-maximizing bundle the market clears. With indivisible items, anonymous equilibrium item pricings may not exist; even when they do, finding such prices is NP-hard [[Deng et al., 2002](#)]). [Mount and Reiter \[1974\]](#) consider the size of the message space needed to compute Walrasian equilibria, and [Nisan and Segal \[2006\]](#) show that polynomial communication is sufficient. This stands in sharp contrast to the situation for submodular buyers, where exponential communication is needed to compute prices which support an efficient allocation [[Nisan and Segal, 2006](#)].³

Our work is also related to the growing area of learning for mechanism design, where a mechanism is selected from some class of mechanisms as a function of sampled buyers. Recent work measures the sample complexity of revenue maximization in the single-parameter [[Blum et al., 2003](#), [Balcan et al., 2005](#), [Blum and Hartline, 2005](#), [Goldberg et al., 2006](#), [Elkind, 2007](#), [Balcan et al., 2007, 2008](#), [Hartline and Roughgarden, 2009](#), [Babaioff et al., 2011](#), [Cesa-Bianchi et al., 2013](#),

³the submodular setting, prices may be over arbitrary bundles and individualized to each buyer, in contrast to the gross substitutes setting, where efficient allocations can be supported by anonymous item pricings.

Dhangwatnotai et al., 2014, Cole and Roughgarden, 2014, Medina and Mohri, 2014, Huang et al., 2015, Roughgarden and Schrijvers, 2015, Morgenstern and Roughgarden, 2015] and multi-parameter settings [Devanur et al., 2011, Dughmi et al., 2014].

In particular, sample complexity results for revenue and welfare maximization using item pricings are known. Most papers consider buyers that make decisions sequentially, avoiding over-demand from uncoordinated resolution of indifferences. In the unlimited supply setting (in which over-demand cannot arise), Balcan et al. [2008] show how to learn approximately revenue-optimal prices with polynomial sample complexity using a covering argument. Extending this work, Balcan et al. [2007] show how to handle large but limited supply settings and to handle welfare maximization. These papers imply that demand concentration will have convergence rates of $\tilde{O}(\frac{k^2}{\epsilon} \ln \frac{1}{\delta})$; our result improves this by a factor of k (where k is the number of distinct types of goods). This percolates into the guarantees regarding welfare concentration, in that our concentration result requires a weaker lower bound on the supply of each good. Our precise convergence rate for welfare concentration is looser than theirs in several parameters, in part because we allow buyers to break indifferences in a worst-case way. Using similar techniques to Balcan et al. [2007], we could recover similar convergence rates for welfare. Devanur and Hayes [2009a] also consider prior-independent pricing for revenue maximization with unit demand buyers in an online setting. Similarly, for the more restricted setting of unit-demand valuations in the online adwords matching setting, Devanur and Hayes [2009b] show a similar generalization guarantee for prices. This is generalized to arbitrary valuations by aga [2014], whose convergence rate depends logarithmically on n and linearly on k ; Molinaro and Ravi [2013] removes this dependence on n but has quadratic dependence on k . Our convergence rate is independent of n and depends only linearly on k , and are therefore faster than any of these previous results. We thank Nikhil Devanur for pointing out this related line of work.

We achieve these tighter convergence rates using compression schemes to derive several uniform convergence results; this compression tool was first used in the context of game theory by Balcan et al. [2014], who use it directly to upper bound the PAC complexity of a learning problem rather than to imply uniform convergence over a class.

Our definition of genericity for unit demand valuations has also been called *independence* by Braverman et al. [2014], Ashlagi et al. [2009]. In fact, our bound of over-demand for generic unit demand valuations can be obtained from results due to Braverman et al. [2014], although over-demand was not the focus of their work. However, our analysis extends to more general valuation functions.

Our genericity results rely on a novel *swap graph* construction, which bears a slight resemblance to the exchangability graph of Murota [1996b]. The exchangability graph also has nodes defined by goods, but without considering a Walrasian allocation and pricing. In contrast, our swap graph is models indifferences at equilibrium.

2 Model

We consider a market with m indivisible goods, where good g has supply $s_g \geq 1$. We will write the set of goods as $[m] = \{1, \dots, m\}$ and denote the bundles of the m goods as $\mathcal{G} = 2^{[m]}$. The market will also have a set N of n buyers, where each buyer demands at most one copy of each good.⁴ For simplicity, we consider valuations defined over subsets of goods rather than arbitrary sets of

⁴More formally, copies of a good beyond the first copy have marginal valuation zero.

copies; see Appendix A for a formal treatment. For each buyer $q \in N$, let $v_q : \mathcal{G} \rightarrow [0, H]$ denote q 's valuation function. We will assume v_q is monotone and $v_q(\emptyset) = 0$. Our goal is to produce a feasible assignments of goods to buyers.

Definition 2.1 (Allocation). *An allocation $\mu : N \rightarrow \mathcal{G}$ assigns each buyer a bundle such that the whole assignment is feasible:*

$$\sum_{q \in N} \mathbf{1}\{g \in \mu(q)\} \leq s_g \quad \text{for all } g \in [m].$$

As is typical, we consider quasi-linear utility functions $u_q : \mathcal{G} \times \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}_{\geq 0}$ defined by

$$u_q(S; p) = v_q(S) - \sum_{g \in S} p_g \quad \text{for all } q \in N \quad \text{and} \quad S \in \mathcal{G}.$$

We will consider prices assigned to goods—each *copy* of a good has the same price. We will write prices as vectors $p = (p_g)_{g \in [m]} \in \mathbb{R}_{\geq 0}^m$ or as functions over bundles $p : \mathcal{G} \rightarrow \mathbb{R}_{\geq 0}$ such that $p(S) = \sum_{g \in S} p_g$.

Definition 2.2 (Demand correspondence). *The demand correspondence for buyer $q \in N$ at prices p is $\mathcal{D}_q(p) = \operatorname{argmax}_{S \in \mathcal{G}} \{u_q(S; p)\}$. We call bundles $S \in \mathcal{D}_q(p)$ demand bundles. Note that $\mathcal{D}_q(p)$ contains only bundles with non-negative utility, since $u_q(\emptyset; p) = 0$ for every p .*

We focus our investigation on *Walrasian equilibria*, defined by a pricing and an allocation.

Definition 2.3 (Walrasian equilibrium). *For valuations $\{v_q\}_{q \in N}$, we say that a pair (p, μ) of prices p and allocation $\mu = (\mu_q)_{q \in N}$ is a Walrasian equilibrium (WE) if both:*

- $\mu_q \in \mathcal{D}_q(p)$ for all $q \in N$; and
- $p_g = 0$ implies $\sum_{q \in N} \mathbf{1}\{g \in \mu(q)\} < s_g$.

We call the price vector p a Walrasian equilibrium price vector. Note that there may be many distinct Walrasian equilibrium price vectors: in fact, the set of all Walrasian prices forms a lattice. The minimum Walrasian equilibrium price vector p is the Walrasian equilibrium price vector that is coordinate-wise minimal amongst all Walrasian equilibrium price vectors.⁵

Likewise, we call the allocation μ a Walrasian allocation. While there may be multiple distinct Walrasian allocations, it is known that all such allocations must maximize welfare.

In general, Walrasian equilibrium prices p are not sufficient to coordinate a corresponding allocation μ because buyers might have indifferences ($|\mathcal{D}_q(p)| > 1$). If buyers choose their bundle arbitrarily, the resulting allocation can violate supply constraints. To measure the amount of violation, we make the following two natural definitions.

Definition 2.4 (Demanders and over-demand). *The set $U(g; p)$ of demanders for a good $g \in [m]$ at price p is the set of buyers that have some demand set containing g :*

$$U(g; p) = \{q \in N : \exists D \in \mathcal{D}_q(p) \quad \text{where } g \in D\}.$$

⁵The fact that the price vectors form a lattice guarantees the existence the minimal price.

Then, the over-demand $OD(g; p)$ for g at prices p is the number of demanders beyond the supply of a good:

$$OD(g; p) = \max\{|U(g; p)| - s_g, 0\}.$$

That is, the over-demand is the worst-case excess demand if bidders break ties in their demand correspondence arbitrarily.

To build intuition, we focus the first part of our paper on unit demand bidders, where

$$v_q(S) = \operatorname{argmax}_{g \in S} \{v_q(g)\} \quad \text{for all } S \in \mathcal{G},$$

and non-empty bundles $S \in \mathcal{D}_q(p)$ are *demand goods*.⁶

3 Lower bound

To build intuition for why tie breaking at equilibrium can lead to infeasibility, we give an example of a market with n buyers in which the over-demand of a good can be $\Omega(n)$ regardless of how buyers break ties, so long as they cannot coordinate with one another after seeing the market instance.⁷

Lemma 3.1. *There exist unit demand valuations such that at the minimal Walrasian prices p , some good has over-demand $\Omega(n)$.*

Proof. Consider a market with n unit demand buyers $N = [n]$ and $m = n$ distinct goods. For a distinguished good $g \in [m]$, every buyer q has valuation $v_q(q) = v_q(g) = 1$ and $v_q(g') = 0$ for all $g' \notin \{q, g\}$. The minimal Walrasian equilibrium prices are $p = \mathbf{0}$, and the unique max-welfare allocation is $\mu_q = q$. At these prices, g is demanded by every buyer. Hence, $OD(g; p) = n - 1$. \square

Note that if buyers resolve indifferences uniformly at random, $n/2$ buyers will attempt to buy good g . While one might hope that the right choice of tie breaking rules would solve this problem, we show that it cannot. Formally, we will suppose that each buyer q has a tie breaking rule e_q^p which can depend on v_q and the price vector p . Then, we consider the over-demand of a good when all bidders use their tie breaking rules.

Definition 3.2 (Tie-breaking over-demand). *Given a set of prices p , buyers N and tie breaking rule e_q^p for buyer q , let the demanders of g be*

$$U^e(g; p) = \{q \in N : g \in \mathcal{D}_q(p) \text{ and } g = e_q^p(\mathcal{D}_q(p))\},$$

and the tie breaking over-demand with respect to e^p be

$$OD^e(g; p) = \max\{|\{q \in U^e(g; p)\}| - s_g, 0\}.$$

Without loss of generality (by the min-max principle), it suffices to consider deterministic tie breaking rules when constructing a randomized lower bound instance.

⁶Any demand bundle with $|S| > 1$ must have some good $g \in S$ for which $v(S) = v(g)$, by the definition of unit demand valuations.

⁷This lack of coordination is formalized by requiring that the tie breaking rule buyer q uses to select among favorite bundles must be independent of the valuations of buyers $q' \neq q$.

Lemma 3.3. *There exists a distribution over unit demand valuations such that for any set of tie breaking rules, the expected tie breaking over-demand from n buyers is $\Omega(n)$.*

Proof. The distribution we construct will contain instances similar to the one from Lemma 3.1, with n unit demand buyers and $m = n$ distinct goods. We choose a permutation σ over the goods $[m]$ and a distinguished good $g^* \in [m]$ uniformly at random. Once σ and g^* are chosen, we define buyer valuations such that for all buyers q , $v_q(\sigma(q)) = v_q(g^*) = 1$ and $v_q(h) = 0$ for all $h \notin \{\sigma(q), g^*\}$. In this market, the minimal Walrasian equilibrium prices is $p = \mathbf{0}$ and the unique max-welfare allocation has $\mu_q = \sigma(q)$ for all $q \in N$. At equilibrium, there is a single buyer q^* for whom $\sigma(q^*) = g^*$ and $\mathcal{D}_{q^*}(p) = \{g^*\}$. For every other buyer, $\sigma(q) \neq g^*$ and q has exactly two goods in their demand correspondence: $\mathcal{D}_q(p) = \{\sigma(q), g^*\}$. Each such buyer will select a good $e_q^p(\{\sigma(q), g^*\}) \in \{\sigma(q), g^*\}$ using their tie breaking rule.

Over the randomness of the instance (σ and g^*) we have $\mathbb{P}[e_q^p(\{\sigma(q), g^*\}) = g^*] = 1/2$. We can then lower bound the expected tie breaking over-demand for the distinguished good as

$$\begin{aligned} & \mathbb{E}_{\sigma, g^*} [OD^e(g^*; p)] \\ &= \sum_{q \in N} \mathbb{P}[\sigma(q) = g^*] \\ &+ \mathbb{P}[\sigma(q) \neq g^*] \cdot \mathbb{P}[e_q^p(\{\sigma(q), g^*\}) = g^*] - \underbrace{1}_{\text{supply}} \\ &= \frac{1}{2} \cdot n(1 - 1/n) = \frac{n-1}{2} \end{aligned}$$

which completes the proof. \square

While this result shows that over-demand can be high without coordination, it seems rather artificial—the buyers’ valuations are extremely similar. We will soon give a simple and natural condition which rules out this example, and more generally ensures that the over-demand for any good is at most 1 at the minimal Walrasian prices regardless of tie breaking. This bound is the best possible: for *any* buyer valuations, minimal Walrasian prices always induce over-demand of at least 1 for every good with positive price.

Lemma 3.4. *Fix any set of buyer valuations $\{v_q : q \in [n]\}$, and let p be a minimal Walrasian equilibrium price vector. For any good g with positive price $p(g) > 0$, we have $OD(g; p) \geq 1$.*

Proof. Let μ be a Walrasian allocation for p . By the Walrasian equilibrium condition, $|q : g \in \mu_q| = s_g$ for any good with $p(g) > 0$. Suppose that $OD(g; p) = 0$, i.e., g is in some bundle in buyer q ’s demand set if and only if $g \in \mu_q$. In this case there exists $\epsilon > 0$ such that if we set $p(g) \leftarrow p(g) - \epsilon$, (p, μ) remains a Walrasian equilibrium—the allocation is unchanged, and every buyer continues to receive an allocation in their demand set. But this contradicts minimality of p . \square

4 Unit demand

Now that we have seen how indifferences at equilibrium can lead to over-demand, we consider whether the over-demand is large for “typical” instances. To build intuition, we start with the special case of unit demand valuations, where $v_q(S) = \operatorname{argmax}_{g \in S} \{v_q(g)\}$. Such valuations can be

encoded with a real number for each pair of buyer $q \in N$ and good $g \in [m]$, denoted $v_q(g)$. Thus, without loss of generality an allocation is a many-to-one matching between buyers and goods—each buyer should be matched to at most 1 good, and each good g should be matched to at most s_g buyers.

We now give conditions on unit demand valuations to ensure that when buyers buy *arbitrary* (singleton) demand sets from their demand correspondence at the minimal Walrasian prices, the resulting allocation has high welfare and low over-demand. Accordingly, we need to reason precisely about how the equilibrium prices depend on the valuations. Getting access to this relation is surprisingly tricky—typical characterizations of Walrasian equilibrium prices are not enough for our needs. For instance, two standard characterizations show that unit demand Walrasian prices i) are dual variables to a particular linear program (the “many-to-one matching linear program”), and ii) are computed from ascending price auction dynamics. The first observation reduces computing prices to an optimization problem, but it does not provide fine-grained information about how the prices depend on the valuations. The second observation is useful for computing prices, but the auction may proceed in a complicated manner, obscuring the relationship between the prices and valuations.

4.1 Swap Graph

Accordingly, we define a graph called the *swap graph* of a Walrasian equilibrium (p, μ) . This graph directly encodes buyer indifferences induced by the equilibrium prices. Furthermore, the swap graph allows us to read off equations involving the prices and valuations. We define the swap graph as follows.

Definition 4.1 (Swap graph). *The swap graph $G = (V, E)$ defined with respect to a Walrasian equilibrium (p, μ) has a node for each good g and an additional null node \perp representing the empty allocation: $V = [m] \cup \{\perp\}$. There is a directed edge $(a, b) \in E$ for $a \neq b, b \neq \perp$ for each buyer q that receives good a in μ but also demands b , i.e. if $\mu_q = a$ and $b \in \mathcal{D}_q(p)$ for some buyer $q \in N$. Note that while there may be parallel edges—representing the same indifferences by different buyers—there are no self loops.*

Since we will phrase our arguments in terms of the swap graph in the remainder of the section, we will first recast Lemma 3.4 using our new language.

Corollary 4.2. *For buyers with unit demand valuations and a minimal Walrasian equilibrium (p, μ) , every node in the swap graph G with in-degree zero has price zero.*

Almost by definition; the over-demand of a good g is its in-degree in the swap graph.

Lemma 4.3. *Let G be the swap graph corresponding to a Walrasian equilibrium. If a node g in G has in-degree d , then $OD(g; p) \leq d$.*

Proof. By construction of the swap graph, a node with in-degree d corresponds to a good g with d buyers with g in their demand correspondence but not in their allocation. Because μ is a feasible allocation, at most s_g buyers can be allocated good g in μ , and since μ is an equilibrium allocation, g is in the demand correspondence for each of these buyers. Thus, there can be at most $s_g + d$ demanders for good g . By definition of over-demand, we have $OD(g; p) \leq d$. \square

So to bound the maximum over-demand of any good, it suffices to bound the in-degree for every good in the swap graph. While the in-degree may be large in the worst case, we can introduce a simple condition on valuations that will rule out these pathological market instances.

4.2 Generic Valuations

Recall that in Section 3, we showed that over-demand can be high at the minimal Walrasian equilibrium prices. So, to provide a better bound on over-demand, we need additional assumptions; ideally, a condition that will hold “typically”. In the lower bound instance from Section 3, over-demand is large because the buyers have valuations that are too similar. Indeed, consider a market with two goods a and b where all buyers have the same *difference* in valuations between a and b . If *some* buyer is indifferent between a and b —by Lemma 3.4, this must be the case at minimal Walrasian prices—*all* buyers are indifferent. This observation motivates our genericity condition.

Definition 4.4 (Generic valuations). *A set of valuations $\{v_q(g) \in \mathbb{R} : q \in N, g \in [m]\}$ is generic if they are linearly independent over $\{-1, 0, 1\}$, i.e.*

$$\sum_{q \in N} \sum_{g \in [m]} \alpha_{q,g} v_q(g) = 0 \quad \text{for } \alpha_{q,g} \in \{-1, 0, 1\} \quad \text{implies } \alpha_{q,g} = 0 \text{ for all } q \in N, g \in [m].$$

Remark 4.5. *Note that this condition holds with probability 1 given any continuous perturbation of a profile of valuation functions, and so for many natural distributions, a profile of valuation functions will “generically” (i.e., with high probability) satisfy our condition. We also show how to discretely perturb a fixed set of valuations to satisfy our condition in Appendix D.*

Remark 4.6. *Our definition of generic in the unit demand setting is also called independence by Ashlagi et al. [2009], Braverman et al. [2014], although over-demand was not the focus in these works. In fact, Braverman et al. [2014] give an alternative way to bound over-demand for generic unit demand valuations (Theorem 4.10). We present the unit-demand case using the swap graph construction, in order to generalize smoothly to broader classes of valuations.*

4.3 Over-Demand

Now, we are ready to present the main technical result of this section: When buyers with generic valuations select an *arbitrary* good in their demand correspondence given minimal Walrasian equilibrium prices, over-demand is low and welfare is high. We will show that the *in-degree* for any node in our swap graph is at most 1. This will imply that no good has over-demand more than 1, regardless of its supply. We proceed via a series of properties of the swap graph. First, under genericity, the swap graph is acyclic.

Lemma 4.7. *The swap graph G , defined with respect to Walrasian equilibrium (p, μ) and generic valuations $\{v_q(g) : q \in N, g \in [m]\}$, is acyclic.*

Proof. Since the null node \perp has no incoming edges by construction, it cannot be part of any cycle. Thus, suppose that there is a cycle $a_0 \rightarrow a_1 \rightarrow \dots \rightarrow a_k \rightarrow a_0$ of non-null nodes. We label the buyers so that buyer q_i is allocated a_i in the allocation μ , but also has good a_{i+1} in their demand correspondence as well (taking the subscript modulo $k + 1$) for $i = 0, \dots, k$.

By construction, all buyers are distinct and $k \geq 1$. Furthermore, each edge represents an indifference relationship for some buyer. In particular,

$$v_q(a_q) - p(a_q) = v_q(a_{q+1}) - p(a_{q+1}) \quad \text{for all } q = 0, \dots, k.$$

Summing these equations and canceling prices, we have

$$\sum_{i=0}^k v_{q_i}(a_i) = \sum_{i=0}^k v_{q_i}(a_{i+1}),$$

contradicting the genericity assumption since all buyers q are distinct. Hence, the swap graph must be acyclic. \square

Because the swap graph is acyclic, we can choose a partial order of the nodes so that all edges go from smaller nodes to larger nodes. For the remainder of the argument, we assume nodes are labeled by such an ordering (i.e. we now have for every edge $(a_i, a_j) \in E$ implies $i < j$). Now, the price of a good can be written in terms of the valuations of smaller goods.

Lemma 4.8. *For every good g , the price $p(g)$ can be written as a linear combination of valuations $v_q(j)$ over $\{-1, 0, 1\}$ for $q \in N$ and $j < g$. Specifically, for every set of goods $g_1 < g_2 < \dots < g_k < g$ such that $g_1 \rightarrow \dots \rightarrow g_k \rightarrow g$ forms a path in the graph and g_1 has in-degree 0, there are buyers q_1, \dots, q_k such that $\mu_{q_i} = g_i$ and*

$$p(g) = \sum_{i=1}^k (v_{q_i}(g_{i+1}) - v_{q_i}(g_i))$$

where $g_{k+1} = g$. If the first node $g_1 = \perp$, we define $v_{q_1}(\perp) = 0$.

Proof. We proceed by induction on the path length k . In the base case, $k = 1$ and we have path $g_1 \rightarrow g$. If $g_1 = \perp$ then buyer q_1 must be indifferent between the empty allocation and good g , so $p(g) = v_{q_1}(g)$. If $g_1 \neq \perp$, then buyer q_1 's indifference between g_1 and g yields

$$v_{q_1}(g_1) - p(g_1) = v_{q_1}(g) - p(g),$$

so $p(g) = v_{q_1}(g) - v_{q_1}(g_1) + p(g_1)$. Note that g_1 has in-degree zero so Corollary 4.2 implies $p(g_1) = 0$. This shows the base case.

For the inductive case, we assume that we can write prices in our desired form for any good with a path of length at most $k - 1$ from a source node. Consider a good g that has path length k from a source node. Because buyer q_k is indifferent between g and g_k ,

$$v_{q_k}(g_k) - p(g_k) = v_{q_k}(g) - p(g)$$

so $p(g) = v_{q_k}(g) - v_{q_k}(g_k) + p(g_k)$. Note that $p(g_k)$ has a path length of at most $k - 1$ from a source node, and so we can apply the induction hypothesis to get complete the induction:

$$p(g) = v_{q_k}(g) - v_{q_k}(g_k) + \sum_{i=1}^{k-1} (v_{q_i}(g_{i+1}) - v_{q_i}(g_i)). \quad \square$$

Finally, we can bound the in-degree of any node under genericity.

Lemma 4.9. *For generic buyer valuations, every node in the swap graph defined with respect to a Walrasian equilibrium (p, μ) with minimal Walrasian prices has in-degree at most 1.*

Proof. Suppose otherwise and let good g be the smallest indexed good with at least two incoming edges. Then there are two sequences of goods

$$g_1 \leq \dots \leq g_k \leq g \quad \text{and} \quad g'_1 \leq \dots \leq g'_{k'} \leq g$$

such that $g_k \neq g'_{k'}$, $g_1 \rightarrow \dots \rightarrow g_k \rightarrow g$ and $g'_1 \rightarrow \dots \rightarrow g'_{k'} \rightarrow g$ form paths in the graph, and g_1 and g'_1 have in-degree 0. We will write $g = g_{k+1} = g'_{k'+1}$. By Lemma 4.8, we can express $p(g)$ in two distinct ways:

$$p(g) = \sum_{i=1}^k (v_{q_i}(g_{i+1}) - v_{q_i}(g_i))$$

and

$$p(g) = \sum_{i=1}^{k'} (v_{q'_i}(g'_{i+1}) - v_{q'_i}(g'_i)),$$

where buyer q_i has $\mu_{q_i} = g_i$ and also has good g_{i+1} in her demand set, and buyer q'_i has $\mu_{q'_i} = g'_i$ and has good g'_{i+1} in her demand set. Taking the difference, we have

$$\sum_{i=1}^k (v_{q_i}(g_{i+1}) - v_{q_i}(g_i)) - \sum_{i=1}^{k'} (v_{q'_i}(g'_{i+1}) - v_{q'_i}(g'_i)) = 0.$$

Since we either have $q_k \neq q_{k'}$ or $g_k \neq g'_{k'}$, the above linear combination is not trivial and contradicts genericity. \square

Finally, by Lemma 4.3 the over-demand for any good is at most its in-degree in the swap graph, so we have bounded over-demand under generic valuations.

Theorem 4.10. *For any set of unit demand buyers with generic valuations and for p the minimal Walrasian equilibrium price vector, the over-demand for any good $g \in [m]$ is at most 1.*

As a result, when generic buyers face minimal⁸ prices p and buy a good in their demand set while resolving indifferences arbitrarily, the excess demand of any good is at most 1.

4.4 Welfare

Now that we have considered over-demand under genericity, what about welfare? If buyers break ties arbitrarily, it is not hard to see that welfare may be very bad: buyers who are indifferent between receiving a good and receiving nothing may all decide to demand nothing, giving zero welfare. However, if we simply rule out this specific kind of indifference, we can show that genericity implies near-optimal welfare.

⁸We show in Appendix B that even with generic valuations, *non-minimal* Walrasian prices can still induce high over-demand, further justifying our focus on minimal Walrasian prices.

If B is not a feasible allocation, we assume $\text{Welfare}_N(B)$ is calculated by resolving over-demand in a worst-case way, as Lemma 4.12 describes below. As an intermediate step in our calculations we bound a relaxed notion of welfare, which *assumes supply is sufficient to satisfy the demand of all buyers*. Given bundles B_1, \dots, B_n , we define the *relaxed* welfare of this (pseudo-)allocation as follows.

Definition 4.11 (Relaxed Welfare). *Given n bundles B_1, \dots, B_n , the relaxed welfare for a market N is $\text{RWelfare}_N(B_1, \dots, B_n) = \sum_{q \in N} v_q(B_q)$.*

Note that because B_1, \dots, B_n may not be a feasible allocation, $\text{RWelfare}_N(B_1, \dots, B_n)$ can in principle be larger than the optimal welfare obtainable over feasible allocations in the market over buyers N , which we denote Opt-Welfare_N . However, subject to our genericity conditions, we need only consider allocations that over-allocate any good by at most 1. For such allocations, the difference between $\text{RWelfare}_N(B)$ and the welfare of a corresponding feasible solution is small (and in general, if the over-demand can be bounded by a small quantity, the gap between welfare and relaxed welfare is small). Furthermore, if B is a feasible allocation then $\text{Welfare}_N(B) = \text{RWelfare}_N(B)$.

Lemma 4.12. *Consider a set of bundles B_1, \dots, B_n such that $|\{q : g \in B_q\}| \leq s_g + d$ for all g and some $d \leq s_g$. We consider two possible ways to convert $\{B_i\}$ into a feasible allocation $\{\widehat{B}_i\}$.*

- *Let $\sigma : N \rightarrow [n]$ be an ordering of the buyers. For any buyer $q \in N$ with bundle B_q , if for some $g \in B_q$ more than s_g buyers q' who precede q in the ordering demand g , let $\widehat{B}_q = \emptyset$. Else, let $\widehat{B}_q = B_q$. (i.e. we let buyers choose bundles in order while supply remains)*

Then, $\widehat{B}_1, \dots, \widehat{B}_n$ is feasible, and

$$\text{RWelfare}_N(B_1, \dots, B_n) \leq \text{Welfare}_N(\widehat{B}_1, \dots, \widehat{B}_n) - d \cdot m \cdot H.$$

- *For each g , randomly select a subset N_g of s_g buyers with $g \in B_q$; let $g \in \widehat{B}_q$ if $q \in N_g$. (i.e. independently for each good, allocate its supply uniformly at randomly to its set of demanders)*

Then, $\widehat{B}_1, \dots, \widehat{B}_n$ is feasible, and if buyers are subadditive:

$$\left(1 - \frac{d}{d + \min_g s_g}\right) \cdot \text{RWelfare}_N(B_1, \dots, B_n) \leq \mathbb{E} \left[\text{Welfare}_N(\widehat{B}_1, \dots, \widehat{B}_n) \right].$$

The proof of Lemma 4.12 can be found in Appendix E.

For the welfare argument only, we assume that a buyer chooses a *nonempty* demand bundle whenever possible. Subject to this restriction, buyers can break ties however they like. Then, we can lower bound the welfare of these buyers N in terms of the optimal welfare Opt-Welfare_N .

Theorem 4.13. *Consider any set of buyers N with generic unit demand valuations bounded in $[0, H]$ and minimal Walrasian equilibrium prices p . For each buyer q , let $b_q \in \mathcal{D}_q(p)$ be some arbitrary set in their demand correspondence, assuming only that $b_q \neq \emptyset$ whenever $|\mathcal{D}_q(p)| > 1$. Then the welfare obtained by the resulting allocation is nearly optimal:*

$$\text{Welfare}_N(b_1, b_2, \dots, b_n) \geq \text{Opt-Welfare}_N - 2 \cdot m \cdot H.$$

Proof. Let μ be a Walrasian allocation for prices p and, thus, a welfare-optimal allocation so that $\text{Opt-Welfare}_N = \sum_{q \in N} v_q(\mu_q)$. We also know that $\mu_q \in \mathcal{D}_q(p)$ for all buyers q , by the properties of Walrasian equilibria. Because we also have $b_q \in \mathcal{D}_q(p)$ for every q , we know

$$\text{Opt-Welfare}_N - \sum_{g \in [m]} p(g) s_g = \sum_{q \in N} [v_q(\mu_q) - p(\mu_q)] = \sum_{q \in N} [v_q(b_q) - p(b_q)].$$

Rearranging,

$$\sum_{q \in N} [v_q(b_q) - p(b_q)] = \sum_{q \in N} v_q(b_q) - \sum_{g \in [m]} p(g) s_g - \sum_{g \in [m]} p(g) \cdot od(g, b) + \sum_{g \in [m]} p(g) \cdot ud(g, b)$$

where $od(g, b)$ and $ud(g, b)$ are defined to be the over-demand and under-demand, respectively, of good g in allocation b compared to the optimal allocation μ :

$$od(g, b) = \max \left\{ \sum_{q \in N} [\mathbf{1}\{g = b_q\} - \mathbf{1}\{g = \mu_q\}], 0 \right\}$$

$$ud(g, b) = \max \left\{ \sum_{q \in N} [\mathbf{1}\{g = \mu_q\} - \mathbf{1}\{g = b_q\}], 0 \right\}$$

Note that $od(g, b) \geq 0$, and that minimal Walrasian prices are at most the maximum value H , and hence we can rearrange the above equations to obtain

$$\text{Opt-Welfare}_N \leq \sum_{q \in N} v_q(b_q) + \sum_{g \in [m]} ud(g, b) \cdot H.$$

Hence, it suffices to upper bound the total under-demand induced over all goods g by b .

Next, note that by assumption, $b_q = \emptyset$ if and only if $\mathcal{D}_q(p) = \{\emptyset\}$, in which case it must also be that $\mu_q = \emptyset$. Hence, we have that $|\{q : b_q = \emptyset\}| \leq |\{q : \mu_q = \emptyset\}|$. We then have

$$\sum_{q \in N} |\mu_q| \leq \sum_{q \in N} |b_q| = \sum_{g \in [m]} (|\mu_q| + od(g, b) - ud(g, b)), \quad \text{and so} \quad \sum_{g \in [m]} ud(g, b) \leq \sum_{g \in [m]} od(g, b).$$

Under genericity, Theorem 4.10 implies $od(g, b) \leq 1$, and hence

$$\sum_{g \in [m]} ud(g, b) \leq \sum_{g \in [m]} od(g, b) \leq m.$$

Thus, we obtain:

$$\text{Opt-Welfare}_N \leq \sum_{q \in N} v_q(b_q) + m \cdot H = \text{RWelfare}_N(b_1, \dots, b_n) + m \cdot H$$

Finally, applying Lemma 4.12 completes the proof. \square

Generalization Results for Unit Demand Buyers We now state our main generalization results specialized to unit demand buyers. Section 6 states the analogous generalization theorems for buyers with arbitrary valuation functions, and provides the proofs of the general statements. Fixing some tie-breaking rule $e : 2^{\mathcal{G}} \times \mathcal{V} \rightarrow \mathcal{G}$ for choosing a demanded set from their demand correspondence;⁹ we will use $C^e(g; p; N) = \sum_{q \in N} \mathbf{1}[g \in e(\mathcal{D}_q(p), v_q)]$ to denote the *number* of copies of good g demanded in market N at prices p , when buyer q uses $e(\cdot, v_q)$ to break ties. Our first theorem bounds the over-demand for g on N' .

Theorem 4.14. *Fix some pricing p and two sampled markets N, N' of unit demand buyers, with $|N| = |N'| = n$. Then, fixing a good g , if $C^e(N; g; p) \leq s_g + 1$, with probability $1 - \delta$,*

$$C^e(g; p; N') - s_g \leq O\left(\sqrt{m \cdot s_g \ln \frac{1}{\delta}} + m \ln \frac{1}{\delta}\right).$$

Thus, for any $0 < \alpha < 4/5$, if $s_g = \Omega(m \ln \frac{1}{\delta} / \alpha^2)$, we have

$$C^e(g; p; N') \leq (1 + \alpha)s_g.$$

Our second theorem guarantees that the exact Walrasian prices p computed on a market N will induce approximately optimal welfare when used on a new set of buyers N' sampled from the same distribution. Let $\text{Welfare}_N(p)$ denote the welfare buyers in N achieve at prices p , assuming over-demand for a good g is resolved in a worst-case way for welfare.

Theorem 4.15. *Fix two sampled markets $N, N' \sim \Pi$ of unit demand buyers for which $|N| = |N'| = n$. For any $\alpha \in (0, 4/5)$, if p are welfare-optimal prices for N and*

$$\text{Opt-Welfare}_N = \Omega\left(\frac{H^3 n^{.5} m^2 \log^4(m) \ln^2(m) \ln^2 \frac{1}{\delta}}{\alpha^2}\right)$$

then, with probability $1 - \delta$,

$$\text{Welfare}_{N'}(p) \geq (1 - \alpha) \text{Opt-Welfare}_{N'}.$$

5 Towards gross substitutes

Now that we have seen our techniques in the unit demand setting, we will generalize our results to more general valuations. Ideally, we would like to extend our results to buyers who have *gross substitutes* valuations.

Definition 5.1 (Gross substitutes). *A valuation function v_q satisfies gross substitutes (GS) if for every price vectors $p' \geq p$ and $S \in \mathcal{D}_q(p)$, there is a bundle $S' \in \mathcal{D}_q(p')$ with*

$$S \cap \{g \in [m] : p(g) = p'(g)\} \subseteq S'.$$

Note that unit demand valuations also satisfy GS.

⁹For this section, assume only that $e(\mathcal{X}) \neq \emptyset$ unless $\mathcal{X} = \{\emptyset\}$.

If all buyers have GS valuations, there always exist Walrasian equilibria, and the minimal equilibrium can be found by following a tatônement procedure [Kelso and Crawford, 1982]. Gul and Stacchetti [1999] show that the class GS is in some sense the most general class of valuations that are guaranteed to have WE.

While GS is an attractive class to target, its definition is axiomatic rather than constructive; Definition 5.1 shows us what properties a GS valuation must satisfy, rather than how to concretely construct GS valuations. This poses a problem for defining genericity: it is not obvious which, if any, valuations in GS satisfy a candidate definition of genericity, so it may be that generic valuations do not even have Walrasian equilibria!

Ultimately, we will prove our results for buyers with *matroid based valuations* (MBV), a subclass of gross substitutes valuations that is conjectured to be equal to all of gross substitutes [Ostrovsky and Paes Leme, 2015]. Such valuations can be explicitly constructed from a set of numeric weights and a matroid, giving us a path to define generic valuations and certify that they are contained in GS.

However, there are still several obstacles to overcome, and the arguments are significantly more involved than for unit demand. Roughly, the central difficulty is establishing a connection between valuations on bundles and valuations on the items in the bundle, since we will require genericity on values for single goods. While this connection is immediate in the case of unit demand—favorite bundles are simply single goods—the situation for MBV is more delicate. For instance, some bundles may contain “irrelevant” goods, which don’t contribute at all to the valuation. Hence, we propose a more complex swap graph to capture the richer structure of the valuations.

Despite the complications, our high-level argument parallels the unit demand case. We will first define a swap graph and connect it to over-demand. Then, we define a generic version of MBV (GMBV) and prove properties about the swap graph for valuations in GMBV. We only restrict to MBV when defining genericity; our arguments in the first step apply to GS valuations.

5.1 Swap Graph with GS Valuations

To define the swap graph for gross substitute valuations, the core idea is to have an edge (a, b) represent a single *swap* of good a for good b in some larger bundle, naturally generalizing our construction for unit demand. The main challenge is ensuring that we faithfully model buyer indifference—which are between *bundles* of goods—via indifferences of single swaps. More precisely, in order to bound over-demand with arbitrary tie-breaking, we must ensure that if a buyer is indifferent between her Walrasian allocation and some other bundle B , then there must be at least one incoming edge to every good in B . If the swap graph satisfies this property, we can describe potential demand for goods in terms of in-degree of nodes like we did for unit demand.

At first, we might hope that if a buyer is indifferent between bundles B_1 and B_2 , every good in B_1 can be exchanged for a good in B_2 while preserving utility; this would ensure that each good in B_2 has an incoming edge. While this property turns out to not be true for general bundles, it is true for the “smallest” bundles in a buyer’s demand correspondence.

Definition 5.2 (Minimal demand correspondence). *For price vector p and buyer $q \in N$, the minimum demand correspondence is*

$$\mathcal{D}_q^*(p) = \{S \in \mathcal{D}_q(p) : T \notin \mathcal{D}_q(p) \text{ for all } T \subsetneq S\}.$$

We call bundles $B \in \mathcal{D}_q^(p)$ minimum demand bundles, or minimum bundles for short.*

Crucially, the minimum demand correspondence for any buyer with a GS valuation forms the bases of a matroid [Ben-Zwi et al., 2013]. (We give a self-contained proof in Appendix C.)

Lemma 5.3 (Ben-Zwi et al. [2013]). *Let $v_q : \mathcal{G} \rightarrow [0, H]$ be a GS valuation. For any price vector p , the minimum demand correspondence $\mathcal{D}_q^*(p)$ forms the set of bases of some matroid.*

By standard facts from matroid theory [Oxley, 1997], all elements of $\mathcal{D}_q^*(p)$ are the same size. More importantly, the bases satisfy the exchange property: for B_1, B_2 in the basis set \mathcal{B} of a matroid and for every $b \in B_1 \setminus B_2$, there exists $b' \in B_2 \setminus B_1$ such that

$$B_1 \cup b' \setminus b \in \mathcal{B} \quad \text{and} \quad B_2 \cup b \setminus b' \in \mathcal{B}.$$

Thus if all bundles in a buyer's demand correspondence have the same size, then we can take \mathcal{B} to be $\mathcal{D}_q^*(p)$ and we have exactly what we need: for any two bundles in the correspondence, every good can be swapped for some other good while remaining in the correspondence. Unfortunately, buyers may prefer bundles of different sizes. We will now introduce the swap graph, and then explain how it models buyers who prefer bundles of different sizes.

Definition 5.4 (Swap graph). *Let buyers have GS valuations, (p, μ) be a WE, and for each buyer $q \in N$ fix a minimum demand bundle $M_q \in \mathcal{D}_q^*(p)$ where $M_q \subseteq \mu_q$. Define the swap graph $G(p, \mu, (M_q)_{q \in N})$ to have a node for every good $g \in [m]$ and an additional node which we refer to as the null node \perp . There is a directed edge from (a, b) for every buyer $q \in N$ such that $a \in M_q$, $b \notin \mu_q$, and there exists $B \in \mathcal{D}_q^*(p)$ with $b \in B$ where*

$$M_q \cup b \setminus a \in \mathcal{D}_q^*(p).$$

Further, the graph contains an edge from \perp to good b for each buyer $q \in N$ with $b \notin \mu(q)$ such that (i) $b \in B$ with $B \in \mathcal{D}_q^(p)$ and $B \setminus b \in \mathcal{D}_q^*(p)$, and (ii) b has positive price.*

There are two main conditions when specifying edges from the null good: (i) involves demand bundles, while (ii) requires positive price. These two conditions address two distinct problems.

The first condition models cases where the demand correspondence has bundles of different sizes. Suppose there is some bundle B strictly larger than the minimum cardinality demand bundle. We need all goods in B to have an incoming edge, to reflect over-demand if a buyer selects B . We cannot ensure these edges via matroid properties, since B is not a minimum cardinality demand bundle and so it is not a basis of the matroid. However, there is always a minimum bundle B^{\min} contained in B , which is a matroid basis, so we have incoming edges to B^{\min} .

The following interpolation lemma shows that for every good $g \in B \setminus B^{\min}$ the bundle $B^{\min} \cup g$ is also a demand bundle, so the swap graph will also have edges from \perp to g , covering all goods in B as desired.

Lemma 5.5. *Consider a buyer with a GS valuation. If we have two bundles B, B' such that $B \subseteq B'$, and B, B' are both bundles in the buyer's demand correspondence at price p , then*

$$u(B; p) = u(B''; p) = u(B'; p)$$

for any B'' with $B \subseteq B'' \subseteq B'$.

Proof. Let B'' satisfy the conditions in the lemma statement and define $\{b_1, \dots, b_k\} = B'' \setminus B$; we proceed by induction on k . Since GS valuations are submodular,

$$u(B'; p) - u(B' \setminus b_1; p) \leq u(B \cup b_1; p) - u(B; p),$$

so

$$u(B'; p) \leq u(B \cup b_1; p) - u(B; p) + u(B' \setminus b_1; p).$$

Since B is most demanded, $u(B \cup b_1; p) - u(B; p) \leq 0$. If $B \cup b_1$ is strictly worse, then $B' \setminus b_1$ must have strictly higher utility than B' , which contradicts B' being a most demanded bundle. So, we must have $u(B \cup b_1; p) = u(B; p)$.

By induction we assume that $B \cup \{b_1, \dots, b_{t-1}\}$ is a most demanded bundle. Using the same argument as in the base case, we can use submodularity to show that

$$u(B \cup \{b_1, \dots, b_{t-1}\}; p) = u(B \cup \{b_1, \dots, b_t\}; p).$$

Hence, B'' is a most demanded bundle at price p . □

The second condition for edges from the null good—giving edges only to goods with positive price—handles a subtle case. Up to this point, we have argued informally that if B is a bundle in $\mathcal{D}_q^*(p)$, then all goods in $B \setminus M_q$ will have an incoming edge. We plan to bound the over-demand by the in-degree, but there is an important corner case: goods with price 0. Such goods can be freely added to *any* buyer's bundle, ruining the bound on over-demand. However, all is not lost: the problem stems from buyers demanding goods with with zero marginal valuation. We call bundles with such goods *degenerate*, and assume (or require) that buyers do not select them.

Definition 5.6 (Non-degenerate). *A bundle S is non-degenerate with respect to S' if goods in S' have non-zero marginal value:*

$$v(S \setminus g) < v(S) \quad \text{for each } g \in S \cap S'.$$

When $S' = [m]$, we will say S is non-degenerate. The non-degenerate correspondence for buyer $q \in N$ at price p is the set of demand bundles defined by

$$\mathcal{D}_q^\bullet(p) = \{S \in \mathcal{D}_q(p) : v_q(S \setminus g) < v_q(S) \quad \forall g \in S\}.$$

Minimum bundles are an example of non-degenerate bundles.

Lemma 5.7. *Let v be a GS valuation function. If S is a minimum demand bundle with respect to this valuation and some set of prices, then S is non-degenerate.*

Proof. Fix a buyer and a price vector p over goods and some minimum demand bundle S . By hypothesis, the buyer's utility $u(S; p)$ must strictly decrease if we eliminate a good from S , i.e. $\forall g \in S$ we have $u(S \setminus g; p) < u(S; p)$. We need to show that $v(S \setminus g) < v(S)$ for every $g \in S$ to prove that S is non-degenerate. Because S is a minimum demand bundle,

$$\begin{aligned} u(S; p) &= v(S) - p(S) > v(S \setminus g) - p(S \setminus g) \\ &> v(S \setminus g) + p(g) \geq v(S \setminus g) \end{aligned}$$

using non-negativity of the prices. □

If we only consider buyers who purchase bundles in $\mathcal{D}_q^\bullet(p)$, there is no problem with zero-price goods g : for buyers with non-degenerate bundles that contain g , that good must automatically be in any minimum bundle for those buyers. Hence, buyers cannot be indifferent between non-degenerate bundle B that contains g and $B \setminus g$. So, we only need edges in the swap graph from \perp to goods with positive price.

We are almost ready to formally connect the swap graph with over-demand, but there is one last wrinkle. Since a buyer is assigned a bundle rather than a single good, they may have two different swaps to the same good: perhaps $M_q \setminus a_1 \cup b$ and $M_q \setminus a_2 \cup b$ are both demanded. We want to count this buyer as causing demand 1 rather than 2, since they will select a bundle of distinct goods. This consideration motivates our definition of *buyer in-degree*.

Definition 5.8 (Buyer in-degree). *Consider buyers with GS valuations $\{v_q\}_{q \in N}$, the minimal Walrasian prices p and allocation μ , and the corresponding swap graph G . The buyer in-degree of node b is the number of distinct buyers with an edge directed to b .*

To prevent buyers from piling on zero-priced goods with zero marginal value, we will require buyers to take non-degenerate bundles. Then, the following definition of over-demand is natural.

Definition 5.9 (Non-degenerate over-demand). *The non-degenerate over-demand $OD^\bullet(g; p)$ for a good $g \in [m]$ at Walrasian prices p is*

$$OD^\bullet(g; p) = \max\{|U^\bullet(g; p)| - s_g, 0\},$$

where

$$U^\bullet(g; p) = \{q \in N : \exists B \in \mathcal{D}_q^\bullet(p) \text{ and } g \in B\}.$$

Like the unit demand case, this is simply the worst-case over-demand assuming that buyers choose an arbitrary non-degenerate bundle from their demand correspondence, without any assumption on how they break ties. Finally, we can relate the swap graph to over-demand.

Lemma 5.10. *Let buyers have GS valuations and let (p, μ) be a Walrasian equilibrium with minimal prices. If a node $g \in [m]$ in the swap graph G has buyer in-degree at most d , then $OD^\bullet(g; p) \leq d$.*

Proof. Consider any good g . We wish to bound the size of $U^\bullet(g; p)$, the set of buyers who have g in a non-degenerate bundle. These buyers fall into two classes: buyers with $g \in M_q$, and buyers with $g \notin M_q$. The number of buyers of the first kind is at most the supply s_g , since M allocates only fewer copies of goods than the Walrasian allocation μ .

To bound the number of buyers of the second kind, we will show that there is an edge to g in the swap graph for each such buyer. Consider any buyer q with $g \notin M_q$ and some non-degenerate bundle $B_q^\bullet \in \mathcal{D}_q^\bullet(p)$ where $g \in B_q^\bullet$. Consider any minimum demand bundle $B_q^* \subseteq B_q^\bullet$. If $g \in B_q^*$, then there is an edge corresponding to q from M_q to g , since both M_q and B_q^* lie in $\mathcal{D}_q^*(p)$ (which forms the set of bases for some matroid by Lemma 5.3).

Otherwise, if $g \notin B_q^*$, Lemma 5.5 shows that $B_q^\bullet(p) \setminus g$ is also a demand bundle:

$$v_q(B_q^\bullet \setminus g) - p(B_q^\bullet \setminus g) = v_q(B_q^\bullet) - p(B_q^\bullet)$$

Since B_q^\bullet is non-degenerate,

$$v_q(B_q^\bullet \setminus g) < v_q(B_q^\bullet)$$

and so $p(g) > 0$. At the same time, Lemma 5.5 also shows that $B_q^* \cup g$ is a demand bundle. Combining these last two facts, there must be an edge corresponding to q from \perp to g in the swap graph. So, the number of demanders of the second kind ($g \notin M_q$) is at most the buyer in-degree of g , which is at most d by assumption.

Now, we count up the demanders of both kinds: at most $s_g + d$ buyers demand g in a non-degenerate bundle. Hence, the non-degenerate over-demand is at most d . \square

5.2 Matroid Based Valuations

Now that we have seen the swap graph, we move on to describe our genericity condition on valuations. As discussed above, it is not clear how to define genericity for general GS valuations due to the axiomatic nature of the definition. We will define genericity for a subclass of GS valuations studied by Ostrovsky and Paes Leme [2015] called *matroid based valuations* (MBV). Such valuations are defined in terms of a simpler class of valuations, known as VIWM.

Definition 5.11 (VIWM). *Let X be some ground set and let w be a function $w : X \rightarrow \mathbb{R}_{\geq 0}$. Consider a matroid $\mathcal{M} = (\mathcal{I}, X)$ and define $\mathcal{M}_S = (\mathcal{I}_S, S)$ as the restricted matroid to set $S \subseteq X$ where $\mathcal{I}_S = \{T \in \mathcal{I} : T \subseteq S\}$. The valuation induced by weighted matroid (\mathcal{M}, w) is*

$$v(S) = \max_{T \in \mathcal{I}_S} \sum_{j \in T} w(j).$$

Some examples of valuations induced by a weighted matroid include unit demand valuations (the matroid has independent sets that are singletons and the empty set) and additive valuations (the matroid has the power set 2^X as its independent sets). For any matroid $\mathcal{M} = (\mathcal{I}, X)$, the valuations induced by \mathcal{M} lie in GS [Fujishige and Yang, 2003, Murota, 1996a, Murota and Shioura, 1999]. However, Ostrovsky and Paes Leme [2015] show that VIWM is a *strict* subset of GS by showing that two operations on valuation functions preserve GS, but do not preserve VIWM.

Definition 5.12 (Merging (convolution)). *For any $S \subseteq X$ and valuation functions $v', v'' : 2^X \rightarrow \mathbb{R}_{\geq 0}$, the merging of v', v'' is the valuation defined by*

$$v^{merg}(S) = \max_{(S', S'')=S} \{v'(S') + v''(S'')\}.$$

Definition 5.13 (Endowment (marginal valuation)). *Suppose $S \subseteq X$ and valuation $v' : 2^{X \cup Y} \rightarrow \mathbb{R}_{\geq 0}$, where $X \cap Y = \emptyset$. Then, the endowment of v' by Y is the valuation defined by*

$$v^{end}(S) = v'(S \cup Y) - v'(Y).$$

Accordingly, Ostrovsky and Paes Leme [2015] propose a new class of valuations closing VIWM under the two operations, called MBV, and conjecture that MBV is equal to all of GS.

Definition 5.14 (Matroid based valuation). *Matroid based valuations (MBV) is the The smallest class of valuations functions containing VIWM closed under finite merges and endowments.*

Conjecture 5.15 (Ostrovsky and Paes Leme [2015]). *The class of MBV is precisely the class of GS valuations: $GS = MBV$.*

5.3 Structural results for MBV

In the case of matchings, valuations of goods are described by a single number. When working with MBV, we will need to get a handle on the value of bundles in terms of the weights on goods in order to reason about integer linear combinations of weights for our genericity argument. In this section, we show that an MBV v evaluated on a non-degenerate set B can be written as the sum of weights corresponding to elements in B plus some integer linear combination of weights corresponding to the additional endowed goods.

Note that if each buyer $q \in N$ has valuations v_q in MBV, then we can decompose v_q into τ_q valuations in VIWM, say $\{v_q^{(t)} : t \in [\tau_q]\}$ where each $v_q^{(t)}$ can be identified with an endowed set $T_q^{(t)}$ of goods disjoint from $[m]$, weight function $w_q^{(t)}$, and independent sets $\mathcal{I}_q^{(t)}$ which make up the matroid $\mathcal{M}_q^{(t)} = (\mathcal{I}_q^{(t)}, [m] \cup T_q^{(t)})$. Note that valuations are defined over bundles of endowed goods and the original goods $[m]$, while the weights are defined over bundles of endowed goods and goods $[m]$.

We will use the notation $W_q(A)$ to denote the set of weights belonging to buyer q that are evaluated on elements in the set A , i.e.

$$W_q(A) = \{w_q^{(t)}(a) : a \in A, t \in [\tau_q]\} \quad \text{and} \quad W(A) = \bigcup_{q \in N} W_q(A).$$

We also will write the weights for goods allocated to each buyer in $\mu = (\mu_q)_{q \in N}$ as $W(\mu) = \bigcup_{q \in N} W_q(\mu_q)$.

We note that each buyer's valuations v_q can be thought of as a tree structure where each node has at most two children, and where the τ_q leaves of the tree are VIWM $v_q^{(t)} : 2^{[m] \cup T_q^{(t)}} \rightarrow \mathbb{R}_{\geq 0}$ for each $t \in [\tau_q]$, where $T_q^{(t)}$ is the union of all goods in endowment operations along the path from leaf t to the root of the tree; such goods are distinct from the goods $[m]$. Internal nodes with two children represent the merge operation applied to the two children. Each internal node with one child represents an endowment operation applied to the child valuation function. Fig. 1 gives an example tree of a valuation v in MBV class.

We let the *depth* of an MBV v be the depth of the corresponding tree, i.e., the length of the longest path from the root to a leaf.

Now, we want to show that if we have a non-degenerate set S with respect to some valuation function v , then $v(S)$ can be written as an integer linear combination of weights from each non-endowed good in S plus some weights from endowed goods. We will denote by $\text{dom}(v)$ the domain of valuation v which in particular always contains \mathcal{G} , i.e. bundles of goods $[m]$.

We first need a useful fact: if S is non-degenerate for an MBV, then every subset $S' \subseteq S$ is also non-degenerate.

Lemma 5.16. *If v is MBV and S is non-degenerate then for every $S' \subseteq S$, S' is also non-degenerate.*

Proof. We know that MBV is contained in GS and every GS valuation is submodular. Hence, if S is non-degenerate but we have a set $S' \subseteq S$ and a good $g \in S' \cap [m]$ such that $v(S' \setminus g) = v(S')$ then we can form the string of inequalities, where the second follows from v being submodular

$$0 < v(S) - v(S \setminus g) \leq v(S') - v(S' \setminus g) = 0.$$

This leads to a contradiction and proves the lemma. □

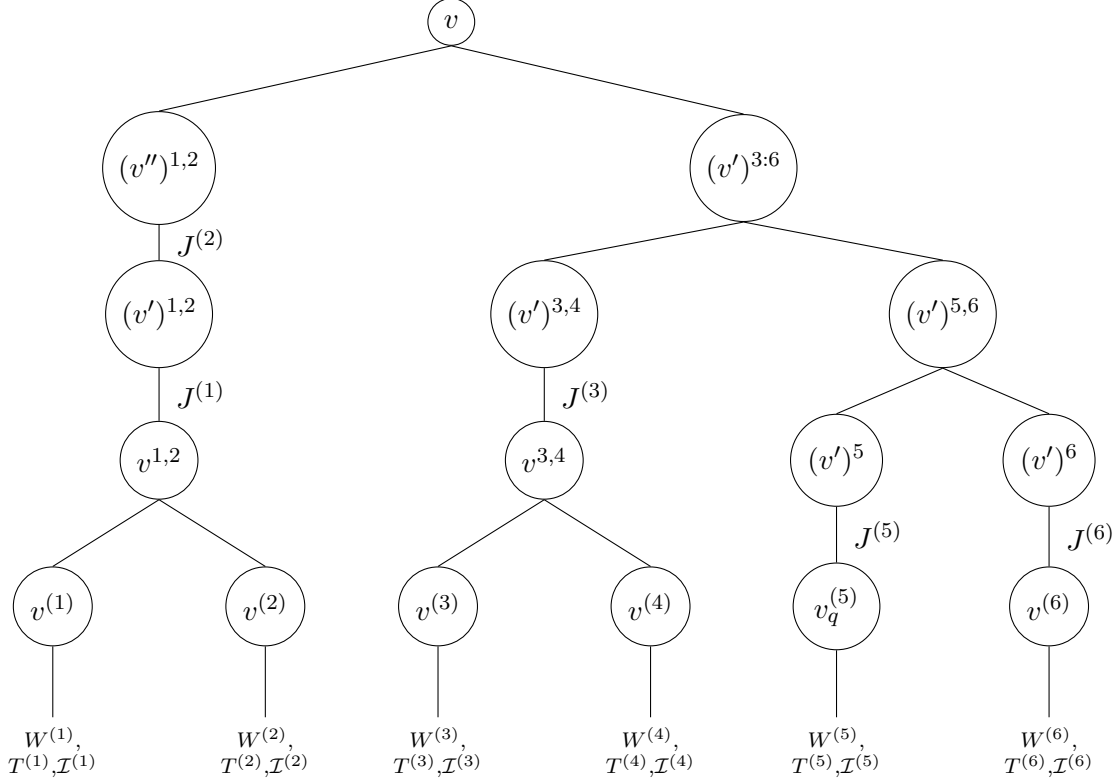


Figure 1: An example for the valuation v that is MBV. The leaf nodes represent VIWMs that have endowed goods $T^{(t)}$, weights $w^{(t)}$, and matroid with independent sets $\mathcal{I}^{(t)}$ with ground set $[m] \cup T^{(t)}$ for $t = 1, \dots, \tau$ where $\tau = 6$ in this case. When there is a parent node with a single child, that represents an endowment operation with the set of endowed goods on the edge label, while if it has two children it is the merge of the two child valuations. Note that in this example, we would require $T^{(1)} = T^{(2)}$ and $T^{(3)} = T^{(4)}$ because a merge operation must take two valuations defined over the same set of goods. Also, v has domain $2^{[m]}$, so that $J^{(1)} \cup J^{(2)} = T^{(1)}$, $J^{(1)} \subseteq T^{(1)}$, $J^{(2)} \subseteq T^{(1)} \setminus J^{(1)}$, $J^{(3)} = T^{(3)} = T^{(4)}$, $J^{(5)} = T^{(5)}$, and $J^{(6)} = T^{(6)}$.

Remark 5.17. *Below, we will frequently manipulate formal integer linear combinations without needing to know precisely what the coefficients are—the important thing is that we are working with some integer linear combination of the given weights. In these cases, we will abuse notation and use variables γ with subscripts to represent arbitrary integer values, where different occurrences of the same variable may not represent the same value, e.g. $v(S) + \sum_{g \in [m]} \gamma_g w(g) = \sum_{g \in [m]} \gamma_g w(g)$ does not mean that γ_g is the same integer on both sides of the equation.*

Now, we can connect the value of non-degenerate bundles to the weights at the leaf VIWM.

Lemma 5.18. *Let $v : \text{dom}(v) \rightarrow \mathbb{R}_{\geq 0}$ be MBV and composed of τ valuations in VIWM with endowed sets $T = \bigcup_{t \in [\tau]} T^{(t)}$ and weight functions $\{w^{(t)} : t \in [\tau]\}$. If $S \in \text{dom}(v)$ is non-degenerate*

with respect to v , then for each $g \in S \cap [m]$ there exists $t_g \in [\tau]$ such that

$$v(S) = \sum_{g \in S \cap [m]} w^{(t_g)}(g) + \sum_{w \in W(T)} \gamma_w w. \quad (1)$$

Proof. The proof is by induction on the depth of the valuation v . In the base case v is a VIWM with matroid $\mathcal{M} = (\mathcal{I}, [m] \cup T)$ and weight function w . By definition of VIWM, we know

$$v(S) = \max_{S' \in \mathcal{I}_S} \sum_{g \in S'} w(g).$$

A basic fact of weighted matroids shows that the maximum can be constructed using the greedy algorithm [Oxley, 1997]. That is, we arrange all elements $g \in S$ in decreasing order of weight w and start with a set $S' = \emptyset$, greedily add elements g to S' as long as $S' \in \mathcal{I}$. We stop when we consider all elements or when we find an element with negative weight. Suppose we have an ordering σ on S in decreasing weight such that the greedy algorithm produces an independent set S^* achieving the maximum.

Since S is non-degenerate, for each $g \in S \cap [m]$ we have $v(S \setminus g) < v(S)$. We claim that $g \in S^*$. If not, suppose we drop good g from the ordering σ . We get an ordering σ_{-g} that orders all goods in $S \setminus g$ in decreasing weight. Then, the greedy algorithm ran on orderings σ and σ_{-g} will produce the same independent set S^* , and hence the same value. But we know that $v(S \setminus g) < v(S)$ strictly, so g must be in the independent set S^* .

Since we fixed S^* before considering $g \in S \cap [m]$, the same argument shows that the whole set $S \cap [m]$ must be contained in S^* , and so

$$v(S) = \sum_{g \in S^*} w(g) = \sum_{g \in S \cap [m]} w(g) + \sum_{h \in S^* \setminus (S \cap [m])} w(h) = \sum_{g \in S \cap [m]} w(g) + \sum_{h \in S^* \cap T} w(h).$$

This gives our base case.

We now assume that our inductive hypothesis (Eq. (1)) holds for any MBV v with depth at most $k-1$ and any non-degenerate bundle S , and we prove Eq. (1) for a MBV v with depth $k > 1$ and any non-degenerate bundle S . Since $k > 1$, v is the result of either a merge or endowment operation. In the first case, let v be the result of an endowment operation from valuation $v' : 2^{[m] \cup T} \rightarrow \mathbb{R}_{\geq 0}$ with endowment set $J \subseteq T$. Let $S \in \text{dom}(v) \subseteq [m] \cup T \setminus J$ be a non-degenerate set, so that $v(S \setminus g) < v(S)$ for every $g \in S \cap [m]$. By definition of endowment,

$$v(S) = v'(S \cup J) - v'(J).$$

We need only show that $S \cup J$ is non-degenerate for valuation v' to apply the induction hypothesis, because v' is a MBV and has depth $k-1$. For contradiction, suppose that $S \cup J$ is degenerate: for some element $g^* \in S \cap [m]$ we have $v'(S \cup J) = v'(S \cup J \setminus g^*)$. Then,

$$v(S) = v'(S \cup J) - v'(J) = v'(S \cup J \setminus g^*) - v'(J) = v(S \setminus g^*)$$

is a contradiction, since S is non-degenerate for valuation v . Hence, $S \cup J$ is non-degenerate and by the inductive hypothesis, for each $g \in S \cap [m]$ we have $t_g \in [\tau]$ such that

$$v'(S \cup J) - v'(J) = \sum_{g \in S \cap [m]} w^{(t_g)}(g) + \sum_{w \in W(T)} \gamma_w w - v'(J) = \sum_{g \in S \cap [m]} w^{(t_g)}(g) + \sum_{w \in W(T)} \gamma_w w.$$

If v is not the result of an endowment operation, it must be the result of a merge operation. In this case, we let v be the merge of two valuations v^1 and v^2 in MBV, both with depth at most $k - 1$. We are given $S \in \text{dom}(v)$ non-degenerate for v , which gives us

$$v(S) = \max_{(S_1, S_2)=S} \{v^1(S_1) + v^2(S_2)\}.$$

Suppose that the maximum is achieved at a partition (S_1^*, S_2^*) of S . We claim that S_1^*, S_2^* are non-degenerate for v^1, v^2 respectively. For contradiction, suppose that S_1^* is degenerate for v^1 . Then, there is a $g \in S_1^* \cap [m]$ such that

$$v(S) = v^1(S_1^*) + v^2(S_2^*) = v^1(S_1^* \setminus g) + v^2(S_2^*) \leq v(S \setminus g),$$

but this is a contradiction since S is non-degenerate for v . Thus, S_1^* is non-degenerate for v^1 ; by symmetry also S_2^* is non-degenerate for v^2 . Therefore, we can apply the induction hypothesis on the two valuations v^1 and v^2 : for each $g \in S \cap [m]$, there is a $t_g \in [\tau]$ such that

$$v^1(S_1^*) + v^2(S_2^*) = \sum_{g \in S_1^* \cap [m]} w^{(t_g)}(g) + \sum_{g \in S_2^* \cap [m]} w^{(t_g)}(g) + \sum_{w \in W(T)} \gamma_w w.$$

Since S_1^* and S_2^* form a partition of S , we can combine the two sums:

$$v(S) = \sum_{g \in S \cap [m]} w^{(t_g)}(g) + \sum_{w \in W(T)} \gamma_w w,$$

concluding the proof. \square

Corollary 5.19. *Let $v : \text{dom}(v) \rightarrow \mathbb{R}_{\geq 0}$ a MBV composed of τ valuations in VIWM with endowed sets $T = \bigcup_{t \in [\tau]} T^{(t)}$ and weight functions $\{w^{(t)} : t \in [\tau]\}$. If $S \in \text{dom}(v)$ and there exists a $g \in S \cap [m]$ such that $v(S \setminus g) < v(S)$, then there exists $t_g \in [\tau]$ such that*

$$v(S) = w^{(t_g)}(g) + \sum_{w \in W(S \cap [m] \setminus g)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \quad (2)$$

Proof. We use the same proof method as in Lemma 5.18. \square

5.4 Generic MBV

Now that we have defined MBV, defining genericity is straightforward: we simply require genericity on the set of weights. Each valuation in MBV can be decomposed as a set of weights over the goods in $[m]$ as well as the goods that have been *endowed*.

Definition 5.20 (Generic matroid based valuations). *We say that a collection of valuations $\{v_q : \mathcal{G} \rightarrow [0, H] : q \in N\}$ are generic matroid based valuations (GMBV) if each v_q is a matroid based valuation and the weights across buyers over all goods $[m]$ and endowments T_q introduced in endowment operations are linearly independent over the integers. We will write the set $A_q = [m] \cup T_q$. Formally, the weights satisfy*

$$\sum_{w \in W(A)} \gamma_w w = 0$$

for $\gamma_w \in \mathbb{Z}$ if and only if the coefficients are all zero.

By Lemma 5.18, if a non-degenerate bundle S for a GMBV v then we know that $v(S)$ is a nontrivial integer linear combination of the weights. This property will be continually used in the results that follow in bounding buyer-in-degree in the swap graph.

5.5 Properties of swap graph with GMBV

To bound over-demand, our plan is the same as the unit demand case: 1) show that any node with no incoming edges must have price zero, as in Corollary 4.2 2) show that the swap graph is acyclic, as in Lemma 4.7, 3) show that we can write the price of a good g as an integer linear combination of the weights for the goods prior to g in a topological sort of the swap graph, as in Lemma 4.8. Once we show these three corresponding results, we can bound *buyer-in-degree* in the swap graph and hence over-demand, due to Lemma 5.10.

As mentioned before, the arguments are more involved than the corresponding arguments for the matching case because we need to work around situations where weights for goods may not end up in the final bundle valuation. A simpler version of this problem can be seen in the matchings section, where proofs frequently treat the null node \perp as a special case—this is because we treat buyers as *all* having value 0 for the null node, so it is clearly non-generic. We require more delicate reasoning for working with GMBV, since there are more opportunities where buyers may have marginal value 0 for some good.

We begin by considering the source nodes in the swap graph. As we had in the unit demand setting in Corollary 4.2, any node in our swap graph with no incoming edge has price zero. In fact, this lemma does not require valuations to be in GMBV; GS valuations suffice. We point out that while the analogous theorem in the unit demand case is an immediate consequence of a fact about the minimal Walrasian equilibrium of GS valuations, we need a bit of care to handle the more complex swap graph in the present case.

Lemma 5.21. *If agents have valuations in GS and (p, μ) is a Walrasian equilibrium with minimal Walrasian prices then any non-null node in the swap graph $G(p, \mu, M)$ with in-degree zero must have price zero, for any choice of minimum demand bundles $M = (M_q)_{q \in N}$ contained in μ .*

Proof. Let the node corresponding to g have no incoming edges. If $p(g) > 0$ we can decrease the price $p(g)$ by a small positive value

$$\epsilon_g = \frac{1}{2} \times \min \left\{ p(g), \min_{\substack{q \in N \\ B, B' \subseteq [m]}} \{u_q(B; p) - u_q(B'; p) : u_q(B; p) - u_q(B'; p) > 0\} \right\}$$

while not modifying the prices p on other goods; let the new price vector be p' .

We claim p' along with the original allocation μ forms a Walrasian equilibrium. Evidently μ continues to clear the market. We need to show that each buyer continues to prefer their allocation under μ . Buyers who are assigned g under μ continue to prefer their allocation, since their utility has only increased as we decrease the price of g .

Now, consider buyers q that are not allocated g in μ . By our choice of ϵ_g , the only bundles B that q might strictly prefer to M_q at the new prices p' contain g and satisfy

$$u_q(B; p) = u_q(M_q; p).$$

By M_q being a minimum bundle, we must have $|B| \geq |M_q|$. If $|B| = |M_q|$ then B is also a minimum bundle at price p , and since such bundles form a Matroid Basis, there must be some $h \in M_q$ such that $M_q \cup g \setminus h$ is also minimal demand under prices p . By construction of the swap graph, g would have an incoming edge—impossible.

If $|B| > |M_q|$, then take any minimum demand bundle $B' \subsetneq B$ at price p . If $g \in B'$ we again get an in-edge to g via a swap from M_q , while if $g \notin B'$ then $B' \cup g$ would be a demand bundle by Lemma 5.5 giving an edge from \perp to g ; in both cases, contradiction. So, buyers who are not allocated g also continue to prefer their allocation under μ .

Thus, the new prices p' and the allocation μ form a Walrasian equilibrium with $p' \leq p$. This is a contradiction: p are minimal Walrasian prices by assumption. \square

Next, we show that if we have valuations that are GMBV, then the corresponding swap graph is acyclic.

Lemma 5.22. *Let (p, μ) be a Walrasian equilibrium for buyers with valuations $\{v_q : q \in N\}$ that are GMBV. Then for any choice of minimum demand bundles $M = (M_q : q \in N)$ where $M_q \subseteq \mu_q$, the corresponding swap graph $G(p, \mu, M)$ is acyclic.*

Proof. Suppose $\mathcal{C} = \{(a_0, a_1), (a_1, a_2), \dots, (a_{k-2}, a_{k-1}), (a_{k-1}, a_0)\}$ are edges in the swap graph G forming a simple cycle. None of these goods can be the null node \perp because \perp has no incoming edges. Let buyer $q_i \in N$ have edge $(a_i, a_{i+1}) \in \mathcal{C}$ (addition in subscript is mod k) so that $a_i \in M_{q_i}$; note that $q_i \neq q_{i+1}$ for all i , since $a_{i+1} \in \mu_{q_{i+1}}$ and $a_{i+1} \notin \mu_{q_i}$ by construction of the swap graph. Because edges in the swap graph indicate indifferences, we have the following set of equalities

$$u_{q_i}(M_{q_i}; p) = u_{q_i}(M'_{q_i}; p) \quad \text{so that} \quad v_{q_i}(M_{q_i}) - p(a_i) = v_{q_i}(M'_{q_i}) - p(a_{i+1}) \quad i = 0, \dots, k-1$$

where $M'_{q_i} = M_{q_i} \cup a_{i+1} \setminus a_i$. Summing this equation over all the buyers in the cycle, the prices cancel and we have

$$\sum_{i=0}^{k-1} v_{q_i}(M_{q_i}) = \sum_{i=0}^{k-1} v_{q_i}(M'_{q_i}). \quad (3)$$

Let a be a good in \mathcal{C} appearing at position i^* (since the cycle is simple, i^* is unique) and q be the corresponding buyer at this position. Let ℓ be the number of weights corresponding to buyer q for good a on the LHS of Eq. (3). Since $a \in M_q$ by construction of the swap graph and M_q is non-degenerate, we have one weight from q for good a each time buyer q shows up in cycle \mathcal{C} by Lemma 5.18, that is

$$\begin{aligned} \sum_{i=0}^{k-1} v_{q_i}(M_{q_i}) &= \sum_{i:q_i=q} v_q(M_q) + \sum_{i:q_i \neq q} v_{q_i}(M_{q_i}) \\ &= \ell \cdot \sum_{g \in M_q} w_q^{(t_g)}(g) + \sum_{i:q_i \neq q} v_{q_i}(M_{q_i}) + \sum_{w \in W(T)} \gamma_w w \\ &= \ell \cdot w_q^{(t_a)}(a) + \ell \cdot \underbrace{\sum_{g \in M_q \setminus a} w_q^{(t_g)}(g)}_{\text{No weights for good } a} + \underbrace{\sum_{i:q_i \neq q} v_{q_i}(M_{q_i})}_{\text{No weights for buyer } q} + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

On the RHS of Eq. (3), we again have a weight from buyer q for good a each time q appears in the cycle, except for position i^* where $a \notin M_q \cup a_{i^*+1} \setminus a$. Thus, there are exactly $\ell - 1$ weights corresponding to q for good a on the RHS of the equation:

$$\begin{aligned}
\sum_{i=0}^{k-1} v_{q_i}(M'_{q_i}) &= \sum_{i:q_i=q} v_q(M_q \cup a_{i+1} \setminus a_i) + \sum_{i:q_i \neq q} v_q(M'_{q_i}) \\
&= v_q(M_q \cup a_{i^*+1} \setminus a) + \sum_{\substack{i:q_i=q \\ i \neq i^*}} v_q(M_q \cup a_{i+1} \setminus a_i) + \sum_{i:q_i \neq q} v_q(M'_{q_i}) + \sum_{w \in W(T)} \gamma_w w \\
&= \underbrace{v_q(M_q \cup a_{i^*+1} \setminus a)}_{\text{No weight for } a} + \underbrace{\sum_{\substack{i:q_i=q \\ i \neq i^*}} \sum_{g \in M_q \cup a_{i+1} \setminus a_i} w_q^{(t'_g)}(g)}_{\ell-1 \text{ weights for good } a \text{ belonging to } q} + \underbrace{\sum_{i:q_i \neq q} v_q(M'_{q_i})}_{\text{No weights for buyer } q} + \sum_{w \in W(T)} \gamma_w w.
\end{aligned}$$

By genericity Eq. (3) cannot hold, so the swap graph must be acyclic. \square

Once we have acyclicity, we can use the swap graph to read off the form of the Walrasian prices. Just like in the unit demand case, we can trace back from g to a source node. The price of the source node is zero, and each hop along the swap graph adds a difference of valuations to the price until we finally arrive at good g .

Because the swap graph $G(p, \mu, M)$ is acyclic for any M , we can choose a partial order on the nodes so that all edges go from nodes earlier in the ordering to nodes later in the ordering. For the next result we fix a good g in the swap graph and a simple path $g_1 \rightarrow g_2 \rightarrow \dots \rightarrow g_k = g$ originating from source node g_1 with no incoming edges. We label the buyers such that the edge from g_{i-1} to g_i belongs to q_i . Note that unlike the unit demand case, this path may visit several goods held by the same buyer (we could have $q_i = q_j$ and $i \neq j$). This feature complicates the genericity analysis, since we need to make sure that a single buyer's weights do not cancel themselves out.

Lemma 5.23. *Given the swap graph G and the simple path defined above, if $g_1 = \perp$ then there exists $B \in \mathcal{D}_{q_2}(p)$ which contains g_2 and $B \setminus g_2 \in \mathcal{D}_{q_2}^*(p)$ such that $p(g_2) = v_{q_2}(B) - v_{q_2}(B \setminus g_2)$ and if $k > 2$ then there exists $t_j \in [\tau_{q_j}]$ for $j = 2, \dots, k$ such that*

$$p(g_k) = w_{q_k}^{(t_k)}(g_k) + v_{q_2}(B) - v_{q_2}(B \setminus g_2) + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \quad (4)$$

Otherwise, if $g_1 \neq \perp$ then there exists $t_j \in [\tau_{q_j}]$ for $j = 1, \dots, k$ such that

$$p(g_k) = w_{q_k}^{(t_k)}(g_k) + \sum_{1 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \quad (5)$$

Proof. By Lemma 5.22, the swap graph is acyclic and defines a partial order on the goods. Note that because buyer q_k has an edge leading to g_k , good g_k cannot be in μ_{q_k} . We will prove the lemma by induction on k , the number of goods in the path. For the base case, we consider $k = 2$. We will first consider the case where the source node $g_1 = \perp$. By construction of the swap graph we know that there exists $B \in \mathcal{D}_{q_2}(p)$ such that $B \setminus g_2 \in \mathcal{D}_{q_2}^*(p)$, thus we have

$$u_{q_2}(B; p) = u_{q_2}(B \setminus g_2; p) \quad \text{and} \quad p(g_2) = v_{q_2}(B) - v_{q_2}(B \setminus g_2).$$

Now, if $g_1 \neq \perp$ and $k = 2$ then we have

$$u_{q_2}(M_{q_2}; p) = u_{q_2}(M_{q_2} \cup g_2 \setminus g_1; p) \quad \text{and} \quad p(g_2) = v_{q_2}(M_{q_2} \cup g_2 \setminus g_1) - v_{q_2}(M_{q_2}) + p(g_1).$$

Note that because g_1 has in-degree 0, we know from Lemma 5.21 that $p(g_1) = 0$. Further, because M_{q_2} and $M_{q_2} \cup g_2 \setminus g_1$ are minimum demand bundles for q_2 we can apply Lemma 5.18: for each $g \in M_{q_2}$ there exists $t_g \in [\tau_{q_2}]$, and for each $g \in M_{q_2} \cup g_2 \setminus g_1$ there exists $t'_g \in [\tau_{q_2}]$ such that

$$\begin{aligned} p(g_2) &= \sum_{g \in M_{q_2} \cup g_2 \setminus g_1} w_{q_2}^{(t'_g)}(g) - \sum_{g \in M_{q_2}} w_{q_2}^{(t_g)}(g) + \sum_{w \in W(T)} \gamma_w w \\ &= w_{q_2}^{(t'_{g_2})}(g_2) + \sum_{w \in W_{q_2}(\mu_{q_2})} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

We now consider the inductive case when the length of the path is $k > 2$. Let us assume that $g_1 = \perp$ and Eq. (4) holds for good g_{k-1} . By construction of the swap graph, we know that buyer q_k is indifferent to receiving M_{q_k} or $M'_{q_k} = M_{q_k} \cup g_k \setminus g_{k-1}$ where $g_k \notin \mu_{q_k}$. We then have,

$$p(g_k) = v_{q_k}(M'_{q_k}) - v_{q_k}(M_{q_k}) + p(g_{k-1}).$$

We then apply the inductive hypothesis on $p(g_{k-1})$ for $g_1 = \perp$. We will use Lemma 5.18 to conclude that for each $g \in M_{q_k}$ there exists $t_g \in [\tau_{q_k}]$, for each $g \in M'_{q_k}$ there is a $t'_g \in [\tau_{q_k}]$, for each $g < k$ there is a $t''_g \in [\tau_{q_g}]$, and for some $B \in \mathcal{D}_{q_2}(p)$ where $B \setminus g_2 \in \mathcal{D}_{q_2}^*(p)$ we have

$$\begin{aligned} p(g_k) &= \sum_{g \in M'_{q_k}} w_{q_k}^{(t'_g)}(g) - \sum_{g \in M_{q_k}} w_{q_k}^{(t_g)}(g) + w_{q_{k-1}}^{(t''_{k-1})}(g_{k-1}) + v_{q_2}(B) - v_{q_2}(B \setminus g_2) \\ &\quad + \sum_{2 < j < k-1} \gamma_j w_{q_j}^{(t''_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \\ &= w_{q_k}^{(t'_k)}(g_k) + v_{q_2}(B) - v_{q_2}(B \setminus g_2) + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t''_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

Proving Eq. (5) for $g_1 \neq \perp$ follows similarly, where there exists $t_j \in [\tau_{q_j}]$ for each $j = 2, \dots, k$ such that

$$\begin{aligned} p(g_k) &= v_{q_k}(M'_{q_k}) - v_{q_k}(M_{q_k}) + p(g_{k-1}) = v_{q_k}(M'_{q_k}) - v_{q_k}(M_{q_k}) + w_{q_{k-1}}^{(t_{k-1})}(g_{k-1}) \\ &\quad + \sum_{1 < j < k-1} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \\ &= w_{q_k}^{(t_k)}(g_k) + \sum_{1 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

□

By Lemma 5.21, goods with zero in-degree must have price 0. Also, goods with positive in-degree must have strictly positive price under genericity. (This fact is also true in the unit demand case, but we did not need it there.)

Corollary 5.24. *Let (p, μ) be a Walrasian equilibrium and $G(p, \mu, M)$ be the corresponding swap graph for any M . Any good with positive in-degree in G will have positive price.*

Proof. Suppose g has positive in-degree. Let $g_1 \rightarrow g_2 \rightarrow \dots \rightarrow g_k = g$ be a simple path originating from source node g_1 to g with the incoming edge to g_i corresponding to buyer q_i . If $g_1 \neq \perp$, then Eq. (5) implies $p(g) > 0$ because no terms could cancel with $w_{q_k}^{(t_k)}(g_k)$ by genericity.

Otherwise, $g_1 = \perp$. If the path is of length 2 ($k = 2$), g has an edge from the null good and must have positive price by construction of the swap graph. So, suppose $k > 2$. By construction of the swap graph, we know that there exists $Q_2 \in \mathcal{D}_{q_2}(p)$ such that if we eliminate g_2 from it, then we will have a minimum demand bundle, i.e. $Q_2 \setminus g_2 \in \mathcal{D}_{q_2}^*(p)$. From Corollary 5.19, we can write the valuation $v_{q_2}(Q_2)$ in terms of the weights corresponding to elements inside of Q_2 that belong to q_2 and an integer linear combination of the weights for the endowed goods T , i.e.

$$v_{q_2}(Q_2) = w_{q_2}^{(t_2)}(g_2) + \sum_{w \in W_{q_2}(Q_2 \setminus g_2)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w.$$

Further, because $Q_2 \setminus g_2 \in \mathcal{D}_{q_2}^*(p)$, we know from Lemma 5.18 that we can write $v_{q_2}(Q_2 \setminus g_2)$ as a sum of weights corresponding to elements in $Q_2 \setminus g_2$ plus an integer linear combination of weights from the endowed goods:

$$v_{q_2}(Q_2 \setminus g_2) = \sum_{a \in Q_2 \setminus g_2} w_{q_2}^{(t_a)}(a) + \sum_{w \in W(T)} \gamma_w w.$$

We then use Eq. (4) to write $p(g)$ as

$$\begin{aligned} p(g) &= w_q^{(t)}(g) + v_{q_2}(Q_2) - v_{q_2}(Q_2 \setminus g_2) + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \\ &= w_q^{(t)}(g) + w_{q_2}^{(t_2)}(g_2) + \sum_{w \in W_{q_2}(Q_2 \setminus g_2)} \gamma_w w + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

Recall that $g \neq g_2, \dots, g_{k-1}$. If $q \neq q_2$, there would be the weight $w_q^{(t)}(g)$ left that would not be able to cancel with any other term, due to genericity. Thus, we consider $q = q_2$. It may be the case that Q_2 contains g and some weight from $W_{q_2}(Q_2 \setminus g_2)$ cancels with $w_q^{(t)}(g)$. However, the weight $w_{q_2}^{(t_2)}(g_2)$ cannot cancel with any other term, due to genericity and all goods along the path are distinct. Thus, in any case there is some weight that remains and so $p(g) \neq 0$ by genericity, so $p(g) > 0$. \square

5.6 Bounding the Buyer In-Degree

We are now ready to prove our final property of the swap graph, by bounding the buyer in-degree under genericity. This result corresponds to Lemma 4.9 from the unit demand case, and crucially uses the fact that the null node only has edges to goods with strictly positive price; if we allowed edges from the null node to price zero goods, every bidder would have such an edge and bidder in-degree could be as large as n , rather than bounded by 1.

Theorem 5.25. *For GMBV buyers and Walrasian equilibrium (p, μ) with minimal Walrasian prices, each node in the swap graph G has buyer in-degree at most 1.*

Proof. We will use Lemma 5.23 to prove this and so we relabel the nodes in the swap graph $G(p, \mu, M)$ using the usual partial ordering—all edges go from a lower indexed node to a higher indexed node. Consider a node g with buyer in-degree more than 1. We know that its price $p(g)$ satisfies Eq. (4) or Eq. (5), depending on whether there is a path originating from a source node g_1 to g in the swap graph with $g_1 = \perp$ or $g_1 \neq \perp$, respectively. We will write T as the set of all endowed goods for all agents.

Suppose good g has at least two incoming edges where two edges belong to different buyers q and r . We consider two paths to g : $g_1, g_2, \dots, g_k = g$ and $h_1, h_2, \dots, h_\ell = g$, where g_1 and h_1 are source nodes, the edge to g_j belongs to q_j and the edge to h_i belongs to r_i with $q_k = q$ and $r_\ell = r$. There are three cases:

- 1) Both paths originate at a source node that is non-null, i.e. $g_1, h_1 \neq \perp$.
- 2) Exactly one of the source nodes along a path is from the null node, i.e. $g_1 = \perp, h_1 \neq \perp$.
- 3) The source nodes along both paths are from the null node, i.e. $g_1, h_1 = \perp$.

For case 1), we have two paths to g and by Lemma 5.23 we can write $p(g)$ in two different ways. For each g_j there exists $t_j \in [\tau_{q_j}]$ for $j = 1, \dots, k$ as well as for each h_j there exists $t'_j \in [\tau_{r_j}]$ for $j = 1, \dots, \ell$ such that

$$p(g) = w_q^{(t_k)}(g) + \sum_{1 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \quad (\text{LHS-1})$$

$$p(g) = w_r^{(t'_\ell)}(h) + \sum_{1 < j < \ell} \gamma_j w_{r_j}^{(t'_j)}(h_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \quad (\text{RHS-1})$$

Consider the weight $w_q^{(t_k)}(g_k)$ on LHS-1. There is no other term in LHS-1 that can cancel with it because the weights are generic and all goods $g_j \neq g$ for $j < k$. Further, good g is not in buyer q 's allocation μ_q , so $w_q^{(t_k)}(g) \notin W(\mu)$ and g is not an endowed good so $w_q^{(t_k)}(g) \notin W(T)$. Also, no term in RHS-1 can cancel with $w_q^{(t_k)}(g)$ because $r \neq q$. This gives a non-trivial linear combination contradicting genericity, completing case 1).

For case 2), the two paths that terminate at g in the swap graph have two different types of source nodes, say $g_1 = \perp$ and $h_1 \neq \perp$. When $k > 2$, we have

$$p(g) = w_q^{(t_k)}(g) + v_{q_2}(Q_2 \cup g_2) - v_{q_2}(Q_2) + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \quad (\text{LHS-2})$$

$$p(g) = w_r^{(t'_\ell)}(g) + \sum_{1 < j < \ell} \gamma_j w_{r_j}^{(t'_j)}(h_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w, \quad (\text{RHS-2})$$

Note that if $g \notin Q_2$ then the term $w_q^{(t_k)}(g)$ for $q \neq r$ could not cancel with any term in LHS-2 due to genericity as we have argued before. Further, this weight would also not cancel with any term in RHS-2 because $q \neq r$ and the weights are generic.

We now consider the case that $g \in Q_2$. Note that by monotonicity we have $v_{q_2}(Q_2 \cup g_2) \geq v_{q_2}(Q_2 \cup g_2 \setminus g)$. If this were an equality we would have the following relation

$$u_{q_2}(Q_2 \cup g_2; p) = u_{q_2}(Q_2 \cup g_2 \setminus g; p) - p(g) < u_{q_2}(Q_2 \cup g_2 \setminus g; p)$$

since g has positive price from Corollary 5.24. This is a contradiction because $Q_2 \cup g_2$ is a demand bundle of q_2 . Hence we must have $v_{q_2}(Q_2 \cup g_2) > v_{q_2}(Q_2 \cup g_2 \setminus g)$ and we can apply Corollary 5.19 to conclude that for some $t^* \in [\tau_{q_2}]$,

$$v_{q_2}(Q_2 \cup g_2) = w_{q_2}^{(t^*)}(g) + \sum_{w \in W_{q_2}(Q_2 \cup g_2 \setminus g)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w.$$

We have $Q_2 \in \mathcal{D}_{q_2}^*(p)$ so we can apply Lemma 5.18 to write $v_{q_2}(Q_2)$ as a sum of weights, where $s^* \in [\tau_{q_2}]$, we can rewrite LHS-2 to get

$$\begin{aligned} p(g) = & w_q^{(t_k)}(g) + w_{q_2}^{(t^*)}(g) - w_{q_2}^{(s^*)}(g) + \sum_{w \in W_{q_2}(Q_2 \cup g_2 \setminus g)} \gamma_w w \\ & + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t'_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

Note that we have three weights corresponding to good g and it may be the case that $t_k = t^* = s^*$. By genericity, two of these weights may cancel (if they are identical) but there will be at least one weight that does not cancel. By genericity this could not cancel with any term on the RHS-2, since those weights are for buyers different from q . Thus, we get a non-trivial linear combination, contradicting genericity.

Note if $k = 2$, we know that $g = g_2$ and $v_q(Q_2 \cup g) - v_q(Q_2) > 0$ (the price of g is positive). Thus, $v_q(Q_2 \cup g)$ must contain a weight for g belonging to q , by Corollary 5.19, that cannot cancel with the weight $w_r^{(t')}(g)$ in the RHS-2 due to $q \neq r$ and genericity. Hence, either way in case 2) we get a contradiction.

We now consider our last case 3), $g_1, h_1 = \perp$. We will need to consider three subcases without loss of generality: a) $k = \ell = 2$; b) $k = 2, \ell > 2$; c) $k, \ell > 2$. For our first subcase a), we can write the price of g in two different ways, according to Lemma 5.23:

$$p(g) = v_q(Q \cup g) - v_q(Q) \quad \text{and} \quad p(g) = v_r(R \cup g) - v_r(R),$$

where Q, R are minimum bundles for bidders q and r respectively, and $Q \cup g, R \cup g$ are demand bundles (but not minimum) for bidders q and r . Crucially, we also know that $p(g) > 0$, by construction of the swap graph. Therefore, $Q \cup g$ and $R \cup g$ must be non-degenerate with respect to g , for instance

$$v_q(Q) - p(Q) = v_q(Q \cup g) - p(Q \cup g) = v_q(Q \cup g) - p(Q) - p(g) < v_q(Q \cup g) - p(Q),$$

so $v_q(Q) < v_q(Q \cup g)$; the same argument applies to $R \cup g$. Furthermore, since Q and R are minimum, they are also non-degenerate by Lemma 5.7. Combined with Lemma 5.18 applied to $Q \cup g$ and $R \cup g$, we have

$$p(g) = v_q(Q \cup g) - v_q(Q) = w_q^{(t)}(g) + \sum_{w \in W_q(Q)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w, \quad (\text{LHS-3a})$$

$$p(g) = v_r(R \cup g) - v_r(R) = w_r^{(t')}(g) + \sum_{w \in W_r(R)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \quad (\text{RHS-3a})$$

However, LHS-3a consists of integer linear combinations of weights for buyer q , while the right hand side RHS-3a only has integer linear combinations of the weights for buyer r , leading to a non-trivial linear combination contradicting genericity.

For subcase b) we can write the prices in two ways:

$$p(g) = v_q(Q \cup g) - v_q(Q) = w_q^{(t)}(g) + \sum_{w \in W_q(Q)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w, \quad (\text{LHS-3b})$$

$$p(g) = w_r^{(t_\ell)}(g) + v_{r_2}(R_2 \cup h_2) - v_{r_2}(R_2) + \sum_{2 < j < \ell} \gamma_j w_{r_j}^{(t_j)}(h_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \quad (\text{RHS-3b})$$

where Q is some minimum demand bundle for q that does not contain g and R_2 is some minimum demand bundle for buyer r_2 that does not contain h_2 . If $r_2 \neq r$ then the term $w_r^{(t_\ell)}(g)$ in RHS-3b cannot cancel with any term in LHS-3b, due to genericity. If $r_2 = r$ then there is no term in RHS-3b that can cancel with $w_q^{(t)}(g)$ on the LHS-3b because of genericity and $q \neq r$.

We then consider our last subcase c), when $k, \ell > 2$. We then use Lemma 5.23 to write the price of good g in two ways, where Q_2 is some minimum demand bundle for q_2 that does not contain g_2 and R_2 is some minimum demand bundle for buyer r_2 that does not contain h_2 :

$$p(g) = w_q^{(t_k)}(g) + v_{q_2}(Q_2 \cup g_2) - v_{q_2}(Q_2) + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \quad (\text{LHS-3c})$$

$$p(g) = w_r^{(t_\ell)}(g) + v_{r_2}(R_2 \cup h_2) - v_{r_2}(R_2) + \sum_{2 < j < \ell} \gamma_j w_{r_j}^{(t_j)}(h_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w \quad (\text{RHS-3c})$$

Note that if $g \notin Q_2$ and $g \notin R_2$, then the term $w_q^{(t_k)}(g)$ for $q \neq r$ could not cancel with any term above due to genericity as we have argued before. Otherwise, without loss of generality $g \in Q_2$. (The case with $g \in R_2$ follows by symmetry.) As we argued above in case 2, we can rewrite the price of g in LHS-3c as

$$\begin{aligned} p(g) = & w_q^{(t_k)}(g) + w_{q_2}^{(t^*)}(g) - w_{q_2}^{(s^*)}(g) + \sum_{w \in W_{q_2}(Q_2 \cup g_2 \setminus g)} \gamma_w w \\ & + \sum_{2 < j < k} \gamma_j w_{q_j}^{(t_j)}(g_j) + \sum_{w \in W(\mu)} \gamma_w w + \sum_{w \in W(T)} \gamma_w w. \end{aligned}$$

Following the same argument as in case 2, we know that there exists at least one weight belonging to buyer q for good g . Similarly, there exists at least one weight (note that it can be the case that $t_k = t^* = s^*$) belonging to buyer r for good g . By genericity these two weights cannot cancel because $q \neq r$.

Hence, the buyer in-degree to any node in the swap graph is at most 1. \square

We now analyze the case where buyers are told to take a most demand bundle that is non-degenerate, i.e. each buyer q takes a bundle from $\mathcal{D}_q^\bullet(p)$. We have already shown that the swap graph has buyer in-degree at most 1 and so we apply Lemma 5.10 to obtain the following result.

Theorem 5.26. *Let (p, μ) be a Walrasian Equilibrium with minimal Walrasian prices p . If buyers have valuations in GMBV and each buyer chooses only non-degenerate bundles at prices p , then the over-demand for any good $g \in [m]$ is at most 1, i.e.*

$$OD^\bullet(g; p) \leq 1, \quad \forall g \in [m].$$

5.7 Bounding welfare

So far, we have shown that over-demand can be no more than one when buyers select any non-degenerate bundle in their demand correspondence at the minimal Walrasian prices. Much like the case for matchings, where we assumed that buyers broke ties in favor of selecting a good in their demand correspondence (rather than selecting the empty set), we will assume that buyers break ties in favor of the *largest cardinality* non-degenerate bundles in their demand correspondence—i.e. we will assume that buyers select arbitrary bundles in what we will call the *max non-degenerate* correspondence $\mathcal{D}_q^{max}(p)$ for buyer q at price p where

$$\mathcal{D}_q^{max}(p) = \{D_q \in \mathcal{D}_q^\bullet(p) : |D_q| \geq |D'_q|, \quad \forall D'_q \in \mathcal{D}_q^\bullet(p)\}. \quad (6)$$

Under this tie-breaking condition, buyers achieve nearly optimal welfare.

Theorem 5.27. *Let buyers have GMBV $\{v_q : \mathcal{G} \rightarrow [0, H]\}_{q \in N}$ in a market in which there are s_g copies of each good $g \in [m]$. At the minimal Walrasian prices p , if each buyer q selects any max non-degenerate bundle $B_q^{max} \in \mathcal{D}_q^{max}(p)$, then welfare will be near optimal:*

$$Welfare_N(B_1^{max}, \dots, B_n^{max}) \geq Opt\text{-Welfare}_N - 2H \cdot m.$$

Proof. Without loss of generality we will assume that the Walrasian allocation μ is non-degenerate; if it were not, we could remove from each allocated bundle all goods with zero marginal valuation while remaining in a buyer's demand set and thus we would obtain the same welfare as the original allocation. Note that the welfare from μ is equal to $Opt\text{-Welfare}_N$.

Let $B^{max} = (B_1^{max}, \dots, B_n^{max})$ denote a collection of max non-degenerate bundles at price p , one per buyer. We know that $|B_q^{max}| \geq |\mu_q|$ by maximality, so B^{max} contains at least as many goods as μ . Note that the utility each buyer receives from the allocation μ is the same as the utility they receives from her bundle in B^{max} . As we did in the proof of Theorem 4.13, we sum over all buyers' utilities to get

$$\begin{aligned} \sum_{q \in N} [v_q(\mu_q) - p(\mu_q)] &= \sum_{q \in N} v_q(\mu_q) - \sum_{g \in [m]} p(g)s_g = \sum_{q \in N} [v_q(B_q^{max}) - p(B_q^{max})] \\ &= \sum_{q \in N} v_q(B_q^{max}) - \sum_{g \in [m]} p(g)s_g + \sum_{g \in [m]} p(g) \cdot ud(g, B^{max}) - \sum_{g \in [m]} p(g) \cdot od(g, B^{max}). \end{aligned}$$

where we define the over-demand (od) and under-demand (ud), respectively, of good g in the tuple of bundles B^{max} compared to the optimal allocation μ as

$$od(g, B^{max}) = \max \left\{ \sum_{q \in N} (\mathbf{1}\{g \in B_q^{max}\} - \mathbf{1}\{g \in \mu_q\}), 0 \right\}$$

$$ud(g, B^{max}) = \max \left\{ \sum_{q \in N} (\mathbf{1}\{g \in \mu_q\} - \mathbf{1}\{g \in B_q^{max}\}), 0 \right\}.$$

Note that the $od(g, B^{max}) \geq 0$ and so we get the following inequality:

$$\text{Opt-Welfare}_N \leq \sum_{q \in N} v_q(B_q^{max}) + \sum_{g \in [m]} p(g) \cdot ud(g, B^{max}) \leq \sum_{q \in N} v_q(B_q^{max}) + H \sum_{g \in [m]} ud(g, B^{max}).$$

To bound the right hand side, we will bound the total under demand of all goods from allocation B^{max} relative to μ .

We now use the fact that B_q^{max} is a max non-degenerate bundle, so we have

$$\sum_{q \in N} |\mu_q| \leq \sum_{q \in N} |B_q^{max}| = \sum_{q \in N} |\mu_q| + \sum_{g \in [m]} (od(g, B^{max}) - ud(g, B^{max}))$$

$$\sum_{g \in [m]} ud(g, B^{max}) \leq \sum_{g \in [m]} od(g, B^{max}).$$

To bound the over demand for each good from B^{max} relative to μ , we again turn to the swap graph. We first bound how many more goods are allocated by B^{max} compared to μ . We consider buyers $q \in N$ and good g where $g \in B_q^{max}$ and $g \notin \mu_q$. Let B_q^* be any minimum demand bundle contained in B_q^{max} . If $g \in B_q^*$ then we know that there will be an edge to g in the swap graph $G(p, \mu, M)$ for any choice of minimum demand bundles $M = (M_1, \dots, M_n)$ contained in μ . Otherwise, if $g \notin B_q^*$ then we know that $B_q^{max} \setminus g$ is also a most demanded bundle from Lemma 5.5. We further know that $p(g) > 0$, due to $v_q(B_q^{max}) > v_q(B_q^{max} \setminus g)$ (B_q^{max} is non-degenerate). So, there is an edge from \perp to g in $G(p, \mu, M)$ for any choice of minimum demand bundles M . Hence, every good that B^{max} allocates over μ has at least one incoming edge.

Because the buyer in-degree in $G(p, \mu, M)$ is at most one for any choice of minimum demand bundles M , we know that for each good g there is at most one buyer q where $g \in B_q^{max}$ but $g \notin \mu_q$. Hence, we get $od(g, B^{max}) \leq 1$ for each $g \in [m]$:

$$\text{Opt-Welfare}_N \leq \sum_{q \in N} v_q(B_q^{max}) + m \cdot H = \text{RWelfare}_N(b_1, \dots, b_n) + m \cdot H.$$

Finally, applying Lemma 4.12 completes the proof. \square

6 Welfare and over-demand generalization

In this section, we show that Walrasian prices *generalize*: the equilibrium prices for a market N of buyers induces similar behavior when presented to a new sample N' of buyers, both in terms of the demand for each good and in terms of welfare. More precisely, a set of prices which minimizes over-demand of each good *and* maximizes welfare when each buyer purchases her most-preferred bundle retains these properties (approximately) on a *new* market N' when each buyer purchases her most-preferred bundle, if buyers in N and N' are drawn independently from the same distribution.

6.1 Results for Arbitrary Valuations

We need a bit more notation to formally state the generalization guarantees for more general valuations. To formalize how over-demand and welfare induced by prices p vary between two markets, we first fix a tie-breaking rule e that buyers use to select amongst demanded bundles for a valuation from set \mathcal{V} . We define classes of *functions* parameterized by pricings mapping valuations to (a) bundles purchased, (b) whether or not a particular good g is purchased, and (c) the value a buyer gets for her purchased bundle.

For technical convenience, we will assume *all* buyers choose bundles from their demand sets in the following systematic way (still independent of other buyers' choices). First, each buyer q will have some set \mathcal{L}_{v_q} of “infeasible” bundles they will *never* choose to buy. For instance, we will prevent buyers from buying items for which they get zero marginal utility: if $v(B) = v(B \cup \{g\})$ for $g \notin B$, having $B \cup \{g\} \in \mathcal{L}_{v_q}$ forces q to avoid the unnecessarily larger bundle (needed for over-demand guarantees). Second, all buyers share some standard “linear” choice rule over the elements of the remaining bundles in their demand set; we use the rule to break ties in favor of larger demanded bundles, needed for welfare guarantees.

Below we formally define an encodable tie-breaking rule, which we will need for our concentration results for demand and welfare.

Definition 6.1 (Encodable tie-breaking). *A function $e : 2^{\mathcal{G}} \times \mathcal{V} \rightarrow \mathcal{G}$ is an encodable rule if*

- $e(\mathcal{X}, v) \in \mathcal{X}$;
- *there exists a vector $y \in (\mathbb{R}_{\geq 0})^m$ (called a separator), and a subset $\mathcal{L}_{v_q} \subseteq 2^{\mathcal{G}}$ (called the infeasible set for buyer q) such that, for all \mathcal{X} , $e(\mathcal{X}, v) \notin \mathcal{L}_{v_q}$, and $\sum_{g \in e(\mathcal{X}, v)} y_g > \sum_{g \in B'} y_g$ for any other $B' \in \mathcal{X} \setminus \mathcal{L}_{v_q}$; and*
- *at any pricing p there is at least one utility-maximizing bundle for each valuation that lies outside the infeasible set \mathcal{L}_{v_q} .*

We show in Lemma F.3 that there exists an encodable tie-breaking rule over demand sets which selects bundles satisfying the properties needed for our over-demand and welfare results to hold (Theorems 4.10 and 4.13) in the case of unit demand buyers, and the properties needed for the analogous results (Theorems 5.26 and 5.27) in the case of GMBV buyers. For the remainder of this section, we fix some encodable tie-breaking rule e and state our results with respect to this rule. We call the bundle that e selects from $\mathcal{D}_q(p)$ the *canonical* bundle for q at p , denoted by $B_q^*(p) = e(\mathcal{D}_q(p), v_q)$. Then, for each good g , let $h_{g,p}(v_q) = \mathbf{1}[g \in B_q^*(p)]$ indicate whether or not g is in q 's canonical bundle at prices p . Let $C^e(g; p; N) = \sum_{q \in N} h_{g,p}(v_q)$, i.e. the number of buyers in N whose canonical bundle at p contains g . For a sample of n buyers $\{v_q\} \stackrel{i.i.d.}{\sim} \Pi$, let $C^e(g; p; \Pi)$ represent the *expected* number of copies of g demanded at prices p if buyers demand canonical bundles. Similarly, let the *welfare* of a pricing p on a market N be $\text{Welfare}_N(p) = \sum_{q \in N} v_q(\widehat{B}_q^*(p))$ for $\widehat{B}_q^*(p)$ a worst-case resolution of the over-demand for each good (See Lemma 4.12). Similarly, let $\text{RWelfare}_p(N) = \sum_{q \in N} v_q(B_q^*(p))$ denote the relaxed welfare of buyers choosing their canonical bundles at prices p ; analogously, for a distribution Π over valuations, define the expected relaxed welfare of buyers choosing canonical bundles at prices p as $\text{RWelfare}_p(\Pi) = n \mathbb{E}_{v \sim \Pi} [v(B_v^*(p))]$. In both cases, these quantities assume there is sufficient supply to satisfy the demand for each good.¹⁰

¹⁰If the supply of each good satisfies the inequality in Theorem 6.2, then each good should have supply which can satisfy at least an $1 - \alpha$ fraction of buyers who choose it. See Lemma 4.12 for a more detailed discussion.

Before presenting the technical details, we show how the behavior induced by prices generalizes for buyers with arbitrary valuations. Our first theorem bounds the over-demand when the Walrasian prices computed for a market are applied to a new market.

Theorem 6.2. *Fix a pricing p and two sampled markets N, N' of buyers with arbitrary valuations, with $|N| = |N'| = n$. Suppose e is encodable. For good g , suppose $C^e(g; p; N) \leq s_g + 1$. Then,*

$$C^e(g; p; N') - s_g = O\left(\sqrt{m \cdot s_g \cdot \ln \frac{1}{\delta}} + m \ln \frac{1}{\delta}\right),$$

with probability $1 - \delta$. Then, for any $\alpha \in (0, \frac{4}{5})$, if $s_g = \Omega(m \ln \frac{1}{\delta} / \alpha^2)$, with probability $1 - \delta$,

$$C^e(g; p; N') \leq (1 + \alpha)s_g.$$

Our second theorem is an analogous generalization result, for welfare instead of over-demand.

Theorem 6.3. *Fix two markets $N, N' \sim \Pi$ for which $|N| = |N'| = n$. Suppose for each good g , the prices p satisfy $C^e(g; p; N) \leq s_g + 1$. With probability $1 - \delta$, if e is encodable, for any $\alpha \in (0, 4/5)$, if p are welfare-optimal for N and*

$$\text{Opt-Welfare}_N = \Omega\left(\frac{H^3 n^{0.5} m^4 \ln^2(m) \ln^2 \frac{1}{\delta}}{\alpha^2}\right)$$

then

$$\text{Welfare}_{N'}(p) \geq (1 - \alpha) \text{Opt-Welfare}_{N'}.$$

The first result (Theorem 6.2), for over-demand, relies on bounding the VC dimension of the class of good g 's demand indicator functions (the class contains a function for each pricing, labeling a valuation v positive if and only if g is in the canonical bundle for v at those prices). To bound the VC dimension, we first prove that the class of *bundle* predictors (the class contains one function for each pricing p , which maps v to v 's canonical bundles at p) is *linearly separable* in a space of $m+1$ dimensions. Using a recent result by Daniely and Shalev-Shwartz [2014], such a class admits a compression scheme of size $m+1$ and can be (ϵ, δ) -PAC learned with $O(\frac{m}{\epsilon} \ln \frac{1}{\delta})$ samples. Since item predictor functions are simply projections of bundle prediction functions, their PAC-complexity of the item predictors is an upper bound on the PAC-complexity of the bundle predictors. Since item predictors are binary valued, we can apply the classical result of Ehrenfeucht et al. [1989] which shows the equivalence of learning and uniform convergence for binary prediction problems, so we can upper bound the class's VC dimension as a function of its PAC sample complexity.

The second result (Theorem 6.3), for welfare, proceeds by bounding the pseudo-dimension of the class of welfare predictor functions (containing, for each pricing a function which map valuations to the *value* of a buyer purchasing her canonical bundle at these prices). The argument uses the existence of a compression scheme to upper bound the number of possible distinct labelings of valuations by bundles. Fixing a bundle labeling of a valuation also fixes the welfare of that valuation; thus, one can upper-bound the size of a “shatterable” set and the pseudo-dimension of the class of welfare predictors.

6.2 Learning Theory Preliminaries

This section reviews several well-known definitions from learning theory. We will show that:

- the demand of a good g from a market N at prices p should be close to the demand for g on a new market N' ; and
- the welfare-optimal prices for N should be approximately welfare-optimal for N'

under appropriate conditions.

Viewed from a learning-theoretic perspective, both of these statements are about the *generalization guarantees* of a class of mechanisms. Learning theory provides many tools to formalize how well properties of prices generalize from one market to the next. In what follows, we will first describe standard tools for measuring the generalizability of binary-valued functions, followed by the analogous tools for understanding the generalizability of real-valued functions. The former will be useful in measuring the concentration of demand for a good g , while the latter will be useful for measuring the concentration of welfare.

Suppose there is a domain \mathcal{V} (for our purposes, the valuation functions), an unknown target function $c : \mathcal{V} \rightarrow \{0, 1\}$, and an unknown distribution Π over \mathcal{V} . We wish to understand how many labeled samples $(v, c(v))$, with $v \sim \Pi$, are necessary and sufficient to be able to compute a \hat{c} which agrees with c almost everywhere with respect to Π . The distribution-independent sample complexity of learning c depends fundamentally on the “complexity” of the set of binary functions \mathcal{F} from which we are choosing \hat{c} . We review two standard complexity measures next.

Let N be a set of n samples from \mathcal{V} . The set N is said to be *shattered* by \mathcal{F} if, for every subset $T \subseteq N$, there is some $c_T \in \mathcal{F}$ such that $c_T(v) = 1$ if $v \in T$ and $c_T(v') = 0$ if $v' \notin T$. That is, ranging over all $c \in \mathcal{F}$ induces all $2^{|N|}$ possible projections onto N . The *VC dimension* of \mathcal{F} , denoted $\text{VC}(\mathcal{F})$, is the size of the largest set S that can be shattered by \mathcal{F} .

Let $\text{err}_N(\hat{c}) = (\sum_{v \in N} |c(v) - \hat{c}(v)|) / |N|$ denote the empirical error of \hat{c} on N , and let $\text{err}(\hat{c}) = \mathbb{E}_{v \sim \Pi}[|c(v) - \hat{c}(v)|]$ denote the *true* expected error of \hat{c} with respect to Π . We say that \mathcal{F} is (ϵ, δ) -PAC learnable with sample complexity n if there exists an algorithm \mathcal{A} such that, for all distributions Π and all target functions $c \in \mathcal{F}$, when \mathcal{A} is given a sample S of size n it produces some $\hat{c} \in \mathcal{F}$ such that $\text{err}(\hat{c}) < \epsilon$ with probability $1 - \delta$ over the choice of the sample. The PAC sample complexity of a class \mathcal{F} can be bounded as a polynomial function of $\text{VC}(\mathcal{F})$, ϵ , and $\ln \frac{1}{\delta}$ [Vapnik and Chervonenkis, 1971]; furthermore, any algorithm which (ϵ, δ) -PAC learns \mathcal{F} over all distributions Π must use nearly as many samples to do so. The following theorem states this well-known result formally.¹¹

Theorem 6.4 (Hanneke [2015], Ehrenfeucht et al. [1989]). *Suppose \mathcal{F} is a class of binary functions. Then, \mathcal{F} can be (ϵ, δ) -PAC learned with a sample of size*

$$n = O\left(\frac{\text{VC}(\mathcal{F}) + \ln \frac{1}{\delta}}{\epsilon}\right)$$

whenever $c \in \mathcal{F}$. Furthermore, any (ϵ, δ) -PAC learning algorithm for \mathcal{F} must have sample complexity

$$n = \Omega\left(\frac{\text{VC}(\mathcal{F}) + \ln \frac{1}{\delta}}{\epsilon}\right).$$

¹¹The upper bound stated here is a quite recent result which removes a $\ln \frac{1}{\epsilon}$ factor from the upper bound; a slightly weaker but long-standing upper bound can be attributed to Vapnik and Kotz [1982].

There is a stronger sense in which a class \mathcal{F} can be learned, called *uniform learnability*. This property implies that with a sufficiently large sample, the error of *every* $c \in \mathcal{F}$ on the sample is close to the true error of c . We say \mathcal{F} is (ϵ, δ) -*uniformly learnable with sample complexity* n if for every distributions Π and a sample N of size n , $|\text{err}_N(c) - \text{err}(c)| < \epsilon$ for every $c \in \mathcal{F}$, with probability $1 - \delta$. Notice that if \mathcal{F} is (ϵ, δ) -uniformly learnable with n samples, then it is also (ϵ, δ) -PAC learnable with n samples. We will use a well-known upper bound on the uniform sample complexity of a class as a function of its VC dimension.

Theorem 6.5 (E.g. [Vapnik and Chervonenkis \[1971\]](#)). *Suppose \mathcal{F} is a class of binary functions. Then, \mathcal{F} can be (ϵ, δ) -uniformly learned with a sample of size*

$$n = O\left(\frac{\text{VC}(\mathcal{F}) \ln \frac{1}{\epsilon} + \ln \frac{1}{\delta}}{\epsilon^2}\right).$$

Both PAC learnability and uniform learnability of binary-valued functions are characterized by the class's VC dimension. When learning real-valued functions, for example, to guarantee convergence of the welfare of pricings), we use a real-valued analogue to VC dimension which will give a sufficient condition for uniform convergence, called the *pseudo-dimension* [[Pollard, 1984](#)]. Formally, let $c : \mathcal{V} \rightarrow [0, H]$ be a real-valued function over \mathcal{V} , and \mathcal{F} be the class we are learning over. Let S be a sample drawn from Π , $|N| = n$, labeled according to c . Both the empirical and true errors of a hypothesis \hat{c} are defined as before, though $|\hat{c}(v) - c(v)|$ is now in the interval $[0, H]$ rather than in $\{0, 1\}$. Let $(r_1, \dots, r_n) \in [0, H]^n$ be a set of *targets* for N . We say (r_1, \dots, r_n) *witnesses* the shattering of N by \mathcal{F} if, for each $T \subseteq N$, there exists some $c_T \in \mathcal{F}$ such that $c_T(v_q) \geq r_q$ for all $v_q \in T$ and $c_T(v_q) < r_q$ for all $v_q \notin T$. If there exist such witnesses, we say N is *shatterable* by \mathcal{F} . The *pseudo-dimension* of \mathcal{F} , denoted $\text{PD}(\mathcal{F})$, is the size of the largest set S which is shatterable by \mathcal{F} . We will derive sample complexity upper bounds from the following theorem, which connects the sample complexity of uniform learning over a class of real-valued functions to the pseudo-dimension of the class.

Theorem 6.6 (E.g. [Anthony and Bartlett \[1999\]](#)). *Let \mathcal{F} be a class of real-valued functions with range in $[0, H]$ and pseudo-dimension $\text{PD}(\mathcal{F})$. For every $\epsilon > 0, \delta \in [0, 1]$, the sample complexity of (ϵ, δ) -uniformly learning the class \mathcal{F} is*

$$n = O\left(\left(\frac{H}{\epsilon}\right)^2 \left(\text{PD}(\mathcal{F}) \ln \frac{H}{\epsilon} + \ln \frac{1}{\delta}\right)\right).$$

Moreover, a conceptually simple algorithm achieves the guarantee in [Theorem 6.6](#): simply output the function $c \in \mathcal{F}$ with the smallest empirical error on the sample. These algorithms are called *empirical risk minimizers*.

6.3 Learning from Pricings and Learning with Compression

We now introduce several other powerful tools from learning theory, recast in the language of mechanism design. Let $(v_1, B_1), \dots, (v_n, B_n), v_q \in \mathcal{V}, B_q \in \mathcal{G}$ represent samples drawn from some distribution Π over \mathcal{V} , labeled by bundles $B_q \in \mathcal{G}$. For each set of prices p , the functions h_p, f_p will map valuations to utility-maximizing bundles and the value the buyer of utility-maximizing bundles, respectively. We prove bounds on the sample complexity of uniform convergence over three classes of functions:

1) the class of functions which map valuations v_q to $v_q(B_q)$ (where B_q is utility-maximizing):

$$\mathcal{H}_V = \{f_p : \mathcal{V} \rightarrow \mathbb{R}, f_p(v_q) = v_q(B^*), B^* = e(\mathcal{D}_q(p))\};$$

2) the class of bundle predictors:

$$\mathcal{H} = \{h_p : \mathcal{V} \rightarrow \mathcal{G}, h_p(v_q) = e(\mathcal{D}_q(p))\};$$

3) the projection of \mathcal{H} to its component *good* predictors:

$$\mathcal{H}_g = \{h_{g,p} : \mathcal{V} \rightarrow \{0, 1\}, h_p(v) = (h_{g,p}(v))_g\}.$$

We will show that we can learn the class \mathcal{H} using a *compression scheme*, a tool for proving sample complexity bounds in the multi-label setting.

Definition 6.7. A compression scheme for $\mathcal{F} : \mathcal{V} \rightarrow X$, of size d consists of

- a compression function

$$\mathbf{compress} : (\mathcal{V} \times X)^n \rightarrow (\mathcal{V} \times X)^d,$$

where $\mathbf{compress}(N) \subseteq N$ and $d \leq n$; and

- a decompression function

$$\mathbf{decompress} : (\mathcal{V} \times X)^d \rightarrow \mathcal{F}.$$

For any $f \in \mathcal{F}$ and any sample $(v_1, f(v_1)), \dots, (v_n, f(v_n))$, the functions satisfy

$$\mathbf{decompress} \circ \mathbf{compress}((v_1, f(v_1)), \dots, (v_n, f(v_n))) = f'$$

where $f'(v_q) = f(v_q)$ for each $q \in [n]$.

Intuitively, a compression function selects a subset of d most relevant points from a sample, and based on these points, the decompression scheme selects a hypothesis. When such a scheme exists, the learning algorithm $\mathbf{decompress} \circ \mathbf{compress}$ is an empirical risk minimizer. Furthermore, this compression-based learning algorithm has sample complexity bounded by a function of d , which plays a role analagous to VC dimension in the sample complexity guarantees.

Theorem 6.8 (Littlestone and Warmuth [1986]). Suppose \mathcal{F} has a compression scheme of size d . Then, the PAC complexity of \mathcal{F} is at most

$$n = O\left(\frac{d \ln \frac{1}{\epsilon} + \ln \frac{1}{\delta}}{\epsilon}\right)$$

While compression schemes imply useful sample complexity bounds, it can be hard to show that a particular hypothesis class admits a compression scheme. One general technique is to show that the class is linearly separable in a higher-dimensional space.

Definition 6.9. A class \mathcal{F} is d -dimensionally linearly separable if there exists a function $\psi : \mathcal{V} \times \mathcal{G} \rightarrow \mathbb{R}^d$ such that for any $f \in \mathcal{F}$, there exists some $w_f \in \mathbb{R}^d$ with $f(v) \in \operatorname{argmax}_y \langle w_f, \psi(v, y) \rangle$ and $|\operatorname{argmax}_y \langle w_f, \psi(v, y) \rangle| = 1$.

It is known that a d -dimensional linearly separable class admits a compression scheme of size d .

Theorem 6.10 (Daniely and Shalev-Shwartz [2014]). Suppose \mathcal{F} has a d -dimensional realizable linear separator ψ . Then, there exists a compression scheme for \mathcal{F} of size d .

6.4 Over-Demand Concentration

In this section, we outline the proof of over-demand concentration (Theorem 6.2). The proof proceeds in two steps. First, we bound the VC dimension of the class of good predictors \mathcal{H}_g by $O(m)$ (Theorem 6.11). For arbitrary valuations, the argument first shows that the class of bundle predictors \mathcal{H} is $(m+1)$ -linearly separable and is thus PAC learnable with a sample size proportional to m . Since $h_{g,p} \in \mathcal{H}_g$ is a projection of $h \in \mathcal{H}$, \mathcal{H}_g must also be PAC learnable with the same number of samples. Since functions in \mathcal{H}_g are binary classifiers and PAC learning is completely characterized by VC dimension in the binary setting, we can bound the VC dimension of \mathcal{H}_g . Then, since \mathcal{H}_g has small VC dimension, it is possible to bound the maximum difference between the sampled demand on N and the new demand on N' for any $h_{g,p} \in \mathcal{H}_g$ by $\tilde{O}(\text{VC}(\mathcal{H}_g) \cdot \sqrt{s_g})$ (Theorem 6.13, whose proof follows standard arguments for bounding PAC sample complexity, using Bernstein's inequality in place of Hoeffding's inequality to achieve an upper bound on error which scales as $\sqrt{s_g}$ rather than \sqrt{n}).

Theorem 6.11. *The VC dimension of \mathcal{H}_g is at most $O(m)$ if the tie-breaking rule e is encodable.*

The proof of Theorem 6.11 relies on the following lemma about the linear separability of the class.

Lemma 6.12. *\mathcal{H} is $(m+1)$ -linearly separable, if the tie-breaking rule e is encodable.*

of Theorem 6.11. Lemma 6.12 states that \mathcal{H} , the class which predicts bundles, is $(m+1)$ -linearly separable. Then, Theorem 6.10 implies that there exists a compression scheme of size $m+1$ for \mathcal{H} . Thus, one can PAC-learn \mathcal{H} with at most $n = O\left(\frac{m \ln \frac{1}{\epsilon} + \ln \frac{1}{\delta}}{\epsilon}\right)$ samples, by Theorem 6.8. Since \mathcal{H}_g is a projection of \mathcal{H} onto its the g th coordinate, PAC-learning \mathcal{H}_g cannot require more samples than PAC-learning \mathcal{H} , so this sample complexity bound on n also applies to \mathcal{H}_g . Then, Theorem 6.4 implies that the VC dimension of the binary class \mathcal{H}_g satisfies

$$c \cdot \left(\frac{m + \ln \frac{1}{\delta}}{\epsilon} \right) \geq \frac{\text{VC}(\mathcal{H}_g) + \ln \frac{1}{\delta}}{\epsilon}$$

for some constant c . The claim follows by basic algebra. \square

We briefly sketch the proof of Lemma 6.12 (the full proof can be found in Appendix F.1). In order to show \mathcal{H} is $(m+1)$ -linearly separable, we must define two things. First, we define a mapping $\Psi : \mathcal{V} \times \mathcal{G} \rightarrow \mathbb{R}^{m+1}$ where $\Psi(v, B)_g$ encodes whether or not $g \in B$ for each good $g \in [m]$, and $\Psi(v, B)_{m+1} = v(B)$ encodes the buyer's value for a bundle B . Second, we define a weight vector $w^p \in \mathbb{R}^{m+1}$ for each price vector p where $w_g^p = -p_g$ encodes the *cost* of good $g \in [m]$ at these prices, while $w_{m+1}^p = 1$. Accordingly, the dot product of $\Psi(v, B)$ and w^p encodes the utility of a buyer v buying a bundle B at prices p : $\Psi(v, B) \cdot w^p = v(B) - p(B)$. Thus, a utility-maximizing bundle B^* maximizes the dot product $B^* \in \text{argmax}_B \Psi(v, B) \cdot w^p$.

Unfortunately, there are two obstacles with this plan of attack. First, the statement of Theorem 6.10 assumes that the maximum is *unique*—if not, prediction is not even well-defined. Second, in order to obtain welfare guarantees, we need some assumptions on how ties are broken, namely, that a buyer buys her canonical bundle. To solve both problems simultaneously, we describe how to perturb Ψ and w^p to ensure that the argmax is unique, and that $B^* = \text{argmax}_B (\Psi(v, B) \cdot w^p) = e(\text{argmax}_B v(B) - p(B))$; namely, that ties are broken appropriately by the linear mapping Ψ .

The second ingredient in proving the demand for good g on N is close to N' is the following theorem, which states that the empirical average of $\ell(f, \cdot)$ on a sample N should be close to the expectation of ℓ on the distribution Π for all functions $f \in \mathcal{F}$, for any function $\ell : \mathcal{F} \times (\mathcal{V}, \mathcal{G}) \rightarrow \{0, 1\}$.

Theorem 6.13. *Consider a hypothesis class \mathcal{F} , and $\ell(f, x) \in \{0, 1\}$. Let $N \sim \Pi$ be a sample such that $|N| = n$. Then, with probability $1 - \delta$, for all $f \in \mathcal{F}$,*

$$\ell_{\Pi}(f) - \ell_N(f) \leq \sqrt{\frac{2\text{VC}(\mathcal{F})\ell_N(f) \ln \frac{1}{\delta}}{n}} + \frac{\text{VC}(\mathcal{F}) \ln \frac{1}{\delta}}{n} \leq \frac{3\text{VC}(\mathcal{F}) \ln \frac{1}{\delta} \sqrt{\ell_N(f)}}{\sqrt{n}}.$$

The bound on the distance depends on the size of the sample N as well as the VC dimension of the class \mathcal{F} . The proof of this theorem and the formal proof of Theorem 6.2 can be found in the appendix; the argument follows straightforwardly from [Shalev-Shwartz and Ben-David \[2014\]](#) (specifically, Theorem 6.11 replacing the bound on the difference in losses for a fixed hypothesis by the bound in Lemma B.10).

6.5 Welfare Concentration

Somewhat unusually, our proof of Theorem 6.2 did not go through a combinatorial shattering argument to prove a bound on VC dimension. In contrast, our proof that the *welfare* of Walrasian prices generalizes (Theorem 6.3) relies on an explicit shattering argument. First, we bound the pseudo-dimension of the class of welfare predictors \mathcal{H}_V (Theorem 6.14). This follows from a combinatorial shattering argument using the linear separator for \mathcal{H} . Then, we show that classes with small pseudo-dimension yield strong concentration (Theorem F.7). Similar to the results for over-demand, the latter statement follows from a standard sample complexity argument using Bernstein's inequality in place of Hoeffding's inequality to get a multiplicative bound. We then show that optimal welfare is well-concentrated (Lemma 6.16). Thus, if the welfare on N' for prices p is close to the optimal welfare for N , then the welfare of N' at p must be nearly optimal for N' .

We begin by proving a bound on the pseudo-dimension of the welfare predictors \mathcal{H}_V .

Theorem 6.14. *The pseudo-dimension of \mathcal{H}_V is $O(m^2 \ln m)$ for general valuations \mathcal{V} and $O(m \ln^3 m)$ for unit demand valuations, if e is encodable.*

The proof of Theorem 6.14 (see Appendix F.3) relies on the following combinatorial lemma.

Lemma 6.15. *\mathcal{H} can induce at most $\binom{n}{m+1} \cdot 2^{m(m+1)}$ bundle labelings of n sampled valuations, and at most $\binom{n}{m+1} \cdot m^{m+1}$ for unit demand buyers, if e is encodable.*

Proof. Let $h_p(N) = (h_p(v_1), \dots, h_p(v_n))$ (by abuse of notation) represent the labeling of a sample N by h_p . Lemma 6.12 shows that \mathcal{H} is $(m+1)$ -linearly separable. By Theorem 6.10, there exists a $(m+1)$ -sized compression scheme Ψ for learning the class \mathcal{H} . Thus, there exists a pair of functions (**decompress**, **compress**) such that for any price vector p , there exists $N_p \subseteq N$ with $|N_p| = m+1$ such that (i) **compress** $(h_p(N)) = h_p(N_p)$, (ii) **decompress** $(\text{compress}(h_p(N))) = \text{decompress}(h_p(N_p)) = f$, and $f(v_q) = B_q = h_p(v_q)$ for every $v_q \in N$.

We will bound the number of distinct labelings \mathcal{H} can induce on a sample $N = (v_1, \dots, v_n)$ of size n . To do so, we use the *decompression* function (which depends only on the separator Ψ) to upper-bound the total number of labelings \mathcal{H} can induce on S .

Our approach will be to upper bound the size of the *range* of **decompress** \circ **compress** when restricted to subsamples of a particular population N by some quantity D . By definition of a compression scheme, for every labeling of buyers in N by bundles that can be induced by some price vector, there is an element of the range of **decompress** \circ **compress** that recovers that labeling on N . So, we can upper bound the total number of distinct labelings that can be induced on N by price vectors by upper bounding the range of **decompress** \circ **compress** on N . The idea is that each *compression* is defined by a subset N' of at most $m + 1$ of the buyers in N , and a labeling of the buyers with bundles. So, bounding the number of such labeled subsets suffices to bound the range of **decompress** \circ **compress**.

We prove such a bound for $D = \binom{n}{m+1} \cdot (2^m)^{m+1}$. Fix any subset of unlabeled examples $N' \subseteq N$, where $|N'| = m + 1$. For any $x \in N'$, there are at most 2^m labelings of x from \mathcal{G} . Thus, there are at most $(2^m)^{m+1} = 2^{m(m+1)}$ labelings of all of N' from \mathcal{G} .¹² Thus, **decompress** \circ **compress** can output at most $\binom{n}{m+1} \cdot 2^{m(m+1)}$ distinct functions based on labeled subsets of N . \square

Theorem 6.14, along with known results for uniform convergence over classes with small pseudo-dimension, implies that the welfare of any set of prices p on a sufficiently large sample is very close to the welfare p induces on a fresh sample of the same size. The next lemma shows that the *optimal* welfare allocation on the training and test samples have very similar welfare, since the optimal welfare is a well-concentrated quantity.

Lemma 6.16. *Recall that Opt-Welfare_N denotes the optimal welfare achievable by an allocation of a fixed set of goods in a market to buyers with valuations $N = (v_1, \dots, v_n)$ bounded in $[0, H]$. Then, the optimal welfare on two samples N, N' of size n drawn i.i.d. from distribution Π satisfies*

$$\Pr[|\text{Opt-Welfare}_N - \text{Opt-Welfare}_{N'}| \geq 4\epsilon] \leq \exp\left(-\frac{2\epsilon^2}{nH^2}\right)$$

for all $\epsilon \in (0, 1)$. In particular, $|\text{Opt-Welfare}_N - \text{Opt-Welfare}_{N'}| \leq H\sqrt{n \ln \frac{1}{\delta}}$ with probability at least $1 - \delta$.

The proof of Lemma 6.16 uses McDiarmid's inequality, and can be found in Appendix F.4. We now prove the main result of this section, which guarantees first that for any p , the welfare induced on N is similar to the welfare induced on N' . In particular, if p is the welfare-optimal pricing for N , then it is also nearly welfare-optimal for N' .

of Theorem 6.3. We wish to show that with probability $1 - \delta$,

$$\text{Opt-Welfare}_{N'} - \text{Welfare}_{N'}(p) \leq \alpha \text{Opt-Welfare}_{N'}.$$

¹²If buyers are unit demand, we have a sharper bound: there are at most m labelings per $x \in N'$, or m^{m+1} labelings in total.

We begin by noting that

$$\begin{aligned}
& \text{Opt-Welfare}_{N'} - \text{Welfare}_{N'}(p) \\
& \leq \text{Opt-Welfare}_N - \text{Welfare}_{N'}(p) + H\sqrt{n \ln \frac{1}{\delta}} \quad (\text{Lemma 6.16}) \\
& \leq \text{Welfare}_N(p) + 2mH - \text{Welfare}_{N'}(p) + \sqrt{nH \ln \frac{1}{\delta}} \quad (\text{Theorem 5.27}) \\
& \leq \text{RWelfare}_N(p) + 3mH - \text{Welfare}_{N'}(p) + H\sqrt{n \ln \frac{1}{\delta}} \quad (\text{Theorem 5.26 and Lemma 4.12}) \\
& \leq \text{RWelfare}_N(p) - \text{RWelfare}_{N'}(p) + mH \left(\max_g \sqrt{s_g m \ln \frac{1}{\delta}} + m \ln \frac{1}{\delta} \right) + H\sqrt{n \ln \frac{1}{\delta}} + 3mH \\
& \leq \text{RWelfare}_N(p) - \text{RWelfare}_{N'}(p) + mH \left(\sqrt{nm \ln \frac{1}{\delta}} + m \ln \frac{1}{\delta} \right) + H\sqrt{n \ln \frac{1}{\delta}} + 3mH \quad (7)
\end{aligned}$$

where the second to last step follows by Theorem 6.2 and Lemma 4.12, and the last step by $s_g \leq n$, otherwise, we can replace s_g by n in the bound in Theorem 6.2, since $C^e(N; p; g) \leq n + 1$.

Applying Theorem F.7 to $\sum_{q \in N} v_q(B_q^*(p)) = n \cdot \frac{\sum_{q \in N} v_q(B_q^*(p))}{n}$ and the analogous term for N' , we know that for all pricings p ,

$$\text{RWelfare}_N(p) - \text{RWelfare}_{N'}(p) \leq \left(H^{3/2} \text{PD}(\mathcal{F}) \ln \frac{1}{\delta} \right) \left(\frac{1}{3} + \sqrt{19 \text{RWelfare}_N(p)} \right).$$

Using Theorem 6.14, this reduces to

$$\begin{aligned}
& \text{RWelfare}_N(p) - \text{RWelfare}_{N'}(p) \\
& \leq \left(H^{3/2} m^2 \ln m \ln \frac{1}{\delta} \right) \left(\frac{1}{3} + \sqrt{19 \text{RWelfare}_N(p)} \right) \quad (\text{Theorem 6.14's bound on } \text{PD}(\mathcal{H}_V)) \\
& \leq \left(H^{3/2} m^2 \ln m \ln \frac{1}{\delta} \right) \left(2\sqrt{19 \text{RWelfare}_N(p)} \right) \\
& \leq \left(H^{3/2} m^2 \ln m \ln \frac{1}{\delta} \right) \left(2\sqrt{19 \text{Welfare}_N(p) + mH} \right) \quad (\text{Theorem 5.26 and Lemma 4.12}) \\
& \leq \left(H^{3/2} m^2 \ln m \ln \frac{1}{\delta} \right) \left(2\sqrt{19 \text{Opt-Welfare}_N + mH} \right) \quad (\text{Welfare}_N(p) \text{ is feasible}) \\
& \leq \left(H^{3/2} m^2 \ln m \ln \frac{1}{\delta} \right) \left(2\sqrt{19 \text{Opt-Welfare}_{N'} + H\sqrt{n \ln \frac{1}{\delta}} + mH} \right). \quad (\text{Lemma 6.16})
\end{aligned}$$

Combining with Eq. (7), we have

$$\begin{aligned}
\text{Opt-Welfare}_{N'} - \text{Welfare}_{N'}(p) & \leq \left(H^{3/2} m^2 \ln m \ln \frac{1}{\delta} \right) \left(2\sqrt{19 \text{Opt-Welfare}_{N'} + H\sqrt{n \ln \frac{1}{\delta}} + mH} \right) \\
& \quad + mH \left(\sqrt{nm \ln \frac{1}{\delta}} + m \ln \frac{1}{\delta} \right) + H\sqrt{n \ln \frac{1}{\delta}} + 3mH.
\end{aligned}$$

If we wish for this to be at most $\alpha \text{Opt-Welfare}_{N'}$, it suffices for

$$\text{Opt-Welfare}_{N'} = \Omega \left(\frac{H^3 n^{0.5} m^4 \ln^2(m) \ln^2 \frac{1}{\delta}}{\alpha^2} \right).$$

With probability $1 - \delta$, we know that $\text{Opt-Welfare}_N - \text{Opt-Welfare}_{N'} \leq H \sqrt{n \ln \frac{1}{\delta}}$, so this is satisfied by $\text{Opt-Welfare}_{N'}$ whenever it is satisfied by Opt-Welfare_N with probability $1 - \delta$, which holds by assumption. \square

sketch of Theorem 4.15. This proof follows the identical calculation from Theorem 6.3, replacing the bound on pseudo-dimension by $m \ln^3 m$, as is implied by Theorem 6.14 when buyers are unit demand. \square

Finally, we sketch the proof of Theorem 6.3. The optimal welfare for the two markets must be close, by Lemma 6.16. The welfare of the pricing p output for the first market is nearly optimal even when over-demand is resolved adversarially, since there is over-demand of at most 1 for each good by assumption. By Theorem 6.14, the pseudo-dimension of the class is at most $m^2 \ln^2 m$ for general valuations and at most $m \ln^3 m$ for unit-demand valuations. A pseudo-dimension analogue of Theorem 6.13 implies that the welfare of applying p to the second market will be close to the welfare of applying p to the initial market.

6.6 Lower bounds for learning

We now show that the VC dimension of \mathcal{H}_g , the class of bundle predictors for a fixed good g (at a pricing p) is at least m , implying our bound on the VC dimension of \mathcal{H}_g from the previous section is tight up to constant factors.

Theorem 6.17. *Let V be the set of unit demand valuations. Then $\text{VC}(\mathcal{H}_g) \geq m$ and $\text{PD}(\mathcal{H}_V) \geq m$, for $\mathcal{H}_g, \mathcal{H}_V$ over valuations \mathcal{V} .*

Proof. Fix a particular good g . We will show a set $N = (v_1, \dots, v_n)$ of unit demand buyer valuations which can be shattered by both \mathcal{H}_g and \mathcal{H}_V . Suppose $n = m$ and label buyers such that $q \in N$ also corresponds to some $q \in [m]$. Then, for all $q \neq g$, define $v_q(\{q\}) = 2$, $v_q(\{g\}) = 1$, and $v_q(\{g'\}) = 0$ for all goods $g' \notin \{q, g\}$. Let $v_g(\{g\}) = \frac{1}{2}$ and $v_g(\{g'\}) = 0$. The intuition for this construction is that buyer q will buy good g if the price of their “individual” good q is sufficiently larger than the price of g ($p_q > p_g + 1$), and buyer g will buy good g if the price is less than her value ($p_g \leq v_g$). Consider an arbitrary set $N' \subseteq N$. We will first show that \mathcal{H}_g can label N' positive and $N \setminus N'$ negative. We define prices p as follows.

If $g \in N'$, then let $p_g = 0$ and $p_{g'} = 1 + 2\epsilon$ if $g' \in N'$ and $p_{g'} = \epsilon$ if $g' \notin N'$. Then, for buyer g , $v_g(\{g\}) = \frac{1}{2} > p_g = 0$, and $v_g(\{g'\}) = 0 < p_{g'}$ for all other g' ; thus, $\{g\}$ is the unique bundle in $\mathcal{D}_g(p)$ (and so $h_{g,p}(v_g) = 1$). Then, for $g' \in N'$, $v_{g'}(\{g'\}) - p_{g'} = 2 - (1 + 2\epsilon) < 1 = v_{g'}(g) - p_g$, thus $\{g\}$ is the unique bundle in $\mathcal{D}_{g'}(p)$ (so, $h_{g,p}(v_{g'}) = 1$). For any $g' \notin N'$, $v_{g'}(\{g'\}) - p_{g'} = 2 - \epsilon > 1 = v_{g'}(\{g\}) - p_g$; thus, $\{g'\}$ is the unique bundle in $\mathcal{D}_{g'}(p)$ and so $h_{g,p}(v_{g'}) = 0$.

If, on the other hand, $g \notin N'$, let $p_g = \frac{1}{2} + \epsilon$, $p_{g'} = \frac{3}{2} + 2\epsilon$ if $g' \in N'$ and $p_{g'} = \epsilon$ if $g' \notin N'$. Then, since $v_g(\{g\}) = \frac{1}{2} < \frac{1}{2} + \epsilon = p_g$, $h_{g,p}(v_g) = 0$. For $g' \in N'$, $v_{g'}(\{g\}) - p_g = 1 - \frac{1}{2} - \epsilon > 2 - \frac{3}{2} - 2\epsilon = v_{g'}(\{g'\}) - p_{g'}$, thus $\{g\}$ is the unique bundle in $\mathcal{D}_{g'}(p)$ and so $h_{g,p}(v_{g'}) = 1$. For

$g' \notin N'$, $v_{g'}(\{g\}) - p_g = 1 - \frac{1}{2} - \epsilon < 2 - \epsilon = v_{g'}(\{g'\}) - p_{g'}$, so $\{g'\}$ is the unique bundle in $\mathcal{D}_{g'}(p)$, and so $h_{g,p}(v_{g'}) = 0$.

Thus, for arbitrary $N' \subseteq N$, we can choose prices p such that $h_{g,p}(v) = 1$ if and only if $v \in N'$, so we have shown how to shatter N with \mathcal{H}_g .

The same N is shatterable by \mathcal{H}_V , as we now argue. Consider (r_1, \dots, r_n) a set of targets, with $r_g = \frac{1}{2}$ and $r_{g'} = \frac{3}{2}$ for all $g' \neq g$. Thus, the target is hit for v_g only when buyer g buys good g (and for $v_{g'}$ when buyer g' buys good g'). For a set of buyers $N' \subseteq N$, set the prices p as above for \mathcal{H}_g , but swap the definition of $p_{g'}$ for $g' \in N'$ with that of $p_{g'}$ for $g' \notin N'$. In the previous definition of p , whenever $g' \in N'$, buyer g' bought g and not g' , and when $g' \notin N'$ buyer g' bought g' and not g . So, when these are swapped, g' will buy g' if and only if $g' \in N'$ (thus, hitting the welfare target of $r_{g'}$ if and only if $g' \in N'$). For buyer g , when $g \in N'$, the prices above cause buyer g to buy good g , hitting welfare target r_g (and, when $g \notin N'$, g didn't buy g , therefore missing welfare target r_g). Thus, for this modified definition of p , all buyers $q \in N'$ will have $f_p(v_q) \geq r_q$ and for all $q \notin N'$, $f_p(v_q) < r_q$. Thus, we can shatter N according to these targets. \square

Acknowledgments We warmly thank Renato Paes Leme, Kazuo Murota, Akiyoshi Shioura, and the participants at the Workshop on Complexity and Simplicity in Economics at the Simons Institute for the Theory of Computing for insightful discussions on earlier versions of this work. We also thank the anonymous reviewers for their helpful comments.

References

- A dynamic near-optimal algorithm for online linear programming. *Operations Research*, 62(4): 876–890, 2014.
- Martin Anthony and Peter L Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.
- Kenneth J Arrow and Gerard Debreu. [Existence of an equilibrium for a competitive economy](#). *Econometrica*, pages 265–290, 1954.
- Itai Ashlagi, Mark Braverman, and Avinatan Hassidim. [Ascending unit demand auctions with budget limits](#). Technical report, 2009.
- Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. [Posting prices with unknown distributions](#). In *Symposium on Innovations in Computer Science, Beijing, China*. Tsinghua University Press, January 2011.
- Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. [Mechanism design via machine learning](#). In *IEEE Symposium on Foundations of Computer Science (FOCS), Pittsburgh, Pennsylvania*, pages 605–614, 2005.
- Maria-Florina Balcan, Nikhil Devanur, Jason D Hartline, and Kunal Talwar. [Random sampling auctions for limited supply](#). Technical report, Carnegie Mellon University, September 2007.
- Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. [Reducing mechanism design to algorithm design via machine learning](#). *Journal of Computer and System Sciences*, 74(8):1245–1270, 2008.
- Maria-Florina Balcan, Amit Daniely, Ruta Mehta, Ruth Urner, and Vijay V Vazirani. [Learning economic parameters from revealed preferences](#). In *Conference on Web and Internet Economics (WINE), Beijing, China*, pages 338–353, 2014.
- Oren Ben-Zwi, Ron Lavi, and Ilan Newman. [Ascending auctions and walrasian equilibrium](#). 2013.
- Avrim Blum and Jason D Hartline. [Near-optimal online auctions](#). In *ACM–SIAM Symposium on Discrete Algorithms (SODA), Vancouver, British Columbia*, pages 1156–1163, 2005.
- Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. [Online learning in online auctions](#). In *ACM–SIAM Symposium on Discrete Algorithms (SODA), Baltimore, Maryland*, 2003.
- Mark Braverman, Jing Chen, and Sampath Kannan. [Optimal provision-after-wait in healthcare](#). In *ACM SIGACT Innovations in Theoretical Computer Science (ITCS), Princeton, New Jersey*, pages 541–542, 2014.
- Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. [Regret minimization for reserve prices in second-price auctions](#). In *ACM–SIAM Symposium on Discrete Algorithms (SODA), New York, New York*, pages 1190–1204, 2013.
- Richard Cole and Tim Roughgarden. [The sample complexity of revenue maximization](#). In *ACM SIGACT Symposium on Theory of Computing (STOC), New York, New York*, pages 243–252, 2014.

- Amit Daniely and Shai Shalev-Shwartz. [Optimal learners for multiclass problems](#). In *Conference on Computational Learning Theory (CoLT), Barcelona, Spain*, pages 287–316, 2014.
- Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. [On the complexity of equilibria](#). In *ACM SIGACT Symposium on Theory of Computing (STOC), Montréal, Québec*, pages 67–71, 2002.
- Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. [Prior-independent multi-parameter mechanism design](#). In *Conference on Web and Internet Economics (WINE), Singapore*, pages 122–133, 2011.
- Nikhil R Devanur and Thomas P Hayes. [The Adwords problem: Online keyword matching with budgeted bidders under random permutations](#). In *ACM SIGecom Conference on Economics and Computation (EC), Stanford, California*, pages 71–78, 2009a.
- Nikhil R. Devanur and Thomas P. Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM Conference on Electronic Commerce, EC '09*, pages 71–78, New York, NY, USA, 2009b. ACM.
- Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. [Revenue maximization with a single sample](#). *Games and Economic Behavior*, 2014.
- Shaddin Dughmi, Li Han, and Noam Nisan. [Sampling and representation complexity of revenue maximization](#). In *Conference on Web and Internet Economics (WINE), Beijing, China*, pages 277–291, 2014.
- Andrzej Ehrenfeucht, David Haussler, Michael Kearns, and Leslie Valiant. [A general lower bound on the number of examples needed for learning](#). *Information and Computation*, 82(3):247–261, 1989.
- Edith Elkind. [Designing and learning optimal finite support auctions](#). In *ACM SIGecom Conference on Economics and Computation (EC), San Diego, California*, pages 736–745, 2007.
- Satoru Fujishige and Zaifu Yang. [A note on Kelso and Crawford’s Gross Substitutes condition](#). *Mathematics of Operations Research*, 28(3):463–469, 2003.
- Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. [Competitive auctions](#). *Games and Economic Behavior*, 55(2):242–269, 2006.
- Faruk Gul and Ennio Stacchetti. [The English auction with differentiated commodities](#). *Journal of Economic Theory*, 92(1):66–95, 2000.
- Faruk Gul and Ennio Stacchetti. [Walrasian equilibrium with gross substitutes](#). *Journal of Economic Theory*, 87(1):95–124, 1999.
- Steve Hanneke. [The optimal sample complexity of PAC learning](#). 2015.
- Jason D Hartline and Tim Roughgarden. [Simple versus optimal mechanisms](#). In *ACM SIGecom Conference on Economics and Computation (EC), Stanford, California*, 2009.
- Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. [Making the most of your samples](#). In *ACM SIGecom Conference on Economics and Computation (EC), Portland, Oregon*, pages 45–60, 2015.

- Alexander Kelso and Vincent Crawford. [Job matching, coalition formation, and gross substitutes](#). *Econometrica*, 50(6):1483–1504, 1982.
- Nick Littlestone and Manfred Warmuth. [Relating data compression and learnability](#). Technical report, University of California, Santa Cruz, 1986.
- Andres Munoz Medina and Mehryar Mohri. [Learning theory and algorithms for revenue optimization in second price auctions with reserve](#). In *International Conference on Machine Learning (ICML), Beijing, China*, pages 262–270, 2014.
- Marco Molinaro and R Ravi. [The geometry of online packing linear programs](#). *Mathematics of Operations Research*, 39(1):46–59, 2013.
- Jamie Morgenstern and Tim Roughgarden. [The pseudo-dimension of nearly-optimal auctions](#). In *Conference on Neural Information Processing Systems (NIPS), Montréal, Québec*, pages 136–144, 2015.
- Kenneth Mount and Stanley Reiter. [The informational size of message spaces](#). *Journal of Economic Theory*, 8(2):161–192, 1974.
- Kazuo Murota. [Convexity and Steinitz’s exchange property](#). In William H Cunningham, S Thomas McCormick, and Maurice Queyranne, editors, *Conference on Integer Programming and Combinatorial Optimization (IPCO), Vancouver, British Columbia*, volume 1084 of *Lecture Notes in Computer Science*, pages 260–274. Springer Berlin Heidelberg, 1996a.
- Kazuo Murota. [Valuated matroid intersection I: Optimality criteria](#). *SIAM Journal on Discrete Mathematics*, 9(4):545–561, 1996b.
- Kazuo Murota and Akiyoshi Shioura. [M-Convex function on generalized polymatroid](#). *Mathematics of Operations Research*, 24(1):95–105, 1999.
- Noam Nisan and Ilya Segal. [The communication requirements of efficient allocations and supporting prices](#). *Journal of Economic Theory*, 129(1):192–224, 2006.
- Michael Ostrovsky and Renato Paes Leme. [Gross substitutes and endowed assignment valuations](#). *Theoretical Economics*, 10(3):853–865, September 2015.
- James G Oxley. *Matroid Theory*. Oxford Graduate Texts in Mathematics. Oxford University Press, 1997.
- David Pollard. *Convergence of stochastic processes*. David Pollard, New Haven, Connecticut, 1984.
- Tim Roughgarden and Okke Schrijvers. [Ironing in the dark](#). Under submission, 2015.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- Akiyoshi Shioura and Akihisa Tamura. [Gross substitutes condition and discrete concavity for multi-unit valuations: a survey](#). *Journal of the Operations Research Society of Japan*, 58(1):61–103, January 2015.

Vladimir N Vapnik and A Ya Chervonenkis. [On the uniform convergence of relative frequencies of events to their probabilities.](#) *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

Vladimir Naumovich Vapnik and Samuel Kotz. *Estimation of dependences based on empirical data.* Springer-Verlag, 1982.

A Valuations over bundles of copies versus bundles of goods

We make clear in this section that although the market that buyers face includes the copies s_g of goods $g \in [m]$, we can instead focus on valuations defined on the market of just the goods $[m]$ when buyers demand at most one copy of each good. We closely follow the work of [Shioura and Tamura \[2015\]](#) to work with multiple copy valuation functions, which are functions on integer lattice points. We define the set of all allocations in the multiple copy market (where each good $g \in [m]$ has s_g many copies) as

$$X = \{0, 1, \dots, s_1\} \times \dots \times \{0, 1, \dots, s_m\}.$$

Thus each buyer $q \in N$ has a valuation $v_q : X \rightarrow [0, H]$. As we do in the main body of the paper, we assume that the valuations are monotone. We will denote $\vec{x} \in X$ where

$$\vec{x} = (\vec{x}(g) : g \in [m]).$$

Note that buyers demand at most one copy of each good, so we define the projection $\pi : X \rightarrow \{0, 1\}^m$ as

$$\pi(\vec{x}) = (\min\{\vec{x}(1), 1\}, \min\{\vec{x}(2), 1\}, \dots, \min\{\vec{x}(m), 1\})$$

The set of valuations we consider can be written as

$$V^X = \{v : X \rightarrow [0, H] : v(\vec{x}) = v(\pi(\vec{x}))\}.$$

We now define the demand correspondence $\mathcal{D}^X(p; v_q)$ and Walrasian equilibrium in this setting where buyers have valuations $v_q \in V^X$ for each $q \in N$.

Definition A.1 (Demand Correspondence). *Let buyer q have valuation $v_q : X \rightarrow [0, H]$. We then define the demand correspondence for any price p as*

$$\mathcal{D}^X(p; v_q) = \operatorname{argmax}_{\vec{x} \in X} \{v_q(\vec{x}) - p^T \vec{x}\}.$$

Definition A.2 (WE-X). *A Walrasian Equilibrium for buyers with valuations $v_q : X \rightarrow [0, H]$ is a tuple $(p, (\vec{x}_q : q \in N))$ where*

- $\vec{x}_q \in \mathcal{D}^X(p; v_q)$
- $\sum_{q \in N} \vec{x}_q(g) \leq s_g$ for any $g \in [m]$, and if the inequality is strict, then the price for that good is zero.

We will write $\chi_S \in \{0, 1\}^m$ as the characteristic vector of $S \subseteq [m]$. Note that if (p, μ) is a WE as we defined in [Definition 2.3](#), then when we write the characteristic function $\vec{x}_q = \chi_{\mu_q}$ for each $q \in N$ we will have $(p, (\vec{x}_q : q \in N))$ is a WE-X as we defined above. Thus, we only need to focus on valuations defined over the boolean cube $\{0, 1\}^m$ or equivalently over the subsets of $[m]$.

B Non-minimal Walrasian prices and genericity

We show here that even in a market that contains only generic unit demand buyers, the over-demand at a non-minimal Walrasian price vector can be high. This is in contrast to our result in Theorem 4.10, in which we showed that the over-demand can be at most 1 at *minimal* Walrasian prices when buyers have generic valuations.

Lemma B.1. *There exists generic valuations such that the over-demand $\max_{g \in [m]} OD_p(g, \mu) = n - 1$ for some Walrasian equilibrium (p, μ) where p is not minimal.*

Proof. We consider a slight variant of valuations given in the proof of Lemma 3.1. Consider n unit demand agents $N = [n]$ with generic valuations and good types $[n]$ where each type has a single copy. We assume that each buyer q 's generic valuation satisfies the following for some special good $g \in [m]$

$$v_q(h) < v_q(g) < v_q(q) \quad \forall h \neq q, g$$

Consider the prices $p_q = v_q(q) - v_q(g)$ for $q \in [n]$. This price vector still satisfies the maximum allocation where every buyer q gets good q , but it now has every buyer q have an indifference between good g and q . Hence, the over-demand for good g at these prices is $n - 1$. We know that p is not minimal due to Theorem 4.10. \square

C Minimal demand bundles form a matroid basis

To establish this fact (Lemma 5.3), we look to Ben-Zwi et al. [2013]. They prove essentially the same lemma, but only when valuations are rational numbers. Since we require genericity, it may not be reasonable to also assume that valuations are rational. Fortunately, it is straightforward to modify the proof by Ben-Zwi et al. [2013] for the general case, which we do here. We begin by showing the following claim which will help in our analysis.

Claim C.1. *If $B, B' \in \mathcal{D}^*(p)$ then $|B| = |B'|$.*

Proof. Let $D_1, D_2 \in \mathcal{D}^*(p)$. Let $p' = p$, except $p'(g) = \infty$ for $g \notin D_1 \cup D_2$. Note that $D_1, D_2 \in \mathcal{D}^*(p')$. We define a small quantity ϵ as

$$\epsilon = \min_{\substack{B, B' \subseteq [m] \\ |B \setminus B'|, |B' \setminus B| \leq 1}} \{u(B; p) - u(B'; p) : u(B; p) - u(B'; p) > 0\}.$$

Let $g_2 \in D_2 \setminus D_1$ and define $p'_2 = p'$ except $p'_2(g_2) = p'(g_2) + \epsilon$. We then know that $D_2 \notin \mathcal{D}(p'_2)$ because $u(D_1; p'_2) = u(D_1; p') = u(D_2; p') > u(D_2; p'_2)$. Due to the single improvement property of GS valuations, we know that there exists D_3 such that both $|D_3 \setminus D_2|, |D_2 \setminus D_3| \leq 1$ and $u(D_3; p'_2) > u(D_2; p'_2)$. We then have

$$\begin{aligned} u(D_3; p'_2) > u(D_2; p'_2) &\implies u(D_3; p) - \epsilon \mathbf{1}\{g_2 \in D_3\} > u(D_2; p) - \epsilon \\ &\implies \epsilon - \epsilon \mathbf{1}\{g_2 \in D_3\} > u(D_2; p) - u(D_3; p). \end{aligned} \tag{8}$$

We know that $D_2 \in \mathcal{D}(p)$ so that $u(D_2; p) - u(D_3; p) \geq 0$, which tells us that $g_2 \notin D_3$. Further, we know that if $u(D_2; p) - u(D_3; p) > 0$ then $u(D_2; p) - u(D_3; p) \geq \epsilon$. However, Eq. (8) tells us that

$u(D_2; p) - u(D_3; p) < \epsilon$, thus a contradiction. We then have $u(D_3; p) = u(D_2; p)$. Hence, we have shown that $g_2 \in D_3$, $D_3 \in \mathcal{D}(p)$ and D_3 is not contained in D_2 .

Thus, $\exists g_1 \notin D_2$ such that $g_1 \in D_3 \subseteq D_1 \cup D_2$ which implies that $g_1 \in D_1 \setminus D_2$. Thus, we can say that $D_3 = D_2 \cup g_1 \setminus g_2$. We continue with induction on $|D_2 \setminus D_1|$ to conclude that $|D_2| = |D_1|$. \square

Proof of Lemma 5.3. We now let $\widehat{\mathcal{D}}(p) = \{S \in \mathcal{D}(p) : |S| \leq |T|, \forall T \in \mathcal{D}(p)\}$, which we know forms the bases of some matroid from Gul and Stacchetti [2000]. We first show that $\mathcal{D}^*(p) \subseteq \widehat{\mathcal{D}}(p)$. If $D \in \mathcal{D}^*(p)$ then we know that $|D| = |D^*|$ for all $D^* \in \mathcal{D}^*(p)$ by the previous claim. Hence, $|D| \leq |B|$ for all $B \in \mathcal{D}(p)$, i.e. $D \in \widehat{\mathcal{D}}(p)$.

We now show that $\widehat{\mathcal{D}}(p) \subseteq \mathcal{D}^*(p)$. Assume that $D \in \widehat{\mathcal{D}}(p)$, but $D \notin \mathcal{D}^*(p)$. We then know that there exists a $D^* \subsetneq D$ where $D^* \in \mathcal{D}^*(p)$. This must mean that $|D^*| < |D|$, but this contradicts the fact that $D \in \widehat{\mathcal{D}}(p)$, i.e. $|D| \leq |B|$ for every $B \in \mathcal{D}(p)$. Thus, we have a contradiction, so that $D \in \mathcal{D}^*(p)$. Further, we have shown that $\mathcal{D}^*(p) = \widehat{\mathcal{D}}(p)$. \square

D Constructing genericity via perturbation

When working with matchings (Section 4), we assumed that valuations are generic in the sense of Definition 4.4. Instead, we can imagine perturbing the valuations slightly to achieving genericity. While “any” continuous perturbation should achieve genericity, we can also give a more concrete perturbation procedure. While it is difficult to guarantee a definition like Definition 4.4, we can limit over-demand using a relaxed version of genericity.

Let’s suppose that bidders start with valuations $v_i(j)$ (possibly non-generic) and minimal Walrasian prices p . Each bidder will perturb their valuation by selecting a uniformly random element of $\mathcal{P} \subseteq \mathbb{R}$, a (finite) set of perturbations. We will continue to assume that the maximum matching μ is unique. We will write Δ for the smallest positive difference social welfare between two allocations—possibly not allocating all goods, or infeasible. For instance, this assumption holds in the typical case where all valuations are given with finite precision, when we can treat all valuations as integer multiples of Δ .

Our goal is to show that with high probability, the perturbed values $\widehat{v_i(j)}$ and corresponding minimal prices \widehat{p} have limited over-demand when bidders select an arbitrary most-demanded good. Our analysis will begin with the original valuations $v_i(j)$ and prices p . As we have done throughout, we consider the swap graph. By identical reasoning as before, the graph is acyclic and the nodes can be partially ordered. We will consider the nodes according to this order, using the principle of deferred decision to successively perturb the valuations. As we go, we will construct a modified set of prices \widehat{p}' that guarantees each good is over-demanded by at most 1. Finally, we will argue that \widehat{p}' is indeed the minimal Walrasian prices for $\widehat{v_i(j)}$, tying the knot.

Let the goods be numbered $1, \dots, k$ according to the partial order in the original swap graph G ; for convenience, we will also label bidders with the same index according to their matched good in the original matching μ . We will construct the modified prices \widehat{p}' inductively. We begin by setting $\widehat{p}_1 = 0$, and sampling bidder 1’s perturbation for good 1 but leaving the other perturbations unsampled, by the principle of deferred decisions.

For the inductive case, we have prices \widehat{p} for goods in $[j]$ and sampled bidder valuations by bidders in $[j]$ for goods in $[j]$. To set the price of good $j + 1$, we first have bidder $j + 1$ sample valuations for all goods in $[j + 1]$. Now, we need to specify the valuations of bidders in $[j]$ for the

new good $j + 1$. We consider a set of candidate price functions r_1, \dots, r_j , where

$$r_l(v) := v - \widehat{v}_l(l) + \sum_{(a,b) \in C} \widehat{v}_a(b) - \widehat{v}_a(a)$$

and C is a directed path from a source to l in \widehat{G} (recall that \widehat{G} is acyclic). Roughly, r_l maps the valuation of bidder l for $j + 1$ to a price for l at which bidder l becomes indifferent between l and $j + 1$. Note that by induction, $l \in [j]$ and we have already fixed all \widehat{v} when defining r_l . We now sample $v_l(\widehat{j + 1})$ for $l \in [j]$, and define

$$p'_{j+1} := \max_l(0, r_l(v_l(\widehat{j + 1}))).$$

That is, if all r_l are negative, we define $p'_{j+1} := 0$.

The key point is that in the new swap graph \widehat{G} , there is an edge from l to $j + 1$ exactly when l is binding (i.e., in the argmax) in the definition of p'_{j+1} . Since $r_l(v)$ is a linear function, if the number of possible perturbations $|\mathcal{P}|$ is large enough, we will be unlikely to have a collision. In other words, with high probability the in-degree of $j + 1$ will either be 1, or zero with price 0.

First, we show that the new swap graph \widehat{G} is a subgraph of the old swap graph \widehat{G} , and that the prices p' are close to p .

Lemma D.1. *For each good j , the following two statements hold:*

1. *If C is a path in \widehat{G} ending at j , then C is also a path in G .*
2. *Suppose that all perturbations in \mathcal{P} are bounded by $\Delta/2k$. Then, $|p'_j - p_j| < \Delta j/k$ for each good j .*

Proof. We prove both points simultaneously by induction on j . The base case $j = 1$ is clear: the good 1 is a source node in both graphs, and has price 0.

For the inductive case, we prove the first point first. If j is a source node in \widehat{G} then the first point is trivial. Otherwise, there is a path C ending at j in \widehat{G} . Suppose the node before j in C is l . By induction, the segment of C ending at l is a path in G . So, we just need to show that there is an edge (l, j) in G .

By construction of p' , we know

$$p'_j = \widehat{v}_l(j) - \widehat{v}_l(l) + p'_l \geq 0.$$

If there is no edge (l, j) , then

$$(v_l(l) - p_l) - (v_l(j) + p_j) > 0$$

We can write the price as a sum of differences of valuations (Lemma 4.8), and collecting the positive and negative parts, the above equation shows that the difference in welfare of two (possibly partial) allocations is non-negative. By assumption, the difference is at least $\Delta > 0$, and hence

$$(v_l(l) - p_l) - (v_l(j) + p_j) > \Delta \quad \text{so} \quad v_l(j) - v_l(l) + p_l < -\Delta - p_j \leq -\Delta.$$

This is a contradiction: by induction, $|p'_l - p_l| \leq \Delta l/k$, and $v_l(l), v_j(l)$ differ from $\widehat{v}_l(l), \widehat{v}_l(j)$ by at most $\Delta/2k$ at most, and $l \leq j - 1 \leq k - 1$.

Now, for the second point. Suppose that j is a source node in \widehat{G} . So, $p'_j = 0$. Suppose that there is a non-empty path C from a source node to j in G (if not, then $p_j = p'_j = 0$), and suppose C hits l before landing at j , so

$$p_j = v_l(j) - v_l(l) + p_l.$$

By induction, we know $|p'_l - p_l| < \Delta l/k$. Furthermore, we know

$$r_l(\widehat{v_l(j)}) := \widehat{v_l(j)} - \widehat{v_l(l)} + p'_l < 0.$$

Since $\widehat{v_l(j)}$ and $\widehat{v_l(l)}$ are at most $\Delta/2k$ away from $v_l(j)$ and $v_l(l)$ respectively, we know $p_j \leq \Delta/2k + \Delta/2k + \Delta l/k \leq \Delta j/k$ by induction, and since $l \leq j - 1$.

Otherwise, j is not a source node in \widehat{G} . Suppose that C is the shortest path from a source node to j in \widehat{G} , with l the node before j in C . We have:

$$p'_j = r_l(\widehat{v_l(j)}) = \widehat{v_l(j)} - \widehat{v_l(l)} + p'_l.$$

We know that C is also a path from a source node to j in G , so

$$p_j = v_l(j) - v_l(l) + p_l.$$

The claim follows by induction. □

Finally, we can show that p' is indeed the minimal Walrasian prices for the perturbed valuations.

Lemma D.2. *Suppose that all perturbations in \mathcal{P} are bounded by $\Delta/2k$. The prices p' are minimal Walrasian prices for valuations \widehat{v} .*

Proof. The new prices p' support μ , a feasible allocation. Consider any other feasible allocation μ' . We have changed each bidder's valuation by at most $\Delta/2k$, so the welfare of μ and μ' each change by at most $\Delta/2$ in the new market. Since the original difference in social welfare between μ and μ' is at least Δ , μ must remain the maximum matching in the perturbed market.

The only thing we need to check is that nodes with in-degree 0 (i.e., goods that are not over-demanded) have price 0. But this is also clear, from the construction of p' . □

Like before, we can bound the worst-case over-demand by the in-degree of any node in \widehat{G} .

Theorem D.3. *Suppose that all perturbations in \mathcal{P} are bounded by $\Delta/2k$. If $|\mathcal{P}| = \Omega(n^2k/\beta)$, then with probability at least $1 - \beta$, each node in \widehat{G} has in-degree at most 1.*

Proof. Note that setting prices for good j can only add edges to j , and leaves the rest of the swap graph \widehat{G} unchanged. So, consider a single good j , and consider any pair of bidders (r, s) . Fixing $\widehat{v_r(j)}$, there is at most one value of $\widehat{v_s(j)}$ that collides (possibly leading to two arcs into j), so the probability of any pair colliding is $1/|\mathcal{P}|$. Taking a union bound over the $O(n^2)$ pairs and k goods, the collision probability in the two graph \widehat{G} is at most $O(n^2k/|\mathcal{P}|) = O(\beta)$. □

E Proof of Lemma 4.12

Proof of Lemma 4.12. We first prove the first condition. We know by assumption that $|\{q : g \in B_q\}| \leq s_g + d$ (at most $s_g + d$ buyers have g in their bundle). Let q_g be the label of the $(s_g + 1)$ st (according to the ordering σ) buyer with g in their bundle, i.e. $g \in B_{q_g}$ and

$$\sum_{q': \sigma(q') < \sigma(q_g)} \mathbf{1}\{g \in B_{q'}\} = s_g.$$

Then, let $S_g = \{q : g \in B_q \text{ and } \sigma(q) \geq \sigma(q_g)\}$ be the set of buyers whose bundle contains g and who, according to the ordering σ , are after the s_g th buyer with $g \in B_q$. For each such buyer q let $\widehat{B}_q = \emptyset$. Otherwise, let $\widehat{B}_q = B_q$. Then,

$$\begin{aligned} \text{RWelfare}_N(B_1, \dots, B_n) &= \sum_q v_q(B_q) = \sum_{q \notin \cup_g S_g} v_q(B_q) + \sum_{q \in \cup_g S_g} v_q(B_q) \\ &= \sum_{q \notin \cup_g S_g} v_q(\widehat{B}_q) + \sum_{q \in \cup_g S_g} v_q(B_q) \leq \sum_{q \notin \cup_g S_g} v_q(\widehat{B}_q) + \sum_g \sum_{q \in S_g} v_q(B_q) \\ &\leq \sum_{q \notin \cup_g S_g} v_q(\widehat{B}_q) + \sum_g \sum_{q \in S_g} H \leq \sum_{q \notin \cup_g S_g} v_q(\widehat{B}_q) + d \cdot m \cdot H \\ &\leq \sum_q v_q(\widehat{B}_q) + d \cdot m \cdot H = \text{RWelfare}_N(\widehat{B}_1, \dots, \widehat{B}_n) + d \cdot m \cdot H \end{aligned}$$

where the third equality follows from the definition of \widehat{B}_q , the first inequality from the fact that $q \in \cup_g S_g$ implies $q \in S_g$ for some g , the second inequality from the bound on all the valuations, the third inequality from the assumption on $|S_g| \leq d$ and there being m goods, and the last inequality from the definition of \widehat{B}_q .

We now prove the second statement. Let N_g be the set of sampled bidders for that demand g , and consider a fixed buyer q . For each good $g \in B_q$, let $p_g = \mathbb{P}_{N_g} [q \in N_g] \geq \frac{s_g}{s_g + d}$ and $\underline{p} = \min_g p_g$. We will now prove, for an arbitrary partition (D_1, D_2) of B_q , that

$$\mathbb{E}_{N_1, \dots, N_m} \left[v_q(\widehat{B}_q \setminus D_2) \right] \geq \underline{p} \cdot v_q(B_q \setminus D_2).$$

Our claim follows when $D_2 = \emptyset$ after summing over all buyers q .

We proceed by induction on $|D_1|$. If $|D_1| = 0$, then $D_2 = B_q$ and $B_q \setminus D_2 = \widehat{B}_q \setminus D_2 = \emptyset$, so

$$\mathbb{E}_{N_1, \dots, N_m} \left[v_q(\widehat{B}_q \setminus D_2) \right] = 0 = \underline{p} \cdot 0 = \underline{p} \cdot v_q(\emptyset).$$

Now, assume for all partitions (D'_1, D'_2) of B_q with $|D'_1| < t$, we have

$$\mathbb{E}_{N_1, \dots, N_m} \left[v_q(\widehat{B}_q \setminus D'_2) \right] \geq \underline{p} \cdot v_q(B_q \setminus D'_2).$$

We wish to show that for a partition (D_1, D_2) of B_q with $|D_1| = t \geq 1$, we have

$$\mathbb{E}_{N_1, \dots, N_m} \left[v_q(\widehat{B}_q \setminus D_2) \right] \geq \underline{p} \cdot v_q(B_q \setminus D_2).$$

For notational cleanliness, let R be the random variable $\widehat{B}_q \setminus D_2$. For any good $g \in B_q \setminus D_2$ and bidder q , we have

$$\begin{aligned}
\mathbb{E}_{N_1, \dots, N_m} [v_q(R)] &= \mathbb{P}[q \in N_g] \cdot \mathbb{E}[v_q(R) \mid q \in N_g] + \mathbb{P}[q \notin N_g] \cdot \mathbb{E}[v_q(R) \mid q \notin N_g] \\
&= p_g \cdot \mathbb{E}[v_q(R) \mid g \in R] + (1 - p_g) \cdot \mathbb{E}[v_q(R) \mid g \notin R] \\
&= p_g \cdot \mathbb{E}[v_q(R) \mid g \in R] + (1 - p_g) \cdot \mathbb{E}[v_q(R \setminus \{g\})] \\
&= p_g \cdot \mathbb{E}[v_q(R \setminus \{g\} \cup \{g\})] + (1 - p_g) \cdot \mathbb{E}[v_q(R \setminus \{g\})] \\
&= p_g \cdot \mathbb{E}[v_q(R \setminus \{g\}) + (v_q(R \setminus \{g\} \cup \{g\}) - v_q(R \setminus \{g\}))] + (1 - p_g) \cdot \mathbb{E}[v_q(R \setminus \{g\})] \\
&= p_g \cdot \mathbb{E}[(v_q(R \setminus \{g\} \cup \{g\}) - v_q(R \setminus \{g\}))] + \mathbb{E}[v_q(R \setminus \{g\})] \\
&\geq p_g \cdot \mathbb{E}[(v_q(D_1 \setminus \{g\} \cup \{g\}) - v_q(D_1 \setminus \{g\}))] + \mathbb{E}[v_q(R \setminus \{g\})] \\
&= p_g \cdot \mathbb{E}[(v_q(D_1) - v_q(D_1 \setminus \{g\}))] + \mathbb{E}[v_q(R \setminus \{g\})] \\
&= p_g \cdot (v_q(D_1) - v_q(D_1 \setminus \{g\})) + \mathbb{E}[v_q(R \setminus \{g\})] \\
&= p_g \cdot (v_q(D_1) - v_q(D_1 \setminus \{g\})) + \mathbb{E}[v_q(\widehat{B}_q \setminus D_2 \setminus \{g\})] \\
&= p_g \cdot (v_q(D_1) - v_q(D_1 \setminus \{g\})) + \mathbb{E}[v_q(\widehat{B}_q \setminus (D_2 \cup \{g\}))] \\
&\geq p_g \cdot (v_q(D_1) - v_q(D_1 \setminus \{g\})) + \underline{p} \cdot v_q(B_q \setminus (D_2 \cup \{g\})) \\
&= p_g \cdot (v_q(D_1) - v_q(D_1 \setminus \{g\})) + \underline{p} \cdot v_q(D_1 \setminus \{g\}) \\
&\geq \underline{p} \cdot ((v_q(D_1) - v_q(D_1 \setminus \{g\})) + v_q(D_1 \setminus \{g\})) \\
&= \underline{p} \cdot v_q(D_1) = \underline{p} \cdot v_q(B_q \setminus D_2)
\end{aligned}$$

where the first inequality follows from the subadditivity of v_q and $R = \widehat{B}_q \setminus D_2 \subseteq B_q \setminus D_2 = D_1$; the following equality follows from $g \in D_1 = B_q \setminus D_2$; the second to last inequality follows from our inductive hypothesis (since $g \in B_1 \setminus D_2$, and $(D_1 \setminus \{g\}, D_2 \cup \{g\})$ form a partition satisfying the induction criterion); the final inequality follows from $\underline{p} \leq p_g$ for all g . \square

F Omitted Learning Proofs

F.1 Proof of Lemma 6.12

The high-level idea for the proof of Lemma 6.12 is as follows. In order to show \mathcal{H} is $m + 1$ -linearly separable, we must define two things. First, we need a mapping $\Psi : \mathcal{V} \times \mathcal{G} \rightarrow \mathbb{R}^{m+1}$, where $\Psi(v, B)_g$ will encode whether or not $g \in B$ for each good $g \in [m]$, and $\Psi(v, B)_{m+1} = v(B)$ encodes the agent's value for a bundle B , and (2), for each price vector p a weight vector $w^p \in \mathbb{R}^{m+1}$ where $w_g^p = -p_g$ encodes the *cost* of good $g \in [m]$ at these prices, while $w_{m+1}^p = 1$. These will have the property that their dot product encodes the utility of an agent v buying a bundle B at prices p : that is, $\Psi(v, B) \cdot w^p = v(B) - p(B)$. Thus, as desired, a utility-maximizing bundle B^* will have the property that $B^* \in \operatorname{argmax}_B \Psi(v, B) \cdot w^p$.

Unfortunately, two related problems remain with this formulation. First, the statement of Theorem 6.10 assumes this argmax is unique.¹³ Second, we want to assume that amongst the demanded bundles, agents with valuation v_q break ties according to the encodable tie-breaking rule

¹³If the argmax is not unique, prediction is not even well defined.

$e = (L_{v_q}, y)$. Fortunately, we can slightly perturb Ψ and the set of w^p s such that they encode the tie-breaking rule e , causing $e(\mathcal{D}_q(p))$ to be the unique bundle which maximizes $\Psi(v, B) \cdot w^p$ over B , solving both problems at once. It will be useful to have notation for the set of bundles which maximize the dot product between these quantities, so we define $M(v, p) = \operatorname{argmax}_B \widehat{\Psi}(v, B) \cdot \widehat{w}^p$.

We now define the perturbed versions of Ψ and w^p . Recall that $y \in \mathbb{R}^m$, so y_g will refer to the g th coordinate of this vector. Let

$$\alpha = \min \left\{ 1, \min_{t, B, B': v_t(B) - v_t(B') \neq 0} |v_t(B) - v_t(B')| \right\}$$

be the smallest gap between any valuation's distinct values for two bundles on the sample, or 1 if the gap is above 1.¹⁴ Then, let $y'_g = \frac{y_g \cdot \alpha}{4m \cdot \max_{g'} y_{g'}}$. Then, we define

$$\widehat{w}_g^p = \begin{cases} -p_g + y'_g & \text{if } g \in [1, m] \\ 1 & \text{if } g = m + 1 \end{cases}$$

and

$$\widehat{\Psi}(v_t, B_t)_g = \begin{cases} \mathbf{1}[[g \in B_t] & \text{if } g \in [1, m] \\ v_t(B_t) & \text{if } g = m + 1 \text{ and } B_t \notin \mathcal{L}_{v_t} \\ -H & \text{if } g = m + 1 \text{ and } B_t \in \mathcal{L}_{v_t} \end{cases}$$

These definitions allow us to prove the following lemmas about the bundles $B' \in M(v, p)$. The first shows that all bundles in $M(v, p)$ are utility-maximizing. The second implies that $|M| = 1$ and that $e(\mathcal{D}_q(p)) \in M(v, p)$.

Lemma F.1. *If $B' \in M(v, p)$, then B' is utility-maximal for v at prices p and $B' \notin \mathcal{L}_v$.*

Proof. We first show that $B \notin \mathcal{L}_v$ for every $B \in M(v, p)$. Suppose $B \in \mathcal{L}_v$. Then,

$$\begin{aligned} \widehat{\Psi}(v, B) \cdot \widehat{w}^p &= -H - \sum_{g \in B} p_g + \sum_{g \in B} y'_g \\ &\leq -H + \sum_{g=1}^m y'_g \\ &= -H + \sum_{g=1}^m \frac{\alpha \cdot y_g}{4m \cdot \max_{g'} y_{g'}} \\ &\leq -H + \sum_{g=1}^m \frac{\alpha}{4m} \\ &\leq 0 \end{aligned}$$

where the final inequality follows from $H \geq 1$ and $\alpha \leq 1$. Thus, any $B' \in \mathcal{L}_v$ will have negative dot-product. Thus, if there is any bundle with positive dot product, B' will not be chosen. If some utility-maximizing bundle $B \notin \mathcal{L}_v$, it is not hard to show that $\widehat{\Psi}(v, B) \cdot \widehat{w}^p \geq 0$. Since e is encodable, we know that some utility-maximizing bundle is not in \mathcal{L}_v . Thus, no losing set will maximize the dot product.

¹⁴The gap might be above one if, for example, the agent's valuation for each bundle is a distinct even number.

We will now show that $B' \in \operatorname{argmax}_B v(B) - p(B)$. Since $v(B) - p(B) \leq \widehat{\Psi}(v, B) \cdot \widehat{w}^p \leq v(B) - p(B) + \frac{\alpha}{4}$ for every $B \notin \mathcal{L}_v$, any B which is utility-maximizing and not in \mathcal{L}_v will have larger dot product than any bundle which is not utility-maximizing, by the definition of α . \square

Lemma F.2. *For any v, p , $|M(v, p)| = 1$ and $\{e(\mathcal{D}_q(p))\} = M(v, p)$.*

Proof. Lemma F.1 implies that any $B' \in M(v, p)$ is not in \mathcal{L}_v and is utility-maximizing, so we know for all $B' \in M(v, p)$ that

$$v(B') - p(B') \leq \widehat{\Psi}(v, B') \cdot \widehat{w}^p(B') = v(B') - p(B') + \sum_{g \in B'} \frac{\alpha \cdot y_g}{4m \cdot \max_{g'} y_{g'}}.$$

For all such B' , the first term is equal (they are all utility-maximizing). Then, let us consider B_1, B_2 both of which are utility-maximizing and not in \mathcal{L}_v . It must be the case that

$$\left(\frac{4m \cdot \max_{g'} y_{g'}}{\alpha} \right) \left(\widehat{\Psi}(v, B_1) \cdot \widehat{w}^p(B_1) - \widehat{\Psi}(v, B_2) \right) \cdot \widehat{w}^p(B_2) = \sum_{g \in B_1} y_g - \sum_{g \in B_2} y_g.$$

Then, since e is encodable, we know there exists a unique utility-maximizing bundle B_1 for which this quantity is strictly positive with respect to each other B_2 , and that $B_1 = e(\mathcal{D}_q(p))$. Thus, it must also be that $|M(v, p)| = 1$. \square

We now prove Lemma 6.12.

Proof of Lemma 6.12. For any sample S of size n , $\widehat{\Psi}$, \widehat{w}^p are defined as above, and for any v_t , M is defined as above. Lemma F.1 implies that any $B \in M(v, p)$ is utility-maximizing, and Lemma F.2 implies that $M(v, p)$ is a singleton and agrees with e . Thus, $\widehat{\Psi}$ and the set of \widehat{w}^p are a linear prediction rule which correctly predicts buyers will buy the utility-maximizing bundle according to tie-breaking rule e . \square

F.2 The existence of an encodable tie-breaking rule for GMBV

Lemma F.3. *There exists an encodable tie-breaking rule $e = (\mathcal{L}_v, y)$ which*

- *Always selects type-minimal bundles,*
- *and always breaks ties in favor of larger type-minimal bundles.*

Proof. We need to define \mathcal{L}_v and y and prove they indeed always select a unique type-minimal bundle of largest size, and also that \mathcal{L}_v never contains all utility-maximizing bundles. Let $y_g = 1 + \frac{4^g}{4m \cdot 4^m}$,¹⁵ and \mathcal{L}_v contain the set of non-type-minimal bundles for buyers with valuation v (that is, if $v(B) = v(B \cup \{g\})$ for some $g \notin B$, then $B \cup \{g\} \in \mathcal{L}_v$). By definition of \mathcal{L}_v , e will never select a bundle which is not type-minimal.

We now show that \mathcal{L}_v never contains all utility-maximizing bundles. If there is some bundle $B \in \mathcal{L}_v$, this implies there is some $B' \subsetneq B$ for which $v(B') = v(B)$, and such that B' is type-minimal (thus, $B' \notin \mathcal{L}_v$). Since $p(B') \leq p(B)$, B' is also a utility-maximizing; thus, \mathcal{L}_v does not contain all utility-maximizing bundles for any pricing p .

¹⁵This corresponds to breaking ties lexicographically over bundles.

Finally, we show that, for any \mathcal{X} , there is always a unique $B \in \mathcal{X}$ which maximizes $Y = \sum_{g \in B} y_g$, and that this unique maximizer is amongst the bundles of the largest size in $\mathcal{X} \setminus \mathcal{L}$. Notice that, for any B , $\sum_{g \in B} y_g = |B| + \sum_{g \in B} \frac{4^g}{4m \cdot 4^m}$; this quantity is strictly above $|B| - 1$ and strictly below $|B| + 1$, so the candidate maximizers of Y are all of maximum size. Then, each of those has a unique value of $\sum_{g \in B} \frac{4^g}{4m \cdot 4^m}$, so there is a unique maximizer of Y . \square

F.3 Shattering arguments

Proof of Theorem 6.14. Consider some sample $S = (v_1, \dots, v_n)$ of arbitrary valuations which can be shattered with targets (r_1, \dots, r_n) . By Lemma 6.15, \mathcal{H} induces at most $\binom{n}{m+1} \cdot 2^{m(m+1)}$ bundle labelings on a sample of size n . Fixing the bundle label of a given v_q fixes the welfare for v_q (if \widehat{B} is purchased by v_q , her welfare is $v_q(\widehat{B})$). So, fixing the bundle labeling of all of S fixes the welfare for all of S . Thus, there are at most $\binom{n}{m+1} \cdot 2^{m(m+1)}$ distinct welfare labelings of all of S ; thus, there must be at most that many binary labelings of S according to the welfare targets (r_1, \dots, r_n) . Since S is shatterable,

$$2^n \leq \binom{n}{m+1} \cdot 2^{m(m+1)}$$

so $n = O(m^2 \log m)$ as claimed.

The proof is identical for unit demand valuations, replacing the 2^m possible labels per valuation with the upper bound of $m+1$ (since, when buyers are unit demand, they will buy one item or no items). \square

F.4 Omitted concentration proofs

For this section, we will be referring to loss functions on samples, $\ell(f, x)$, as well as loss of those functions on entire samples $\ell_N(f)$ and distributions $\ell_\Pi(f)$. In the latter case, let $\ell_N(f) = \frac{1}{|N|} \sum_{x \in N} \ell(f, x)$ denote the *empirical* loss of the function f on a sample N ; in the latter case, let $\ell_\Pi(f) = \mathbb{E}[x \sim \Pi] \ell(f, x)$ denote the *true expected* loss of f with respect to Π .

Theorem F.4 (McDiarmid's Inequality). *Suppose X^1, \dots, X^k are independent, and assume f has the property that, for any $i, x^1, \dots, x^k, \widehat{x}^i$,*

$$|f(x^1, \dots, x^k) - f(x^1, \dots, \widehat{x}^i, \dots, x^k)| \leq c_i.$$

Then,

$$\Pr[|\mathbb{E}[f(X^1, \dots, X^k)] - f(X^1, \dots, X^k)| \geq \epsilon] \leq 2e^{-\frac{2\epsilon^2}{\sum_{i=1}^k c_i^2}}.$$

Proof of Lemma 6.16. This follows from a simple application of McDiarmid's inequality, noting that welfare is a smooth function in each argument (i.e., the welfare cannot change by more than H by changing any one of the sampled valuations). \square

Theorem F.5. *Fix a hypothesis $f \in \mathcal{F}$. Let ℓ be some function defined on samples such that $\ell(f, x) \in \{0, 1\}$, $\ell_N(f) = \frac{1}{|N|} \sum_{q \in N} \ell(f, q)$, and $\ell_\Pi(f) = \mathbb{E}[q \sim \Pi] \ell(f, q)$. Then, with probability $1 - \delta$, if $|N| = n$, for a sample $N \sim \Pi$,*

$$\ell_\Pi(f) - \ell_N(f) \leq \sqrt{\frac{2\ell_N(f) \log \frac{1}{\delta}}{n}} + \frac{4 \log \frac{1}{\delta}}{n}.$$

Proof of Theorem F.5. This follows from Lemma B.10 (pg. 427) by [Shalev-Shwartz and Ben-David \[2014\]](#). \square

We now prove Theorem 6.13, which is a direct corollary of Theorem F.5.

Proof of Theorem 6.13. Applying Theorem F.5 and taking a union bound over all possible loss values on the sample (that is, using Sauer's Lemma in conjunction with Theorem F.5 rather than the standard Hoeffding bound). \square

Proof of Theorem 6.2. This follows from Theorem 6.13 along with the bound of $s_g + 1$ on $\ell_S(h_{g,p})$ and the bound on $\text{VC}_{\mathcal{H}_g}$ given in Theorem 6.11. Thus, we have proved the first statement of the theorem. The second follows from the first, with basic algebra. \square

Theorem F.6. Fix a hypothesis $f \in \mathcal{F}$. Let ℓ be some loss function defined on samples N drawn from Π , such that $\ell(f, x) \in [0, H]$. Then, with probability $1 - \delta$, if $|N| = n$,

$$\ell_N(f) - \ell_\Pi(f) \leq \left(\frac{H^{3/2} \log \frac{1}{\delta}}{n} \right) \left(\frac{1}{3} + \sqrt{19n\ell_N(f)} \right)$$

Proof. Let $\alpha_i = \frac{\ell_\Pi(f) - \ell(f, x_i)}{\sqrt{H}}$ be the (normalized) difference between the true error of f and the error of f on sample x_i . Then, applying Bernstein's inequality to the negative of the α_i s, we have that for any $t < 0$,

$$\begin{aligned} \mathbb{P} \left[\sum_{i=1}^m \alpha_i < t \right] &\leq \exp \left(- \frac{t^2/2}{\sum_{g=1}^n \mathbb{E} [\alpha_g^2] + H \cdot t/3} \right) \\ &\leq \exp \left(- \frac{t^2/2}{\frac{n}{\sqrt{H}} \cdot \ell_N(f) + H \cdot t/3} \right) \\ &= \exp \left(- \frac{\sqrt{H}t^2/2}{n \cdot \ell_N(f) + H^{3/2} \cdot t/3} \right). \end{aligned}$$

Setting the right-hand side to equal δ , and solving for t , we have

$$\begin{aligned} \log \frac{1}{\delta} &= \frac{\sqrt{H}t^2/2}{n\ell_N(f) + \frac{H^{3/2} \cdot t}{3}} \\ \Rightarrow t &\leq \frac{H}{3} \log \frac{1}{\delta} + \sqrt{H^2 \log^2 \frac{1}{\delta} + \frac{18\ell_N(f)n \log \frac{1}{\delta}}{\sqrt{H}}} \\ &\leq \frac{H}{3} \log \frac{1}{\delta} + \sqrt{19 \cdot H^2 \log^2 \frac{1}{\delta} \ell_N(f)n} \\ &\leq \frac{H}{3} \log \frac{1}{\delta} + H \log \frac{1}{\delta} \sqrt{19\ell_N(f)n} \\ &\leq \frac{H}{3} \log \frac{1}{\delta} + H \log \frac{1}{\delta} \sqrt{19\ell_N(f)n}. \end{aligned}$$

As $\frac{\sqrt{H}}{n} \sum_{i=1}^n \alpha_i = \ell_{\Pi}(f) - \ell_N(f)$, this implies that with probability $1 - \delta$,

$$\ell_N(f) - \ell_{\Pi}(f) \leq \left(\frac{H^{3/2} \log \frac{1}{\delta}}{n} \right) \left(\frac{1}{3} + \sqrt{19\ell_N(f)n} \right)$$

which proves the claim. \square

Theorem F.7. Consider a hypothesis class \mathcal{F} , and $\ell(f, x) \in [0, H]$. Then, with probability $1 - \delta$, if $|N| = n$, for all $f \in \mathcal{F}$,

$$\ell_N(f) - \ell_{\Pi}(f) \leq \left(\frac{H^{3/2} \text{PD}_{\mathcal{F}} \log \frac{1}{\delta}}{n} \right) \left(\frac{1}{3} + \sqrt{19\ell_N(f)n} \right).$$

Proof of Theorem F.7. Applying Theorem F.6 and taking a union bound (that is, using Sauer's Lemma in conjunction with Theorem F.6 rather than the standard Hoeffding bound). \square