Price Doubling and Item Halving: Robust Revenue Guarantees for Item Pricing

ELLIOT ANSHELEVICH, Rensselaer Polytechnic Institute
SHREYAS SEKAR, Rensselaer Polytechnic Institute

We study approximation algorithms for revenue maximization based on static *item pricing*, where a seller chooses prices for various goods in the market, and then the buyers purchase utility-maximizing bundles at these given prices. We formulate two somewhat general techniques for designing good pricing algorithms for this setting: *Price Doubling* and *Item Halving*. Using these techniques, we unify many of the existing results in the item pricing literature under a common framework, as well as provide several new bicriteria algorithms for approximating both revenue and social welfare simultaneously. More specifically, for a variety of settings with item pricing, we show that it is possible to deterministically obtain a log-approximation for revenue and a constant-approximation for social welfare *simultaneously*; thus one need not sacrifice revenue if the goal is to still have decent welfare guarantees.

The main technical contribution of this paper is a $O((\log m + \log k)^2)$ -approximation algorithm for revenue maximization based on the item halving technique, for settings where buyers have XoS valuations, where m is the number of goods and k is the average supply. Surprisingly, ours is the first known item pricing algorithm with polylogarithmic approximation for such general classes of valuations, and partially resolves an important open question from the algorithmic pricing literature about the existence of item pricing algorithms with logarithmic factors for general valuations [4]. We also use the item halving framework to form envy-free item pricing mechanisms for the popular setting of multi-unit markets, providing a log-approximation to revenue in this case.

Additional Key Words and Phrases: Item Pricing; Revenue Maximization; Non-Discriminatory Pricing; Combinatorial Valuations;

1 INTRODUCTION

In this paper, we tackle the following fundamental problem in algorithmic economics:

Given a market with a set $\mathcal{I}$ of m goods and buyers having complex valuations, compute a set of prices (one per good) that maximize revenue once buyers purchase utility maximizing bundles.

The *item pricing* formulation described above is an abstraction that captures many types of economic interactions between buyers and sellers. Arguably, the appeal of item pricing stems from its simplicity and decentralized nature: the seller posts prices on the goods in advance, and each buyer purchases the most desirable set of available goods at the given prices. Moreover, although

This work was partially supported by NSF awards CNS-1218374 and CCF-1527497.

Author's addresses: E. Anshelevich and S. Sekar, Computer Science Department, Rensselaer Polytechnic Institute, 110, 8th Street, Troy, NY-12180. Email: eanshel@cs.rpi.edu, shreyas.sekar@gmail.com.

The full version of this paper containing proofs of all theorems can be accessed at <https://arxiv.org/abs/1611.02442>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'17, June 26–30, 2017, Cambridge, Massachusetts, USA. ACM ISBN 978-1-4503-4527-9/17/06...\$15.00.

© 2017 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

DOI: <https://doi.org/10.1145/3033274.3085117>

there are exceptions, a seller often cannot choose to charge different prices for different customers, i.e., the item prices are *non-discriminatory*. The lack of price discrimination is the crucial feature that distinguishes our model from the body of work on general posted pricing in mechanism design, where prices are charged per buyer, and not per good as we do in this work. To illustrate the subtleties involved here, for unit-demand [12] and even submodular valuations [9], one can compute per-buyer posted pricing schemes that yield constant factor approximations for revenue, whereas, the item pricing problem that we consider is known to be $O(\log^{1-\epsilon} m)$ -hard to approximate for unit-demand valuations [11].

Item pricing (to maximize revenue) occupies a central place in the study of markets (see for example [7, 22]). Despite the substantial body of work in this domain, surprisingly little is known about maximizing revenue for complex buyer valuations such as submodular or XoS; in fact, a long-standing open question [4, 10] is ‘whether static item prices can be used to achieve non-trivial revenue guarantees for submodular and more general valuations?’. We attempt to fill this conspicuous void in the literature by presenting *the first item pricing algorithm with a polylogarithmic approximation for revenue* for such valuations, partially resolving the open question. Previously, logarithmic revenue guarantees using non-discriminative pricing strategies were only known for simpler settings such as markets with unit-demand buyers [22]. In the process, we also develop a general framework for item pricing via black-box reductions that unifies many of the previous revenue-maximization results in the literature and helps us transform pricing mechanisms for revenue maximization into those that yield bicriteria (revenue,welfare) guarantees.

Informal Model. For the majority of this work, we study item pricing mechanisms in which the seller posts a price p_i for each good $i \in \mathcal{I}$, followed by a consumption stage in which buyers arrive in some order and purchase a utility maximizing subset of the available goods at the fixed prices. Each buyer j has a monotone valuation function $v_j : 2^{\mathcal{I}} \rightarrow \mathbb{R}^+$ and a utility-maximizing bundle refers to a set $S \subseteq \mathcal{I}$ maximizing $v_j(S) - \sum_{i \in S} p_i$. We refer to this as the *sequential mechanism*: in this model, buyers still see the same prices on items, but once the supply of a good is exhausted, it is no longer available for sale to future buyers. Specifically, we assume that some k_i units of good $i \in \mathcal{I}$ are available to be sold. In Section 4, we also consider the stronger class of *simultaneous mechanisms*, where buyers arrive at the same time and consume utility-maximizing bundles in parallel, with the mechanism ensuring that the supply constraints are not violated. In decentralized markets with limited supply where buyers consume goods in an asynchronous fashion, it is reasonable to assume that the sequential mechanism is a more natural model as compared to the simultaneous one [27]. Finally, we will use the term *item pricing* to encompass all mechanisms that only post a single price per good, including both of the mechanisms mentioned above.

Landscape of Item Pricing Results

To obtain a better understanding of our results, it is important to place them in the context of the larger body of work on item pricing for revenue maximization. The problem of revenue maximization is considered to be well understood for simpler valuations such as *unit-demand* [22] or *single-minded* [14], which admit $O(\log m)$ and $O(\sqrt{m} \log k_{\max})$ -approximations respectively via simultaneous mechanisms. On other hand, for submodular and XoS valuations, our understanding of this regime is largely incomplete owing to the super-logarithmic gap between the known lower bound of $\Omega(\log m)$ for submodular valuations, and the upper bound of $2^{\tilde{O}(\sqrt{\log m})}$ [4], corresponding to the current best approximation algorithm for revenue maximization with sequential mechanisms¹. While recent work [18] has yielded sequential item pricing mechanisms with $O(1)$ -approximations

¹The result assumes a unit-supply environment and hence, the bound depends only on the number of goods m

for social welfare, those techniques do not necessarily lead to any direct insights for the revenue problem.

To summarize, for the primary problem studied in this work, i.e., *sequential item pricing for revenue maximization with XoS valuations*: (i) the state of the art is a $2^{\tilde{O}(\sqrt{\log m})}$ -approximation algorithm, (ii) no algorithm can achieve revenue that is within a $\Omega(\log m)$ -factor of the optimum welfare for all instances [4], (iii) no poly-time algorithm can achieve a $\log^{1-\epsilon} m$ -approximation for any $\epsilon > 0$ [11]. Against this backdrop, we present an algorithm whose revenue is within a $O(\log^2 m)$ -factor of the optimum welfare for unit-supply settings, significantly narrowing the gap between the upper and lower bounds. Our result generalizes the logarithmic revenue guarantees that were previously obtained only for restricted valuations such as unit-demand functions [22] to the more general XoS class. Prior to this, logarithmic guarantees for XoS valuations were only known for more complex mechanisms involving bundle pricing [20] or price discrimination [10].

Finally, our work is motivated by concerns that are somewhat orthogonal to those prevalent in fields such as mechanism design. Specifically, we are less interested in constructing novel combinatorial auctions for revenue maximization, and more so in gaining a fundamental understanding of how markets with item pricing can operate. Towards this end, as is common in the item pricing literature [14, 20, 22], we eschew concerns such as incentive compatibility, and consider a full information model, where the seller is aware of each buyer's valuation function.

Revenue, Welfare, and Black-box Reductions. A noteworthy aspect of this work is that, as in much of the previous literature, all of our approximation factors are derived by comparing the revenue obtained by our algorithm to the social welfare of the optimum allocation. Measuring revenue in terms of social welfare allows for our results to have some useful implications [2]. For instance, since the optimum welfare is an upper bound on the revenue obtained by any 'individually rational mechanism', our main result implies that no pricing mechanism can improve upon item pricing (specifically, our algorithm's revenue) by a factor larger than $O(\log^2 m)$ for unit-supply settings. This includes, for instance, an algorithm that updates prices dynamically every round or one that charges per-buyer posted prices. Finally, all of the results in this paper are presented in the form of black-box reductions that convert *any given input allocation* to an item pricing solution, whose revenue guarantees depend on the social welfare of the original allocation. Such black-box reductions are highly sought after in the pricing literature since they allow us to apply the extensive techniques developed for the purely algorithmic resource allocation problem to a more game-theoretic pricing problem where the seller can only indirectly control what items are allocated to the buyers.

1.1 Summary of Contributions

In this paper, we consider a typical setting consisting of a set of m goods in limited supply, with k being the supply averaged over all of the goods in the market, and $k_{\max}$ being the maximum supply of any good. There are N buyers in the market having combinatorial valuation functions, and under a given set of prices, each buyer purchases a subset of the available goods that maximizes her quasilinear utility, i.e., value derived minus price paid.

Unifying Charging Framework. At a fundamental level, our main contribution is a simple, but general charging framework that provides a recipe for revenue maximization by *charging* welfare differentials to revenue. In combination with the two complementary algorithmic paradigms of *price doubling* and *item halving*, the framework can be used to derive black-box reductions from 'item pricing' to 'item allocations'. Price doubling, or multiplicative pricing has been previously employed to derive pricing mechanisms with logarithmic guarantees in the literature [22], whereas item

halving is a new technique for pricing that we propose in this work. In some senses, item halving can be viewed as a more granular version of the supply-limiting approach in mechanism design [26]: our generalization allows us to apply this technique for general combinatorial valuations. The charging framework serves as a focal point for unifying a number of results in the pricing literature (see Table 1); we add value to these results by transforming them into bicriteria approximation algorithms that simultaneously guarantee good revenue and welfare. Secondly, we leverage our framework to design new approximations for revenue maximization in other settings.

Item Halving + Charging Framework: Sequential Mechanism for XoS Buyers. The main technical contribution of this paper is a new approximation algorithm for revenue maximization based on a novel algorithmic approach that we term item halving. The sequential mechanism that we design achieves the stated revenue guarantee for two models of buyer arrival: (i) when the arrival is adversarial within a polynomially bounded set of buyer orders, and (ii) when buyers arrive according to an arbitrary probability distribution. For the completely adversarial case, we present a non-polynomial-time algorithm with the same guarantee, that serves as an existence result.

(CENTERPIECE RESULT). *For settings where buyers have XoS valuations, we can compute item prices in poly-time such that the sequential mechanism with these prices results in a $O((\log m + \log k)^2)$ -approximation to the optimum revenue for random and polynomially adversarial buyer arrival orders.*

Note that XoS valuations are a strict generalization of submodular functions. The above result assumes access to Demand and XoS oracles for accessing buyer valuations (see Section 2). In unit-supply settings ($k = 1$), our approximation guarantee is $O(\log^2 m)$.

Key Features of our Result: The above theorem provides a positive answer to the long standing open question of whether static item pricing could be used to obtain (poly-)logarithmic approximation factors for complex valuations, and in the process, significantly improves upon the previous upper bound of $2^{O(\sqrt{\log(mk) \log \log(mk)})}$ [4]. Prior to this, similar approximation factors were only known for more general pricing schemes such as bundle pricing [20] or discriminatory pricing [10]. In contrast, ours is the first known poly-logarithmic approximation algorithm based on non-discriminatory item pricing. Moreover, we come close to closing the gap between the upper and lower bounds for this setting, the latter being a $\Omega(\log m)$ factor for submodular valuations with unit supply.

Other Implications of our Framework After proving our centerpiece result, we devote the rest of this paper to a more detailed exploration of our charging framework resulting in new item pricing algorithms for other settings. All of the below results are obtained via the simultaneous mechanism, which is stronger than the sequential mechanism (see Section 2). Thereby, we make the case that our charging framework is not specifically tied to any specific mechanism and bears strong implications for future work involving item pricing.

- **Price Doubling and Bicriteria Approximations** By employing the price doubling approach within the confines of our framework, we identify a ‘generic algorithm for item pricing’ that unifies many of the existing results under a common umbrella. Not only are we able to re-derive guarantees from the item pricing literature, we actually strengthen *all of these results* by transforming them to bicriteria approximation algorithms, whose guarantees are listed in Table 1. Specifically, we present a black-box transformation from revenue maximization to welfare maximization for simultaneous item pricing mechanisms: “*given any α -approximate simultaneous mechanism for welfare maximization, we present a poly-time reduction to compute item prices that ensure a $O(\alpha(\log m + \log k + \log \alpha))$ -approximation to the optimal revenue, and simultaneously a 3α -approximation for welfare*”. The result holds for arbitrary combinatorial valuations as long as the black-box algorithm satisfies a mild requirement known as ‘local welfare maximization’ (See Section 4.1).

- **Item Halving for Multi-Unit Valuations** Finally, in Section 4.2 we apply the same item halving machinery developed for XoS valuations to obtain a simultaneous item pricing mechanism with a $O(\log m)$ -revenue guarantee for the popular class of *multi-unit valuations*, where a buyer's value depends only on the number of items that she receives. Previously, the best known result for this setting was a $O(\log m + \log N)$ -approximation based on sequential mechanisms [4].

Setting	Previous Results (Revenue)	Our Results {Revenue,Welfare}
Unit-Demand	$O(\log m + \log k)$ [22]	$\{O(\log m + \log k), 3\}$
Gross Substitutes	-	$\{O(\log m + \log k), 3\}$
Single-Minded	$O(\sqrt{m} \log k_{\max})$ [14]	$\{O(\sqrt{m} \log k_{\max}), O(\sqrt{m})\}$
Unlimited Supply	$O(\log m + \log N)$ [4]	$\{O(\log m + \log N), 2\}$
Bundle Pricing	$O(\alpha^* \log N)$ [20]	$\{O(\alpha^* \log N), 6\alpha^*\}$

Table 1. Bicriteria approximation factors obtained for various settings. For a given valuation class, α^* denotes the best known approximation factor for the allocation problem.

Discussion of Bicriteria Results Simultaneous bicriteria (revenue,welfare) approximations have featured previously in single-item auctions [15], and item pricing for special cases of unit-demand and multi-minded buyers [2]. Our result, on the other hand, has a far greater scope since it is not tied to any class of valuation (see Table 1). We also remark that while many revenue maximization algorithms in the literature do lead to implicit bicriteria factors, the bounds obtained for welfare are trivial since any α -approximation for revenue in terms of the optimum social welfare implies an α -approximation for social welfare as well. On the contrary, our framework allows us to derive explicit bicriteria approximations, where both the revenue and social welfare are comparable to their respective optimal lower bounds. For example, while the results of [22] immediately imply a $(O(\log m), O(\log m))$ -approximation for (revenue,welfare) for unit-demand buyers, our framework provides an improved $(O(\log m), O(1))$ -bicriteria approximation for the unit-supply setting. The main message behind our result is: ‘why only maximize revenue when we can get near-optimal welfare at no cost’.

1.2 Related Work

In terms of similarity, the current paper is closest to the extensive body of work on item pricing and more general envy-free schemes [6, 17, 21] for welfare and revenue maximization. Item pricing for maximizing only welfare has traditionally been a sought after area of research as it is equivalent to computing a Walrasian equilibrium minus the market clearing constraints; since our focus is primarily on revenue, we refer the reader to [19, 25] for more recent algorithmic perspectives on the subject. On the other hand, item pricing for revenue maximization has been gaining traction with respect to both sequential and simultaneous mechanisms. A steady stream of research has yielded near-optimal approximation algorithms for a variety of settings including unit demand [7, 13, 22], single minded [3, 14], multi-unit markets [6, 17], and unlimited supply settings [4]. One of the main contributions of this paper is a general framework that captures many of the above results, and provides a recipe for converting them into bicriteria approximations.

Despite the tremendous body of work on revenue maximization, ours is the first known poly-logarithmic approximation algorithm based on static, item pricing for complex valuations such as submodular and XoS functions. Partial exceptions include the $O(\log m)$ -approximation for ‘simple

submodular functions' in [4] and the $O(\log^3 m)$ -approximation in [10] for settings consisting of a large number of buyers having the exact same valuation function. Although our central result improves upon the upper bound of $2^{\tilde{O}(\sqrt{\log m})}$ from [4] for XoS functions, it is pertinent to mention that their result holds for a more general information model where the seller is not aware of the buyer valuations. For such a general model, it is reasonable to expect that the only possible strategy would be to price items uniformly (common price), as is done in [4, 10]. The presence of a matching lower bound of $2^{\Omega(\sqrt{\log m})}$ [10] for uniform pricing motivates the need for pricing different goods differently. For such a scheme, we argue that knowledge of buyer valuations is necessary since it allows us to quantify each good's relative value in the market.

As mentioned previously, this work differs considerably from the literature on *sequential posted pricing* in mechanism design (e.g., [5, 8, 12]). With the notable exception of [18], most of the pricing schemes in this area are non-anonymous, i.e., based on offering the same good to different buyers at different prices. On the contrary, our prices are static and do not change across rounds.

2 MODEL AND PRELIMINARIES

We consider a market setting with a set $\mathcal{N}$ of N buyers and set $\mathcal{I}$ of m goods. Each buyer j has a monotone valuation function $v_j : 2^{\mathcal{I}} \rightarrow \mathbb{R}^+$, whereas the seller is bound by a supply constraint on each good, i.e., at most k_i units of good $i \in \mathcal{I}$ are available to be sold. We use $k := \frac{\sum_{i \in \mathcal{I}} k_i}{m}$ to denote the average supply of goods in the market. Given an allocation $\vec{S} = (S_1, S_2, \dots, S_N)$ of goods to the buyers, we will use $SW(\vec{S}) = \sum_{j \in \mathcal{N}} v_j(S_j)$ to represent the social welfare of this allocation. For the same allocation, let $N_i(\vec{S})$ denote the set of buyers who received a copy of good i and $k_i(\vec{S}) = |N_i(\vec{S})|$. An allocation is *feasible* only if it does not violate the supply constraints i.e., for each $i \in \mathcal{I}$, no more k_i buyers receive a copy of this good. Finally, let $\theta(\vec{S}) = \sum_{i \in \mathcal{I}} k_i(\vec{S})$ denote the total number of items allocated to buyers in the given allocation $\vec{S}$.

In this work, we will primarily deal with the XoS class consisting of *fractionally subadditive* valuation functions; this class includes, for example, all submodular functions. A valuation function v is said to belong to this class if there exists a set of additive clauses or functions $(a_1, \dots, a_r)$ such that for any $T \subseteq \mathcal{I}$, $v(T) = \max_{j=1}^r a_j(T)$. Each additive function a_j has a single value $a_j(i)$ for each $i \in \mathcal{N}$ so that for a set T of agents, $a_j(T) = \sum_{i \in T} a_j(i)$. We also provide algorithms for the class of *multi-unit valuations*, where a buyer's valuation depends only on the number of items that she receives, i.e., for all $S, T \subseteq \mathcal{I}$ such that $|S| = |T|$, $v(S) = v(T)$. Some of our secondary results are shown for other standard classes of valuations including *unit-demand*, *gross substitutes*, *submodular*, and *single-minded* valuations, all of whose definitions can be found in the full version of this work or a standard text such as [24].

Finally, as is common in the literature, we assume that we have black-box access to certain oracles that allow us to query the set functions. Specifically, we consider the following types of oracles with respect to a valuation v : (i) *Value Oracle*, that when given a set $S \subseteq \mathcal{I}$ returns the value of $v(S)$, (ii) *Demand Oracle* that when queried with a vector of item prices $\vec{p}$ returns a set S that maximizes the quantity $v(S) - \sum_{i \in S} p_i$, and (iii) an *XoS oracle* that for an XoS function v and a set $T \subseteq \mathcal{I}$ returns the additive function a_T that maximizes $a_T(T)$. We refer the reader to [16] for a detailed discussion on the power and limitations of different types of oracles.

2.1 Pricing Mechanisms

In the current work, we restrict our attention to two types of mechanisms based on *non-discriminatory item pricing*, where the seller fixes a single price per good in advance, and buyers purchase utility-maximizing bundles at the given prices. The price paid by each buyer equals the sum of item prices of the goods that she receives.

- (1) **Simultaneous Mechanism.** In this mechanism, the seller posts a price p_i per good, buyers arrive simultaneously, and purchase a utility-maximizing subset $S_j \subseteq \mathcal{I}$ under the given prices. A pair $(\vec{p}, \vec{S})$ is said to be a *valid outcome* of the simultaneous mechanism if the following constraints are satisfied: (i) for every buyer $j \in \mathcal{N}$, S_j maximizes $v_j(S_j) - \sum_{i \in S_j} p_i$, and (ii) at most k_i buyers are allocated a copy of good i .
- (2) **Sequential Mechanism.** The sequential mechanism is very similar to the previous mechanism except that buyers arrive sequentially and once the supply of a good is exhausted, it is no longer available to future buyers. Specifically, this mechanism proceeds as follows:
 - (a) Initially, all of the goods are available.
 - (b) For each good $i \in \mathcal{I}$, the seller posts a single price p_i and decides on the quantity $0 \leq q_i \leq k_i$ of this good to supply.
 - (c) The buyers arrive in some order, and each buyer purchases her utility maximizing bundle from the set of available goods.
 - (d) At any stage of the mechanism, if q_i copies of good i are sold (i.e, the good is sold out), then it is marked as unavailable.

A triple $(\vec{p}, \vec{q}, \vec{S})$ is said to be a valid outcome of the sequential mechanism if by setting the prices p_i and supply constraints q_i , the above mechanism can result in the allocation $\vec{S}$ for some arrival order of the buyers. In fact, we will usually omit $\vec{q}$ from this and simply denote an outcome by $(\vec{p}, \vec{S})$, since if such an outcome is possible, then we can always set $q_i = k_i(\vec{S})$ to obtain the same outcome.

The sequential mechanism is a relaxation of the simultaneous one since any valid outcome of the latter mechanism is also a valid outcome for the sequential mechanism. Moreover, the outcome of the simultaneous mechanism is always *envy-free* as each buyer receives her utility-maximizing bundle and does not envy other buyers, whereas this may not be the case for the sequential mechanism. Finally, we consider three arrival models for buyers in the sequential mechanism: (i) *partially adversarial*, where the buyer arrival order comes from a polynomially bounded set of arrival orders, (ii) *random arrival*, where buyers arrive according to an arbitrary distribution D^Π over the arrival orders, and finally (iii) *fully adversarial*, where any of the $N!$ arrival orders are possible. For the objective of revenue maximization, we are able to provide efficient computational mechanisms for the first two arrival models, and a constructive existence result for the fully adversarial case.

Revenue and Surplus Given a valid outcome $(\vec{p}, \vec{S})$ of either item pricing mechanism, we define $Rev(\vec{p}, \vec{S}) := \sum_{i \in \mathcal{I}} p_i k_i(\vec{S}) = \sum_{j \in \mathcal{N}} \sum_{i \in S_j} p_i$ to be the resulting revenue, and $Surp(\vec{p}, \vec{S}) := \sum_{j \in \mathcal{N}} \{v_j(S_j) - \sum_{i \in S_j} p_i\}$ to be the aggregate surplus or buyer utility. It is not hard to see that $SW(\vec{S}) = Rev(\vec{p}, \vec{S}) + Surp(\vec{p}, \vec{S})$.

2.2 Charging Framework and Black-box Reductions

As a first step towards deriving black-box reductions from item pricing (for revenue maximization) to ‘item allocations’, we present a simple, but general charging framework for identifying pricing solutions with good revenue. The framework is not constructive and merely provides sufficient conditions for revenue approximation. The computationally challenging aspect is to actually design algorithms that compute prices obeying the properties spelled out below, which we do using

Price Doubling and Item Halving. Following a treatment of the charging framework, we briefly introduce the item halving technique that together with the framework yields our central result for fractionally subadditive valuations. In Section 4, we return to a more detailed examination of the charging framework and show how it can be used to unify existing results when applied in conjunction with the price doubling technique.

Revenue-Welfare Charging Property. Our charging framework is composed of Definition 2.1 and Claims 2.2 and 4.1. The driving force behind this framework is the following idea: suppose that we have a sequence of γ benchmark allocations that are provided as input, and we are able to design item pricing mechanisms such that the difference in welfare between successive allocations can be ‘charged’ to the mechanism’s revenue. Then, one of these mechanisms must provide revenue that is comparable to the social welfare of the ‘best benchmark allocation’. Finally, we remark that similar techniques based on charging welfare differentials to revenue have been used previously in the literature to compute revenue maximizing item prices for restricted classes of valuations [14, 22]. However, the framework that we present is considerably more general as it is not tied to any specific pricing strategy or class of buyer valuations.

Definition 2.1. Consider a sequence of benchmark allocations $\vec{B}^{(1)}, \vec{B}^{(2)}, \dots, \vec{B}^{(\gamma)}$, and pricing solutions $(\vec{p}^{(t)}, \vec{s}^{(t)})_{t=1}^{\gamma}$. The sequence of solutions is said to satisfy the *charging property* if $\exists \alpha, \beta \geq 1$, such that (i) for all $1 \leq t \leq \gamma - 1$, $SW(\vec{B}^{(t)}) - SW(\vec{B}^{(t+1)}) \leq \alpha Rev(\vec{p}^{(t)}, \vec{s}^{(t)}) + \frac{SW(\vec{B}^{(t)})}{\beta}$; (ii) $SW(\vec{B}^{(\gamma)}) \leq \alpha Rev(\vec{p}^{(\gamma)}, \vec{s}^{(\gamma)})$.

In most of the applications, γ and α will be $O(\log m)$, and β will be sufficiently large. We now show how to leverage the property to obtain revenue that is within a $O(\gamma\alpha)$ -factor of the social welfare of the first benchmark allocation.

CLAIM 2.2. Consider a sequence of benchmark allocations $\vec{B}^{(1)}, \dots, \vec{B}^{(\gamma)}$, and pricing solutions $(\vec{p}^{(t)}, \vec{s}^{(t)})_{t=1}^{\gamma}$ that satisfy the charging property as per Definition 2.1. Then, there exists $1 \leq \ell \leq \gamma$ such that $Rev(\vec{p}^{(\ell)}, \vec{s}^{(\ell)}) \geq \frac{1}{\gamma\alpha} SW(\vec{B}^{(1)}) \left(1 - \frac{\gamma-1}{\beta}\right)$.

PROOF. Suppose that for some $r \geq 1$, $\vec{B}^{(r)}$ represents the allocation with the maximum social welfare among the sequence of benchmark allocations, i.e., $r = \arg \max_{t=1}^{\gamma} SW(\vec{B}^{(t)})$. Define ℓ to be the integer index in the range $[1, \gamma]$ at which the pricing solution obtains the maximum revenue. So, for all $1 \leq t \leq \gamma$, $Rev(\vec{p}^{(\ell)}, \vec{s}^{(\ell)}) \geq Rev(\vec{p}^{(t)}, \vec{s}^{(t)})$. As per Definition 2.1, we know that for any $r \leq t \leq \gamma - 1$, the following property holds: $SW(\vec{B}^{(t)}) - SW(\vec{B}^{(t+1)}) \leq \alpha Rev(\vec{p}^{(t)}, \vec{s}^{(t)}) + \frac{SW(\vec{B}^{(t)})}{\beta}$.

Summing up the above sequence of inequalities from $t = r$ to $t = \gamma - 1$ along with the inequality $SW(\vec{B}^{(\gamma)}) \leq \alpha Rev(\vec{p}^{(\gamma)}, \vec{s}^{(\gamma)})$, we get that

$$\begin{aligned} SW(\vec{B}^{(r)}) &\leq \alpha \sum_{t=1}^{\gamma} Rev(\vec{p}^{(t)}, \vec{s}^{(t)}) + \sum_{t=1}^{\gamma-1} \frac{SW(\vec{B}^{(t)})}{\beta} \\ &\leq \gamma\alpha Rev(\vec{p}^{(\ell)}, \vec{s}^{(\ell)}) + (\gamma - 1) \frac{SW(\vec{B}^{(r)})}{\beta}. \end{aligned}$$

Therefore, we have that $Rev(\vec{p}^{(\ell)}, \vec{s}^{(\ell)}) \geq \frac{1}{\gamma\alpha} SW(\vec{B}^{(r)}) \left(1 - \frac{\gamma-1}{\beta}\right) \geq \frac{1}{\gamma\alpha} SW(\vec{B}^{(1)}) \left(1 - \frac{\gamma-1}{\beta}\right)$. $\square$

Price Doubling, Item Halving, and Revenue Maximization for XoS Buyers

According to our framework, if we are provided with a sequence of benchmark allocations and pricing outcomes that satisfy Definition 2.1, we have essentially solved the revenue maximization

problem. In order to obtain a full-fledged black-box reduction from pricing to allocation, we need to generate these solutions using a single input allocation, which we do using the price doubling and item halving approaches. Here, we describe the latter approach in the context of XoS valuations. This provides a natural segue to the central result of this paper: a $O((\log m + \log k)^2)$ -revenue approximation for XoS buyers. We defer the treatment of price doubling to Section 4.

Item Halving The main idea behind this technique is that the number of items allocated to buyers in successive allocations decreases by a multiplicative factor of two. Towards this end, each allocation starting with the input allocation $\tilde{B}^{(1)}$ is intrinsically tied to a pricing vector so that running either mechanism in conjunction with those prices results in the successive allocation.

Item Halving for XoS Buyers: Challenges and Techniques. Putting aside our framework for the time being, let us consider a fundamental question: how do we price items in unit supply for buyers having complex valuations? Almost all of the literature has relied on price doubling to obtain item prices that approximately maximize revenue. Specifically, powerful approximation algorithms have resulted from the simple observation made in [4] that offering a ‘single buyer’ a random price (on all goods) from the set $\{2^t \frac{H}{m^2}\}_{t=1}^{1+\log^2 m}$, for a suitable choice of H , is sufficient to ensure good revenue in expectation. Unfortunately, it is known [10] that static pricing based on this technique could be rather sub-optimal ($2^{\Omega(\sqrt{\log m})}$) even for two buyers, since the earlier buyer could have large valuations for low revenue goods intended for the later buyer. The natural next step is to then consider pricing items differently based on the value they generate for buyers in some allocation, for e.g., as in [18]. However, this need not result in good revenue either for similar reasons. Therefore, we infer that current pricing techniques including price doubling are ineffectual in a static setting since they do not take into account the dependencies introduced by limited supply.

The crucial observation that enables our centerpiece result is inspired by Definition 2.1. Suppose that we compute prices for each good based on the average value that it generates for buyers in some benchmark allocation multiplied by a conservative scaling parameter, and that the sequential mechanism based on these prices results in poor revenue. If the prices truly reflect each item’s value and the scaling parameter is judiciously chosen so that the items do not get sold out too early, then we can infer that every utility-maximizing buyer would purchase a high-value bundle with only a few items. More specifically, our result is powered by the following *dichotomous black-box reduction*: given a benchmark allocation, either the mechanism with the above prices results in good revenue, or we can identify a dense sub-allocation that only uses half the number of items as the original benchmark allocation, but whose social welfare is close to the benchmark. Recursively applying this reduction, we either obtain a pricing strategy with good revenue or we simply end up with one good, for which revenue maximization is trivial. We refer to this technique as item halving as each successive application of our black-box mechanism decreases the number of allocated items by a multiplicative factor of two.

3 ITEM PRICING FOR REVENUE MAXIMIZATION WITH XOS BUYERS

In this section, we present the central result of this paper: a computationally efficient, sequential posted pricing mechanism that achieves a $O((\log m + \log k)^2)$ -approximation to the optimum revenue. In fact, as discussed previously, our approximation factor is obtained with respect to the social welfare maximizing allocation. Depending on the arrival order, our $O((\log m + \log k)^2)$ sequential mechanism comes in two flavors: the partially adversarial model, where the arrival order is adversarial within a bounded set of arrival orders, and the random model, where the buyers arrive according to an arbitrary distribution D^Π .

Notation We will implicitly assume that we are dealing with XoS valuations for the rest of the section. Given an allocation $\vec{S} = (S_1, S_2, \dots, S_N)$, we use $x_j^{S_j}$ to denote the XoS clause that maximizes buyer j 's valuation for the set S_j of goods, i.e., $x_j^{S_j}(S_j) = v_j(S_j)$. Then, we can define the total value derived by buyers from any given good $i \in \mathcal{I}$ with respect to an allocation $\vec{S}$ as $U_i(\vec{S}) = \sum_{j \in N_i(\vec{S})} x_j^{S_j}(i)$. Additionally, for any integer $1 \leq r \leq k_i(\vec{S})$, we define $Top_i(r, \vec{S}) \subseteq N_i(\vec{S})$ to be the (r) buyers in $N_i(\vec{S})$ with the r highest values of $x_j^{S_j}(i)$.

Now, we are ready to define our first black-box algorithm that takes as input an allocation $\vec{A}$ and a parameter $\gamma \geq 1$, and computes a series of benchmark allocations $(\vec{B}^{(1)}, \dots, \vec{B}^{(\gamma)})$ as well as pricing solutions $((\vec{p}^{(1)}, \vec{S}^{(1)}) \dots, (\vec{p}^{(\gamma)}, \vec{S}^{(\gamma)}))$ that are in accordance with Definition 2.1. We present our black-box reduction (the core algorithm) in two parts, where the first part handles the main reduction and the computation of successive benchmark allocations, and the second part handles the boundary or tail condition, where a buyer has a large valuation for a single item. Note that the purpose of our core algorithm is to simply compute prices (and supply limits) for our eventual poly-logarithmic sequential mechanism.

Algorithm 1: Core Algorithm

```

input : Allocation  $\vec{A}$ , parameter  $\gamma \geq 1$ .
1 Define  $\vec{B}^{(1)} = \vec{A}$ ;
2 for  $t \leftarrow 1$  to  $\gamma - 1$  do
3   Set  $\vec{p}^{(t)} = \text{Prices}(\vec{B}^{(t)}, \gamma)$  [See Figure 1];
4   Let  $\pi$  be an arrival order chosen by an arbitrary function  $\Gamma(\vec{p}^{(t)}, \vec{B}^{(t)})$  (to be fixed later);
5   Let  $\vec{S}^{(t)}$  be an allocation obtained upon running the sequential mechanism with prices
       $\vec{p}^{(t)}$ , supply  $\vec{k}(\vec{B}^{(t)})$ , and buyer arrival order  $\pi$ ;
      /* Case I: less than half the goods from  $\vec{B}^{(t)}$  are sold in  $\vec{S}^{(t)}$  */
6   if  $2\theta(\vec{S}^{(t)}) \leq \theta(\vec{B}^{(t)})$  then
7     | Set  $\vec{B}^{(t+1)} = \vec{S}^{(t)}$ ;
8   end
      /* Case II: more than half the goods from  $\vec{B}^{(t)}$  are sold in  $\vec{S}^{(t)}$  */
9   else Set  $\vec{B}^{(t+1)} = \text{Alloc-Unsold}(\vec{B}^{(t)}, \vec{S}^{(t)})$  [See Figure 2];
10 end

/* Tail Conditions */
11 Define  $j^*, i^*$  to be a buyer, good pair that maximizes  $v_j(\{i\})$ ;
12 Set  $\vec{p}_{i^*}^{(\gamma)} = v_{j^*}(\{i^*\}) - \epsilon$ , for some sufficiently small  $\epsilon > 0$ , and  $\vec{p}_i^{(\gamma)} = \infty$ , for  $i \neq i^*$ ;
13 Let  $\vec{S}^{(\gamma)}$  be the allocation in which  $j^*$  receives the item  $i^*$ , and no one else receives anything;

```

Informal Description of Core Algorithm. The core algorithm presented here (Algorithm 1) is the main workhorse behind all of our results. The purpose of this algorithm is to take an arbitrary allocation $\vec{A}$ as input along with an arrival order determining function Γ and return a sequence of benchmark allocations and pricing solutions that satisfy the charging property. By varying the formulation of the function Γ , we can use the same core procedure to derive revenue maximization

algorithms for all three arrival orders discussed previously (see Section 3.1). Beginning with the input allocation $\vec{A}$, in each iteration, we use the function $Prices(\vec{S}, \gamma)$ from Figure 1 to compute a price for each good that depends on the average value it generates in allocation $\vec{S} = \vec{B}^{(t)}$ scaled down by a factor of 2γ . Following this, we run the sequential mechanism with the computed prices for the arrival order π determined by Γ and obtain the pricing solution $(\vec{p}^{(t)}, \vec{S}^{(t)})$. Now, if the number of items allocated in $\vec{S}^{(t)}$ is at most half the number allocated in $\vec{B}^{(t)}$, then we set the benchmark allocation for the next iteration $\vec{B}^{(t+1)}$ to be $\vec{S}^{(t)}$. If the condition does not hold, then we invoke the function $Alloc-Unsold$ and set $\vec{B}^{(t+1)}$ to be the allocation comprising of all the items which were allocated in $\vec{B}^{(t)}$ but *unsold* in $\vec{S}^{(t)}$. Thus, at each iteration the number of items allocated in the benchmark solution reduces by at least a factor of two. Finally, the *tail conditions* provide a means for revenue maximization when the benchmark allocation consists of a single good.

Fig. 1. Function $Prices(\vec{S}, \gamma)$ Fig. 2. Function $Alloc-Unsold(\vec{B}, \vec{S})$

We now prove the main theorem of this section, which allows us to use the charging property for the allocations generated by the above algorithm. In Section 3.1, we show how to apply this algorithm and the charging property to actually produce mechanisms with high revenue. The modular nature of Algorithm 1 allows us to use the same core algorithm to obtain revenue maximizing prices for all three arrival orders: partially adversarial, random arrival, and fully adversarial (existence result) by simply varying the formulation of the function Γ .

THEOREM 3.1. *For $\gamma = \log m + \log k$, the benchmark allocations $\{\vec{B}^{(t)}\}_{t=1}^{t=\gamma}$ and pricing solutions $\{(\vec{p}^{(t)}, \vec{S}^{(t)})\}_{t=1}^{t=\gamma}$ generated by the above algorithm for any input allocation $\vec{A}$ satisfy the charging property (Definition 2.1) with $\alpha = \beta = 2\gamma$, and thus for some ℓ we have that*

$$SW(\vec{A}) \leq 4(\log m + \log k)^2 Rev(\vec{p}^{(\ell)}, \vec{S}^{(\ell)}).$$

Note that although this theorem immediately implies that at least one solution $(\vec{p}^{(\ell)}, \vec{S}^{(\ell)})$ has high revenue compared to the welfare of the starting allocation $\vec{A}$, this does not yet give us the result we need since the allocation $\vec{S}^{(\ell)}$ is specific to one particular arrival order of buyers, whereas the actual arrival order is either random or adversarial.

Proof of Theorem 3.1. By Claim 2.2, it suffices if we show that for every $1 \leq t \leq \gamma - 1$: (i) $SW(\vec{B}^{(t)}) - SW(\vec{B}^{(t+1)}) \leq 2\gamma Rev(\vec{p}^{(t)}, \vec{S}^{(t)}) + \frac{SW(\vec{B}^{(t)})}{2\gamma}$ and (ii) $SW(\vec{B}^{(\gamma)}) \leq 2\gamma Rev(\vec{p}^{(\gamma)}, \vec{S}^{(\gamma)})$. We assume without loss of generality that the number of items allocated in $\vec{B}^{(t+1)}$ is exactly half the number of copies allocated in the previous benchmark allocation $\vec{B}^{(t)}$ for all $1 \leq t \leq \gamma - 1$. We begin with an outline of the proof.

- (1) Consider some iteration t of our core algorithm where the benchmark solution is $\vec{B}^{(t)}$. As mentioned in the algorithm, let $\vec{S}^{(t)}$ denote the allocation obtained upon running the sequential posted pricing mechanism with prices that are a scaled down version of each good's average value in $\vec{B}^{(t)}$, (as in $Prices(\vec{B}^{(t)}, \gamma)$). Then, we claim and prove that $SW(\vec{B}^{(t)}) - SW(\vec{S}^{(t)})$ can be bounded in terms of the mechanism's revenue and the welfare of $\vec{B}^{(t)}$, analogous to Definition 2.1. This has obvious implications in showing that our solutions satisfy Definition 2.1 for a given value of t as long as the core algorithm chooses $\vec{B}^{(t+1)}$ to be the output of the sequential mechanism.
- (2) Given a benchmark allocation vector $\vec{B}^{(t)}$, once again consider running the sequential mechanism with scaled down prices, and let $\vec{T}$ denote a sub-allocation of $\vec{B}^{(t)}$ obtained by allocating the 'unsold' items back to the original buyers (as in $Alloc-Unsold(\vec{B}^{(t)}, \vec{S}^{(t)})$). Then, $SW(\vec{B}^{(t)}) - SW(\vec{T})$ can be bounded purely in terms of the revenue of the above sequential mechanism. Therefore, our solutions satisfy Definition 2.1 in the second case where $\vec{B}^{(t+1)}$ is defined to be the output of $Alloc-Unsold$.
- (3) In the final act, we prove that for each t , $\vec{B}^{(t+1)}$ allocates at most half the total number of items as $\vec{B}^{(t)}$, and so at most two items are allocated in $\vec{B}^{(\gamma)}$. In this case, the tail part of our core algorithm, where exactly one good is sold, results in a revenue that is only a constant factor away from the social welfare of $\vec{B}^{(\gamma)}$.

General Claims. We now show two general claims that can be used to bound the difference in welfare of any two consecutive benchmark solutions in our core algorithm.

CLAIM 3.2. *Given any allocation $\vec{B}$ and some parameter $\alpha \geq 1$, consider the sequential posted pricing mechanism with arbitrary buyer arrival order for supply constraint $q_i = k_i(\vec{B})$, and price $p_i = \frac{U_i(\vec{B})}{\alpha q_i}$ for all $i \in \mathcal{I}$. Suppose that $\vec{S}$ denotes the output allocation of this mechanism. Then, $SW(\vec{B}) - SW(\vec{S}) \leq \alpha Rev(\vec{p}, \vec{S}) + \frac{SW(\vec{B})}{\alpha}$.*

The next claim is useful in showing that Definition 2.1 holds in the (second) case when our core algorithm uses the function $Alloc-Unsold(\vec{B}^{(t)}, \vec{S}^{(t)})$ to select the next benchmark allocation.

CLAIM 3.3. *Given any allocation $\vec{B}$ and some parameter $\alpha \geq 1$, consider the sequential posted pricing mechanism with arbitrary buyer arrival order for supply constraint $q_i = k_i(\vec{B})$, and price $p_i = \frac{U_i(\vec{B})}{\alpha q_i}$ for all $i \in \mathcal{I}$. Suppose that $\vec{S}^M$ denotes the output allocation of this mechanism, and $\vec{S}$ is the allocation returned by $Alloc-Unsold(\vec{B}, \vec{S}^M)$. Then, $SW(\vec{B}) - SW(\vec{S}) \leq \alpha Rev(\vec{p}, \vec{S}^M)$.*

Tail Conditions. Applying Claims 3.2 and 3.3 in the context of our core algorithm, we can prove that the first half of the charging property is obeyed by our solutions. It only remains for us to

show that $SW(\vec{B}^{(\gamma)}) \leq \alpha Rev(\vec{p}^{(\gamma)}, \vec{S}^{(\gamma)})$. We begin by highlighting the ‘item halving’ nature of our algorithm. Note that in the initial allocation to our core algorithm $\vec{A}$ (and hence $\vec{B}^{(1)}$), at most mk copies of goods are allocated since k denotes the average supply. In every successive round, the number of items in the (next) benchmark solution is at least halved. Therefore, after $\log(mk) - 1$ rounds, i.e., in $\vec{B}^{(\gamma)}$, at most 2 items are allocated to the buyers. The next lemma follows from this reasoning.

LEMMA 3.4. *Define $(i^*, j^*) := \arg \max_{i \in I, j \in N} v_j(i)$. The social welfare of the allocation $\vec{B}^{(\gamma)}$ is at most $2v_{j^*}(i^*)$.*

Thus, since the revenue obtained by $(\vec{p}^{(\gamma)}, \vec{S}^{(\gamma)})$ is $v_{j^*}(i^*) - \epsilon$, we know that for a small enough epsilon it must be that $SW(\vec{B}^{(\gamma)}) \leq 2\gamma Rev(\vec{p}^{(\gamma)}, \vec{S}^{(\gamma)})$, as desired.

Wrapping up the Proof. Equipped with the various pieces, we can now complete the proof of our black-box reduction by showing that our solutions obey the conditions outlined in Definition 2.1. Fix the inputs to the core algorithm, i.e., allocation $\vec{A}$, and γ as defined previously. Consider any iteration $t \leq \gamma - 1$ of our algorithm. Let us proceed in two cases depending on the algorithm’s choice of $\vec{B}^{(t+1)}$. In the first case, suppose that our algorithm identifies $\theta(\vec{B}^{(t)}) \geq 2\theta(\vec{S}^{(t)})$ and defines $\vec{B}^{(t+1)} = \vec{S}^{(t)}$. Then, we apply Claim 3.2 with $\vec{B} = \vec{B}^{(t)}$, $\vec{S} = \vec{S}^{(t)} = \vec{B}^{(t+1)}$, and $\alpha = 2\gamma$ and get that: $SW(\vec{B}^{(t)}) - SW(\vec{B}^{(t+1)}) \leq 2\gamma Rev(\vec{p}^{(t)}, \vec{S}^{(t)}) + \frac{SW(\vec{B}^{(t)})}{2\gamma}$.

In the second case, we have that $\theta(\vec{B}^{(t)}) < 2\theta(\vec{S}^{(t)})$ and so, we set $\vec{B}^{(t+1)} = \text{Alloc-Unsold}(\vec{B}^{(t)}, \vec{S}^{(t)})$. For this case, we refer to Claim 3.3 with $\vec{B} = \vec{B}^{(t)}$, $\vec{S}^M = \vec{S}^{(t)}$, and $\alpha = 2\gamma$ to obtain the same result as above. Finally, the last condition that $SW(\vec{B}^{(\gamma)}) \leq 2Rev(\vec{p}^{(\gamma)}, \vec{S}^{(\gamma)})$ comes from Lemma 3.4. Therefore, our solutions satisfy Definition 2.1 and as a consequence, we can apply Claim 2.2 with $\alpha = \beta = 2\gamma$ to complete the proof. $\square$

3.1 Applying the Item Halving Framework to Form Pricing with High Revenue

We are now in a position to leverage Theorem 3.1 as a black-box mechanism to obtain computationally efficient posted price mechanisms with good revenue guarantees for different arrival orders. Specifically, for instances having XoS valuations, we can use the $\frac{e}{e-1}$ -approximation algorithm for the allocation problem [16] to form an initial allocation $\vec{A}$ with high welfare. The item-halving algorithm gives us a pricing and allocation $(\vec{p}, \vec{S})$ whose revenue is a $O((\log m + \log k)^2)$ -approximation to the optimum social welfare. However, $\vec{S}$ is a specific allocation and only arises for some buyer arrival order. By choosing the function Γ appropriately, we can make sure that the revenue resulting from all other arrival orders is no worse than that of $\vec{S}$.

3.1.1 Adversarial Buyer Arrival.

THEOREM 3.5. (Computational Result) *Given a set of arrival orders Π , we can compute in time polynomial in N , m , and $|\Pi|$, prices $\vec{p}$ and supply constraints $\vec{q}$ such that the revenue guaranteed by the sequential posted pricing mechanism with these parameters for all $\pi \in \Pi$ is a $\frac{4e}{e-1}(\log m + \log k)^2$ -approximation to the optimum welfare, and thus to the optimum revenue as well.*

For the case where the buyer arrival order is adversarial with $|\Pi|$ being of polynomial size, Theorem 3.5 immediately gives us an efficient algorithm for forming a sequential pricing mechanism which achieves a $O((\log m + \log k)^2)$ approximation for maximum revenue. For the case where Π consists of all possible arrival orders (truly adversarial), this still gives us an interesting existence result: for every instance, there exist *fixed* prices and supply constraints such that, no matter what

order the buyers arrive, the revenue of the outcome is a $O((\log m + \log k)^2)$ approximation to the maximum possible welfare.

(*Description of Mechanism for Partially Adversarial Case*) Set $\vec{A}$ to be the allocation which yields an $\frac{e}{e-1}$ -approximation to maximum social welfare; this can be obtained from [16]. Set $\Gamma(\vec{p}, \vec{B})$ to be as follows. For each $\pi \in \Pi$, run the sequential posted price mechanism with prices $\vec{p}$ and supply constraints $\vec{q} = k(\vec{B})$, and choose π to be the arrival order that results in the smallest revenue. Set $\Gamma(\vec{p}, \vec{B})$ to be this arrival order π . We can now run the core algorithm with $\vec{A}$, Γ as defined previously and $\gamma = \log m + \log k$. $\square$

3.1.2 Random Arrival Orders. We now consider a more realistic model [23] of buyer arrival that is determined by an arbitrary distribution D^Π over the set of arrival orders. We make the standard assumption that we have access to an oracle from which we can draw samples corresponding to the distribution D^Π . Building on our machinery from the previous sections, we design a mechanism that with high probability achieves a $O((\log m + \log k)^2)$ -approximation for revenue, which is the same asymptotic factor that was obtained for the partially adversarial model.

THEOREM 3.6. *Given oracle access to a distribution D^Π over the buyer arrival orders, we can design in poly-time a sequential posted pricing mechanism $\mathcal{M}$ such that with high probability,*

$$SW(OPT) \leq \frac{8e}{e-1} (\log m + \log k)^2 E_{\pi \sim D^\Pi} [Rev(\mathcal{M}, \pi)],$$

where $Rev(\mathcal{M}, \pi)$ is the revenue achieved by the mechanism $\mathcal{M}$ when the buyer arrival order is given by π .

We remark here that while the revenue achieved by our mechanism is in expectation over the arrival orders, the ‘high probability’ refers to the randomization of the algorithm itself. The core algorithm that achieves the desired approximation guarantee is the same as before; the only difference is in the function Γ , which is shown in Figure 3.

Fig. 3. Function $\Gamma(\vec{p}, \vec{B})$ specifically for random arrival orders

Instead of iterating over a finite set of arrival orders Π as we did previously, the new sequential mechanism function draws a logarithmic (in N, m) number of arrival orders according to the input distribution, and for each of these orders simulates the sequential posted pricing mechanism with the same set of prices and supply constraints. Finally, the function returns the outcome of the posted pricing mechanism (corresponding to some arrival order), which has the worst revenue.

(*Description of Mechanism for Random Arrival Case*) Run the core algorithm with $\vec{A}$ being the allocation obtained via the algorithm in [16], $\Gamma(\vec{p}, \vec{B})$ as defined in Figure 3 and $\gamma = \log m + \log k$.

4 IMPLICATIONS OF THE CHARGING FRAMEWORK

In this section, we return to the charging framework that we defined in Section 2. Previously, we only defined the parts of the framework that were necessary to prove our $O((\log m + \log k)^2)$ -revenue result for XoS valuations. Here, we will show some additional properties associated with this framework and show how it can be leveraged to unify many of the results in the item pricing literature and generate new results for other classes of valuations.

In contrast to the previous section, all of the results from this section are obtained using the simultaneous mechanism, in which all buyers purchase bundles in parallel. Recall that simultaneous mechanisms are harder to design owing to the fact that the prices need to be chosen in a manner that ensures that every buyer receives a utility-maximizing bundle without violating the supply constraints. We reiterate that for any given instance, a valid solution $(\vec{p}, \vec{S})$ of the simultaneous mechanism is also a valid solution for the sequential mechanism.

4.1 Price Doubling and Bicriteria Approximations via Simultaneous Mechanisms

The field of item pricing has been witness to a rich algorithmic literature with near-optimal guarantees for revenue. A few of the notable contributions in this landscape include the logarithmic approximations to revenue for unit-demand [22] and unlimited-supply markets [4], and a $O(\sqrt{m} \log k_{\max})$ -approximation for single-minded valuations [14]. Somewhat surprisingly, we observe that our charging framework is found lurking beneath the hood of all of the above results. Applying the price doubling technique along with our framework, we present a single black-box revenue maximization theorem for a general class of market: *all of the results* mentioned above are obtained as corollaries. Further, we strengthen these results by also extracting welfare guarantees without any loss in the asymptotic factors for revenue. We begin by informally describing the price doubling technique and then show how the charging framework can be leveraged to obtain multi-objective guarantees.

Price Doubling Also known as multiplicative pricing, this technique is based on generating pricing solutions by running the mechanism on a sequence of pricing vectors such that the prices across successive solutions differ by a multiplicative factor (usually two). Following this, we charge the difference in social welfare between the allocations (obtained via the mechanism) to the mechanism's revenue so that we can apply the charging property.

CLAIM 4.1. *Consider a sequence of benchmark allocations $\vec{B}^{(1)}, \vec{B}^{(2)}, \dots, \vec{B}^{(\gamma)}$, and pricing solutions $(\vec{p}^{(t)}, \vec{S}^{(t)})_{t=1}^{\gamma}$ that satisfy the charging property. Further, suppose that for all $1 \leq t \leq \gamma$, $\vec{B}^{(t)} = \vec{S}^{(t)}$. Then, for any target welfare parameter $c \in [1, \gamma\alpha]$, there exists an $1 \leq \ell \leq \gamma$ such that*

$$SW(\vec{S}^{(\ell)}) \geq \frac{SW(\vec{B}^{(1)})}{c} \left(1 - \frac{\gamma - 1}{\beta}\right) \text{ and } Rev(\vec{p}^{(\ell)}, \vec{S}^{(\ell)}) \geq \left(1 - \frac{1}{c}\right) \frac{SW(\vec{B}^{(1)})}{\gamma\alpha} \left(1 - \frac{\gamma - 1}{\beta}\right).$$

Essentially, the claim states that as long as the benchmark solution coincides with the allocation obtained via the pricing mechanism, one can obtain excellent bicriteria (revenue, welfare) guarantees.

Black-Box Reduction for Item Pricing. Our main result in this section involves a reduction from revenue to welfare maximization via a simultaneous mechanism. Informally, the result indicates that in order to design item pricing mechanisms with good revenue, it is sufficient to design item pricing mechanisms that approximate welfare; this is true even for arbitrary combinatorial valuations. We now define some pertinent concepts. We say that a buyer has a (monotone) *single-item valuation* if and only if $v(\{i\}) > 0$ for at most one good $i \in \mathcal{I}$. Secondly, an allocation, and by extension an algorithm, is said to be *locally welfare maximizing* if for any buyer j who is not allocated any good, and any set of goods T that are not *sold out*, $v(T) = 0$. For instance, observe that any mechanism

that outputs a socially optimal solution trivially satisfies this property. The following theorem takes as input a black-box algorithm for a simultaneous item pricing mechanism whose social welfare is an α -approximation to $SW(OPT)$ and creates an item pricing algorithm for revenue with a log factor increase in the approximation factor and a small loss in welfare.

THEOREM 4.2. *Let $\mathcal{V}$ be a class of buyer valuations that includes all single-item valuations. Suppose that we are given black-box access to a locally welfare maximizing algorithm Alg , that provides an α -approximate pricing solution for the problem of designing a simultaneous mechanism for welfare maximization for instances where buyer valuations belong to $\mathcal{V}$. Then, we can efficiently compute prices such that the simultaneous mechanism with these prices provides a $O(\alpha(\log(\alpha) + \log(m) + \log(k)))$ -approximation to revenue and a (3α) -approximation to social welfare.*

We remark that the logarithmic gap between welfare and revenue is tight even when all buyers have single item valuations [18]. Further, even though the above result applies to XoS valuations, it does not replace Theorem 3.5 since simultaneous mechanisms cannot achieve good approximations to welfare for this class [19]. Hence, we need sequential mechanisms and the item halving technique to obtain good revenue for XoS buyers.

Black-box Reduction: For a fixed instance G and ‘reserve price’ $r \geq 0$, suppose that $(\vec{p}(r), \vec{S}(r))$ denotes the outcome of the item pricing mechanism Alg for this instance augmented by a set of $m(k+1)$ dummy buyers. Specifically for each good $i \in I$, there are $k_i + 1$ dummy buyers all of whom have single item valuations with $v(\{i\}) = r$. Let $Rev'(\vec{p}(r), \vec{S}(r))$ be the revenue obtained by the mechanism only from the original set of buyers. Now, we define our black-box reduction for $\gamma = 1 + \log(\alpha mk)$, and $SW^0 = SW(Alg(G))$, which is the social welfare of the allocation returned by Alg for the original instance G .

- (1) For $t = 1$ to $t = \gamma$
- (2) Set reserve price $r_t = 2^{t-1} \frac{SW^0}{2mk}$.
- (3) Return the smallest index $1 \leq \ell \leq \gamma$ at which $Rev'(\vec{p}^{(\ell)}, \vec{S}^{(\ell)}) \geq \frac{SW(\vec{S}^{(1)})}{6\gamma}$.

It suffices to show that the sequence of solutions $(\vec{p}^{(t)}, \vec{S}^{(t)})_{t=1}^\gamma$ satisfy Claim 4.1.

Unifying Previous Results in Item Pricing. The next theorem of this paper is a collection of bicriteria approximation algorithms that use simultaneous item pricing mechanisms based on Theorem 4.2 to maximize both revenue and welfare for a number of settings. In the full version, we show how the theorem can be applied to derive bicriteria bounds even for pricing mechanisms that do not fall under the purview of item pricing, for e.g., bundle pricing [20].

THEOREM 4.3. *We can compute the following {revenue, welfare}-approximations in poly-time*

- (1) A $\{O(\log m + \log k), 3\}$ -approximation when buyers have gross substitutes valuation.
- (2) A $\{O(\sqrt{m} \log(k_{max})), O(\sqrt{m})\}$ -approximation for single-minded valuations.
- (3) A $\{O(\log(N) + \log m), 3\}$ -approximation for arbitrary combinatorial valuations (that includes single-item valuations) and unlimited supply.

Here k_{max} refers to the maximum available supply of any good.

Recall that any c -approximation algorithm for revenue obtained in terms of the optimum social welfare is trivially a (c, c) -bicriteria approximation. What makes Theorem 4.3 appealing is that the bound for welfare is much better than c , i.e., there is no asymptotic loss in the approximation factors for both revenue and welfare from the state of the art.

4.2 Simultaneous Mechanism: Item Halving applied to Multi-unit Markets

Previously, we used the item halving technique to derive a sequential mechanism with polylogarithmic revenue guarantees for fractionally subadditive buyers. We now highlight the versatility of this technique by showing that it leads to good approximations even when buyer valuations may not belong to the XoS class. Specifically, we consider the popular multi-unit market setting [17] where buyers only care about the number of items they receive, and show that a small variant of our previously developed algorithm yields a $O(\log m)$ -approximation for revenue. Previously [4] provided a sequential mechanism with a $O(\log m + \log(N))$ guarantee for revenue; not only do we remove the dependence on N , but we provide a stronger (simultaneous) mechanism with the given approximation factor.

A multi-unit market consists of a single good and each buyer's valuation depends only on the number of copies of this good that she receives. For the purpose of notational consistency, we assume that the market consists of a set $\mathcal{I}$ of m goods with one copy each and every buyer's valuation depends only on the cardinality of the set of goods, i.e., for any $j \in \mathcal{N}$, and any $S, T \subseteq \mathcal{I}$, $v_j(T) = v_j(S)$ if $|S| = |T|$; thus we will sometimes overload notation and refer to $v_j(S)$ for $|S| = q$ as $v_j(q)$. Note that such valuations can exhibit complementarities, and therefore do not fall under the XoS class. Additionally, we make the fairly standard assumption (e.g., see [4]) that no single buyer wishes to buy out more than half of the entire supply of goods in the market.

Definition 4.4. We say that an instance with multi-unit valuations satisfies the *no overwhelming buyer assumption* if for every buyer j , $v_j(\frac{m}{2} + q) = v_j(\frac{m}{2})$ for any non-negative integer q .

THEOREM 4.5. *Given any instance of a multi-unit market that satisfies the no overwhelming buyer assumption and an input allocation $\vec{A}$, we can compute in poly-time a single price $\tilde{p}$ such that the simultaneous mechanism with price $\tilde{p}$ on all copies of the good achieves a $O(\log m)$ -revenue approximation to the optimum social welfare (and hence, optimum revenue) as long as $m \geq 2$.*

5 CONCLUSION AND OPEN QUESTIONS

The contributions made by this work can be evaluated at two levels. The fundamental contribution is a general framework for revenue maximization that captures the essence of a number of revenue maximization algorithms, previously developed in an ad-hoc fashion for specific settings and transforms all of these results into bicriteria approximation algorithms with only a constant factor loss in revenue. Our second contribution is more technical: we present the first known algorithm with a log-squared approximation factor for revenue with XoS buyers. In doing so, we partially resolve a decade long open question [4, 10] about the existence of polylogarithmic approximations for revenue maximization based on static item pricing. Our result is surprisingly powerful: it implies that even for quite general combinatorial valuations, we can use arguably the simplest pricing scheme known to us (one number per good) and approximate an all-encompassing benchmark, that is the optimum welfare.

An immediate open question is whether our central result for XoS buyers can be generalized to obtain a bicriteria bound. We make some progress towards answering this question in the full version. A more fundamental question is whether our results extend to a Bayesian information setting, where the seller is only aware of the distribution from which buyers' valuations are drawn. Unfortunately, it is known that no posted price obtains a reasonable revenue in comparison to the optimum welfare [1] in Bayesian environments even for unit-demand. One way to circumvent this would be to develop better benchmarks as in [8]; alternatively, we could assume some natural properties on the distributions such as monotone hazard rate.

REFERENCES

- [1] Saeed Alaei, Jason D. Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. 2015. Optimal Auctions vs. Anonymous Pricing. In *Proceedings of FOCS 2015*.
- [2] Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. 2016. Pricing to Maximize Revenue and Welfare Simultaneously in Large Markets. In *Proceedings of WINE 2016*.
- [3] Maria-Florina Balcan and Avrim Blum. 2007. Approximation Algorithms and Online Mechanisms for Item Pricing. *Theory of Computing* 3, 1 (2007), 179–195. <https://doi.org/10.4086/toc.2007.v003a009>
- [4] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. 2008. Item pricing for revenue maximization. In *Proceedings of ACM EC 2008*.
- [5] Avrim Blum, Anupam Gupta, Yishay Mansour, and Ankit Sharma. 2011. Welfare and Profit Maximization with Production Costs. In *Proceedings of FOCS 2011*. 77–86.
- [6] Simina Brânzei, Aris Filos-Ratsikas, Peter Bro Miltersen, and Yulong Zeng. 2016. Envy-Free Pricing in Multi-unit Markets. *CoRR* abs/1602.08719 (2016). <http://arxiv.org/abs/1602.08719>
- [7] Patrick Briest and Piotr Krysta. 2011. Buying Cheap Is Expensive: Approximability of Combinatorial Pricing Problems. *SIAM J. Comput.* 40, 6 (2011), 1554–1586.
- [8] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. 2016. A duality based unified approach to Bayesian mechanism design. In *Proceedings of the ACM STOC 2016*.
- [9] Yang Cai and Mingfei Zhao. 2017. Simple Mechanisms for Subadditive Buyers via Duality. In *Proceedings of STOC 2017*.
- [10] Tanmoy Chakraborty, Zhiyi Huang, and Sanjeev Khanna. 2013. Dynamic and Nonuniform Pricing Strategies for Revenue Maximization. *SIAM J. Comput.* 42, 6 (2013), 2424–2451.
- [11] Parinya Chalermsook, Julia Chuzhoy, Sampath Kannan, and Sanjeev Khanna. 2012. Improved Hardness Results for Profit Maximization Pricing Problems with Unlimited Supply. In *Proceedings of APPROX-RANDOM 2012*.
- [12] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of STOC 2010*. 311–320.
- [13] Ning Chen and Xiaotie Deng. 2014. Envy-free pricing in multi-item markets. *ACM Trans. Algorithms* 10, 2 (2014), 7.
- [14] Maurice Cheung and Chaitanya Swamy. 2008. Approximation Algorithms for Single-minded Envy-free Profit-maximization Problems with Limited Supply. In *Proceedings of FOCS 2008*. 35–44.
- [15] Ilias Diakonikolas, Christos H. Papadimitriou, George Pierrakos, and Yaron Singer. 2012. Efficiency-Revenue Trade-Offs in Auctions. In *Proceedings of ICALP 2012*. 488–499.
- [16] Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2010. Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders. *Math. Oper. Res.* 35, 1 (2010), 1–13. <https://doi.org/10.1287/moor.1090.0436>
- [17] Michal Feldman, Amos Fiat, Stefano Leonardi, and Piotr Sankowski. 2012. Revenue maximizing envy-free multi-unit auctions with budgets. In *Proceedings of ACM EC 2012*.
- [18] Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. Combinatorial Auctions via Posted Prices. In *Proceedings of SODA 2015*. 123–135.
- [19] Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. On Welfare Approximation and Stable Pricing. *arXiv* abs/1511.02399 (2015). <http://arxiv.org/abs/1511.02399>
- [20] Michal Feldman, Nick Gravin, and Brendan Lucier. 2016. Combinatorial Walrasian Equilibrium. *SIAM J. Comput.* 45, 1 (2016), 29–48. <https://doi.org/10.1137/13094339X>
- [21] Amos Fiat and Amiram Wingarten. 2009. Envy, Multi Envy, and Revenue Maximization. In *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*. 498–504.
- [22] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. 2005. On profit-maximizing envy-free pricing. In *Proceedings of SODA 2005*.
- [23] Thomas Kesselheim, Robert D. Kleinberg, and Rad Niazadeh. 2015. Secretary Problems with Non-Uniform Arrival Order. In *Proceedings of STOC 2015*. 879–888. <https://doi.org/10.1145/2746539.2746602>
- [24] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic game theory*. Vol. 1. Cambridge University Press Cambridge.
- [25] Tim Roughgarden and Inbal Talgam-Cohen. 2015. Why Prices Need Algorithms. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 19–36.
- [26] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. 2012. Supply-limiting mechanisms. In *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*. 844–861.
- [27] Asher Wolinsky. 1988. Dynamic markets with competitive bidding. *The Review of Economic Studies* 55, 1 (1988), 71–84.