

# Posted Pricing sans Discrimination

Shreyas Sekar

Rensselaer Polytechnic Institute, Troy, NY-12180

shreyas.sekar@gmail.com

## Abstract

In the quest for market mechanisms that are easy to implement, yet close to optimal, few seem as viable as posted pricing. Despite the growing body of impressive results, the performance of most posted price mechanisms however, rely crucially on “price discrimination” in the presence of large supply or non-linear production costs. With this in mind, we study the problem of *designing non-discriminatory pricing mechanisms for social welfare maximization* in a Bayesian setting where the seller only has stochastic information regarding buyer valuations. Our main contribution is a framework for static item pricing in the face of production costs, i.e., the seller posts one price per good and buyers arrive sequentially and purchase utility-maximizing bundles. The framework yields constant factor approximations to the optimum welfare when buyer valuations are fractionally subadditive, and extends to settings where the seller is completely oblivious to buyer valuations. Our results indicate that even in markets with complex buyer valuations and non-linear costs, it is possible to obtain good guarantees without price discrimination, i.e., charging buyers differently for the same good.

## 1 Introduction

In the quest for market mechanisms that are *simple*, yet *approximately optimal* [Hartline and Roughgarden, 2009], few candidates seem as appealing as *posted price mechanisms*, where the seller simply posts prices on the items and buyers consume utility-maximizing sets of goods. This optimism towards posted prices is not without cause: a long line of research has established that in addition to a slew of favorable properties, mechanisms based on posted pricing provide near-optimal performance guarantees in a number of diverse applications [Blumrosen and Holenstein, 2008; Feldman *et al.*, 2015; Maehara *et al.*, ]. More recently, some of these results have been extended to settings with general buyer valuations, where sellers only possess distributional information regarding the same [Feldman *et al.*, 2015; Hsu *et al.*, 2016]; such results are a crucial first step towards

the larger goal of understanding the performance of posted prices in more realistic, and general markets.

Real markets, however, often feature sellers who possess multiple copies of goods or more generally, costs associated with producing or transporting each copy of good. Here, our picture of posted pricing mechanisms is far from complete. Papers dealing with multi-unit supply in combinatorial auctions have often resorted to *price discrimination*, i.e., charging buyers differing prices for the same good, in order to extract good performance bounds [Blum *et al.*, 2011; Chakraborty *et al.*, 2010; Cohen-Addad *et al.*, 2016]. In the literature, one also encounters discriminative pricing policies under more benign labels such as *dynamic pricing* [Chakraborty *et al.*, 2009], *sequential posted pricing* [Chawla *et al.*, 2010] or *non-anonymous reserve prices* [Daskalakis and Pierrakos, 2011]. Looking beyond multi-unit supply to the more general case of production costs, little is known about the performance of posted price mechanisms when the underlying costs are non-linear, even less so in Bayesian settings.

*Is discriminatory pricing essential for good performance in the presence of multi-unit supply?  
Do posted prices give desirable guarantees in a Bayesian setting involving production costs?*

These are the questions that we seek to answer in this work. Our work is motivated by both theory and practice; price discrimination, while accepted practice in some industries (airline, taxi), is inherently unfair to the buyers and can lead to adverse effects in the long-run [Anderson *et al.*, 2010]. On the other hand, extending the state-of-the-art in *Bayesian Mechanism Design* to markets with general buyer valuation functions and production costs often necessitates blurring the line between good and copy, resulting in mechanisms that price each copy of a single good as if it were a distinct good.

Our central result answers both the questions posed above for Bayesian settings with fractionally subadditive (XoS) buyer valuations and convex production costs. Specifically, we design an incentive compatible, non-discriminatory mechanism based on static posted prices that extracts a constant fraction of the optimum social welfare. At a high level, our work presents a strong case for treating buyers fairly: by bounding the gains obtained via discriminatory pricing, we argue that the incentives for discrimination may not offset the

negative impact of treating buyers unfairly.

**Market Model** We consider a market with a set  $\mathcal{I}$  of  $m$  goods and  $N$  buyers having combinatorial valuation functions:  $v : 2^{\mathcal{I}} \rightarrow \mathbb{R}^+ \cup \{0\}$ , i.e.,  $v_j(S)$  denotes buyer  $j$ 's value for a set  $S \subseteq \mathcal{I}$  of goods. The seller faces a convex production cost function with non-decreasing marginal costs for each good, i.e., for every  $i \in \mathcal{I}$ , the seller incurs a cost of  $C_i(\ell)$  for producing  $\ell$  units of this good. Such convex costs strictly generalize settings with limited supply. For the majority of this work, we assume a Bayesian setting with prior  $\mathcal{F} = F_1 \times F_2 \dots \times F_N$ : buyer  $i$ 's private valuation is drawn independently from distribution  $F_i$ , which is known to the seller. Our objective is to maximize the social welfare, which is defined as the total value derived by the buyers minus the production cost incurred by the seller. Towards this goal, we propose two types of posted pricing mechanisms depending on whether or not the seller has to commit to producing an exact amount of each good in advance. Both these mechanisms follow a simple template: (i) the seller posts a single price per good applicable to all buyers, and (ii) buyers arrive sequentially in some arbitrary, possibly *adversarial* order and purchase utility maximizing subsets of the available goods.

1. **On the Fly Mechanisms** (OTF): The seller fixes an upper bound on the number of copies of each good that he is willing to produce and incurs a production cost only for the items that are actually sold. Once the number of copies of a good sold reaches the upper bound, the item becomes unavailable for future buyers.
2. **Commitment Mechanisms**: The seller commits to producing a certain quantity of each good and incurs production costs whether or not those items are sold.

When the seller has limited supply, the two classes of mechanisms are equivalent. However, for strictly non-linear production costs, on the fly mechanisms are a relaxation of commitment mechanisms. Finally, all of the mechanisms in this work are trivially *dominant strategy incentive compatible*.

### Comparisons to Existing Work: Dynamic Pricing and Unit Supply

To the best of our knowledge, there are no previous results in Bayesian mechanism design dealing with convex production costs (specifically with a view towards non-discriminatory pricing). This is particularly surprising given that increasing marginal costs are a natural consideration in many markets since the ‘‘limited supply assumption is too pessimistic and additional sources can often be found at higher cost’’ [Blum *et al.*, 2011].

At the same time, the current work acts as a bridge between the closely related streams of literature pertaining to *multi-unit combinatorial auctions* [Bartal *et al.*, 2003; Blum *et al.*, 2011; Huang and Kim, 2015] and *posted pricing with unit-supply* [Feldman *et al.*, 2015]. The literature on multi-unit combinatorial auctions has yielded constant-factor approximations for maximizing social welfare in the presence of limited supply and production costs, albeit by offering different posted prices to different buyers. In this work, we study the same setting and approximate social welfare via the more

demanding class of static, non-discriminatory pricing mechanisms instead of dynamic pricing as in previous work<sup>1</sup>. On the other hand, in [Feldman *et al.*, 2015], the authors use non-discriminatory posted prices to maximize welfare when the seller has exactly one unit supply of each good. We extend their result to settings with non-linear production costs on the goods.

## 1.1 Summary of Contributions

### Central Result: 4-approximate On the Fly Mechanism

We begin with the class of OTF mechanisms for buyers having XoS valuations, which generalize submodular functions. We present a black-box reduction that transforms any approximation algorithm  $Alg$  for the *allocation problem*: ‘the problem of allocating goods to buyers to maximize social welfare’, into a non-discriminatory posted price mechanism that ensures a social welfare of  $\frac{1}{2} E_{\vec{v} \sim \mathcal{F}}[SW(Alg(\vec{v}))]$ , where  $SW(Alg(\vec{v}))$  is the welfare guaranteed by  $Alg$  on input valuations  $\vec{v}$ . The black-box result makes use of both a demand oracle and an XoS oracle for the buyer valuations. To supplement the reduction, we also present a poly-time 2-approximation algorithm ( $Alg$ ) for the allocation problem with XoS buyers and convex costs. Together, they yield our main computational result: *a mechanism that posts a single price per good and extracts one-fourth of the optimum welfare for arbitrary buyer arrival orders and convex costs.*

**(Significance of the Result)** Previously, posted pricing mechanisms with good welfare approximations were only known for restrictive settings with one unit supply of each good and no production cost [Feldman *et al.*, 2015]. Although their result could arguably be extended to markets with increased supply, this would require posting different prices for different copies of the same item, i.e., by treating identical copies of a good as distinct goods, any problem with multi-unit supply can be reduced to one with unit supply. On the contrary, our mechanism offers the same price on all copies of a good and makes no assumption on the cost function other than convexity, which is standard in market design [Blum *et al.*, 2011].

**Commitment Mechanism and XoS valuations:** For the challenging class of commitment mechanisms, our framework yields welfare guarantees that are sensitive to the *structure of the allocation* returned by the algorithm. Formally, we define an instance to be  $\alpha$ -bounded (as in [Feige *et al.*, 2013]) with respect to an allocation algorithm  $Alg$ , if in the solution returned by the algorithm, the expected value derived by the buyers is at least a factor  $\alpha$  times the expected production cost incurred by the seller. Then, our mechanism obtains ‘desirable welfare guarantees’ when  $\alpha$  is large, i.e., for instances where the algorithm returns a solution whose expected social welfare is large compared to the cost incurred.

More specifically, *we present a non-discriminatory commitment mechanism whose approximation factor to the optimum welfare is  $4 \frac{\alpha-1}{\alpha-2}$  for XoS buyers ( $\alpha > 2$ ).* If the cost functions are ‘sufficiently convex’, then  $\alpha$  is large irrespective

<sup>1</sup>It is noteworthy that the results in [Blum *et al.*, 2011] were obtained for a prior-free setting; we obtain stronger approximation guarantees under a distributional assumption.

of the buyer valuations, and our mechanism provides good bounds.

**Sampling Buyer Distributions:** Our mechanisms for the Bayesian setting exhibit a runtime polynomial in the support size of the product distribution  $\mathcal{F}$ . However, one can convert these mechanisms into computationally efficient procedures in a straightforward manner by repeatedly sampling valuation profiles and computing allocations for the sampled valuations. Specifically, applying the sampling techniques in [Feldman *et al.*, 2015], we obtain mechanisms that run in time polynomial in  $N, m$ , and  $\frac{1}{\epsilon}$  with an additive welfare loss of  $\epsilon$  from the stated bound. In order to clearly illustrate our technical ideas and avoid messy notation, we suppress the sampling aspect of our mechanisms for the rest of this work.

## 1.2 Other Related Work

This paper also bridges the gap between two closely related fields: auction theory, and envy-free item pricing. Although the mechanisms in this work technically fall under the purview of Bayesian mechanism design [Chawla *et al.*, 2010], they are not auctions in the traditional sense as they do not elicit any input from the buyers; hence, they are trivially dominant strategy incentive compatible (DSIC). Our work is closer in spirit to envy-free pricing algorithms for computing Walrasian equilibrium [Roughgarden and Talgam-Cohen, 2015]. Since Walrasian equilibria do not exist [Gul and Stacchetti, 1999] for complex valuations such as XoS, one could view the pricing mechanisms in this work as relaxations that sacrifice envy-freeness but guarantee existence and welfare.

To the best of our knowledge, there are no mechanisms in the literature that meet all of the desiderata fulfilled by our framework: (i) general buyer valuations (XoS), (ii) (convex) production cost functions, (iii) non-discriminatory pricing, (iv) non-identically distributed buyers, and (v) dominant strategy incentive compatibility. However, if we relax some of these constraints, we stumble upon a number of closely related works (see for example [Bei and Huang, 2011; Brânzei *et al.*, 2014]).

Owing to their popularity in real-world applications, there has been a surge in the theoretical investigation of *dynamic pricing mechanisms*. Specifically, [Blum *et al.*, 2011; Huang and Kim, 2015] look at an OTF-like mechanism where the seller offers different prices in each round and prove welfare bounds that depend on the nature of the production cost function. On the other hand, our mechanism guarantees fairness as the prices remain static throughout. Moreover, our OTF mechanism obtains a 4-approximation to the optimum welfare for arbitrary convex production costs.

## 2 Model and Preliminaries

As defined previously, we consider a set  $\mathcal{I}$  of  $m$  goods: each  $i \in \mathcal{I}$  has a convex production cost, i.e., for any  $\ell$ , the seller incurs a (marginal) cost of  $c_i(\ell)$  for producing the  $\ell^{\text{th}}$  copy of this good, and  $c_i(\ell)$  is non-decreasing as  $\ell$  increases. The seller can choose to produce any number of goods for which he faces an aggregate production cost of  $C_i(\ell) = \sum_{r=1}^{\ell} c_i(r)$ . Such convex costs strictly generalize limited supply. For instance,  $c_i(\ell) = 0$  for  $\ell \leq k$  and

$c_i(\ell) = \infty$  otherwise, indicates a fixed supply of  $k$  units.

Given an allocation  $\vec{S} = (S_1, S_2, \dots, S_N)$  where  $S_j$  is the set of goods allocated to buyer  $j$ , suppose that  $\forall i \in \mathcal{I}$ ,  $k_i$  represents the number of units of this good allocated to the buyers. Then, the *social welfare* of this allocation is  $SW(\vec{S}) = \sum_{j=1}^N v_j(S_j) - \sum_{i \in \mathcal{I}} C_i(k_i)$ . For the rest of this work, we will assume that all buyers have monotone non-decreasing XoS valuations, which strictly generalize many well-studied functions such as unit-demand, budget additive, and submodular valuations [Dobzinski *et al.*, 2010].

**Fractionally Subadditive (or ‘XoS’) Valuations**  $\exists$  a set of additive functions  $(a_1, \dots, a_r)$  such that for any  $T \subseteq \mathcal{I}$ ,  $v(T) = \max_{j=1}^r a_j(T)$ . An additive function (also known as a *clause*)  $a_j$  has a single value  $a_j(i)$  for each  $i \in \mathcal{I}$  so that for a set  $T$  of goods,  $a_j(T) = \sum_{i \in T} a_j(i)$ .

**Oracle Access** The standard approach in the literature while dealing with set functions (where the input representation is exponential in size) is to assume the presence of an oracle that allows indirect query access to the valuation. All of our results in Sections 3 and 4 assume black-box access to two types of oracles: (i) a *demand oracle* that when queried with a vector of prices  $\vec{p}$  returns a set  $S$  that maximizes the quantity  $v(S) - \sum_{i \in S} p_i$ , and (ii) an *XoS oracle* that for an XoS function  $v$  and a set  $T \subseteq \mathcal{I}$  returns the additive clause  $a_i$  that maximizes  $a_i(T)$ . We refer the reader to [Dobzinski *et al.*, 2010] for a thorough comparison of these two oracles.

### Seller’s Mechanism: Static Posted Prices

A posted pricing mechanism with item prices is said to be *static* or *non-discriminatory* if no two buyers are offered the same good at different prices. We now formally define the two types of mechanisms considered in this work. Note that both our mechanisms do not elicit any input from the buyers.

**On the fly mechanism** In an OTF mechanism, a seller incurs production costs only for the goods that are actually sold to the buyers.

1. the seller posts a price vector  $\vec{p}$  where  $p_i$  is the price of good  $i \in \mathcal{I}$ , and fixes an upper bound  $k_i$  on the number of copies of each good  $i$  that he is willing to produce.
2. buyers arrive in some arbitrary order; each buyer looks at the set of available goods and purchases her utility maximizing bundle.
3. an item becomes unavailable once the number of copies purchased equals the upper bound.

The upper bound can be thought of as the seller preparing his machinery for producing a certain number of goods, but stopping short of actually producing them.

**Commitment Mechanism** In a commitment mechanism, the seller produces the goods before the buyers arrive and thus incurs an additional cost for any goods that are not sold. From an algorithmic perspective, these mechanisms are much harder than on the fly mechanisms.

1. the seller posts a single price per good and produces a fixed quantity of each good in advance.
2. buyers arrive in some arbitrary order, and each buyer purchases her utility maximizing bundle from the remaining items.

### 3 On the fly mechanism for XoS buyers

In this section, we present our main result: an OTF mechanism that obtains a 4-approximation to the optimum welfare. The mechanism comprises of two independent components: (i) a black-box reduction for transforming any algorithm for the allocation problem into a posted price mechanism that guarantees the half of the social welfare of the algorithm (ii) a 2-approximation algorithm for the allocation problem. We first state the black-box result formally.

**Claim 3.1.** *Given a product distribution  $\mathcal{F}$ , an algorithm  $Alg$ , demand and XoS oracles, we can compute in poly-time a price vector  $\vec{p}^*$  and supply upper bound  $\vec{k}^*$  such that the OTF mechanism with these parameters provides a welfare guarantee (in expectation) of  $\frac{1}{2}E_{\vec{v} \sim \mathcal{F}}[SW(Alg(\vec{v}))]$ .*

Such black-box reductions (from pricing to the allocation problem) have recently gained traction in mechanism design [Bei and Huang, 2011; Dughmi *et al.*, 2017; Feldman *et al.*, 2015] as they allow us to reuse the extensive techniques developed for computing welfare maximizing allocations in order to derive prices that indirectly result in desirable allocations. However, the black-box techniques developed in [Bei and Huang, 2011; Dughmi *et al.*, 2017] only yield Bayesian incentive compatible (BIC) mechanisms, whereas the above claim trivially indicates dominant strategy incentive compatibility.

Claim 3.1 strictly generalizes the black-box result in [Feldman *et al.*, 2015] for deriving welfare-maximizing prices for settings with unit supply and our proof is inspired by the approach developed in that paper. However, owing to the non-linear nature of the production costs, the pricing scheme that we propose is more involved than the one used by [Feldman *et al.*, 2015]. More specifically, when dealing with unit supply, one can achieve good social welfare by simply pricing each item at a fraction of its (expected) contribution towards the social welfare of some allocation. However, this approach does not yield desirable guarantees with production costs. Instead, we generalize the single-item pricing scheme from [Feldman *et al.*, 2015] in a non-trivial direction by pricing each item at the average of its contribution towards the buyer welfare and its contribution towards the production cost, i.e., for every copy of the item produced, we create a dummy buyer whose valuation equals the marginal production cost and then price the item at its average contribution towards social welfare.

We reiterate the important point that when the representation of  $\mathcal{F}$  is not polynomial, using standard sampling techniques [Feldman *et al.*, 2015], we can get a welfare guarantee of  $\frac{1}{2}E_{\vec{v} \sim \mathcal{F}}[SW(Alg(\vec{v}))] - \epsilon$  in poly-time. Before proving the theorem, we state the concomitant ‘allocation algorithm’.

#### Algorithm for Allocating Goods to XoS buyers

The secondary algorithmic contribution of this section is a 2-approximation algorithm for the following allocation problem: *given a set  $\mathcal{I}$  of goods with convex production costs and deterministic XoS valuations, compute an allocation  $\vec{S} = (S_1, \dots, S_N)$  that maximizes social welfare.* Note that the allocation algorithm is only applicable for a single profile of buyer valuations, i.e., the deterministic setting. Our pricing

mechanism uses this algorithm as a black-box for different valuations drawn from  $\mathcal{F}$ .

**Claim 3.2.** *(Allocation Algorithm) For any given input with XoS buyer valuations, Algorithm 1 returns an allocation of goods to buyers such that the resulting social welfare is at least one-half of that of the welfare maximizing allocation.*

---

#### Algorithm 1 Allocation Algorithm for XoS buyers

---

**Input:** Access to buyer valuations in the form of Demand, XoS oracles

**Initialization Phase**

- 1:  $S_1, S_2, \dots, S_N \leftarrow \emptyset$  // Initial Allocation
- 2: Define  $\vec{p}$  s.t.  $p_i = c_i(1) \forall i \in \mathcal{I}$  // Internal Prices

**Allocation Phase:** For each buyer  $j \in 1, 2, \dots, N$

- 3: Let  $S'_j$  be the utility-maximizing bundle demanded by buyer  $j$  at prices  $\vec{p}$ ; set  $S_j = S'_j$
  - 4: Let  $(a_r^j)$  denote the maximizing XoS clause for buyer  $j$  and set  $S_j$ . Set  $q_j(i) = a_r^j(i)$  for all  $i \in S_j$ .
  - 5: **Update:** For each item  $i \in S_j$ , let  $X_i$  be the set of buyers who have received a copy of  $i$  so far.
  - 6: If  $p_i$  corresponds to  $q_j(i)$  for some  $j \in X_i$ , remove  $j$  from  $X_i$  and remove item  $i$  from  $j$ 's allocation.
  - 7: Update  $p_i = \min\{c_i(|X_i| + 1), \min_{j \in X_i} q_j(i)\}$ .
- 

The functioning of Algorithm 1 can be interpreted as follows: the market consists of  $N + 1$  buyers (buyer 0 is the seller) who arrive sequentially in some arbitrary order. Each new buyer can acquire their maximum utility bundle by purchasing each item from either the seller (at the current marginal production) or from previous buyers (at a price derived from the XoS clause corresponding to that buyer). Note that while Algorithm 1 uses prices as an internal device, its final output is simply an allocation and does not involve any pricing. Our algorithm generalizes the 2-approximation algorithm for XoS buyers and unit-supply first studied in [Dobzinski *et al.*, 2010] to settings with production costs.

The main implication of Claims 3.1 and 3.2 is a computationally efficient posted pricing mechanism whose welfare guarantee is one-fourth that of the optimum allocation.

**Theorem 3.3.** *We can compute in poly-time a price vector  $\vec{p}^*$  and supply vector  $\vec{k}^*$  such that for any arrival order, the expected welfare of the posted price mechanism with these prices is at least one-fourth of the optimum social welfare.*

All of the heavy-lifting for the above theorem is done by Claim 3.1, and it suffices to prove the claim.

**Proof of Claim 3.1** We now present a general framework for non-discriminatory pricing in the presence of convex production costs. The framework consists of two ingredients: first, we provide a technique to derive prices for each good separately. Second, we provide some sufficient conditions for obtaining pricing mechanisms with good welfare for arbitrary valuations.

#### Framework Part I: Profit-Surplus Equivalence

We now restrict our attention to markets with a single good  $i$  and show how to derive prices that obtain good welfare.

Suppose we allocate  $k$  units of this good to buyers whose valuations for the good are  $v_1 \geq v_2 \geq \dots \geq v_k$ . The total welfare due to this particular allocation is  $\sum_{\ell=1}^k [v_\ell - c_i(\ell)]$ . We are not interested in negative welfare, so we assume that  $\sum_{\ell=1}^k v_\ell \geq \sum_{\ell=1}^k c_i(\ell)$ . Define the following functions: (i)

**Buyer Surplus:**  $f_i(p) = \sum_{\ell=1}^k [v_\ell - p]$  (this denotes the total utility derived by the buyers at price  $p$ ) and (ii) **Seller Profit:**  $\pi_i(p) = pk - \sum_{\ell=1}^k c_i(\ell)$ .

The purpose of this section is to identify a convenient price where profit and surplus coincide (see Lemma 3.4); such a price is desirable for both the buyers and the seller. The choice of this price is justified in Lemma 3.5. Observe that both  $\pi_i(p)$  and  $f_i(p)$  are continuous and monotone in  $p$  with  $f_i(0) \geq \pi_i(0)$  whereas  $f_i(v_1) \leq \pi_i(v_1)$ . So, there must exist some price  $p_i^* \in [0, v_1]$  satisfying the following condition.

$$\text{(Profit-Surplus Equivalence)} f_i(p_i^*) = \pi_i(p_i^*). \quad (1)$$

**Lemma 3.4.** *The total profit if all  $k$  items are sold at price  $p_i^*$  is exactly half the social welfare due to  $i$ , i.e.,  $\pi_i(p_i^*) = \frac{1}{2} \left( \sum_{\ell=1}^k [v_\ell - c_i(\ell)] \right)$ .*

The lemma follows from Equation 1 using simple algebra.

### Framework Part 2: Sufficient Conditions

We now identify sufficient conditions that lead to OTF mechanisms having good welfare guarantees for XoS valuations. Specifically, suppose that we are given as input an *allocation algorithm*  $Alg$ . We propose three conditions that a price vector must satisfy so that the OTF mechanism with these prices provides an expected welfare that is comparable to that of the allocation returned by  $Alg$  in expectation over  $\mathcal{F}$ .

**Notation:** Let  $Alg(\vec{v}) = (A_1(\vec{v}), \dots, A_N(\vec{v}))$  denote the allocation output by  $Alg$  for valuation profile  $\vec{v}$ . Suppose that under this allocation,  $N_i(\vec{v})$  represents the set of buyers to whom good  $i$  is allocated, and  $k_i(\vec{v}) = |N_i(\vec{v})|$ . Fixing  $\vec{v}$ , let  $o_j$  be the XoS clause corresponding to buyer  $j$  that maximizes  $o_j(A_j(\vec{v}))$ , i.e.,  $o_j(A_j(\vec{v})) = v_j(A_j(\vec{v}))$ . Then, we define the total value derived from good  $i$ ,  $V_i(\vec{v}) = \sum_{j \in N_i(\vec{v})} o_j(i)$ , and also  $V_i = E_{\vec{v} \sim \mathcal{F}}[V_i(\vec{v})]$ . Similarly, we can also define the quantity representative of social welfare  $SW_i(\vec{v}) = \sum_{j \in N_i(\vec{v})} o_j(i) - \sum_{\ell=1}^{k_i(\vec{v})} c_i(\ell)$ , and its expected value  $SW_i$ .

**Lemma 3.5.** *Given an approximation algorithm  $Alg$ , suppose that there exists a price vector  $\vec{p}^*$  and an upper bound supply vector  $\vec{k}^*$ , satisfying the following conditions*

1. for every good  $i$ ,  $p_i^* k_i^* - C_i(k_i^*) \geq 0$ ;
2. for every good  $i$ ,  $V_i - p_i^* E_{\vec{v} \sim \mathcal{F}}[k_i(\vec{v})] \geq \frac{1}{\alpha} SW_i$ ;
3. for every good  $i$ ,  $p_i^* k_i^* - C_i(k_i^*) \geq \frac{1}{\alpha} SW_i$ .

Then, an on the fly mechanism with parameters  $(\vec{p}^*, \vec{k}^*)$  results in a welfare of  $\frac{E_{\vec{v} \sim \mathcal{F}}[SW(Alg(\vec{v}))]}{\alpha}$ .

*Proof.* Fix some valuation profile  $\vec{v} \sim \mathcal{F}$ , and a buyer  $j$ . Suppose that for this valuation,  $Sold_j(\vec{v})$  is the set of goods that are sold out (supply limit reached) when buyer  $j$  arrives. Consider another valuation profile  $\vec{a}_{-j}$  for buyers other than  $j$ , so that this valuation is independent of  $\vec{v}$ . Define  $\vec{a} = (v_j, \vec{a}_{-j})$ ,

and define  $o_j$  to be the XoS clause that maximizes  $A_j(\vec{a})$ , i.e.,  $v_j(A_j(\vec{a})) = o_j(A_j(\vec{a})) = \sum_{i \in A_j(\vec{a})} o_j(i)$ .

Under the fixed valuation  $\vec{v}$  and for each  $\vec{a}_{-j}$ , buyer  $j$  could have purchased the set of goods in  $A_j(\vec{a}) \setminus Sold_j(\vec{v})$ . Therefore, the utility of this buyer under the fixed valuation is at least  $E_{\vec{a}_{-j}}[\sum_{i \in \mathcal{I}} \mathbb{1}(j \in N_i(\vec{a})) \cdot \mathbb{1}(i \notin Sold_j(\vec{v})) \cdot o_j(i) - p_i^*]$ .

Observe that  $Sold_j(\vec{v})$  does not depend on  $\vec{a}_{-j}$  whereas  $N_i(\vec{a})$  and  $o_j(i)$  depend only on  $\vec{a}$ . Therefore, summing up the above quantity for all buyers and taking the expectation over  $\vec{v}$ , we get that the expected total surplus (buyer utility) is at least  $\sum_{i \in \mathcal{I}} \sum_{j=1}^N E_{\vec{v}}[\mathbb{1}(i \notin Sold_j(\vec{v}))] E_{\vec{a}}[\mathbb{1}(j \in N_i(\vec{a})) \cdot o_j(i) - p_i^*]$ . Applying the definition of  $V_i$ , this quantity can be rewritten as  $\sum_{i \in \mathcal{I}} Pr(i \text{ is not sold out}) E_{\vec{v}}[V_i(\vec{v}) - k_i(\vec{v}) p_i^*]$ . Finally, from Condition (2) of the lemma, we get our final lower bound on the aggregate buyer utility (surplus).

$$E_{\vec{v} \sim \mathcal{F}} \left[ \sum_{j=1}^N u_j(\vec{v}) \right] \geq \frac{1}{\alpha} \sum_{i \in \mathcal{I}} Pr(i \text{ is not sold out}) SW_i.$$

Next, for any fixed valuation  $\vec{v}$ , let  $S(\vec{v})$  and  $U(\vec{v})$  be the set of items that are sold out (all  $k_i^*$  copies), and the ones that are not respectively. Suppose that for any  $\vec{v}$ ,  $t_i(\vec{v})$  is the number of copies of good  $i$  sold by the mechanism. Then, the seller's profit (in expectation) is exactly  $E_{\vec{v} \sim \mathcal{F}}[\sum_{i \in S(\vec{v})} (k_i^* p_i^* - C_i(k_i^*)) + \sum_{i \in U(\vec{v})} (t_i(\vec{v}) p_i^* - C_i(t_i(\vec{v})))]$ .

Moreover, it is not hard to deduce that the term  $t_i(\vec{v}) p_i^* - \sum_{\ell=1}^{t_i(\vec{v})} c_i(\ell)$  is non-negative for all  $i$  as per our choice of  $p_i^*$ . Therefore, ignoring the profit due to the unsold items, we can apply the third condition in the lemma statement and get the following lower bound on the expected profit:  $\frac{1}{\alpha} \sum_{i \in \mathcal{I}} Pr(i \text{ is sold out}) SW_i$ .

By definition, social welfare is simply the expected buyer surplus plus the expected seller profit. Adding up our two lower bounds, we can complete the proof of the lemma.  $\square$

### Final leg: Showing the Black-Box Result

Armed with our framework from the previous sections, we are now ready to prove the black-box reduction in Claim 3.1. Specifically, we will compute prices based on the first part of the framework and then show that a mechanism with these prices satisfies the conditions of Lemma 3.5 for  $\alpha = 2$ .

- **Pricing Scheme:** For every good  $i$ , set (i)  $p_i^* = \frac{V_i}{2k_i^*} + \frac{E_{\vec{v} \sim \mathcal{F}}[C_i(k_i(\vec{v}))]}{2k_i^*}$  and (ii)  $k_i^* = E_{\vec{v} \sim \mathcal{F}}[k_i(\vec{v})]$ .

Note that  $p_i^*$  denotes the price that results in profit-surplus equivalence, taken in expectation over  $\mathcal{F}$ . We assume without loss of generality that  $k_i^*$  is integral for all  $i$ . It only remains to prove that these prices and supply constraints satisfy the conditions of Lemma 3.5 with  $\alpha = 2$ .

1. **Condition (1):**  $p_i^* k_i^* - C_i(k_i^*) \geq 0$ : By definition,  $p_i^* k_i^* - C_i(k_i^*) = \frac{V_i}{2} + \frac{1}{2} E_{\vec{v} \sim \mathcal{F}}[C_i(k_i(\vec{v}))] - C_i(k_i^*)$ . Since the cost functions are convex, we can apply Jensen's inequality and get that  $E_{\vec{v} \sim \mathcal{F}}[C_i(k_i(\vec{v}))] \geq C_i(E_{\vec{v} \sim \mathcal{F}}[k_i(\vec{v})]) = C_i(k_i^*)$ . Therefore, we get that  $p_i^* k_i^* - C_i(k_i^*) \geq \frac{V_i}{2} - \frac{E_{\vec{v} \sim \mathcal{F}}[C_i(k_i(\vec{v}))]}{2} = \frac{1}{2} SW_i \geq 0$ .

2. **Condition (2):**  $V_i - p_i^* E_{\vec{v} \sim \mathcal{F}}[k_i(\vec{v})] \geq \frac{1}{\alpha} SW_i$ : As per the definition of  $p_i^*$ , we have that  $p_i^* k_i^* = \frac{1}{2}(V_i + E_{\vec{v} \sim \mathcal{F}}[C_i(k_i(\vec{v}))])$ . Taking the negative of this equation and adding  $V_i$  to both sides gives us  $V_i - p_i^* k_i^* = \frac{1}{2}(V_i - E_{\vec{v} \sim \mathcal{F}}[C_i(k_i(\vec{v}))]) = \frac{1}{2} SW_i$ .
3. **Condition (3):**  $p_i^* k_i^* - C_i(k_i^*) \geq \frac{1}{\alpha} SW_i$ : This follows from the proof of Condition (1) for  $\alpha = 2$ .  $\square$

## 4 Commitment Mechanisms

We move on to the much harder class of mechanisms, where the seller has to initially commit to producing a fixed amount of each good under uncertainty, and therefore, incurs a loss on every item that is not sold. Given an instance of our problem and an algorithm  $Alg$ , using the same notation as in the proof of Claim 3.1, the expected social welfare provided by the algorithm is  $E_{\vec{v} \sim \mathcal{F}}[SW(Alg(\vec{v}))] = E_{\vec{v} \sim \mathcal{F}}[\sum_{j=1}^N v_j(A_j(\vec{v})) - \sum_{i \in \mathcal{I}} C_i(k_i(\vec{v}))]$ . Then, the instance is said to be  $\alpha$ -bounded with respect to  $Alg$  for  $\alpha \geq 1$  if  $E_{\vec{v} \sim \mathcal{F}}[\sum_{j=1}^N v_j(A_j(\vec{v}))] \geq \alpha E_{\vec{v} \sim \mathcal{F}}[\sum_{i \in \mathcal{I}} C_i(k_i(\vec{v}))]$ .

Our main result in this section is a commitment mechanism that obtains a  $\frac{4\alpha-1}{\alpha-2}$  approximation to the optimum welfare: for large enough values of  $\alpha$ , we obtain a constant factor approximation and as  $\alpha \rightarrow \infty$ , the performance of the commitment mechanism approaches that of the OTF mechanism from Theorem 3.3.  $\alpha$ -boundedness captures the *recoverable welfare* in terms of the initial investment. In previous work [Feige *et al.*, 2013],  $\alpha$ -boundedness was defined in terms of the optimum solution for a given instance; on the other hand, our definition depends both on the instance and the benchmark algorithm.

**Theorem 4.1.** 1. *Given an allocation algorithm  $Alg$ , for every instance with fractionally subadditive buyers that is  $\alpha$ -bounded with respect to  $Alg$  for  $\alpha \in [2, \infty]$ , there exists a commitment mechanism that extracts a welfare of  $E_{\vec{v} \sim \mathcal{F}}[\frac{SW(Alg(\vec{v}))}{2}] \frac{\alpha-2}{\alpha-1}$ .*

2. *For any given instance with fractionally subadditive buyers, we can compute posted prices in poly-time such that the resulting commitment mechanism always guarantees a social welfare within a factor  $4\frac{\alpha-1}{\alpha-2}$  of the optimum allocation, as long as the instance is  $\alpha$ -bounded with respect to the algorithm from Claim 3.2.*

**Discussion:** Essentially the parameter  $\alpha$  allows us to distinguish between *good* and *bad* instances of our problem. The performance of our commitment mechanism improves as  $\alpha$  increases. In many markets, it is reasonable to expect that the cost of producing the goods is not too large in comparison to the social value [Feige *et al.*, 2013]. Alternatively, if the production cost function is sufficiently convex, then the solution returned by our algorithm from Claim 3.2 is always  $\alpha$ -bounded for a large enough value of  $\alpha$  to ensure good welfare. For example, if  $c_i(\ell) = \ell^2$  for all  $i \in \mathcal{I}$  and  $Alg$  allocates at least 3-copies of every good, then the resulting allocation is  $\alpha$ -bounded for  $\alpha = 2.5$  and our commitment mechanism provides a welfare guarantee of  $\frac{1}{6} E_{\vec{v}}[SW(Alg(\vec{v}))]$ .

(*Proof Sketch*) The proof relies on the same framework as the OTF mechanism from Section 3, so we only sketch the

key differences. In the proof of Claim 3.1, we defined profit and surplus functions for every  $i \in \mathcal{I}$  and selected a price  $p_i^*$  so that the profit due to the sold items is exactly the surplus due to the unsold items, and the profit due to the unsold items is non-negative. However, this is no longer true for commitment mechanisms as the profit due to the unsold items may be negative due to the incurred production costs, thereby leading to poor welfare. To compensate for the negative profits when items are unsold, the main idea is to reduce prices so that the higher surplus offsets the increased production costs when the items are unsold. Using the same notation as in the proof of Claim 3.1, we define a reduced surplus function for each good  $i$ ,  $h_i(p) = V_i - p E_{\vec{v} \sim \mathcal{F}}[k_i(\vec{v})] - E_{\vec{v}}[C_i(k_i(\vec{v}))] = f_i(p) - E_{\vec{v}}[C_i(k_i(\vec{v}))]$ . For every good  $i \in \mathcal{I}$ , we select its price  $p_i^*$  to be the point where  $h_i(p) = \pi_i(p)$ , i.e., where reduced surplus and profit meet, and quantity to produce as  $k_i^* = E_{\vec{v} \sim \mathcal{F}}[k_i(\vec{v})]$ . Using the same sequence of ideas as in the proof of Claim 3.1, the theorem follows.  $\square$

## 5 Conclusions and Future Directions

In this work, we considered the problem of maximizing social welfare in the presence of distributional information regarding buyer valuations. Our main contribution is a non-discriminatory posted pricing mechanism that achieves one-fourth of the optimum welfare even when the seller faces convex production costs on the goods. Convex costs are a natural generalization of fixed supply and capture the *diseconomies of scale* that occur when expensive resources are utilized to achieve higher levels of production, for example, when manufacturing plants have a standard capacity up to which they incur constant marginal costs but the marginal cost increases beyond this threshold due to overtime, new equipment, etc. [Hochbaum and Wagner, 2015]. Other applications of convex costs include the energy cost incurred by a data center as a function of utilization [Buchbinder *et al.*, 2011]. A natural question that arises is whether our results can be extended to settings with concave production costs or economies of scale. As a first step in this direction, we can show good approximations for social welfare when the derivative of the concave function is not too volatile.

All of our results assume that we can access the buyer valuation functions by means of oracle queries, which is a standard assumption in the literature. However, it is natural to envisage prior-free scenarios where the seller has absolutely no information about buyer valuations. Can we design posted pricing mechanisms in such an extreme case? In the full version of this paper (available on arXiv), we answer this question in the affirmative by presenting an OTF mechanism for markets with XoS buyers and convex costs that achieves a  $O(\log m \log N)$ -approximation to the optimum social welfare. Previously, good approximations were only known for dynamic pricing mechanisms [Blum *et al.*, 2011].

## Acknowledgements

The author wishes to express his gratitude to Elliot Anshelevich and Brendan Lucier for providing useful feedback on an earlier version of this paper.

## References

- [Anderson *et al.*, 2010] Eric T Anderson, Duncan I Simester, et al. Price stickiness and customer antagonism. *The Quarterly Journal of Economics*, 125(2):729–765, 2010.
- [Bartal *et al.*, 2003] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of TARK*, pages 72–87, 2003.
- [Bei and Huang, 2011] Xiaohui Bei and Zhiyi Huang. Bayesian incentive compatibility via fractional assignments. In *Proceedings of SODA 2011*, pages 720–733, 2011.
- [Blum *et al.*, 2011] Avrim Blum, Anupam Gupta, Yishay Mansour, and Ankit Sharma. Welfare and profit maximization with production costs. In *Proceedings of FOCS 2011*, pages 77–86, 2011.
- [Blumrosen and Holenstein, 2008] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *Proceedings of ACM EC*, page 49, 2008.
- [Brânzei *et al.*, 2014] Simina Brânzei, Yiling Chen, Xiaotie Deng, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. The fisher market game: Equilibrium and welfare. In *Proceedings of AAAI 2014*, pages 587–593, 2014.
- [Buchbinder *et al.*, 2011] Niv Buchbinder, Navendu Jain, and Ishai Menache. Online job-migration for reducing the electricity bill in the cloud. In *Proceedings of NETWORKING*, pages 172–185, 2011.
- [Chakraborty *et al.*, 2009] Tanmoy Chakraborty, Zhiyi Huang, and Sanjeev Khanna. Dynamic and non-uniform pricing strategies for revenue maximization. In *Proceedings of FOCS 2009*, pages 495–504, 2009.
- [Chakraborty *et al.*, 2010] Tanmoy Chakraborty, Eyal Even-Dar, Sudipto Guha, Yishay Mansour, and S. Muthukrishnan. Approximation schemes for sequential posted pricing in multi-unit auctions. In *Proceedings of WINE 2010*, pages 158–169, 2010.
- [Chawla *et al.*, 2010] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of STOC 2010*, pages 311–320, 2010.
- [Cohen-Addad *et al.*, 2016] Vincent Cohen-Addad, Alon Eden, Michal Feldman, and Amos Fiat. The invisible hand of dynamic market pricing. In *Proceedings of ACM EC 2016*, pages 383–400, 2016.
- [Daskalakis and Pierrakos, 2011] Constantinos Daskalakis and George Pierrakos. Simple, optimal and efficient auctions. In *Proceedings of WINE*, pages 109–121, 2011.
- [Dobzinski *et al.*, 2010] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010.
- [Dughmi *et al.*, 2017] Shaddin Dughmi, Jason D. Hartline, Robert D. Kleinberg, and Rad Niazadeh. Bernoulli factories and black-box reductions in mechanism design. In *Proceedings of STOC*, 2017.
- [Feige *et al.*, 2013] Uriel Feige, Nicole Immorlica, Vahab S. Mirrokni, and Hamid Nazerzadeh. PASS approximation: A framework for analyzing and designing heuristics. *Algorithmica*, 66(2):450–478, 2013.
- [Feldman *et al.*, 2015] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of SODA*, pages 123–135, 2015.
- [Gul and Stacchetti, 1999] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic theory*, 87(1):95–124, 1999.
- [Hartline and Roughgarden, 2009] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of ACM EC 2009*, pages 225–234, 2009.
- [Hochbaum and Wagner, 2015] Dorit S Hochbaum and Michael R Wagner. Production cost functions and demand uncertainty effects in price-only contracts. *IIE Transactions*, 47(2):190–202, 2015.
- [Hsu *et al.*, 2016] Justin Hsu, Jamie Morgenstern, Ryan Rogers, Aaron Roth, and Rakesh Vohra. Do prices coordinate markets? In *Proceedings of STOC 2016*, 2016.
- [Huang and Kim, 2015] Zhiyi Huang and Anthony Kim. Welfare maximization with production costs: a primal dual approach. In *Proceedings of SODA 2015*, pages 59–72, 2015.
- [Maehara *et al.*,] Takanori Maehara, Yasushi Kawase, Hanna Sumita, Katsuya Tono, and Ken-ichi Kawarabayashi. Optimal pricing for submodular valuations with bounded curvature. In *Proceedings of AAAI 2017*.
- [Roughgarden and Talgam-Cohen, 2015] Tim Roughgarden and Inbal Talgam-Cohen. Why prices need algorithms. In *Proceedings of ACM EC 2015*, pages 19–36, 2015.