Unit-Demand Pricing for Min-Buying Customers

Titing Cui, Michael L. Hamilton, Jourdain Lamperski, Kasra Ganjizadeh Tari
University of Pittsburgh, Pittsburgh, PA 15213,
tic54@pitt.edu, mih140@pitt.edu, lamperski@pitt.edu, kag339@pitt.edu

We study the unit-demand pricing problem under *min-buying* behavior, where customers purchase the lowest-priced item that yields non-negative utility. In contrast to the quasi-linear model, we show that uniform pricing becomes substantially more powerful under min-buying. Specifically, in the well-studied case of unit-demand pricing for quasi-linear utility-maximizing customers, two classical lines of work show that: (i) when valuations are IID for n items, an optimally chosen uniform price is a two-approximation of the optimal discriminatory (non-uniform) pricing (Chawla et al. 2007, 2010), and (ii) when valuations are independent but not identical, and satisfy a regularity condition, a uniform price is a ≈ 2.62 approximation of the optimal differentiated pricing (Alaei et al. 2019, Jin et al. 2019). In this paper, we consider the analogous problems and their extensions for min-buying customers.

First, we study symmetric environments. We show that for min-buying customers with IID valuations, a single uniform price across all items is optimal. We then extend this optimality result to a broad class of valuations drawn from exchangeable and affiliated distributions. However, uniform pricing is not always optimal in symmetric environments: for min-buying customers we construct an exchangeable valuation distribution for which the gap between uniform pricing and optimal item pricing is logarithmic in the number of items which is tight and matches the worst-case gap in the quasi-linear setting. Next, we consider asymmetric environments in which valuations are independent but not identically distributed. When these valuations are regular and customers are min-buying, we show that a uniform price achieves $\approx 69.8\%$ of the optimal item-pricing revenue and this is tight. Overall, our results highlight that the value of non-uniform pricing substantially diminishes, and often disappears entirely, under more conservative demand models in which customers engage in satisficing behaviors such as min-buying.

Key words: unit-demand pricing, min-buying, algorithmic pricing, price discrimination
History: First Version: March 14, 2025. This Version: August 31, 2026

1. Introduction

In a unit-demand pricing problem (UDP), the task is to find prices for n items that maximize a seller's expected revenue given some distribution over customer valuations. Once prices are set, each unit-demand customer draws a vector of valuations, and purchases the item they prefer most (or nothing) according to some decision rule. The most commonly studied decision rule is when customers are *quasi-linear utility-maximizers* who choose the item that maximizes the difference between their valuation and the price, or no item when all differences are negative. The next most

common decision rule is called *min-buying*, as popularized by Rusmevichientong et al. (2006). Under min-buying, customers purchase the lowest priced item that gives them non-negative utility.

Both quasi-linear utility-maximization and min-buying are useful modeling choices in different settings. While quasi-linear utility-maximization is an intuitive and natural way to model rational customer behavior, it requires assumptions about how customers make choices that do not hold in all settings. Namely, it requires that customers translate their preferences into a common numerical scale and then choose based on utility comparisons between all pairs of items. Research in behavioral economics provides many examples where, due to limited attention, cognitive costs, and/or other biases, customers fail to maximize their quasi-linear utility in this way and instead are better modeled as following simple heuristics to make purchasing decisions (for a survey of results in this area, see Kahneman and Thaler (2006)). Inspired by such settings, when customers are not modeled as fully rational, a common alternative is to suppose that they engage in *satisficing* behavior. Under this behavior, customers screen items by whether they are "good enough" (i.e., yield non-negative utility) and then pick the cheapest acceptable option (Kahneman 2003). This is equivalent to following a min-buying rule.

Beyond concerns of rationality, min-buying rules are also commonly used to model how customers make decisions under incomplete information. For instance, in settings with opaque products (i.e., products where some attribute of the good, such as color or style, is hidden until after purchase), a common assumption is to model customers as *risk-averse*, meaning that they value the opaque product according to the minimum-valued possible allocation (Elmachtoub and Hamilton 2021, Chen et al. 2025, Housni et al. 2025). UDP problems with min-buying customers are also structurally identical to so-called oblivious sequential posted price (OPP) problems where arriving customers are shown a sequence of prices in ascending order (Chawla et al. 2010), as well as certain ranked assortment optimization problems (Aouad et al. 2018, Berbeglia and Joret 2020).

For either choice of decision rule, UDPs of the above form are foundational models for understanding the value of *price discrimination*, which concerns setting different prices for horizontally differentiated items, and *price differentiation*, which concerns setting prices for vertically differentiated items. Much recent work on UDPs has focused on quantifying the value of price discrimination or differentiation relative to the strategy of offering a single uniform price for all items. For quasi-linear customers, however, the story is mixed. On the one hand, in the natural setting of IID valuations and quasi-linear utility-maximizing customers, a uniform price need not be optimal (Thanassoulis 2004), and optimal discriminatory pricing can earn up to twice as much revenue (Chawla et al. 2010, Dütting et al. 2016). (Note pricing in this setting is referred to as *discriminatory* since, when valuations are IID, every permutation of a price vector yields identical revenue, thus the assignment of relatively high and low prices to particular items is arbitrary). On the other

hand, a complementary line of work shows that uniform pricing can be remarkably effective even in asymmetric (i.e., vertically differentiated) settings. For quasi-linear utility-maximizing customers with independent but non-identical item valuations satisfying a regularity condition, an optimally chosen uniform price always earns at least 38% (i.e., approximately 1/2.62) of the optimal non-uniform pricing revenue (Jin et al. 2019). This guarantee is surprisingly strong given how different the item distributions may be, potentially reflecting stark differences in item quality, and highlights the general effectiveness of uniform pricing.

Viewed through this lens, much of the literature on UDPs can be understood as quantifying when non-uniform prices are necessary. Because uniform prices are fair, easy to implement, and easy to optimize, they are attractive to sellers so long as restricting attention to a uniform price does not result in too much revenue loss. For quasi-linear utility maximizers, the performance of uniform pricing relative to optimal non-uniform pricing is well understood: prior work provides tight or nearly tight bounds in canonical regimes, including IID valuations, independent regular but non-identical valuations, and general worst-case distributions. However, the practical and managerial implications of these results are less clear. Further, as these results are closely tied to the assumption that customers maximize quasi-linear utility which may not hold in many practical settings, it is then natural to ask how these classical conclusions change when customers are min-buying rather than quasi-linear utility maximizers. In this paper, we answer this question by proving analogues of these foundational results under min-buying behavior. In contrast to the mixed results in the quasi-linear case, we find that uniform pricing is optimal in many symmetric environments and, more generally, achieves substantially stronger guarantees.

1.1. Our Contributions

A summary of our main contributions is as follows:

1. First, we study the value of discriminatory pricing in symmetric valuation environments when customers are min-buying. We consider three classes of symmetric valuations: IID distributions, affiliated and exchangeable distributions, and general exchangeable distributions. We show that when customers are min-buying and valuations are IID, a uniform price is always optimal (cf. Theorem 1) and thus discriminatory pricing has no additional power in this setting. We then extend this result to the more general class of affiliated and exchangeable distributions and show uniform pricing continues to be optimal (cf. Theorem 2). Finally, we complement these positive results with a lower-bounding construction for general exchangeable valuations that demonstrates, both in the min-buying and quasi-linear settings, that discriminatory pricing can earn logarithmically more revenue than uniform pricing (cf. Theorem 3) matching the known upper-bound.

2. Second, we study the value of differentiated pricing when valuations are not drawn symmetrically across items. Our main result in this setting is to show that for min-buying customers, when valuations are drawn from independent and (not necessarily identical) regular distributions, a uniform price is guaranteed to earn $\approx 69.8\%$ of the revenue of the optimal non-uniform pricing and this is tight (cf. Theorem 4). This resulting approximation substantially improves over the $\approx 38\%$ that can be guaranteed in the quasi-linear case, further emphasizing the differences between these two models.

Overall, our work highlights how the unit-demand pricing problem changes when customers are min-buyers versus quasi-linear utility-maximizers. We situate our results within the broader literature on unit-demand pricing and show that, both for price discrimination among horizontally differentiated items and for price differentiation across (potentially) vertically quality-differentiated items, the gains from non-uniform pricing depend critically on the customer decision rule, and are markedly smaller under min-buying (often vanishing altogether). En route to our results, we also develop techniques for how to analyze pricing in general symmetric settings and how to exploit the structure of min-buying behavior which may be of independent interest.

1.2. Literature Review

The literature on unit-demand pricing is vast and interdisciplinary, with contributions from economics, operations research/management, marketing, and computer science. In this subsection, we describe some of the most relevant streams of work on min-buying rules and various unit-demand pricing problems, and connect these results with our findings.

Models of Min-Buying Min-buying models of customer behavior have a long history, going back at least to Simon (1955), who highlights the potential value of alternative models of customer decision-making, including satisficing rules based on simple min/max criteria. More recently, behavioral evidence that quasi-linear maximization can, in some settings, provide a poor explanation of consumer choice has spurred renewed interest in min-buying models (Gigerenzer and Goldstein 1996, Kahneman 2003). These min-buying-style models are also supported by empirical research in several contexts. For example, Stüttgen et al. (2012) use eye-tracking data to show that many consumers engage in satisficing behavior (rather than comparing across all available items) and Karle et al. (2015) provide theoretical and empirical evidence that customers may exhibit loss aversion and gravitate toward lower-priced items more strongly than would be expected under a fully rational, quasi-linear utility-maximizing model.

Building on their use as a model of bounded rationality, a growing body of work studies pricing and revenue-management problems under min-buying-type assumptions (e.g., Goyal et al. (2016), Cao et al. (2022), Bhardwaj et al. (2025)). Min-buying rules are also commonly used to model

customer decisions under incomplete information, such as arises in settings with *opaque products*, where information about the good is hidden until after purchase. In models of opaque selling, it is common to assume that customers are *risk-averse*, meaning that they value an opaque product according to the minimum-valued possible allocation (Elmachtoub and Hamilton 2021, Chen et al. 2025, Housni et al. 2025). Furthermore, unit-demand pricing under min-buying is closely connected to certain assortment-optimization models (Aouad et al. 2018). In particular, Berbeglia and Joret (2020) establish an exact correspondence between min-buying UDPs and revenue-ordered assortments, allowing our results to translate directly into new guarantees for specific instances of their model.

UDPs in Symmetric Settings There is a deep body of work studying the value of price discrimination in symmetric settings. When valuations are drawn IID and customers are quasi-linear utility maximizers, a stream of papers (Thanassoulis 2004, Chawla et al. 2007, 2010, Dütting et al. 2016) shows that non-uniform discriminatory prices can strictly outperform uniform pricing. For n items, the ratio between the revenue of the optimal item pricing and that of the optimal uniform pricing can be as large as $2 - \frac{1}{n}$, and this bound is tight. Beyond their direct implications for pricing, these results have also been influential in the algorithmic game theory literature because of their close connections to prophet inequalities, which have spurred substantial follow-up work; see, for example, Lucier (2017), Correa et al. (2019).

Beyond IID valuations, prior work has studied other symmetric settings in which valuations are correlated across items and exhibit some form of positive dependence (Anunrojwong et al. 2022), as well as the general setting in which valuations are exchangeable. In particular, for general exchangeable valuations and quasi-linear utility-maximizing customers, Elmachtoub and Hamilton (2021) show that discriminatory pricing can earn at most a logarithmic factor more revenue than uniform pricing and ask whether this bound is tight. We answer this question in the affirmative by constructing an exchangeable instance in which, under min-buying, discriminatory pricing earns a $\Theta(\log n)$ factor more revenue than uniform pricing, where n is the number of items. This lower bound also applies to quasi-linear customers as, for any fixed price vector, revenue is weakly higher under quasi-linear utility maximization than under min-buying, whereas uniform-pricing revenue is identical under the two decision rules. Conversely, every upper bound for quasi-linear customers applies to min-buying UDPs. Together, these observations establish that the logarithmic gap is tight up to constant factors under both decision rules.

In parallel with work characterizing approximation guarantees for uniform pricing, a stream of papers studies the algorithmic problem of solving UDPs, sometimes referred to as *item pricing*, in symmetric settings. For quasi-linear utility-maximizing customers with IID valuations, Chen

et al. (2014, 2018) show that finding the optimal item pricing is NP-hard even when valuations are restricted to three support points. On the positive side, when valuations are IID and bounded, Cai and Daskalakis (2015) provide an additive PTAS, while Bowers et al. (2025) give a multiplicative PTAS. In this paper, we study the analogous problem for min-buying customers and show that uniform pricing is always optimal when valuations are IID. This result simultaneously characterizes the performance of uniform pricing and reduces the corresponding item-pricing problem to a one-dimensional optimization.

UDPs Beyond Symmetric Settings Outside symmetric settings, there is also a vast body of work studying the value of price differentiation. The most relevant stream is that of Alaei et al. (2019), Jin et al. (2019, 2020), which show that when customers are quasi-linear utility maximizers and valuations are drawn independently but non-identically from regular distributions, the revenue of an optimal uniform price is always within a factor of approximately 2.62 of the optimal differentiated-pricing revenue, and that this factor is tight. In our work, we study the analogous problem for min-buying customers and characterize the exact worst-case approximation factor, showing that an optimal uniform price is within a factor of approximately 1.43 of the optimal differentiated pricing. Equivalently, our result characterizes the tight revenue gap between anonymous pricing and optimal order-oblivious posted pricing for independent regular valuations (and alternatively, the result of Jin et al. (2019) can be cast as giving the tight approximation factor between differentiated pricing and uniform pricing when customers are *max-buying*, meaning they purchase highest priced item for which they obtain non-negative utility). On the algorithmic side, there is a significant literature on solving or approximating the optimal item-pricing problem when valuations are non-identical and customers are quasi-linear utility maximizers; see Aggarwal et al. (2004), Guruswami et al. (2005), Balcan et al. (2008), Gallego and Berbeglia (2021). Several of these techniques and guarantees also extend to related min-buying formulations. However, relative to this literature, our focus is on establishing tight constant-factor approximation guarantees for the simpler policy class of uniform pricing.

1.3. Paper Outline

The remainder of the paper is organized as follows. In Section 2, we describe the model and give some preliminary results about unit-demand pricing for min-buying customers. In Section 3, we study the performance gap between uniform pricing and discriminatory pricing when valuations are drawn symmetrically. In Section 4, we consider when items are valued asymmetrically and bound the gap between uniform pricing and differentiated pricing. Finally, in Section 5 we conclude and describe avenues for future work. All omitted examples and proofs referenced in the main body can be found in Sections A and B in the Appendix.

2. Model and Preliminaries

We consider a seller offering n items, described by the index set $[n] := \{1, 2, \ldots, n\}$, to customers with non-negative, random valuations $\mathbf{V} = (V_1, \ldots, V_n)$. For each item i, we will denote the marginal valuation distribution by $V_i \sim \mathrm{F}_i$. We suppose each marginal distribution F_i admits a continuous density f_i[1]. We denote the survival function (cCDF) of a distribution F_i by $\overline{\mathrm{F}}_i(x) := 1 - \mathrm{F}_i(x)$. Assuming knowledge of the valuation distribution, the seller sets prices, $\mathbf{p} = (p_1, p_2, \ldots, p_n)$ for the n items, where throughout we will assume the items are indexed so that $p_1 \leq p_2 \leq \ldots \leq p_n$.

Facing item prices $\mathbf{p}$, we suppose customers are unit-demand and min-buying, and thus purchase the minimum priced item such that their valuation exceeds the price, or nothing if all prices exceed their valuations. Since items are indexed to be in ascending price order we can equivalently say a customer with valuation vector $\mathbf{v} = (v_1, v_2, \ldots, v_n)$ purchases the item $i^* := \min\{i | v_i \geq p_i\}$, or no item if $v_i < p_i$ for all $i \in [n]$. We will denote the expected revenue of a seller using prices $\mathbf{p}$ by $\mathcal{R}_{DP}(\mathbf{p})$, and let $\mathcal{R}_{DP} := \max_{\mathbf{p}} \mathcal{R}_{DP}(\mathbf{p})$ with the argument suppressed to denote the maximum achievable expected revenue. When $\mathbf{p}$ is such that the prices for the items are all the same, i.e., $\mathbf{p} = p \cdot \mathbf{e}_n$ for some $p \in \mathbb{R}$ and where $\mathbf{e}_n := (1, 1, \ldots, 1) \in \mathbb{R}^n$ is the vector of 1's, we will refer to this as a **u**niform **p**ricing and denote the revenue by $\mathcal{R}_{UP}(p)$. When referring to the expected revenue of the optimal uniform pricing we again omit the argument and denote it by $\mathcal{R}_{UP} := \max_p \mathcal{R}_{UP}(p)$.

In this paper, we will compare $\mathcal{R}_{DP}$ and $\mathcal{R}_{UP}$ for several important classes of valuation distributions. In the remainder of this section, we will describe some preliminary results about these revenue functions that will aid us in our subsequent analysis.

2.1. Preliminaries

In this subsection, we will describe some preliminary results related to the structure of unit-demand pricing for min-buying customers. First, note that when valuations are independent, the revenue of a uniform price p can be expressed simply as:

$$\mathcal{R}_{UP}(p) = p\left(1 - \prod_i \mathrm{F}_i(p)\right), \tag{1}$$

which is the price p multiplied by the probability that at least one valuation exceeds that price. The revenue of uniform pricing is the same regardless of whether customers are min-buying, quasi-linear utility maximizing, or max-buying.

The min-buying behavior becomes relevant for non-uniform price vectors. Specially, for the some prices $\mathbf{p}$, and when items are re-indexed so that $p_i \leq p_{i+1}$ for all $i \in [n]$, the min-buying assumption allows us to express the revenue generated by $\mathbf{p}$ in a simple product/sum form:

$$\mathcal{R}_{DP}(\mathbf{p}) := \sum_{i=1}^{n} p_i \Pr\left(i = \arg\min_{j \in [n]} \{j : V_j \geq p_j\}\right) = \sum_{i=1}^{n} p_i \Pr(V_i \geq p_i \text{ and } V_j < p_j,\ \forall j \in [i-1])$$

$$= \sum_{i=1}^{n} p_i \Pr(V_i \geq p_i) \prod_{j=1}^{i-1} \Pr(V_j < p_j)$$
$$= \sum_{i=1}^{n} p_i \overline{\mathrm{F}}_i(p_i) \prod_{j=1}^{i-1} \mathrm{F}_j(p_j). \tag{2}$$

Often we will reference how a pricing policy operates over a restricted subset of items $S \subset [n]$. We denote the sub-vector of a price vector $\mathbf{p}$ over an index set S by the subscript $\mathbf{p}_S$, i.e., $\mathbf{p}_S = (p_j)_{j \in S}$, and denote the revenue of a price vector over the restricted set S by the superscript as:

$$\mathcal{R}_{DP}^{S}(\mathbf{p}_S) := \sum_{i \in S} p_i \Pr\left(i = \arg\min_{j \in S}\{j : V_j \geq p_j\}\right). \tag{3}$$

For unit-demand pricing with min-buying customers, there is a natural recursive formula for the expected revenue of an item pricing relative to a restricted set of the items. Specifically, for any $k \in [n]$, let $S_k := \{1, \ldots, k\}$ be the set of the first k items, and define $\bar{S}_k := [n] \setminus S_k$ to be the complement set of the last $n-k$ items. Then for any k, Eq. (2) can alternatively be expressed over S_k and $\overline{S}_k$ as:

$$\mathcal{R}_{DP}(\mathbf{p}) = \sum_{i=1}^{k} p_i \overline{\mathrm{F}}_i(p_i) \prod_{j=1}^{i-1} \mathrm{F}_j(p_j) + \sum_{i=k+1}^{n} p_i \overline{\mathrm{F}}_i(p_i) \prod_{j=1}^{i-1} \mathrm{F}_j(p_j)$$
$$= \mathcal{R}_{DP}^{S_k}(\mathbf{p}_{S_k}) + \left(\sum_{i=k+1}^{n} p_i \overline{\mathrm{F}}_i(p_i) \prod_{j=k+1}^{i-1} \mathrm{F}_j(p_j) \right) \prod_{j=1}^{k} \mathrm{F}_j(p_j)$$
$$= \mathcal{R}_{DP}^{S_k}(\mathbf{p}_{S_k}) + \mathcal{R}_{DP}^{\bar{S}_k}(\mathbf{p}_{\overline{S}_k}) \prod_{j=1}^{k} \mathrm{F}_j(p_j). \tag{4}$$

Eq. (4) captures a form of separability in revenue between the k lowest priced items and the remaining $n-k$. If apply this property to S_1 we have:

$$\mathcal{R}_{DP}(\mathbf{p}) = \mathcal{R}_{DP}^{S_1}(p_{S_1}) + \mathcal{R}_{DP}^{\bar{S}_1}(\mathbf{p}_{\overline{S}_1}) \mathrm{F}(p_{S_1}), \tag{5}$$

which neatly separates the n item pricing problem for min-buying customers into a single item pricing problem in the first term (the term $\mathcal{R}_{DP}^{S_1}(p_{S_1})$ equals $p_{S_1} \overline{\mathrm{F}}_1(p_{S_1})$ which is the revenue of the scalar price p_{S_1} in a standard monopoly pricing problem), and an $n-1$ item pricing problem in the second term. The two terms are tied together only by the constraint that the prices are ascending. Specifically, from Eq. (5) we observe that if $\mathbf{p}^*$ are the optimal item prices, then $\mathbf{p}^*_{\overline{S}_1}$ must be the optimal item pricing over $\overline{S}_1$ subject to the constraint that the prices are increasing, i.e., that $\mathbf{p}_{\overline{S}_1} \geq p_{S_1} \mathbf{e}_{n-1}$ component-wise.

This is a key way in which the revenue function of item pricing for min-buying customers differs from item pricing for quasi-linear utility maximizers. In the quasi-linear case, items "compete"

through their net utilities $v_j - p_j$, so the demand for any one item depends on the full price vector. In contrast, min-buying induces a strong form of cannibalization where lower-priced items screen customers away from higher-priced items. Armed with this insight, we are now ready to study UDP for min-buying customers when valuations are symmetric.

3. Uniform Pricing vs. Discriminatory Pricing

In this section, we will study unit-demand pricing for min-buying customers when valuations are drawn symmetrically. This setting models discriminatory pricing where a seller offers a substitutable set products of similar quality, for instance various colors of the same item, at different prices. As mentioned in the introduction, a surprising and influential line of work has shown that, even under symmetric valuations, discriminatory pricing can substantially outperform uniform pricing when customers are quasi-linear utility maximizers. Here, we consider when this advantage persists for min-buying customers by studying the performance of uniform pricing in three increasingly general symmetric environments: IID valuations, exchangeable and affiliated valuations, and general exchangeable valuations.

3.1. Uniform Pricing Guarantees for IID Valuations

The case of IID valuations is the most commonly studied symmetric setting of the unit-demand pricing problem. For quasi-linear utility-maximizing customers with IID valuations, Chawla et al. (2007), Dütting et al. (2016) show that discriminatory pricing can earn up to twice as much revenue as uniform pricing, and that this factor is tight as the number of items grows large. Our first result shows that this phenomenon disappears for min-buying customers: when valuations are IID, discriminatory pricing offers no advantage over uniform pricing.

THEOREM 1 **(IID Valuations)**. *Suppose there are n items, customers are min-buying, and they draw valuations IID from some positive-valued distribution $V_i \sim F$ for $i \in [n]$. Then a uniform price is revenue-maximizing, i.e., $\mathcal{R}_{DP} = \mathcal{R}_{UP}$.*

Theorem 1 shows that, for min-buying customers and IID valuations, the natural intuition that symmetrically valued products should be priced symmetrically holds true. Theorem 1 also implies that the algorithmic problem of setting prices in this setting is straightforward. Whereas, in the case of quasi-linear utility-maximizing customers, the unit-demand pricing problem is NP-hard to solve even when valuations are IID (Chen et al. 2018), when customers are min-buying the corresponding algorithmic problem simplifies to a one-dimensional optimization to determine the optimal uniform price.

Before we formally prove Theorem 1, to build intuition it is helpful to consider how non-uniform prices can capture revenue in symmetric settings. Specifically, a non-uniform price vector creates

high- and low-price "tiers". The high-price tiers can increase revenue in cases where a customer both has strong preferences and those preferences happen to be for high-priced items. This benefit is offset by a cannibalization effect where the low-priced items can draw demand away from higher-priced items. For min-buying customers, this cannibalization effect is much stronger than in the quasi-linear case. A min-buying customer who finds any lower-priced item acceptable (i.e. would get non-negative utility from purchase) never purchases a higher-priced acceptable item, regardless of how much more they value the higher-priced item. Thus the revenue lift for discriminatory pricing in the quasi-linear case is driven by customers who find multiple items acceptable and nevertheless opt for the more expensive of the acceptable items. We formalize this intuition next.

Proof of Theorem 1. We will prove that revenue is maximized by uniform pricing by induction on the number of items, n. In fact, we will prove that stronger statement that, for every price vector $\mathbf{p} = (p_1, \dots, p_n)$, there exists a price $\hat{p} \in \{p_1, \dots, p_n\}$ such that $\mathcal{R}_{UP}(\hat{p}) \geq \mathcal{R}_{DP}(p_1, \dots, p_n)$. The base case of $n = 1$ follows immediately, so let $n \geq 2$ and suppose the claim holds for $n-1$ items. Fix an arbitrary price vector $\mathbf{p} = (p_1, \dots, p_n)$ and suppose the items are re-ordered so that $p_1 \leq \dots \leq p_n$ (by symmetry this is WLOG).

It follows from Eq. (5) and the induction hypothesis that there exists a price $\hat{p} \in \{p_2, \dots, p_n\}$ (and thus $\hat{p} \geq p_1$), such that:

$$\begin{aligned}
\mathcal{R}_{DP}(\mathbf{p}) &= p_1\big(1 - \mathrm{F}(p_1)\big) + \mathrm{F}(p_1)\mathcal{R}^{\overline{S}_1}_{DP}(p_2, \dots, p_n) \\
&\leq p_1\big(1 - \mathrm{F}(p_1)\big) + \mathrm{F}(p_1)\mathcal{R}^{\overline{S}_1}_{UP}(\hat{p}) \\
&= \mathcal{R}_{DP}\big(p_1, \hat{p}\mathbf{e}_{n-1}\big), \qquad (6)
\end{aligned}$$

i.e., that the revenue $\mathcal{R}_{DP}(\mathbf{p})$ is dominated by a discriminatory pricing $\tilde{\mathbf{p}} := (p_1, \hat{p}, \dots, \hat{p})$.

Now we will show $\mathcal{R}_{DP}\big(p_1, \hat{p}\mathbf{e}_{n-1}\big)$ is dominated by one of two possible uniform prices, either p_1 or $\hat{p}$. To that end, first note that if $\mathrm{F}(\hat{p}) = 1$, then

$$\mathcal{R}_{DP}\big(p_1, \hat{p}\mathbf{e}_{n-1}\big) = p_1\big(1 - \mathrm{F}(p_1)\big) \leq \mathcal{R}_{UP}(p_1),$$

and the result follows. Thus suppose that $\mathrm{F}(\hat{p}) < 1$ and note since $p_1 \leq \hat{p}$, this also implies that $\mathrm{F}(p_1) < 1$. Define $\lambda := \frac{1-\mathrm{F}(p_1)}{1-\mathrm{F}(p_1)^n}$, and consider:

$$\begin{aligned}
\max\{\mathcal{R}_{UP}(p_1), \mathcal{R}_{UP}(\hat{p})\} &\geq \lambda\mathcal{R}_{UP}(p_1) + (1-\lambda)\mathcal{R}_{UP}(\hat{p}) \\
&= p_1\big(1 - \mathrm{F}(p_1)\big) + \hat{p}\mathrm{F}(p_1)\left(1 - \mathrm{F}(\hat{p})^{n-1}\right)\frac{1 - \mathrm{F}(\hat{p})^n}{1 - \mathrm{F}(\hat{p})^{n-1}}\frac{1 - \mathrm{F}(p_1)^{n-1}}{1 - \mathrm{F}(p_1)^n} \\
&\geq p_1(1 - \mathrm{F}(p_1)) + \mathrm{F}(p_1)\hat{p}(1 - \mathrm{F}(\hat{p})^{n-1}) \\
&= \mathcal{R}_{DP}(\tilde{\mathbf{p}}), \qquad (7)
\end{aligned}$$

where the first inequality follows from the fact that $0 < \lambda \leq 1$, and the second inequality follows because (i) the function $\frac{1-x^n}{1-x^{n-1}}$ is increasing for $x \in [0,1)$, and (ii) since $p_1 \leq \hat{p}$ we have that $\mathrm{F}(p_1) \leq \mathrm{F}(\hat{p})$ and thus,

$$\frac{1-\mathrm{F}(\hat{p})^n}{1-\mathrm{F}(\hat{p})^{n-1}} \geq \frac{1-\mathrm{F}(p_1)^n}{1-\mathrm{F}(p_1)^{n-1}}.$$

Finally, the last equality follows from Eq. (5). Thus from Eq. (6) and Eq. (7), we see that a uniform pricing always dominates any discriminatory pricing. By the principle of mathematical induction, uniform pricing is revenue maximizing for all n. □

We note that the proof of Theorem 1 requires no regularity assumptions on the valuation distribution, nor does it require a detour through mechanism design or prophet inequalities as optimality proofs often do in the quasi-linear setting. In this sense, it is a purely pricing-based proof that proceeds by exploiting the special structure generated by min-buying behavior and IID valuations. The proof of Theorem 1 also serves a useful template for more general proofs of uniform pricing optimality: proceeding by induction on the number of items, reducing to the two price case, and then showing that discriminatory pricing is upper-bounded by a convex combination of two uniform-pricing revenues. We will extend this argument in the next subsection.

3.2. Uniform Pricing Guarantees for Exchangeable and Affiliated Valuations

Although IID valuations are commonly studied, they are not the most general model of symmetric valuations. In the remainder of this section, we consider the broader symmetric setting in which valuations are drawn from an exchangeable distribution. Exchangeability requires that the joint distribution of the valuation vectors be invariant to any permutation of the items. Thus, while exchangeability preserves symmetry across items, it still allows for item valuations to be correlated. This correlation makes exchangeable distributions expressive enough to capture a variety of market characteristics that IID valuations cannot. For example, an exchangeable distribution can represent variation in breadth of preference by controlling how many items a customer is likely to value highly. It can also capture heterogeneity in the strength of preferences by allowing for mixtures of high- and low-valuation customer types (as we do in Example EC.1).

Unfortunately, uniform pricing is not optimal for all exchangeable distributions, even when customers are min-buying. To illustrate, we give a simple example where discriminatory pricing out performs uniform pricing in Example 1 next.

Example 1 ($\mathcal{R}_{DP} > \mathcal{R}_{UP}$ for an Exchangeable Distribution). Suppose there are two items, customers are min-buying, and their valuations are drawn according to the exchangeable three-point distribution:

$$(V_1, V_2) = \begin{cases} (1,1), & \text{with probability } \frac{1}{2}, \\ (2,0), & \text{with probability } \frac{1}{4}, \\ (0,2), & \text{with probability } \frac{1}{4}. \end{cases}$$

For this instance, a uniform price of 1 and uniform price of 2 both earn revenue 1 and thus $\mathcal{R}_{UP} = 1$. Now consider the discriminatory pricing $\mathbf{p} = (1,2)$. Customers with a valuation vector of $(1,1)$ purchase the first item at price 1, customers with a valuation vector of $(2,0)$ also purchase the first item at price 1, and customers with a valuation vector $(0,2)$ purchase the second item at price 2. Therefore,

$$\mathcal{R}_{DP}(1,2) = \frac{1}{2}\cdot 1 + \frac{1}{4}\cdot 1 + \frac{1}{4}\cdot 2 = \frac{5}{4}.$$

Thus, differentiated pricing earns strictly more revenue than uniform pricing, even with only two items and an exchangeable distribution supported on three valuation vectors. □

Example 1 presents a barrier to generalizing Theorem 1 to all exchangeable valuation distributions. However, note that the distribution in Example 1 encodes a form of negative dependence among item valuations: conditional on a customer valuing one item highly, at 2, their valuation for the other item is 0. Although such negative dependence may be plausible in some settings, when items are valued symmetrically because they are substitutes within a common product category, it may also be natural to suppose that valuations are positively associated. That is, customers who value some items highly may also tend to value the remaining items highly. In the mechanism-design literature, this form of positive dependence is commonly modeled through a so-called *affiliation* condition (Milgrom 1981, Milgrom and Weber 1982, Anunrojwong et al. 2022). We give a formal definition next.

Definition 1 (Affiliation). A joint distribution over valuations $V_1, \ldots, V_n$ is **affiliated** if it has continuous CDF F and density f such that, for all valuation vectors $\mathbf{v}$ and $\mathbf{v}'$,

$$f(\mathbf{v} \vee \mathbf{v}')f(\mathbf{v} \wedge \mathbf{v}') \geq f(\mathbf{v})f(\mathbf{v}'), \tag{8}$$

where $\mathbf{v} \vee \mathbf{v}'$ and $\mathbf{v} \wedge \mathbf{v}'$ denote the component-wise maximum and minimum of $\mathbf{v}$ and $\mathbf{v}'$, respectively.

Unpacking Eq. (8), it encodes positive dependence by enforcing that aligned valuations where the vector is all high (the pairwise maxes) or all low (the pairwise minimums) are more likely than mixes of high and low valuations. This is easiest to see for two items, where affiliation is equivalent to requiring that $f(x_1,x_2)f(y_1,y_2) \geq f(x_1,y_2)f(y_1,x_2)$ for all $x_1 < y_1$ and $x_2 < y_2$.

As IID valuations are trivially affiliated, the setting of exchangeable and affiliated valuations is strictly more general than the IID setting studied in Theorem 1. Moreover, the distribution in Example 1 is not affiliated. Thus we may consider whether uniform pricing remains optimal when valuations are both exchangeable and affiliated. The challenge in handling this more general setting is that the positive dependence among item valuations can be quite intricate. Furthermore, for quasi-linear customers, there exist exchangeable and affiliated distributions under which discriminatory pricing earns more than twice the revenue of the optimal uniform pricing; we provide

a simple construction in Example EC.1. Thus, relative to the IID setting, discriminatory pricing can be even more powerful for quasi-linear customers when valuations are exchangeable and affiliated. Nevertheless, imposing affiliation in addition to exchangeability is sufficient to recover the optimality of uniform pricing for min-buying customers, as we show next.

THEOREM 2 **(Exchangeable and Affiliated Valuations)**. *Suppose there are n items, customers are min-buying, and their valuations are drawn according to an exchangeable and affiliated joint distribution. Then a uniform price is revenue-maximizing, i.e., $\mathcal{R}_{DP} = \mathcal{R}_{UP}$.*

The proof of Theorem 2, while more involved, follows the same general structure as the proof of Theorem 1. First, we use the induction hypothesis, together with standard closure properties of exchangeable and affiliated distributions, to reduce from considering arbitrary discriminatory price vectors to just two-level pricings in which the first item has a low price (p_1) and all remaining items have a common higher price ($\hat{p} > p_1$). We then bound the revenue of this two-level pricing by the maximum of $\mathcal{R}_{UP}(p_1)$ and $\mathcal{R}_{UP}(\hat{p})$. However an additional issue arise; under the two-level pricing, either the first item is purchased at price p_1, or it is rejected, meaning $V_1 < p_1$, and the remaining items are considered for purchase. Unlike in the IID setting, for exchangeable distributions rejection of the first item may change the joint distribution of the remaining valuations, so we must control how this conditioning affects the probability of a future purchase. Intuitively, it is here where affiliation rescues the proof as it implies that customers whose valuation for the first item is low also tend to also have lower valuations for the higher-priced items, lowering the revenue potential of discriminatory pricing.

Formally, by following same convex-combination argument used in the proof of Theorem 1, we show that to gain the necessary control over the revenue after conditioning in the exchangeable and affiliated case, it is enough to establish the following inequality. Letting $M := \max_i V_i$, for all $a < b$ we require that

$$\Pr\left(V_1 < a \mid M \geq a\right) \geq \Pr\left(V_1 < a \mid M \geq b\right). \tag{9}$$

In words, Eq. (9) states that conditional on the customer's maximum valuation exceeding a higher threshold, any particular valuation is less likely to remain below the low-price threshold a. Unfortunately, we can not just directly apply the affiliation condition (i.e. Eq. (8)) directly to the pair (V_1, M) as this is not justified by the standard closure properties of affiliated distributions. The maximum M is an aggregation of the valuation vector which includes V_1, and affiliation is not generally preserved under arbitrary increasing aggregations. Positive association alone is also insufficient: it compares fixed increasing events, whereas Eq. (9) requires a monotone comparison across different conditioning thresholds.

We resolve this difficulty by passing to the full vector of order statistics. Let $V^{(k)}$ denote the k^{th} order statistic and note $M = V^{(n)}$. By exchangeability, the order-statistic vector has density $g(y_1, \ldots, y_n) := n! f(y_1, \ldots, y_n) \mathbf{1}\{y_1 \leq \cdots \leq y_n\}$. Now note that the ordered cone $\{y_1 \leq \ldots \leq y_n\}$ is a sub-lattice. Its indicator is therefore affiliated, and multiplying the original affiliated density by this indicator preserves affiliation (Karlin and Rinott 1981). Hence the order statistics are affiliated. Marginalizing the order-statistic vector then shows that each pair $(V^{(k)}, M)$, for $k < n$ is affiliated. This affiliation then implies that, for every fixed a and $k < n$, the function $\phi_{k,a}(t) := \Pr(V^{(k)} < a \mid M \geq t)$ is non-increasing. Finally, using exchangeability and an averaging argument, we translate these order-statistic comparisons back to an individual coordinate to show that, conditional on the symmetric event $\{M \geq t\}$, the probability that $V_1 < a$ equals the average probability that each order statistic lies below a. This establishes Eq. (9) and completes the comparison between the two-level pricing and the two corresponding uniform prices.

In the next subsection, we consider the general symmetric setting of exchangeable distributions and show that, without additional restrictions on the valuations, the worst-case performance of uniform pricing in the min-buying and quasi-linear settings.

3.3. Uniform Pricing for General Exchangeable Distributions

In this final subsection, we complement our results demonstrating the optimality of uniform pricing for min-buying customers in the IID and affiliated and exchangeable settings with a negative result showing that, in the most general symmetric settings, even under min-buying, uniform pricing may earn logarithmically less than optimal discriminatory pricing. For general exchangeable valuations and quasi-linear utility-maximizing customers, Elmachtoub and Hamilton (2021) show that uniform pricing is a $1 + \log(n)$ approximation to the revenue of optimal discriminatory pricing, and this guarantee applies to min-buying customers as well since the revenue of any pricing with quasi-linear customers upper-bounds the revenue with min-buying ones (for uniform pricing they are of course identical). Elmachtoub and Hamilton (2021) also conjecture that their bound is not tight and that a constant-factor approximation may be possible. In the following theorem, we show that this conjecture is incorrect and that, in fact, for both the min-buying and quasi-linear settings, the logarithmic approximation is best possible up to constant factors.

THEOREM 3 **(Exchangeable Valuations)**. *Suppose there are n items and customers are min-buying. There exists an exchangeable distribution such that $\mathcal{R}_{DP} \geq \frac{3}{32} \log(n) \mathcal{R}_{UP}$ which implies that uniform pricing cannot guarantee better than $\Theta(\log(n))$ of the optimal discriminatory pricing revenue.*

The proof of Theorem 3 follows by explicitly constructing an exchangeable distribution over multiple customer types, where each type i is associated with a valuation level, v_i, and a number of

items, $k_i \in [n]$, that the customer values. Conditional on a type, a subset of items of the prescribed size is chosen uniformly at random, and the customer assigns the same positive valuation to every item in that subset, valuing all remaining items as zero. As the valued subset is chosen uniformly at random for each customer, the resulting joint distribution is exchangeable. The valuation levels and type probabilities are then chosen in an equal-revenue fashion so that no single uniform price can extract more than a constant amount of revenue.

The key feature of the construction is that customers with higher valuation levels value fewer items, creating a strong form of negative dependence in the valuations. Consequently, under a discriminatory price vector with geometrically decreasing prices, a type with a high valuation has a non-negligible probability of being directed toward an appropriately high-priced item, whereas types with lower valuations are more likely to purchase at one of the lower prices. Each valuation level therefore contributes on the order of the reciprocal of the number of remaining price levels to discriminatory-pricing revenue. Summing these contributions yields a harmonic, and hence logarithmic, revenue advantage. We defer the details of the construction and analysis to the Appendix.

Combined with Theorem 3 of Elmachtoub and Hamilton (2021), Theorem 3 implies that, for general exchangeable valuations, the worst-case approximation factor of uniform pricing is $\Theta(\log n)$ for both min-buying and quasi-linear utility-maximizing customers. Thus overall, uniform pricing under min-buying provides much stronger guarantees than in the quasi-linear case when valuations across items are symmetric and exhibit some form of positive associated (such as affiliation). However, when the items are symmetric but their valuations may be arbitrarily negatively correlated, uniform pricing can perform similarly poorly under the two decision rules.

In the next section, we move beyond symmetric environments and study the performance of uniform pricing when valuations across items are independent but not necessarily identically distributed across items.

4. Uniform Pricing vs. Differentiated Pricing

In this section, we consider unit-demand pricing problems in which item valuations are independent but not identically distributed. These problems model differentiated pricing, where a seller offers an array of items that need not be perfect substitutes. For instance, standard, professional, and premium versions of software or electronics where customers will purchase at most one of the options, but where the differences in valuations for the goods can reflect differences in features and not just idiosyncratic preference. We refer to pricing in this setting as *differentiated*, rather than discriminatory, because differences in prices may now be justified by genuine differences in item quality reflected in the valuations. In settings with independent but non-identical valuations, we no longer expect uniform pricing to be optimal. Nevertheless, a surprising line of work shows that,

when customers are quasi-linear utility maximizers and distributions are suitably light-tailed, a uniform price can often achieve a significant fraction of the optimal differentiated-pricing revenue (Alaei et al. 2019, Jin et al. 2019). Here, we study the analogous comparison between uniform and differentiated pricing for min-buying customers.

To begin, consider the general case in which valuations for the n items are drawn independently but not identically according to $V_i \sim F_i$, $i \in [n]$. In this setting, an optimal uniform price always earns at least a $1/n$ fraction of the revenue of the optimal differentiated pricing, since uniform pricing can replicate the largest revenue contribution obtained from any one of the n items. Unfortunately, without additional assumptions this trivial guarantee is tight even when customers are min-buying. We demonstrate this for min-buying customers in Example EC.2 in the appendix using a construction similar to one of Hartline and Roughgarden (2009). As no non-trivial approximation of differentiated pricing is possible in the general independent case, in what follows we will study the effectiveness of uniform pricing under the assumption that the valuation distributions are regular.

4.1. Independent Regular Valuations

The instance in Example EC.2 creates a large revenue gap between uniform and differentiated pricing revenue by placing small probabilities on increasingly high item valuations. Different items therefore generate revenue at widely different price scales, and no single uniform price can capture all of these opportunities simultaneously. To rule out such pathological constructions, it is common to assume that each item's valuation distribution is *regular* (Milgrom and Weber 1982, Alaei et al. 2019, Jin et al. 2019, 2020). Specifically, for a distribution F, let $\mathcal{R}_F(q) := qF^{-1}(1-q)$ for $q \in [0,1]$ denote F's revenue-quantile function, where F^{-1} is the generalized quantile function. We call F *regular* if $\mathcal{R}_F$ is concave. Many common distribution are regular, such as the Normal, Uniform, and Exponential distributions. The regularity condition is also implied by many well-known assumptions, including the monotone hazard rate (MHR) condition and the assumption that the density function is log-concave.

For the remainder of this section, we will assume that each item's valuation is drawn independently from a regular distribution. Our main result will be to characterize the worst-case guarantee for the revenue of uniform pricing relative to optimal differentiated pricing. In order to state our main result precisely, we first must define two functions. Namely, define

$$A(x) := -\frac{\log(1-x^2)}{x}, \qquad x \in (0,1), \tag{10}$$

where $A(0) := 0$, and let $c(x)$ solve the ODE:

$$c'(x) = A'(x)\big(1 - xc(x)\big), \tag{11}$$

with boundary condition $c(0) = 1$. Let $x^\star$ be the unique solution to the equation $xc(x) = 1$ (we will prove uniqueness formally in Section 4.2), and define $\Gamma^\star := \frac{1}{x^\star} \approx 1.4314$. We will show that when valuations are independent and regular, the ratio between the revenue of the optimal differentiated pricing and the revenue of the optimal uniform pricing never exceeds $\Gamma^\star$.

THEOREM 4 **(Independent and Regular Valuations)**. *Suppose customers are min-buying and their valuations are drawn independently from continuous regular distributions $V_i \sim F_i$ for $i \in [n]$. Then, the optimal differentiated and uniform pricing revenues satisfies:*

$$\frac{\mathcal{R}_{DP}}{\mathcal{R}_{UP}} \leq \Gamma^\star \approx 1.4314, \tag{12}$$

and this factor is tight.

Theorem 4 is the min-buying analogue of the tight result of Jin et al. (2019) for quasi-linear customers. Through the equivalence between min-buying UDPs and order-oblivious posted pricing, Theorem 4 also characterizes the tight approximation ratio between anonymous pricing and optimal order-oblivious posted pricing. Note, Theorem 4 provides a substantially stronger guarantee for uniform pricing under min-buying than is possible in the quasi-linear setting, improving the guaranteed fraction of optimal differentiated-pricing revenue from approximately 38% to nearly 70%. Thus, although uniform pricing is not generally optimal in asymmetric valuation environments, it remains a remarkably effective policy for the broad class of independent, regular valuations. In the remainder of this section, we prove Theorem 4.

4.2. Proof of Theorem 4

At a high-level, our proof of Theorem 4 follows the broad extremal-reduction philosophy developed in (Alaei et al. 2019, Jin et al. 2019, 2020), successively narrowing types of valuation distributions we must consider until the problem reduces to a one-dimensional optimization. Specifically, our proof has four steps. First, we reduce the space of valuation distributions that we must consider from all regular distributions to so-called triangular-revenue distributions, whose revenue-quantile functions are triangular (cf. Lemma 1). For every instance, this reduction preserves the differentiated pricing revenue while weakly reducing the uniform pricing revenue. Second, we further reduce from all triangular-revenue distributions to a canonical subclass of triangular distributions (cf. Lemma 2). Third, we optimize over these canonical triangular-revenue distributions and prove that $\mathcal{R}_{DP}/\mathcal{R}_{UP}$ is bounded above by $\Gamma^\star$ (cf. Lemma 3). Finally fourth, we construct a sequence of instances for which the ratio between differentiated pricing and uniform pricing revenue converges to $\Gamma^\star$, establishing the tightness of the Theorem 4 (cf. Lemma 4). Overall, our proof follows a similar approach to Jin et al. (2019), adapting their framework to min-buying customers and deriving a direct bound that we then optimize.

4.2.1. Reduction to Triangular-Revenue Distributions Following the reduction approach in Alaei et al. (2019), Jin et al. (2019, 2020), we first identify the regular distributions that are least favorable to uniform pricing relative to differentiated pricing. Specifically, for parameters $a, v > 0$, we say that a random variable V follows a *triangular-revenue distribution*, denoted by $V \sim \mathrm{Tri}(a, v)$, if its cumulative distribution function (CDF) satisfies:

$$\Pr(V < t) = \frac{t}{t+a}, \qquad 0 \le t < v, \tag{13}$$

and the remaining probability mass placed at v, i.e., $\Pr(V = v) = \frac{a}{a+v}$. As v is the maximum valuation, we will often refer to it as the *cap* of the valuation. For such a distribution, the revenue-quantile curve is triangular, i.e.,

$$\mathcal{R}^{\mathrm{Tri}(a,v)}(q) := \begin{cases} vq, & 0 \le q \le \dfrac{a}{a+v}, \\ a(1-q), & \dfrac{a}{a+v} \le q \le 1. \end{cases} \tag{14}$$

where $\mathcal{R}^{\mathrm{Tri}(a,v)}(q)$ is the quantile-revenue function. The two branches of the revenue function meet at $q = a/(a+v)$, where the slope changes from v to $-a$, and so the revenue is maximized at a price of v or, equivalently, at the quantile $q = a/(a+v)$. In the following lemma we demonstrate that the worst-case analysis of any independent, regular instance facing min-buying customers can be reduced to triangular-revenue functions.

Lemma 1 **(Reduction to triangular revenue-quantile curves)**. *Consider any instance where customers are min-buying and their valuations are drawn independently from continuous, regular distributions, $V_i \sim F_i$ for $i \in [n]$. There exists a n independent, triangular-revenue distributions $\{\widetilde{F}_i\}_{i=1}^n$ such that:*

$$\widetilde{\mathcal{R}}_{DP} \ge \mathcal{R}_{DP} \text{ and } \widetilde{\mathcal{R}}_{UP} \le \mathcal{R}_{UP},$$

where $\widetilde{\mathcal{R}}_{DP}$ and $\widetilde{\mathcal{R}}_{UP}$ are the optimal revenues of differentiated and uniform pricing on the transformed instance of triangular-revenue distributions, respectively.

The intuition of Lemma 1 is straightforward and identical to the quasi-linear case where similar reductions follow. Namely, once an item's revenue at its displayed differentiated price is fixed, regularity forces a minimum amount of demand at every lower price. Replacing the original distribution by the corresponding triangular-revenue distribution retains the relevant differentiated-price contribution but makes the item as unattractive as regularity permits at lower uniform prices. Triangular-revenue distributions are therefore extremal for the revenue ratio.

After the reduction, we no longer need to optimize over arbitrary concave revenue curves. Each item is characterized by only two parameters: its cap v, which determines its relevant price scale,

and its parameter a, which determines how much the probability of purchasing it contributes when the price is below its cap. The remaining question is now to understand how large the gap can be between the revenue of differentiated pricing and uniform pricing when items' valuations follow the triangular-revenue distributions.

4.2.2. Further Reduction to Canonical Triangular-Revenue Distributions Lemma 1 reduces any instance of independent regular valuation distributions to one consisting of independent two-parameter triangular-revenue distributions. However, while these transformed instances are simpler, analyzing the worst-case ratio over arbitrary triangular-revenue instances remains challenging. In particular, even for triangular-revenue distributions, the optimal differentiated pricing may set the prices of some items below their caps, leaving an additional set of price variables to optimize over. In the next step of the proof, we further reduce the class of triangular-revenue distributions that we need to consider to one where we may assume the differentiated pricing is at each distributions cap.

The reduction is driven by a simple property of min-buying. Specifically, define the *continuation revenue* from a collection of more expensive items as the expected revenue they generate conditional on the customer reaching them, or equivalently, conditional on no cheaper item being purchased. Suppose the continuation revenue from all items with prices above some value v is C, and consider adding an item $V_j \sim \mathrm{Tri}(a, v)$ priced at its cap, i.e., $p_j = v$, to the instance. The new item is purchased with probability $\frac{a}{a+v}$, and otherwise the customer rejects it and proceeds to the more expensive items. Hence, the new continuation revenue is:

$$T_{a,v}(C) = \frac{a}{a+v}v + \frac{v}{a+v}C = \frac{a+C}{1+a/v}, \tag{15}$$

where $T_{a,v}(C)$ denotes the continuation revenue after adding the item $V_j \sim \mathrm{Tri}(a, v)$ to the instance. From Eq. (15), we see that, once the items are ordered by price, the entire collection of more expensive items affects every cheaper item only through the scalar continuation value C. The number of items, the item's individual distributions, and the individual prices no longer matter for differentiated-pricing revenue once C is fixed. We emphasize that this scalar-state property is specific to the min-buying choice rule and is what allows us to compress a complicated triangular-revenue instance into what we refer to as a *canonical* triangular-revenue instance, following the terminology of Alaei et al. (2019). Our reduction proceeds by successively removing information that is irrelevant to the continuation recursion while ensuring that uniform pricing does not become more profitable.

Formally, a canonical instance consists of a terminal item $V_0 \sim \mathrm{Tri}(1, P)$ and a finite sequence of N triangular-revenue distributions $V_k \sim \mathrm{Tri}(a_k, v_k)$ whose caps are indexed in decreasing order,

$P > v_1 > \ldots > v_N > 1$. We evaluate a canonical instance under the pricing which sets each items price at its cap, i.e., with $p_0 = P$ and $p_k = v_k$ for each $k \in [N]$. For a canonical instance, let C_k denote the continuation revenue obtained after starting with the terminal item and incorporating the first k finite items in decreasing-cap order. By Eq. (15), C_k satisfies

$$C_0 = \frac{P}{P+1}, \qquad C_k = \frac{a_k + C_{k-1}}{1 + a_k/v_k}, \qquad k \in [N].$$

The importance of the canonical representation is that, instead of optimizing over arbitrary triangular-revenue distributions and arbitrary differentiated prices, we now optimize over an ordered sequence of caps $(v_1, \ldots, v_N)$ and an associated sequence of nonnegative triangular parameters $(a_1, \ldots, a_N)$, subject to cap-pricing revenue governed by the scalar continuation recursion above.

In the next lemma, we show that every triangular-revenue instance can be transformed into a canonical instance while weakly decreasing optimal uniform-pricing revenue and producing a cap-pricing whose revenue is arbitrarily close to the original optimal differentiated-pricing revenue.

LEMMA 2 **(Reduction to Canonical Triangular-Revenue Distributions)**. *Consider an instance $\mathcal{I}$ of n independent triangular-revenue valuations $V_i \sim F_i$, $i \in [n]$, normalized so that $\mathcal{R}_{UP} = 1$. For every $\epsilon > 0$, there exist an integer $N \in [n]$, a value $P > 1$, and a canonical instance $\widetilde{\mathcal{I}}$ consisting of independent triangular-revenue valuations $\widetilde{V}_0 \sim Tri(1, P)$ and $\widetilde{V}_k \sim Tri(a_k, v_k)$ for $k \in [N]$ whose caps satisfy $P > v_1 > v_2 > \ldots > v_N > 1$, such that, under the differentiated pricing $\widetilde{\mathbf{p}} = (\widetilde{p}_0, \widetilde{p}_1, \ldots, \widetilde{p}_k) = (p_0, v_1, \ldots, v_k)$ we have*

$$\widetilde{\mathcal{R}}_{DP}(\widetilde{\mathbf{p}}) \geq \mathcal{R}_{DP} - \epsilon \text{ and } \widetilde{\mathcal{R}}_{UP} \leq 1,$$

where $\widetilde{\mathcal{R}}_{DP}(\widetilde{\mathbf{p}})$ and $\widetilde{\mathcal{R}}_{UP}$ are the revenue of differentiated pricing optimal uniform pricing on the transformed instance $\widetilde{\mathcal{I}}$ of canonical triangular-revenue distributions, respectively.

The proof of Lemma 2 consists of three conceptually simple transformations. First, if an item with a triangular-revenue distribution is not priced at its cap under differentiated pricing, then we can truncate the distribution at its displayed differentiated price without changing the item's purchase probability at that price, while weakly reducing its demand under uniform pricing. Second, items with a common cap can be aggregated into a single triangular-revenue item. Third, after indexing the remaining items in decreasing order of their caps, we compress the first several items in this ordering, which have the largest caps and therefore the highest differentiated prices, into a terminal equal-revenue component and at most one boundary item. Similar aggregation and ordered-triangular reductions have been used in Alaei et al. (2019), Jin et al. (2020); however, here our transformations are chosen specifically to preserve the min-buying continuation-revenue

recursion. In the proof, we also show that removing all items with caps less than 1 does not decrease differentiated-pricing revenue.

Lemma 2 preserves the two quantities needed for the worst-case analysis. The canonical instance has uniform-pricing revenue no larger than the normalized benchmark, while its designated cap pricing earns revenue arbitrarily close to the original differentiated-pricing optimum. It is therefore sufficient to bound the cap-pricing revenue of a canonical instance of triangular-revenue distributions.

4.2.3. Bound for Canonical Triangular-Revenue Distributions The canonical reduction in Lemma 2 leaves us with a substantially simpler extremal problem for bounding the maximal ratio between $\mathcal{R}_{DP}$ and $\mathcal{R}_{UP}$. Starting from the terminal continuation value $C_0 = \frac{P}{P+1}$, which approaches one as P becomes large, we need to answer how large the continuation recursion $C_k = \frac{a_k + C_{k-1}}{1 + a_k/v_k}$ can become when the optimal uniform-pricing revenue is at most one.

To answer this question, we must handle the tradeoff between accumulating differentiated-pricing revenue and increasing the opportunity for uniform pricing. To illustrate this, consider some item k with valuation distribution $\mathrm{Tri}(a_k, v_k)$ priced at its cap v_k. Increasing a_k raises the probability that item k sells at its cap v_k, thereby creating a larger revenue opportunity at the price v_k. The same increase, however, also raises the item's acceptance probability $a_k/(a_k + p)$ at every lower uniform price $p \le v_k$. Thus, creating more differentiated-pricing revenue at one price simultaneously increases demand under every lower uniform price. Revenue opportunities at different prices therefore cannot be created independently.

This coupling can be seen directly from the uniform-pricing constraint. Consider the uniform price $p = v_k$. Since $v_1 > v_2 \cdots > v_k$, the first k items, together with the terminal item, are active at this price. The terminal item $V_0 \sim \mathrm{Tri}(1, P)$ rejects $p = v_k$ with probability $v_k/(v_k + 1)$, while item $i \le k$ rejects with probability $v_k/(v_k + a_i)$. Hence,

$$\mathcal{R}_{UP}(p) = v_k \left(1 - \frac{v_k}{v_k + 1} \prod_{i=1}^{k} \frac{v_k}{v_k + a_i}\right) \le 1,$$

since instance is normalized so that $\mathcal{R}_{UP} \le 1$. Let $x_k := \frac{1}{v_k}$ be the reciprocal cap value. Rearranging, and and taking logarithms gives us

$$\sum_{i=1}^{k} \log(1 + a_i x_k) \le -\log(1 - x_k^2).$$

Finally, the concavity of $\log(1 + a_i x)$ implies $\log(1 + a_i x_k) \ge x_k \log(1 + a_i)$, and therefore

$$\sum_{i=1}^{k} \log(1 + a_i) \le -\frac{\log(1 - x_k^2)}{x_k} = A(x_k).$$

We interpret $A(x_k)$ as the cumulative uniform-pricing budget available up to reciprocal price $x_k = 1/v_k$. Thus, the parameters $a_1, \ldots, a_k$ cannot be chosen independently: increasing one of the parameters a_i may generate additional differentiated-pricing revenue at its own price, but it consumes part of the same budget and therefore leaves less room for profitable revenue opportunities at other prices. In this sense, regularity links the revenue opportunities created at different prices through a common uniform-pricing constraint.

This perspective is closely related to normalized common-price programs in the anonymous-pricing literature. Alaei et al. (2019) normalize anonymous-pricing revenue to one and impose the corresponding constraint simultaneously at all prices. Jin et al. (2020) exploit the triangular structure to obtain tractable cumulative constraints at the relevant price boundaries. Our analysis builds on this general perspective, but the objective that must be controlled is different. Under min-buying, allocating triangular mass across price scales affects differentiated revenue through the continuation recursion $C_k = \frac{a_k + C_{k-1}}{1 + a_k x_k}$. The remaining question is therefore how large this recursion can become subject to the cumulative budget constraints above.

The key step is to compare the discrete canonical problem with a continuous relaxation. Rather than optimizing the parameters a_k directly, we first move each reciprocal cap to the smallest value consistent with its cumulative budget, so that the budget constraint binds at every step. This change can only increase continuation revenue. We then show that the resulting relaxed recursion is bounded the continuous benchmark $c(x)$, defined by

$$c'(x) = A'(x)(1 - xc(x)), \qquad c(0) = 1.$$

The differential equation is the infinitesimal counterpart of the discrete continuation recursion. We also prove that the largest value attained by this benchmark is

$$c(x^\star) = \Gamma^\star.$$

It follows that no canonical instance can generate differentiated pricing revenue greater than $\Gamma^\star$ when its uniform-pricing revenue is at most one. We summarize the resulting bound in the following lemma.

LEMMA 3 **(Independent and Canonical Triangular-Revenue Distributions)**. *For every finite canonical instance $\mathcal{I}^{\text{can}}$ with $\mathcal{R}_{UP} \le 1$, the revenue of cap pricing satisfies*

$$\mathcal{R}_{DP}(\mathbf{v}) \le \Gamma^\star. \tag{16}$$

The use of triangular reductions, common-price constraints, and continuous limiting instances has close precedents in Alaei et al. (2019), Jin et al. (2019, 2020). The distinctive feature of our

analysis is what happens after the triangular reduction. Specifically, the min-buying choice rule yields a one-dimensional continuation recursion, and the interaction between this recursion and the uniform-pricing budget leads to the continuous envelope introduced in Eq. (11). The resulting envelope comparison is therefore specific to min-buying and is what produces the constant $\Gamma^\star$.

Lemma 3 completes the upper-bound argument in Theorem 4. Starting from an arbitrary instance of independent regular valuations, Lemma 1 reduces the instance to triangular-revenue distributions without decreasing the relevant differentiated-pricing revenue or increasing uniform-pricing revenue. Lemma 2 reduces the triangular instance, up to an arbitrarily small loss, to the canonical form analyzed above. Lemma 3 then bounds the resulting canonical revenue by $\Gamma^\star$. Consequently, for every $\varepsilon > 0$

$$\mathcal{R}_{DP} - \varepsilon \le \widetilde{\mathcal{R}}_{DP}(\mathbf{v}) \le \Gamma^\star \, \widetilde{\mathcal{R}}_{UP} \le \Gamma^\star \, \mathcal{R}_{UP},$$

Taking $\varepsilon \to 0$, we obtain the upper bound result in Theorem 4, which can be equivalently expressed in terms of the optimal uniform-pricing guarantee as:

$$\mathcal{R}_{UP} \ge \alpha^\star \mathcal{R}_{DP}.$$

Finally, it remains to show that the bound from studying continuous envelope can actually be achieved by a finite instance.

4.2.4. Tightness of Theorem 4 In this subsection, we show that the factor $\Gamma^\star$ is tight. Specifically, the tightness construction reverses the relaxation used in Lemma 3. This is analogous in spirit to the lower-bound constructions of Alaei et al. (2019), Jin et al. (2020), where the continuous or relaxed extremal solution is approximated by increasingly fine, finite triangular-revenue instances.

The construction follows a simple principle. First fix an endpoint $X \in (0,1)$. We then divide the available budget $[0, A(X)]$ into many small pieces and associate one triangular-revenue item with each piece. The cap of each item is selected so that the corresponding uniform-pricing constraint binds exactly. Hence, the construction leaves essentially none of the available uniform-pricing budget unused. As the number of items increases, the discrete continuation recursion then becomes a first-order discretization of the continuous envelope. Thus, it can get arbitrarily close to the maximum $\Gamma^\star$. The following lemma formalizes the resulting lower bound.

LEMMA 4. *For every $\varepsilon > 0$, there exists a finite canonical instance $\mathcal{I}$ such that*

$$\mathcal{R}_{UP} = 1, \qquad \mathcal{R}_{DP} > \Gamma^\star - \varepsilon.$$

Lemma 4 therefore shows that the continuous-envelope bound is asymptotically attainable. The same uniform-pricing constraints that generate the upper bound can be made to bind on an increasingly dense collection of prices, while the corresponding continuation recursion converges to the envelope itself. Hence the constant $\Gamma^\star$ reflects the true worst-case economics of the problem.

One final technical issue remains. The canonical triangular-revenue distributions contain atoms at their caps, whereas Theorem 4 is stated for continuous regular valuations. We resolve this issue in Lemma EC.5 by smoothing each finite triangular instance into an independent continuous regular instance. The smoothing preserves the revenue of the designated cap-pricing while changing optimal uniform-pricing revenue by an arbitrarily small amount. Applying Lemma EC.5 to the sequence of canonical instances in Lemma 4 therefore gives, for every $\varepsilon > 0$, there exists an independent continuous regular instance satisfying

$$\frac{\mathcal{R}_{DP}}{\mathcal{R}_{UP}} > \Gamma^\star - \varepsilon.$$

As the details are standard we defer the statement and proof to Section B.3 in the appendix. Together with the upper bound, this establishes the tightness of Theorem 4.

5. Conclusion

In this paper, we examined unit-demand pricing problems where customers are min-buying and established optimality and approximation results for uniform pricing policies. Our findings highlight key differences between pricing for min-buying customers versus pricing for quasi-linear utility maximizers. Notably, when valuations are symmetric and positively correlated, as is the case for IID valuations and their generalization to exchangeable and affiliated valuations, we show that discriminatory pricing provides no additional benefit in min-buying models. Further, even in cases where valuations are not IID, we demonstrate that the advantages of differentiated pricing are substantially reduced for min-buying customers relative to the quasi-linear case.

Our work thus underscores the fragility of previous analyses of discriminatory and differentiated pricing in the quasi-linear setting. When valuations are IID for instance, our result imply that any additional revenue generated by discriminatory pricing over uniform pricing is driven primarily by the behavioral assumption that customers choose the item that maximizes their utility over other acceptable items. However, this assumption is tenuous and may fail in practice. Research in behavioral economics provides numerous examples where customers do not behave rationally and/or fail to maximize their linear utility (Kahneman and Thaler 2006). In setting where customers are modeled as boundedly rational, risk-averse customers, min-buying may be just as faithful a model of customer's purchasing behavior as a quasi-linear assumption, however unlike the quasi-linear assumption, min-buying strongly supports using a well chosen uniform price in most cases.

Further, our work provides a set of tools to enable further research on pricing for min-buying customers. For instance, one interesting direction is to extend our analysis to richer feasibility environments, such as multi-unit combinatorial settings, and to characterize when uniform (or otherwise simple) pricing policies are optimal or provide strong guarantees for suitability risk-averse customers. It is also would be interesting to further explore the implications of our results for oblivious posted-price mechanisms (which recall are mathematically identically to UDP for min-buying customers) and study mechanism-design questions concerning the relative value of anonymity in pricing, price discrimination, and control over the order of offers. Finally, it would also be useful to study behavioral models that interpolate between min-buying and quasi-linear utility maximization, for example when customers choose the cheapest acceptable item only with some probability or compare only to a limited subset of products. Such models could clarify how sensitive the value of non-uniform pricing is to the precise customer choice rule.

Notes

1. Continuity is not strictly necessary for our results. Unless otherwise stated, we will assume continuity for ease of exposition as it allows us to avoid complications due to tie-breaking between items.

References

Aggarwal, Gagan, Tomás Feder, Rajeev Motwani, An Zhu. 2004. Algorithms for multi-product pricing. *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings 31*. Springer, 72–83. 6

Alaei, Saeed, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis, Yang Yuan. 2019. Optimal auctions vs. anonymous pricing. *Games and Economic Behavior* **118** 494–510. 1, 6, 16, 17, 18, 19, 20, 22, 23, ec8

Anunrojwong, Jerry, Santiago Balseiro, Omar Besbes. 2022. On the robustness of second-price auctions in prior-independent mechanism design. *Proceedings of the 23rd ACM Conference on Economics and Computation*. 151–152. 5, 12

Aouad, Ali, Vivek Farias, Retsef Levi, Danny Segev. 2018. The approximability of assortment optimization under ranking preferences. *Operations Research* **66**(6) 1661–1669. 2, 5

Balcan, Maria-Florina, Avrim Blum, Yishay Mansour. 2008. Item pricing for revenue maximization. *Proceedings of the 9th ACM Conference on Electronic Commerce*. 50–59. 6

Berbeglia, Gerardo, Gwenaël Joret. 2020. Assortment optimisation under a general discrete choice model: A tight analysis of revenue-ordered assortments. *Algorithmica* **82**(4) 681–720. 2, 5

Bhardwaj, Aaditya, Ben Black, Trivikram Dokka, Christopher Kirkbride. 2025. Algorithms for min-buying in networks. *arXiv preprint arXiv:2502.14459* . 4

Bowers, Robin, Marius Garbea, Emmanouil Pountourakis, Samuel Taggart. 2025. Polynomial-time approximation schemes via utility alignment: Unit-demand pricing and more. *arXiv preprint arXiv:2506.20030* . 6

Cai, Yang, Constantinos Daskalakis. 2015. Extreme value theorems for optimal multidimensional pricing. *Games and Economic Behavior* **92** 266–305. 6

Cao, Yufeng, Anton J Kleywegt, He Wang. 2022. Network revenue management under a spiked multinomial logit choice model. *Operations Research* **70**(4) 2237–2253. 4

Chawla, Shuchi, Jason D Hartline, Robert Kleinberg. 2007. Algorithmic pricing via virtual valuations. *Proceedings of the 8th ACM Conference on Electronic Commerce*. 243–251. 1, 5, 9

Chawla, Shuchi, Jason D Hartline, David L Malec, Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. *Proceedings of the forty-second ACM symposium on Theory of computing*. 311–320. 1, 2, 5

Chen, Xi, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, Mihalis Yannakakis. 2014. The complexity of optimal multidimensional pricing. *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1319–1328. 5

Chen, Xi, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, Mihalis Yannakakis. 2018. The complexity of optimal multidimensional pricing for a unit-demand buyer. *Games and Economic Behavior* **110** 139–164. 6, 9

Chen, Xingxing, Levi DeValve, Jacob Feldman, Yukai Huang, Danny Segev. 2025. Approximation schemes for dynamic pricing with opaque products. *Available at SSRN 5400432* . 2, 5

Correa, José, Patricio Foncea, Dana Pizarro, Victor Verdugo. 2019. From pricing to prophets, and back! *Operations Research Letters* **47**(1) 25–29. 5

Dütting, Paul, Felix A Fischer, Max Klimm. 2016. Revenue gaps for discriminatory and anonymous sequential posted pricing. *CoRR, abs/1607.07105* . 2, 5, 9

Elmachtoub, Adam N, Michael L Hamilton. 2021. The power of opaque products in pricing. *Management Science* **67**(8) 4686–4702. 2, 5, 14, 15, ec5

Gallego, Guillermo, Gerardo Berbeglia. 2021. Bounds and heuristics for multi-product personalized pricing. *arXiv preprint arXiv:2102.03038* . 6

Gigerenzer, Gerd, Daniel G Goldstein. 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychological review* **103**(4) 650. 4

Goyal, Vineet, Retsef Levi, Danny Segev. 2016. Near-optimal algorithms for the assortment planning problem under dynamic substitution and stochastic demand. *Operations Research* **64**(1) 219–235. 4

Guruswami, Venkatesan, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, Frank McSherry. 2005. On profit-maximizing envy-free pricing. *SODA*, vol. 5. 1164–1173. 6

Hartline, Jason D, Tim Roughgarden. 2009. Simple versus optimal mechanisms. *Proceedings of the 10th ACM conference on Electronic commerce*. 225–234. 16, ec2

Housni, Omar El, Adam N Elmachtoub, Harsh Sheth, Jiaqi Shi. 2025. Price and assortment optimization under the multinomial logit model with opaque products. *arXiv preprint arXiv:2502.08124* . 2, 5

Jin, Yaonan, Pinyan Lu, Qi Qi, Zhihao Gavin Tang, Tao Xiao. 2019. Tight approximation ratio of anonymous pricing. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 674–685. 1, 3, 6, 16, 17, 18, 22

Jin, Yaonan, Pinyan Lu, Zhihao Gavin Tang, Tao Xiao. 2020. Tight revenue gaps among simple mechanisms. *SIAM Journal on Computing* **49**(5) 927–958. 6, 16, 17, 18, 20, 22, 23, ec8, ec9

Kahneman, Daniel. 2003. Maps of bounded rationality: Psychology for behavioral economics. *American economic review* **93**(5) 1449–1475. 2, 4

Kahneman, Daniel, Richard H Thaler. 2006. Anomalies: Utility maximization and experienced utility. *Journal of economic perspectives* **20**(1) 221–234. 2, 24

Karle, Heiko, Georg Kirchsteiger, Martin Peitz. 2015. Loss aversion and consumption choice: Theory and experimental evidence. *American Economic Journal: Microeconomics* **7**(2) 101–120. 4

Karlin, Samuel, Yosef Rinott. 1980. Classes of orderings of measures and related correlation inequalities. i. multivariate totally positive distributions. *Journal of Multivariate Analysis* **10**(4) 467–498. ec3

Karlin, Samuel, Yosef Rinott. 1981. Total positivity properties of absolute value multinormal variables with applications to confidence interval estimates and related probabilistic inequalities. *The Annals of Statistics* 1035–1049. 14

Lucier, Brendan. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* **16**(1) 24–47. 5

Milgrom, Paul R. 1981. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics* 380–391. 12

Milgrom, Paul R, Robert J Weber. 1982. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* 1089–1122. 12, 16

Rusmevichientong, Paat, Benjamin Van Roy, Peter W Glynn. 2006. A nonparametric approach to multiproduct pricing. *Operations Research* **54**(1) 82–98. 2

Simon, Herbert A. 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics* 99–118. 4

Stüttgen, Peter, Peter Boatwright, Robert T Monroe. 2012. A satisficing choice model. *Marketing Science* **31**(6) 878–899. 4

Thanassoulis, John. 2004. Haggling over substitutes. *Journal of Economic theory* **117**(2) 217–245. 2, 5

Appendix A: Omitted Examples

In this section, we include all examples which were omitted from the main body.

Example EC.1 ($\frac{\mathcal{R}_{DP}}{\mathcal{R}_{UP}} > 2.23$ for Exchangeable & Affiliated Valuations, Quasi-Linear Case). In this example, we give an explicit construction and analysis of an instance where valuations are drawn from an exchangeable and affiliated distribution, and for quasi-linear customers discriminatory pricing earns more than twice that of the optimal uniform pricing. Thus, the separation between the min-buying and quasi-linear case is strictly greater for exchangeable and affiliated distributions than it is in the IID setting.

Suppose there are $n = 56$ items, partitioned into three groups of items G_1, G_2, and G_3 where $|G_1| = 1$, $|G_2| = 5$, and $|G_3| = 50$. Let customers be quasi-linear utility maximizers who form valuations by first drawing a latent type $T \in \{1, 2, 3\}$ according to $\Pr(T = 1) = \frac{99}{100}$, $\Pr(T = 2) = \frac{99}{10000}$, and $\Pr(T = 3) = \frac{1}{10000}$. Conditional on T, the customers drawn valuations IID according to one of three distributions. For a type-1 customer, valuations are always equal to one, i.e., $V_i = 1$ with probability 1. For a type-2 customer, valuations are drawn according to a binary distribution such that $V_i = 1$ with probability $\frac{7}{10}$ and otherwise $V_i = 100$ with probability $\frac{3}{10}$. Finally, for a type-3 customer, valuations are drawn according to:

$$V_i = \begin{cases} 10000 & \text{with probability } \frac{1}{25}, \\ 100 & \text{with probability } \frac{36}{125}, \\ 1 & \text{with probability } \frac{84}{125}. \end{cases}$$

Thus overall, customers draw valuations according to a mixture of IID distributions, and the mixture family satisfies a monotone likelihood ratio property and hence is exchangeable and affiliated (this is easy to verify directly, we omit the computation for brevity).

Now, over this instance we will first bound the revenue of uniform pricing. If $p \le 1$, then $\mathcal{R}_{UP}(p) = p \le 1$ since valuations are always at least one, and this is tight when $p = 1$. If $p \in (1, 100]$, then only type-2 and type-3 customers can purchase and since these types have total probability $1/100$, $\mathcal{R}_{UP}(p) \le \frac{p}{100} \le 1$. Finally if $p \in (100, 10000]$, then only type-3 customers can purchase $\mathcal{R}_{UP}(p) \le \frac{p}{10000} \le 1$. As prices over 10000 earn no revenue, we conclude that $\mathcal{R}_{UP} = 1$

Now consider the feasible discriminatory pricing that charges $p_1 = .99$ for items in group G_1, $p_2 = 99.89$ for items in group G_2, and $p_3 = 9900.89$ for items in group G_3. A type-1 customer always purchases the item in G_1, contributing expected revenue of $\frac{99}{100}(0.99) = 0.9801$. A type-2 customer purchases an item in G_2 whenever at least one of the items in G_2 is valued at 100 and the item in G_1 is valued at 1, otherwise they purchase an item in G_1. The expected revenue from a type-2 customer is:

$$\frac{99}{10000}\left((99.89)\frac{7}{10}\left(1-\left(\frac{7}{10}\right)^5\right)+.99\left(1-\frac{7}{10}\left(1-\left(\frac{7}{10}\right)^5\right)\right)\right) \ge .5799.$$

Finally, a type-3 customer purchases an item in G_3 whenever none of the six items in $G_1 \cup G_2$ are valued at 10000 and at least one item in G_3 is valued at 10000. A lower-bound on expected revenue from type-3 customers, considering only this purchasing event or otherwise assuming they buy an item in G_1 is:

$$\frac{1}{10000}\left(9900.89\left(\frac{24}{25}\right)^6\left(1-\left(\frac{24}{25}\right)^{50}\right)+.99\left(1-\left(\frac{24}{25}\right)^6\left(1-\left(\frac{24}{25}\right)^{50}\right)\right)\right) \ge .6743$$

Combining the revenue from the three types, we have that $\mathcal{R}_{DP} \geq .9801 + .5799 + .6743 = 2.2343$. Thus discriminatory pricing can earn more than twice as much revenue as uniform pricing even when valuations are exchangeable and affiliated. Further, for simplicity, we have limited this construction to three customer types but extending the construction allows us to achieve any level of gap between uniform pricing and discriminatory pricing revenue. □

EXAMPLE EC.2 ($\mathcal{R}_{DP}/\mathcal{R}_{UP} = n$ FOR GENERAL INDEPENDENT VALUATIONS). In this example, we give an explicit construction and analysis of a unit-demand pricing problem with min-buying customers and independent valuations where a single price can guarantee no better than $\frac{1}{n}$ of the revenue of the optima; differentiated pricing. We note that a similar construction appears in Section 5 of Hartline and Roughgarden (2009), albeit in a slightly different context.

Fix any $\epsilon \in (0,1)$ and choose some $B > 2/\varepsilon$. Suppose customers are min-buying and draw valuations V_i for items $i \in [n]$ independently according to two-point distributions where $V_i = 0$ with probability $1 - B^{-i}$ and $V_i = B^i$ with probability B^{-i}. For this instance, the optimal uniform price must be in the support of the valuations. To bound the optimal uniform pricing revenue, consider the price $p = B^i$ for some $i \in [n]$. This uniform price earns:

$$\mathcal{R}_{UP}(B^i) = B^i \Pr\left(\max_j V_j \geq B^i\right) = B^i\left(1 - \prod_{j=i}^{n}(1 - B^{-j})\right) \leq B^i \sum_{j=i}^{n} B^{-j} = \sum_{j=0}^{n-i} B^{-j} \leq \frac{B}{B-1},$$

where the first inequality follows from the union bound and the second by adding terms and summing the geometric series. Thus for any price in the support $\mathcal{R}_{UP}(p) \leq \mathcal{R}_{UP}(B^i) \leq \frac{B}{B-1}$.

Now consider the differentiated pricing which uses prices $\mathbf{p} := (B, B^2, B^3, \ldots, B^n)$. Under this pricing, a min-buying customer buys item i with probability $B^{-i}\prod_{j=1}^{i-1}(1 - B^{-j})$, and the revenue of differentiated pricing under $\mathbf{p}$ can be bounded as

$$\begin{aligned}
\mathcal{R}_{DP}(\mathbf{p}) &= \sum_{i=1}^{n} B^i B^{-i} \prod_{j=1}^{i-1}(1 - B^{-j}) \\
&= \sum_{i=1}^{n}\prod_{j=1}^{i-1}(1 - B^{-j}) \\
&\geq \sum_{i=1}^{n}\prod_{j=1}^{\infty}(1 - B^{-j}) \\
&\geq \sum_{i=1}^{n}(1 - \sum_{j=1}^{\infty} B^{-j}) \\
&= \sum_{i=1}^{n}(1 - \frac{1}{B-1}) = n\frac{B-2}{B-1}
\end{aligned}$$

where the second inequality follows from the fact that $\prod_j (1 - x_j) \geq 1 - \sum_j x_j$ for any collection $\{x_j\} \subset [0,1]$. Combining the upper-bound for uniform pricing with the lower bound for differentiated pricing we see $\frac{\mathcal{R}_{DP}}{\mathcal{R}_{UP}} \geq n\frac{B}{B-2} = \frac{n}{1-\epsilon}$ as desired. Further, taking $\epsilon \to 0$, we obtain worst-case approximation factor of uniform pricing for general independent valuations is exactly n. □

Appendix B: Omitted Proofs

In this section, we include proofs of all results which were omitted from the main body.

B.1. Omitted Proofs from Section 3

Proof of Theorem 2. As in the proof of Theorem 1, our proof will proceed by induction on the number of items, n, to prove the stronger statement that, for every exchangeable and affiliated joint distribution F over valuations $V_1, \dots, V_n$, and for every price vector $\mathbf{p} = (p_1, \dots, p_n)$, there exists a price $\hat{p} \in \{p_1, \dots, p_n\}$ such that $\mathcal{R}_{UP}(\hat{p}) \geq \mathcal{R}_{DP}(\mathbf{p})$. The base case $n = 1$ follows immediately, so let $n \geq 2$ and suppose the claim holds for $n-1$ items under every exchangeable and affiliated distribution. Fix some arbitrary price vector $\mathbf{p} = (p_1, \dots, p_n)$ and recall since the valuation distribution is exchangeable, we may relabel the items so that $p_1 \leq p_2 \leq \ldots \leq p_n$ without loss of generality.

Now, even though valuations are no longer independent, we can still give an recursive analog of Eq. (4) as:

$$\mathcal{R}_{DP}(\mathbf{p}; F) = p_1 \Pr(V_1 \geq p_1) + \Pr(V_1 < p_1)\mathcal{R}^{\overline{S}_1}_{DP}\left((p_2, \dots, p_n); F^{V_1 < p_1}\right), \tag{EC.1}$$

where we will use the notation $\mathcal{R}_{DP}(\cdot\,; G)$ to emphasize that the revenue is taken as an expectation over valuations $\mathbf{V} \sim G$, and where $F^{V_1 < p_1}$ denotes the conditional distribution of $(V_2, \dots, V_n)$ after observing the event $V_1 < p_1$. We may also suppose that $\Pr(V_1 < p_1) > 0$ without loss of generality since, if $\Pr(V_1 < p_1) = 0$, then the lowest priced first item is always purchased and $\mathcal{R}_{DP}(\mathbf{p}) = \mathcal{R}_{UP}(p_1)$.

Now, in order to apply our induction hypothesis in Eq. (EC.1), we must show that $F^{V_1 < p_1}$ is also exchangeable and affiliated. To see this, recall that affiliation is preserved under products, conditioning on sublattice events, and marginalization (Karlin and Rinott 1980), and note that, up to normalization, the density of $F^{V_1 < p_1}$ is

$$f^{V_1 < p_1}(v_2, \dots, v_n) \propto \int_0^{p_1} f(x, v_2, \dots, v_n) dx.$$

The set $\{(v_1, \dots, v_n) | v_1 < p_1\}$ is a sublattice, hence its indicator satisfies the affiliation inequality Eq. (8), and thus $f(v_1, \dots, v_n)\mathbb{I}\{v_1 < p_1\}$ is affiliated. Marginalizing this function over v_1 shows that the unnormalized density $f^{V_1 < p_1}$ is affiliated. Normalization preserves this property, and hence the conditional distribution is affiliated. Moreover, because the original density f is exchangeable, the conditional density remains exchangeable in the remaining coordinates $2, \dots, n$.

We may therefore apply the induction hypothesis to the conditional distribution $F^{V_1 < p_1}$ and the continuation price vector $(p_2, \dots, p_n)$. It follows that there exists a price $\hat{p} \in \{p_2, \dots, p_n\}$ such that:

$$\mathcal{R}^{\overline{S}_1}_{DP}\left((p_2, \dots, p_n); F^{V_1 < p_1}\right) \leq \mathcal{R}^{\overline{S}_1}_{UP}\left(\hat{p}; F^{V_1 < p_1}\right).$$

Since $\hat{p}$ is one of the continuation prices, we have $\hat{p} \geq p_1$. By the induction hypothesis and Eq. (EC.1), we have:

$$\begin{aligned}\mathcal{R}_{DP}(\mathbf{p}; F) &\leq p_1 \Pr(V_1 \geq p_1) + \Pr(V_1 < p_1)\mathcal{R}^{\overline{S}_1}_{UP}\left(\hat{p}; F^{V_1 < p_1}\right) \\ &= \mathcal{R}_{DP}(\tilde{\mathbf{p}}; F),\end{aligned} \tag{EC.2}$$

where $\tilde{\mathbf{p}} := (p_1, \hat{p}\mathbf{e}_{n-1})$. Thus, the original price vector $\mathbf{p}$ is weakly dominated by the two-tier price vector $\tilde{\mathbf{p}}$. It remains to show that

$$\mathcal{R}_{DP}(\tilde{\mathbf{p}}; F) \leq \max\left\{\mathcal{R}_{UP}(p_1; F), \mathcal{R}_{UP}(\hat{p}; F)\right\}. \tag{EC.3}$$

To this end, define the following four events:

$$A_1 := \left\{ \max_{i\in[n]} V_i \geq p_1 \right\}, \qquad A_2 := \left\{ \max_{i\in[n]} V_i \geq \hat{p} \right\},$$
$$B_1 := \left\{ V_1 < p_1, \max_{i\in[2,n]} V_i \geq p_1 \right\}, \qquad B_2 := \left\{ V_1 < p_1, \max_{i\in[2,n]} V_i \geq \hat{p} \right\}.$$

A_1 and A_2 are the events in which an item is purchased under uniform prices p_1 and $\hat{p}$, respectively. B_1 and B_2 are the events that occur when the first item is rejected at price p_1 and another item is acceptable at price p_1 or $\hat{p}$, respectively. Note we may suppose $\Pr(A_2) > 0$ since, if not, the high-price tier never sells implying that

$$\mathcal{R}_{DP}(\tilde{\mathbf{p}}; F) = p_1 \Pr(V_1 \geq p_1) \leq p_1 \Pr(A_1) = \mathcal{R}_{UP}(p_1; F),$$

and also $\Pr(A_2) > 0$ implies that $\Pr(A_1) > 0$ because $A_2 \subseteq A_1$.

Now, following the proof of Theorem 1, define $\lambda := \frac{\Pr(V_1 \geq p_1)}{\Pr(A_1)}$. Since A_1 is the disjoint union of $\{V_1 \geq p_1\}$ and B_1, we can express $1 - \lambda = \frac{\Pr(B_1)}{\Pr(A_1)}$. Consequently,

$$\begin{aligned} \max\{\mathcal{R}_{UP}(p_1; F), \mathcal{R}_{UP}(\hat{p}; F)\} &\geq \lambda \mathcal{R}_{UP}(p_1; F) + (1-\lambda)\mathcal{R}_{UP}(\hat{p}; F) \\ &= p_1 \Pr(V_1 \geq p_1) + \hat{p} \frac{\Pr(B_1)}{\Pr(A_1)} \Pr(A_2). \end{aligned} \tag{EC.4}$$

To complete the proof, it suffices to show that $\frac{\Pr(B_1)}{\Pr(A_1)} \Pr(A_2) \geq \Pr(B_2)$ since $\mathcal{R}_{DP}(\tilde{\mathbf{p}}; F) = p_1 \Pr(V_1 \geq p_1) + \hat{p} \Pr(B_2)$. By rearranging and conditioning, this inequality is equivalent to showing:

$$\Pr(V_1 < p_1 \mid A_1) \geq \Pr(V_1 < p_1 \mid A_2). \tag{EC.5}$$

The remainder of the proof will be to show Eq. (EC.5). Before proceeding, let us pause to note that Eq. (EC.5) is intuitive for exchangeable and affiliated valuation distributions. It states that the higher the lower bound imposed on the customer's maximum valuation (here, $\hat{p}$ rather than p_1), the more the conditional marginal distribution of each item valuation should shift upward. In particular, V_1 should become less likely to lie below the low-price threshold p_1. To formalize this intuition, we will combine exchangeability with the positive-dependence structure implied by affiliation by passing to the joint distribution of the order statistics.

Let $V^{(i)}$ denote the i^{th} order statistic of the valuations, and let $M := V^{(n)}$ denote the maximum. Since the original density f is exchangeable, the joint density of the order statistics is

$$g(y_1, \ldots, y_n) = n! f(y_1, \ldots, y_n) \mathbb{I}\{y_1 \leq \ldots \leq y_n\}. \tag{EC.6}$$

The ordered cone $\mathcal{C}_n := \{(y_1, \ldots, y_n) | y_1 \leq \ldots \leq y_n\}$ is a sublattice. Therefore its indicator function satisfies Eq. (8), and $g(y_1, \ldots, y_n)$ is an affiliated density. Further, again since affiliation is preserved under marginalization, for every $k \in [n-1]$, the pair $(V^{(k)}, M)$ is affiliated.

Our strategy is now to establish an analogue of Eq. (EC.5) for each pair $(V^{(k)}, M)$ and then use an averaging argument to obtain the desired inequality for (V_1, M). To that end, for every fixed a and every $k \in [n-1]$, define $\phi_{k,a}(t) := \Pr(V^{(k)} < a \mid M \geq t)$ for all t such that $\Pr(M \geq t) > 0$. We will show that $\phi_{k,a}(t)$ is non-increasing in t. Fix $s < t$, and define the events

$$L := \{V^{(k)} < a\}, \qquad H := \{V^{(k)} \geq a\},$$
$$C_1 := \{s \leq M < t\}, \qquad C_2 := \{M \geq t\},$$

Let $h_k(x,z)$ denote the joint density of $(V^{(k)}, M)$. Since h_k is affiliated, for every $x < a \le y$ and $s \le z < t \le w$, we have

$$h_k(x,z)h_k(y,w) \ge h_k(x,w)h_k(y,z).$$

Integrating this inequality over $x < a$, $y \ge a$, $s \le z < t$, and $w \ge t$ gives:

$$\int_{x<a}\int_{y\ge a}\int_{s}^{t}\int_{t}^{\infty} h_k(x,z)h_k(y,w)\,dw\,dz\,dy\,dx \ge \int_{x<a}\int_{y\ge a}\int_{s}^{t}\int_{t}^{\infty} h_k(x,w)h_k(y,z)\,dw\,dz\,dy\,dx.$$

The left-hand integral factors over (x,z) and (y,w), while the right-hand integral factors over (x,w) and (y,z). Therefore, this inequality is equivalent to:

$$\Pr(L\cap C_1)\Pr(H\cap C_2) \ge \Pr(L\cap C_2)\Pr(H\cap C_1), \tag{EC.7}$$

which is equivalent to $\Pr(L \mid C_1) \ge \Pr(L \mid C_2)$ by conditioning. Moreover, note $\{M \ge s\}$ is the disjoint union of C_1 and C_2. Thus $\Pr(L \mid M \ge s)$ is a convex combination of $\Pr(L \mid C_1)$ and $\Pr(L \mid C_2)$. It follows that $\Pr(L \mid M \ge s) \ge \Pr(L \mid C_2) = \Pr(L \mid M \ge t)$. Thus, $\phi_{k,a}(s) \ge \phi_{k,a}(t)$, and we conclude that $\phi_{k,a}$ is non-increasing.

Now we will translate these order-statistic comparisons back to an individual valuation. By exchangeability, conditional on the symmetric event $\{M \ge t\}$, every coordinate has the same probability of lying below a. Therefore,

$$\Pr(V_1 < a \mid M \ge t) = \frac{1}{n}\mathbb{E}\left[\sum_{i=1}^{n}\mathbf{1}\{V_i < a\}\,\middle|\,M \ge t\right] = \frac{1}{n}\sum_{k=1}^{n}\phi_{k,a}(t) = \frac{1}{n}\sum_{k=1}^{n-1}\phi_{k,a}(t), \tag{EC.8}$$

where the final equality follows from noting that when $t \ge a$, the event $\{M \ge t\}$ implies that $M \ge a$, and the final term is zero. As every term on the right-hand side is non-increasing in t, it follows that the probability $\Pr(V_1 < a \mid M \ge t)$ is also non-increasing for $t \ge a$. Taking $a = p_1$ and $t \in \{p_1, \hat{p}\}$ gives $\Pr(V_1 < p_1 \mid A_1) \ge \Pr(V_1 < p_1 \mid A_2)$ which proves Eq. (EC.5). Substituting this inequality into Eq. (EC.4) and combining it with Eq. (EC.2), gives that $\mathcal{R}_{DP}(\mathbf{p}; F) \le \max\{\mathcal{R}_{UP}(p_1; F), \mathcal{R}_{UP}(\hat{p}; F)\}$ which completes the induction.

Thus every discriminatory price vector is weakly dominated by a uniform price chosen from one of its own price levels, and this implies that $\mathcal{R}_{DP} = \mathcal{R}_{UP}$ as desired. □

Proof of Theorem 3. As the upper-bound is handled by a result of Elmachtoub and Hamilton (2021), here we will construct a lower-bounding instance where the gap in revenue between discriminatory pricing and uniform pricing is logarithmic. To that end, suppose there are n items and for each $i \in [n]$, define $r_i := 4^{n-i}$, $\lambda_i = \frac{3}{8r_i}$, $m_i = n - i + 1$, and $k_i = \max\{1, \left\lfloor \frac{n}{2m_i} \right\rfloor\}$. Consider the distribution of valuations $\mathbf{V} = (V_1, \ldots, V_n) \sim F$ such that there are n valuation types. Each type $i \in [n]$ occurs with probability λ_i and a type-i valuation vector is constructed by choosing a subset $S \subseteq [n]$ of cardinality k_i uniformly at random among all subsets, then setting $V_j = 2r_i$ if $j \in S$ and $V_j = 0$ if $j \notin S$. With the remaining probability $1 - \sum_i \lambda_i$, let $\mathbf{V} = \mathbf{0}$. Note this distribution over valuations is clearly exchangeable and that this is a valid probability distribution since,

$$\sum_{i=1}^{n}\lambda_i = \frac{3}{8}\sum_{i=1}^{n}4^{i-n} = \frac{3}{8}\sum_{i=0}^{n-1}4^{-i} = \frac{1}{2}\left(1 - 4^{-n}\right) < 1.$$

Now to upper-bound the revenue of an optimal uniform pricing, consider any uniform price $p \in (0, 2r_1]$ (since if $p > 2r_1$ no buyer purchases). Let $i \in [n]$ be the unique index such that $2r_{i+1} < p \le 2r_i$ (letting $r_{n+1} = 0$). For any type-t valuation vector where $t > i$, we have

$$\max_{j \in [n]} V_j = 2r_t \le 2r_{i+1} < p$$

and thus they will not purchase any item at price p. Consequently,

$$\mathcal{R}_{UP}(p) \le p \sum_{j=1}^{i} \lambda_j \le 2r_i \sum_{j=1}^{i} \frac{3}{8r_j} = \frac{3}{4} \sum_{j=1}^{i} \frac{r_i}{r_j} = \frac{3}{4} \sum_{j=0}^{i-1} 4^{-j} = 1 - 4^{-i} < 1,$$

which implies that $\max_p \mathcal{R}_{UP}(p) \le 1$.

Now we will lower-bound the revenue of optimal discriminatory pricing by considering the policy which uses prices $\mathbf{p} = (r_1, r_2, \ldots, r_n)$[2]. Consider a type-$i$ valuation vector (i.e., the valuation vector where $|S| = k_i$ items have valuation $2r_i$, chosen u.a.r) and consider the event where the highest index item with non-zero valuation is i, denoted by E_i. Under event E_i, the item i is valued as $2r_i$ which exceeds its price of r_i, and it is the cheapest item which gives non-negative utility thus i is purchased. Further, note that since S is a uniformly random size k_i subset of $[n]$,

$$\Pr(E_i \mid \text{Valuation of type } i) = \frac{\binom{i-1}{k_i - 1}}{\binom{n}{k_i}}. \tag{EC.9}$$

Now, if we can prove that that $\Pr(E_i \mid \text{Valuation of type } i) \ge \frac{1}{4m_i}$, it will imply our desired result since $\mathcal{R}_{UP} \le 1$ and

$$\begin{aligned}
\mathcal{R}_{DP} \ge \mathcal{R}_{DP}(r_1, \ldots, r_n) &\ge \sum_{i=1}^{n} r_i \Pr(E_i \mid \text{Valuation of type } i) \lambda_i \\
&\ge \sum_{i=1}^{n} r_i \frac{1}{4m_i} \frac{3}{8r_i} \\
&= \frac{3}{32} \sum_{i=1}^{n} \frac{1}{n - i + 1} \\
&= \frac{3}{32} H_n,
\end{aligned}$$

where H_n is the n^{th} harmonic number. Thus, in the remainder of this example we will focus on lower-bounding the expression in Eq. (EC.9).

We will consider two cases. First suppose that i is such that $k_i = 1$. In this case, $\Pr(E_i \mid \text{Valuation of type } i) = \frac{1}{n}$ by Eq. (EC.9), and $k_i = 1$ implies that $\frac{n}{2m_i} < 2$ by its definition. Thus $\Pr(E_i \mid \text{Valuation of type } i) = \frac{1}{n} > \frac{1}{4m_i}$ as desired. Now suppose i is such that $k_i \ge 2$. Since $k_i = \left\lfloor \frac{n}{2m_i} \right\rfloor$, we have $k_i \le \frac{n}{2m_i} < k_i + 1$. Using $i = n - m_i + 1$, by Eq. (EC.9) we obtain:

$$\begin{aligned}
\Pr(E_i \mid \text{Valuation of type } i) &= \frac{k_i}{n} \prod_{j=1}^{k_i - 1} \frac{i - j}{n - j} \\
&= \frac{k_i}{n} \prod_{j=1}^{k_i - 1} \left(1 - \frac{m_i - 1}{n - j}\right)
\end{aligned}$$

$$
\begin{aligned}
&\geq \frac{k_i}{n}\left(1-\sum_{j=1}^{k_i-1}\frac{m_i-1}{n-j}\right)\\
&\geq \frac{k_i}{n}\left(1-\frac{(k_i-1)(m_i-1)}{n-k_i+1}\right)\\
&\geq \frac{k_i}{n}\left(1-\frac{k_i-1}{2k_i}\right)\\
&\geq \frac{k_i}{2m_i(k_i+1)}\left(1-\frac{k_i-1}{2k_i}\right)\\
&\geq \frac{1}{4m_i}
\end{aligned}
$$

where the first inequality follows from the inequality $\prod_j(1-a_j)\geq 1-\sum_j a_j$ for any numbers $\{a_j\}\subset[0,1]$, the third and fourth inequalities follows from recalling i is such that $2k_i(m_i-1)\leq n-2k_i\leq n-k_i+1$ which implies that $\frac{(k_i-1)(m_i-1)}{n-k_i+1}\leq\frac{k_i-1}{2k_i}$ and $\frac{k_im_i}{n}\geq\frac{k_i}{2(k_i+1)}$, and the final inequality is by rearranging. Thus $\Pr(E_i\,|\,\text{Valuation of type } i)\geq\frac{1}{4m_i}$ for all i, as desired. □

B.2. Omitted Proofs from Section 4

Proof of Lemma 1. Suppose $\mathbf{p}=(p_1,\ldots,p_n)$ is the differentiated price vector given $V_i\sim F_i$ where $p_1\leq p_2\leq\cdots\leq p_n$ and F_i is regular for $i\in[n]$. For each item, define $q_i:=\Pr(V_i\geq p_i)$ as the acceptance probability at the differentiated prices. Without loss of generality, we assume $q_i\in(0,1)$ for all i, since otherwise the item generates no revenue and can be omitted. The revenue function for each item i under price p_i can therefore be written as $\mathcal{R}_{F_i}(q_i):=p_iq_i$. In this proof, we construct an instance with triangular revenue distributions that preserves the revenue of differentiated pricing while weakly decreasing the revenue of uniform pricing.

To construct the triangular revenue distribution instance, we first define

$$a_i:=\frac{p_iq_i}{1-q_i},\tag{EC.10}$$

and replace V_i by $\widetilde{V}_i\sim\mathrm{Tri}(a_i,p_i)$. By the definition of $\widetilde{V}_i$, we have

$$\Pr(\widetilde{V}_i\geq p_i)=\frac{a_i}{a_i+p_i}=q_i,$$

where the equality follows by plugging in the definition of a_i. Recall that the revenue of differentiated pricing is

$$\mathcal{R}_{DP}(\mathbf{p})=\sum_{i=1}^{n}p_iq_i\prod_{j=1}^{i-1}(1-q_j).$$

The replacement preserves every differentiated price acceptance probability, and hence

$$\widetilde{\mathcal{R}}_{DP}\geq\widetilde{\mathcal{R}}_{DP}(\mathbf{p})=\mathcal{R}_{DP}.$$

It remains to show that uniform pricing revenue weakly decreases after the replacement. Note that, the revenue-quantile function of the triangular revenue distribution satisfies

$$R_{F_i}(0)=\mathcal{R}_{\mathrm{Tri}(a_i,p_i)}(0)=0,\qquad R_{F_i}(q_i)=\mathcal{R}_{\mathrm{Tri}(a_i,p_i)}(q_i).$$

Since F_i is regular, by the concavity of the revenue-quantile function, we have

$$\mathcal{R}_{F_i}(x)=xF_i^{-1}(1-x)\geq\mathcal{R}_{\mathrm{Tri}(a_i,p_i)}(x)=xF^{-1}_{\mathrm{Tri}(a_i,p_i)}(1-x)$$

for any $x \in (q_i, 1)$, equivalently, $F_i^{-1}(1-x) \geq F^{-1}_{\mathrm{Tri}(a_i,p_i)}(1-x)$. By the monotonicity of F_i^{-1}, we immediately obtain that, for any $\hat{p} \leq p_i$,

$$\Pr(V_i \geq \hat{p}) = 1 - F_i(\hat{p}) \geq 1 - F_{\mathrm{Tri}(a_i,p_i)}(\hat{p}) = \Pr(\widetilde{V}_i \geq \hat{p}).$$

Further, when $\hat{p} > p_i$, the replacement never accepts. Hence,

$$\Pr(\widetilde{V}_i \geq \hat{p}) \leq \Pr(V_i \geq \hat{p}) \qquad \text{for every } \hat{p}.$$

By independence, the probability that all items reject a common price is weakly larger under the replacement. Thus,

$$\widetilde{\mathcal{R}}_{UP}(p) \leq \mathcal{R}_{UP}(p) \qquad \text{for every } p,$$

and maximizing over p yields

$$\widetilde{\mathcal{R}}_{UP} \leq \mathcal{R}_{UP},$$

which completes the proof.

□

Proof of Lemma 2. For any triangular revenue instance $\mathcal{I}$ and $\varepsilon > 0$, suppose there exists a differentiated price vector $\mathbf{p} = (p_1, \ldots, p_n)$ that satisfies $\mathcal{R}_{DP}(\mathbf{p}; \mathcal{I}) \geq \mathcal{R}_{DP}(\mathcal{I}) - \frac{\varepsilon}{2}$, where, for the simplicity of notations, we use $\mathcal{R}_{DP}(\mathbf{p}; \mathcal{I})$ to denote the revenue of differentiated pricing with price $\mathbf{p}$ under $\mathcal{I}$, and $\mathcal{R}_{DP}(\mathcal{I})$ and $\mathcal{R}_{UP}(\mathcal{I})$ to denote the optimal revenue of differentiated pricing and uniform pricing under $\mathcal{I}$, respectively. In the following, we show that $\mathcal{I}$ can be transformed into the desired canonical instance, which preserves the revenue of differentiated pricing arbitrarily close and weakly decreases the revenue of uniform pricing. Some of the steps in the proof follow similar analyzes of triangular revenue distributions in Alaei et al. (2019), Jin et al. (2020). We include the details because the differentiated pricing revenue in our model is governed by the min-buying continuation recursion.

Consider an item whose valuation distribution is $V_i \sim \mathrm{Tri}(a_i, v_i)$ and is displayed at price p_i. An item priced above its cap v_i never sells and may be removed, so we assume $p_i \leq v_i$. Now, if we replace V_i by $\widehat{V}_i \sim \mathrm{Tri}(a_i, p_i)$, at the displayed price, we have

$$\Pr\left(\widehat{V}_i \geq p_i\right) = \frac{a_i}{a_i + p_i} = \Pr(V_i \geq p_i).$$

The replacement therefore preserves the item's acceptance probability at its displayed price. Since differentiated pricing revenue under min-buying depends only on the displayed prices and their acceptance probabilities, the revenue of the selected price vector is unchanged. Further, we show that the replacement also weakly reduces uniform-price demand. For $t \leq p_i$, the two distributions have the same acceptance probability: $\Pr(\widehat{V}_i \geq t) = \Pr(V_i \geq t) = \frac{a_i}{a_i + t}$. For $t > p_i$, the truncated item never accepts, whereas the original item may still accept. Hence,

$$\Pr(\widehat{V}_i \geq t) \leq \Pr(V_i \geq t) \qquad \text{for every } t.$$

By repeating this replacement for every item, we can therefore assume each item is priced at its cap under the differentiated pricing.

Next, we show that items with a common cap can be aggregated. Suppose m items share a common cap v, with parameters $a_1, a_2 \ldots, a_m$. Now, we replace them by one item with valuation distribution $\mathrm{Tri}(a, v)$, where a is chosen so that $\frac{v}{v+a} = \prod_{j=1}^{m} \frac{v}{v+a_j}$. Equivalently, $a = v\left(\prod_{j=1}^{m}\left(1+\frac{a_j}{v}\right)-1\right)$. Let z denote the continuation revenue from all more expensive items. The original instance generates continuation revenue $v\left(1-\prod_{j=1}^{m}\frac{v}{v+a_j}\right)+z\prod_{j=1}^{m}\frac{v}{v+a_j}$, which remains the same after replacement. We next verify that aggregation weakly reduces demand at every lower uniform price. For $t \le v$, the original instance and the aggregated item have rejection probabilities

$$\prod_{j=1}^{m}\frac{t}{t+a_j} \qquad \text{and} \qquad \frac{t}{t+a},$$

respectively. Let $\lambda := \frac{v}{t} \ge 1$ and $b_j := \frac{a_j}{v}$. The desired inequality $\frac{t}{t+a} \ge \prod_{j=1}^{m}\frac{t}{t+a_j}$ is equivalent to

$$1+\lambda\left(\prod_{j=1}^{m}(1+b_j)-1\right) \le \prod_{j=1}^{m}(1+\lambda b_j),$$

which follows by plugging a and rearranging. To see why this is true, expand the products on both sides. Every term of degree $d \ge 1$ is multiplied by λ on the left and by λ^d on the right. Since $\lambda^d \ge \lambda$, the inequality follows.

From now on, we assume that the caps of the triangular revenue distributions are distinct and ordered as $v_1 > v_2 > \cdots v_n$. Define the continuation revenue recursively by:

$$C_0 := 0, \qquad C_k := \frac{a_k + C_{k-1}}{1 + a_k/v_k}, \qquad k \in [n]. \tag{EC.11}$$

The quantity C_k is the revenue generated by the first k items under min-buying, where items are incorporated from the highest cap to the lowest. In particular, C_n is the revenue of the selected differentiated prices. If $C_n \le 1$, choose P sufficiently large that $\frac{P}{P+1} \ge C_n - \frac{\varepsilon}{2}$, and retain only the item $\mathrm{Tri}(1, P)$. Its uniform-pricing revenue is strictly below one, while $\frac{P}{P+1} \ge \mathcal{R}_{DP} - \varepsilon$. The result follows in this case. Henceforth, we assume that $C_n > 1$.

We then compress the expensive prefix. Following the first-crossing organization used in triangular-instance reductions for sequential posted pricing by (Jin et al. 2020), let r be the first index at which the continuation revenue reaches one:

$$C_{r-1} < 1 \qquad \text{and} \qquad C_r \ge 1. \tag{EC.12}$$

Note that the continuation revenue is

$$C_r = \frac{a_r + C_{r-1}}{1 + a_r/v_r} = \frac{a_r}{a_r + v_r} v_r + \frac{v_r}{a_r + v_r} C_{r-1},$$

which is a convex combination of $C_{r-1} < 1$ and v_r. Therefore, we have

$$1 \le C_r < v_r, \qquad \text{and} \qquad v_r > 1.$$

We replace the first r items by an ideal equal-revenue continuation of one followed by at most one boundary item with cap v_r. Define

$$\bar{a} := \frac{v_r(C_r - 1)}{v_r - C_r}. \tag{EC.13}$$

Because $1 \le C_r < v_r$, we have $\bar{a} \ge 0$. If $C_r = 1$, then $\bar{a} = 0$, and the boundary item may be omitted. We next verify that this replacement has the two desired properties. First, we show that the replacement preserves differentiated pricing revenue. By the definition of $\bar{a}$,

$$\frac{\bar{a}+1}{1+\bar{a}/v_r} = C_r. \tag{EC.14}$$

Hence, starting from a continuation equal to one and adding the boundary item $\mathrm{Tri}(\bar{a}, v_r)$ produces exactly the same continuation C_r as the original first r items. Since all remaining continuation updates depend on the prefix only through C_r, this replacement leaves the differentiated revenue generated by the remaining items unchanged.

Second, we show the replacement weakly reduces uniform-price demand. Let $Q(t) := \prod_{i=1}^{r-1} \frac{t}{t+a_i}$ be the probability that the first $r-1$ items all reject a uniform price $t \le v_r$. Their continuation revenue is C_{r-1}, so Lemma EC.1 gives:

$$Q(t) \le \frac{t}{t+C_{r-1}}, \qquad 0 < t \le v_r. \tag{EC.15}$$

We first derive a useful bound on C_r. At the uniform price v_r, the first r items reject with probability $Q(v_r)\frac{v_r}{v_r+a_r}$. Since the optimal uniform-pricing revenue is at most one,

$$v_r\left(1 - Q(v_r)\frac{v_r}{v_r+a_r}\right) \le 1.$$

Together with Eq. (EC.15), this implies that:

$$1 - \frac{1}{v_r} \le \frac{v_r}{v_r+C_{r-1}}\frac{v_r}{v_r+a_r}. \tag{EC.16}$$

Solving the continuation equation for a_r gives $a_r = \frac{v_r(C_r-C_{r-1})}{v_r-C_r}$, and therefore $\frac{v_r}{v_r+a_r} = \frac{v_r-C_r}{v_r-C_{r-1}}$. Substituting this expression into Eq. (EC.16) yields $\frac{v_r-1}{v_r} \le \frac{v_r(v_r-C_r)}{v_r^2-C_{r-1}^2}$. Rearranging gives:

$$C_r \le 1 + \frac{(v_r-1)C_{r-1}^2}{v_r^2} \le 1 + C_{r-1}. \tag{EC.17}$$

We can now compare the rejection probabilities of the original and replacement prefixes. At a uniform price $0 < t \le v_r$, the original first r items reject with probability $Q(t)\frac{t}{t+a_r}$, whereas the replacement rejects with probability $\frac{t}{t+1}\frac{t}{t+\bar{a}}$. By Eq. (EC.15), $Q(t)\frac{t}{t+a_r} \le \frac{t}{t+C_{r-1}}\frac{t}{t+a_r}$, and substituting

$$a_r = \frac{v_r(C_r-C_{r-1})}{v_r-C_r}, \qquad \bar{a} = \frac{v_r(C_r-1)}{v_r-C_r},$$

gives

$$(t+C_{r-1})(t+a_r) - (t+1)(t+\bar{a}) = \frac{1-C_{r-1}}{v_r-C_r}\left[C_r t + v_r(1+C_{r-1}-C_r)\right].$$

The right-hand side is nonnegative because

$$C_{r-1} < 1, \qquad C_r < v_r, \qquad C_r \le 1 + C_{r-1},$$

where the last inequality follows from Eq. (EC.17). Hence,

$$\frac{t}{t+C_{r-1}}\frac{t}{t+a_r} \le \frac{t}{t+1}\frac{t}{t+\bar{a}}.$$

Combining the last two inequalities,

$$Q(t)\frac{t}{t+a_r} \le \frac{t}{t+1}\frac{t}{t+\bar{a}}, \qquad 0 < t \le v_r. \tag{EC.18}$$

Thus, the replacement preserves the continuation C_r while weakly increasing the rejection probability at every relevant uniform price.

It remains to replace the ideal continuation of one by a finite triangular-revenue item. Let $V_0 \sim \mathrm{Tri}(1, P)$, $P > v_r$. Its continuation revenue is $\frac{P}{P+1}$, which converges to one as $P \to \infty$. Keeping the boundary parameter $\bar{a}$ fixed, the continuation after the boundary item becomes $\widehat{C}_r(P) := \frac{\bar{a}+\frac{P}{P+1}}{1+\bar{a}/v_r}$, so that

$$\widehat{C}_r(P) \uparrow C_r \qquad \text{as } P \to \infty. \tag{EC.19}$$

For every $t \le v_r$, the terminal item rejects with probability $t/(t+1)$, independently of P. Therefore, Eq. (EC.18) remains valid for every finite $P > v_r$.

We next remove all items whose caps are at most one. Starting from $C_r \ge 1$, every subsequent item with cap greater than one preserves a continuation value of at least one, because its continuation update is a convex combination of the incoming continuation and its cap. In contrast, if the incoming continuation is $y \ge 1$ and an item has cap $v_k \le 1$, then

$$\frac{a_k + y}{1 + a_k/v_k} - y = \frac{a_k(1 - y/v_k)}{1 + a_k/v_k} \le 0.$$

Moreover, since the updated continuation value is a convex combination of y and v_k, it remains at least $v_k \ge v_{k+1}$. Therefore, the same argument applies inductively to every subsequent item in the non-increasing-cap suffix. Hence, once the caps fall to one or below, the remaining items cannot increase differentiated revenue and may be deleted. Their deletion also weakly reduces uniform-price demand.

Let $L \ge r$ be the largest index such that $v_L > 1$, and retain the items $r+1, \ldots, L$. Let $\widehat{C}_L(P)$ denote the differentiated revenue obtained by starting from $\widehat{C}_r(P)$ and applying the retained continuation maps. Each map

$$y \longmapsto \frac{a_k + y}{1 + a_k/v_k}$$

is increasing and has derivative $\frac{1}{1+a_k/v_k} \le 1$. Consequently,

$$0 \le C_L - \widehat{C}_L(P) \le C_r - \widehat{C}_r(P) \longrightarrow 0.$$

Since deleting the low-cap tail weakly increases differentiated revenue, $C_L \ge C_n$. We may therefore choose P sufficiently large that $\widehat{C}_L(P) \ge C_L - \frac{\varepsilon}{2}$. It follows that

$$\widehat{C}_L(P) \ge C_L - \frac{\varepsilon}{2} \ge C_n - \frac{\varepsilon}{2} \ge \mathcal{R}_{DP}(\mathcal{I}) - \varepsilon.$$

Finally, we verify the uniform-pricing bound. For $t \le v_r$, the new prefix has weakly higher rejection probability than the original prefix by Eq. (EC.18). The retained cheaper items are unchanged, while deleting the low-cap tail only further increases the probability that all items reject. Thus, uniform revenue at every $t \le v_r$ is no greater than in the original instance and is therefore at most one.

For $v_r < t \le P$, all finite retained items are inactive, and only the terminal item can sell. Its revenue is

$$\frac{t}{t+1} < 1.$$

Prices above P generate zero revenue. Hence,

$$\widetilde{\mathcal{R}}_{UP} \le 1.$$

After relabeling the boundary item and the retained suffix, the canonical instance consists of the terminal item $\mathrm{Tri}(1,P)$ and finite items

$$\mathrm{Tri}(a_k, v_k), \qquad P > v_1 > \cdots > v_N > 1.$$

If C_k denotes the continuation revenue in this canonical instance, then

$$C_0 = \frac{P}{P+1}, \qquad C_k = \frac{a_k + C_{k-1}}{1 + a_k/v_k}.$$

Writing $x_k = 1/v_k$ gives

$$C_k = \frac{a_k + C_{k-1}}{1 + a_k x_k},$$

which is the desired canonical representation. □

Proof of Lemma 3. Consider a finite canonical instance satisfying $\mathcal{R}_{UP} \le 1$. Let $V_0 = \mathrm{Tri}(1,P)$ denote the terminal item and $V_i \sim \mathrm{Tri}(a_i, v_i)$ denote the remaining items, where $P > v_1 > v_2 > \cdots > v_n > 1$. For simplicity of notation, we define $x_i := \frac{1}{v_i}$. Therefore, we have $0 < x_1 < x_2 < \cdots < x_n < 1$. Let C_k denote the continuation revenue after incorporating the first k items. Then, by Lemma 2, we have $C_0 = \frac{P}{P+1} < 1$ and $C_k = \frac{a_k + C_{k-1}}{1 + a_k x_k}$ for $k \in [n]$. In the following, we show that the revenue of cap priced canonical instance is upper bounded by $\Gamma^\star$.

We first use the uniform-pricing restriction to bound the cumulative triangular-revenue mass. For $k \in [n]$, consider the uniform price $p = v_k$. At price p, the terminal item is rejected with probability $p/(p+1)$, and exactly the first k finite items are active. Since item $i \le k$ is rejected with probability $p/(p+a_i)$, the normalization $\mathcal{R}_{UP} \le 1$ implies

$$p\left(1 - \frac{p}{p+1}\prod_{i=1}^{k}\frac{p}{p+a_i}\right) \le 1.$$

Rearranging the inequality and taking logarithms for both sides gives

$$\sum_{i=1}^{k} \log\left(1 + \frac{a_i}{p}\right) \le -\log(1 - p^{-2}).$$

Plugging $p = 1/x_k$ in, we obtain

$$\sum_{i=1}^{k} \log(1 + a_i x_k) \le -\log(1 - x_k^2). \tag{EC.20}$$

For every $a \ge 0$, the concavity of $\log(1+ax)$ and the identity $\log(1 + a\cdot 0) = 0$ imply $\log(1+ax) \ge x\log(1+a)$ when $x \in [0,1]$. Applying this inequality term by term in Eq. (EC.20) and rearranging yields

$$\sum_{i=1}^{k} \log(1 + a_i) \le -\frac{\log(1 - x_k^2)}{x_k} = A(x_k).$$

For simplicity, let $s_0 := 0$ and $s_k := \sum_{i=1}^{k} \log(1+a_i)$ for $k \in [n]$. Because $A(x)$ is strictly increasing on $[0,1)$, define the budget-binding reciprocal caps $\widehat{x}_k := A^{-1}(s_k)$ for $k \in [n]$. In particular, $\widehat{x}_0 = 0$. By Eq. (EC.20) and the monotonicity of $A(x)$, we have $\widehat{x}_k \le x_k$. After discarding any item with $a_k = 0$, we may assume $0 = \widehat{x}_0 < \widehat{x}_1 < \cdots < \widehat{x}_n < 1$.

Now we define a relaxed continuation recursion by

$$U_0 := 1, \qquad U_k := \frac{a_k + U_{k-1}}{1 + a_k \widehat{x}_k}, \qquad k \in [n]. \tag{EC.21}$$

Note that the update function $T_{a,x}(z) := \frac{a+z}{1+ax}$ is increasing in z and decreasing in x. Since $C_0 = \frac{P}{P+1} < 1 = U_0$ and $\widehat{x}_k \le x_k$, the induction gives $C_k = T_{a_k,x_k}(C_{k-1}) \le T_{a_k,\widehat{x}_k}(U_{k-1}) = U_k$ for any $k \in [n]$. In the recursion of U_k, the uniform-price budget binds exactly, i.e., $\log(1+a_k) = s_k - s_{k-1} = A(\widehat{x}_k) - A(\widehat{x}_{k-1})$. Therefore, the choice of a_k satisfies

$$a_k = \exp(A(\widehat{x}_k) - A(\widehat{x}_{k-1})) - 1.$$

We next show by induction that $U_k \le c(\widehat{x}_k)$. Recall that $c(x)$ is the solution for

$$c'(x) = A'(x)\big(1 - xc(x)\big), \qquad c(0) = 1.$$

For $k = 0$, by definition we have $U_0 = 1 = c(0) = c(\widehat{x}_0)$. Suppose $U_{k-1} \le c(\widehat{x}_{k-1})$. Since $T_{a,x}(z)$ is increasing in z, $U_k = T_{a_k,\widehat{x}_k}(U_{k-1}) \le T_{a_k,\widehat{x}_k}(c(\widehat{x}_{k-1}))$. Because $a_k = \exp(A(\widehat{x}_k) - A(\widehat{x}_{k-1})) - 1$, we can then apply the result in Lemma EC.2 with $y = \widehat{x}_{k-1}$, $x = \widehat{x}_k$ and $z = c(\widehat{x}_{k-1})$, since by Lemma EC.3 we have $c(x) \ge 1$ for all $x \in (0,1)$. Therefore, $U_k \le T_{a_k,\widehat{x}_k}(c(\widehat{x}_{k-1})) \le c(\widehat{x}_k)$.

Furthermore, we prove that the maximum of $c(x)$ is attained at the unique $x^\star$ satisfying $x^\star c(x^\star) = 1$ in Lemma EC.3. Therefore, combining all the results above, we obtain

$$\mathcal{R}_{DP}(\mathbf{v}) = C_n \le U_n \le c(\widehat{x}_n) \le \max_{x \in [0,1)} c(x) \le \Gamma^\star,$$

which completes the proof. □

Proof of Lemma 4 Fix $X \in (0,1)$. We construct a sequence of finite canonical instances whose uniform-pricing revenue is exactly one and whose differentiated pricing revenue converges to $c(X)$. The result then follows by taking $X = x^\star$.

For $m \ge 1$, let $h := \frac{A(X)}{m}$. For each $k \in [m]$, define $G_k(x) := (1-x^2)(1+hx)^k$, for $x \in [0,1)$, and let $x_k \in (0,1)$ be the unique positive solution to

$$(1 - x_k^2)(1 + hx_k)^k = 1. \tag{EC.22}$$

The role of this equation will become clear below: it makes the uniform-pricing constraint bind exactly at the k-th item boundary.

We first verify that the roots are well defined and ordered. The logarithmic derivative of $G_k(x)$ is

$$\frac{d}{dx} \log G_k(x) = -\frac{2x}{1-x^2} + \frac{kh}{1+hx},$$

and its derivative is strictly negative on $(0,1)$. Thus, $G_k(x)$ is strictly unimodal. Since

$$G_k(0) = 1, \qquad G_k'(0) = kh > 0, \qquad \lim_{x \uparrow 1} G_k(x) = 0,$$

the function first rises above one and then decreases below one. Therefore, it has a unique positive root x_k. Moreover, for every $x>0$,

$$G_{k+1}(x)=G_k(x)(1+hx)>G_k(x).$$

In particular, $G_{k+1}(x_k)>1$, so the positive root of G_{k+1} lies strictly to the right of that of G_k. Hence,

$$0<x_1<x_2<\cdots<x_m<1.$$

Set $v_k:=\frac{1}{x_k}$ for $k\in[m]$, therefore $v_1>v_2\cdots v_m>1$. For some $P>v_1$, consider the canonical instance consisting of $V_0\sim\mathrm{Tri}(1,P)$ and $V_k\sim\mathrm{Tri}(h,v_k)$ for $k\in[m]$. We first show that its optimal uniform-pricing revenue is one. Consider a uniform price $p\in[1,v_1]$, let $x:=1/p$, and suppose exactly the first k finite items are active. The terminal item rejects with probability $1/(1+x)$, and each active item rejects with probability $1/(1+hx)$. Hence,

$$R_{UP}(p)=\frac{1}{x}\left(1-\frac{1}{(1+x)(1+hx)^k}\right).$$

A direct rearrangement gives

$$R_{UP}(p)\le 1 \quad\Longleftrightarrow\quad (1-x^2)(1+hx)^k\le 1.$$

Because exactly k items are active, $x\ge x_k$. By the definition of x_k and the fact that $G_k(x)$ is below one to the right of its positive root, $R_{UP}(p)\le 1$. At the boundary price $p=v_k$, we have $x=x_k$, so Eq. (EC.22) gives $R_{UP}(v_k)=1$. For $v_1<p\le P$, only the terminal item is active and its revenue is $p/(p+1)<1$; prices below one earn at most one, and prices above P earn zero. Therefore, for the constructed canonical instance, we always have

$$\mathcal{R}_{UP}=1. \tag{EC.23}$$

We next consider differentiated pricing and price every item at its cap. Let $C_k^{(m,P)}$ denote the continuation revenue after the terminal item and the first k items have been incorporated. Then

$$C_0^{(m,P)}=\frac{P}{P+1},\qquad C_k^{(m,P)}=\frac{h+C_{k-1}^{(m,P)}}{1+hx_k}.$$

Let $C_k^{(m)}$ denote the same recursion with $C_0^{(m)}=1$. Since each continuation function is increasing and 1-Lipschitz, $0\le C_m^{(m)}-C_m^{(m,P)}\le\frac{1}{P+1}$. Thus, for each fixed m, $C_m^{(m,P)}\longrightarrow C_m^{(m)}$ as $P\to\infty$. By Lemma EC.4,

$$C_m^{(m)}\longrightarrow c(X)\qquad\text{as } m\to\infty.$$

Now choose $X=x^\star$. By Lemma EC.3, $c(x^\star)=\Gamma^\star$. Hence, for any $\varepsilon>0$, we may first choose m sufficiently large and then $P>v_1$ sufficiently large so that $C_m^{(m,P)}>\Gamma^\star-\varepsilon$. Since cap pricing is a feasible differentiated price vector,

$$\mathcal{R}_{DP}\ge C_m^{(m,P)}>\Gamma^\star-\varepsilon.$$

Combining this with Eq. (EC.23) proves there exist a canonical instance such that

$$\mathcal{R}_{UP}=1,\qquad \mathcal{R}_{DP}>\Gamma^\star-\varepsilon.$$

□

B.3. Additional Omitted Proofs

LEMMA EC.1 **(Prefix Rejection Bound)**. *Consider an ordered collection of independent triangular-revenue valuations*

$$V_i \sim \mathit{Tri}(a_i, v_i), \qquad i = 1, \dots, m,$$

whose caps satisfy $v_i \geq t$ for some $t > 0$. Suppose each item is priced at its cap, and define the continuation revenue recursively by

$$Z_0 := 0, \qquad Z_i := \frac{a_i + Z_{i-1}}{1 + a_i/v_i}, \qquad i = 1, \dots, m.$$

Let

$$Q_m(t) := \prod_{i=1}^{m} \frac{t}{t + a_i}$$

be the probability that every item rejects the uniform price t. Then

$$Q_m(t) \leq \frac{t}{t + Z_m}. \tag{EC.24}$$

Proof of Lemma EC.1 We proceed by induction on the number of items. The result is immediate for $m = 0$, because $Z_0 = 0$ and $Q_0(t) = 1$. Suppose the result holds for the first $m-1$ items. After adding the final item, the continuation revenue becomes $Z_m = \frac{a_m + Z_{m-1}}{1 + a_m/v_m}$, whereas the rejection probability becomes $Q_m(t) = Q_{m-1}(t)\frac{t}{t+a_m}$. By the induction hypothesis, $Q_m(t) \leq \frac{t}{t+Z_{m-1}}\frac{t}{t+a_m}$. It therefore remains to show that

$$\frac{t}{t + Z_{m-1}} \frac{t}{t + a_m} \leq \frac{t}{t + Z_m}.$$

Because all denominators are positive, this inequality is equivalent to $t(t + Z_m) \leq (t + Z_{m-1})(t + a_m)$. Using the definition of Z_m, the difference between the right- and left-hand sides is

$$\begin{aligned}(t + Z_{m-1})(t + a_m) - t(t + Z_m) &= (t + Z_{m-1})(t + a_m) - t\left(t + \frac{a_m + Z_{m-1}}{1 + a_m/v_m}\right) \\ &= \frac{a_m\big(a_m t + a_m Z_{m-1} + t Z_{m-1} + v_m Z_{m-1}\big)}{a_m + v_m} \geq 0.\end{aligned}$$

Hence, $Q_m(t) \leq \frac{t}{t+Z_m}$, which completes the induction. □

LEMMA EC.2 **(Continuous Envelope Dominance)**. *Suppose $0 \leq y < 1$ and the continuation value at reciprocal cap y is $z \geq 1$. Let $c : [y, 1) \to \mathbb{R}$ be the unique solution to*

$$c'(s) = A'(s)\big(1 - s c(s)\big), \qquad c(y) = z.$$

For any $x \in [y, 1)$, define

$$a(x) := \exp\big(A(x) - A(y)\big) - 1.$$

Let $d(x)$ be the continuation value when incorporating another item with a valuation distribution $\mathit{Tri}(a(x), 1/x)$ where $y \leq x < 1$. Specifically, $d(x)$ is given by

$$d(x) := \frac{a(x) + z}{1 + x a(x)}.$$

Then, the discrete continuation update $d(x)$ is upper-bounded by the continuous envelope $c(x)$, namely,

$$d(x) \leq c(x) \qquad \text{for every } x \in [y, 1). \tag{EC.25}$$

Proof of Lemma EC.2. In this proof, we compare the discrete continuation value $d(s)$ with the continuous envelope $c(s)$ as the reciprocal cap x increases from y. Specifically, we show that an appropriately weighted difference between them is nondecreasing, which implies that the continuous envelope remains weakly above the discrete continuation.

For $x \in [y,1)$, the definition of $a(x)$ gives

$$a'(x) = A'(x)\big(1+a(x)\big).$$

Note that the direct derivative of $c(x)-d(x)$ contains a linear term proportional to this difference. We absorb the term by defining $\mu(x) := \exp\left(\int_y^x tA'(t)\,dt\right)$ and the weighted gap

$$h(x) := \mu(x)\big(c(x)-d(x)\big).$$

Since $\mu(x) \geq 0$ for any $x \in [y,1)$, to show $d(x)$ is always upper bounded by $c(x)$, we only need to show $h(x) \geq 0$. By the definition of $a(x)$, we immediately obtain $a(y) = \exp\left(A(y)-A(y)\right)-1 = 0$, and hence $h(y) = 0$.

We next show that $h(x)$ is nondecreasing from y. Since $\mu'(x) = xA'(x)\mu(x)$ and $c'(x) = A'(x)\big(1-xc(x)\big)$, taking the derivative of $h(x)$ yields

$$\begin{aligned} h'(x) &= \mu(x)\left(c'(x)-d'(x)\right) + \mu'(x)\left(c(x)-d(x)\right) \\ &= \mu(x)\left[c'(x)-d'(x)+xA'(x)\big(c(x)-d(x)\big)\right] \\ &= \mu(x)\left[A'(x)\big(1-xd(x)\big)-d'(x)\right] \end{aligned} \tag{EC.26}$$

where the second equality follows from plugging $\mu'(x)$ and the third equality follows from simplifying. Substituting the definition of $d(x)$ and its derivative into Eq. (EC.26) and simplifying gives us

$$h'(x) = \mu(x)\frac{a(x)}{\left(1+xa(x)\right)^2}\left(a(x)+z-A'(x)(1-x)(1-zx)\right). \tag{EC.27}$$

It remains to verify that the right-hand side of Eq. (EC.27) is nonnegative. We first establish an auxiliary bound

$$A'(x)(1-x)^2 \leq 1, \qquad 0<x<1. \tag{EC.28}$$

Recall that $A(x) = -\frac{\log(1-x^2)}{x}$, the derivative of $A(x)$ is

$$A'(x) = \frac{2}{1-x^2} + \frac{\log(1-x^2)}{x^2} \leq \frac{2}{1-x^2} - 1.$$

where the inequality follows since $\log(1-x^2) \leq -x^2$ for any $x \in (0,1)$. Consequently, multiplying by $(1-x)^2$ yields

$$A'(x)(1-x)^2 \leq (1-x)^2\left(\frac{2}{1-x^2}-1\right) = \frac{(1-x)(1+x^2)}{1+x} \leq \frac{(1-x)(1+x)}{1+x} \leq 1,$$

where the first equality follows by simplifying and the second inequality follows since $1+x^2 \leq 1+x$ for any $x \in (0,1)$. This proves Eq. (EC.28).

Further, because $z \geq 1$, we have $1-zx \leq 1-x$ for $x \in (0,1)$. Note that the power-series representation $A(x)$ is $A(x) = -\frac{\log(1-x^2)}{x} = \sum_{i=1}^{\infty}\frac{x^{2m-1}}{m}$, which implies $A'(x) \geq 0$. Thus, we obtain

$$A'(x)(1-x)(1-zx) \leq A'(x)(1-x)^2 \leq 1 \leq a(x)+z,$$

where the last inequality follows since $z \geq 1$ and $a(x) \geq 0$. Since by definition $\mu(x) > 0$, $a(x) \geq 0$, $1 + xa(x) > 0$, we conclude that $h'(x) \geq 0$, for every $x \in [y, 1)$. Therefore, h is nondecreasing and $h(y) = 0$, which implies

$$h(x) \geq 0 \qquad \text{for every } x \in [y, 1).$$

Because $\mu(x) > 0$, this is equivalent to

$$d(x) \leq c(x) \qquad \text{for every } x \in [y, 1),$$

as claimed. □

LEMMA EC.3 **(Maximum of the Continuous Envelope).** *Let $c : [0, 1) \to \mathbb{R}$ be the unique solution to*

$$c'(x) = A'(x)\big(1 - xc(x)\big), \qquad c(0) = 1.$$

Then there exists a unique $x^\star \in (0, 1)$ satisfying

$$x^\star c(x^\star) = 1. \tag{EC.29}$$

Moreover, $c(x)$ is strictly increasing on $[0, x^\star]$ and strictly decreasing on $([x^\star, 1)$. Consequently,

$$\max_{0 \leq x < 1} c(x) = c(x^\star) = \frac{1}{x^\star} = \Gamma^\star. \tag{EC.30}$$

Proof of Lemma EC.3 Note that $c'(x) = A'(x)\big(1 - xc(x)\big)$ and we have shown $A'(x) > 0$ for $x \in (0, 1)$ in the proof of Lemma EC.2. The essence of the proof is therefore to show $1 - xc(x)$ crosses zero exactly once. For simplicity, define $g(x) := 1 - xc(x)$.

We first prove the existence of the crossing point x^*. Suppose that $g(x)$ has no zero on $(0, 1)$. Since $g(0) = 1$ and g is continuous, this would imply

$$g(x) > 0 \qquad \text{for every } x \in [0, 1).$$

Consequently, $c'(x) > 0$ for every $x \in (0, 1)$, so $c(x)$ is strictly increasing. Fix any $x_0 \in (0, 1)$. Since $c(0) = 1$, the strict monotonicity implies $c(x_0) > 1$. We can therefore choose $x \in (x_0, 1)$ such that $x > \frac{1}{c(x_0)}$. Because $c(x)$ is increasing, we have

$$xc(x) \geq xc(x_0) > 1,$$

which contradicts with the assumption $g(x) = 1 - xc(x) > 0$. Hence, there exists $x^\star \in (0, 1)$ such that $g(x^\star) = 0$, or equivalently, $x^\star c(x^\star) = 1$.

We next prove the uniqueness of the crossing point x^*. Taking the derivative of $g(x)$ gives

$$g'(x) = -c(x) - xc'(x).$$

At any zero point x^* of $g(x)$, we have $g'(x^*) = -c(x^*) = -\frac{1}{x^*} < 0$. Thus, every zero of $g(x)$ is a strict crossing from positive to negative, which implies the crossing point x^* is unique since $g(x)$ is continuous. If x_1 were the first zero, then $g(x) < 0$ just to its right. If another zero $x_2 > x_1$ exists, take the first such zero. Continuity of $g(x)$ gives $g(x) \leq 0$ immediately to its left, whereas $g'(x_2) < 0$ implies $g(x) > 0$ immediately to its left. Therefore, it contradicts the definition of x_2. Hence, $x^\star$ is the unique zero of $g(x)$. It follows that

$$g(x) > 0 \quad \text{for } 0 \leq x < x^\star, \qquad g(x) < 0 \quad \text{for } x^\star < x < 1.$$

Consequently, the function $c(x)$ is strictly increasing before $x^\star$ and strictly decreasing afterward. Since $c(x)$ is increasing on $[0, x^\star]$ and $c(0) = 1$, we obtain $c(x) \geq 1$ for $x \in [0, x^\star]$. For any $x > x^\star$, we have $g(x) < 0$ which implies $xc(x) > 1$. Hence, $c(x) \geq 1$ for all $x \in [0, 1]$. Further, $c(x)$ attains its unique global maximum at $x^\star$. Finally, the first order condition $x^\star c(x^\star) = 1$ implies

$$\max_{0 \leq x < 1} c(x) = c(x^\star) = \frac{1}{x^\star} = \Gamma^\star,$$

which completes the proof. □

LEMMA EC.4 **(Convergence of the Canonical Recursion)**. *Fix $X \in (0,1)$. For each $m \geq 1$, let $h := \frac{A(X)}{m}$, and, for each $k \in [m]$, let $x_k \in (0,1)$ be the unique positive solution to*

$$(1 - x_k^2)(1 + hx_k)^k = 1. \tag{EC.31}$$

Define the discrete continuation sequence by

$$C_0^{(m)} := 1, \qquad C_k^{(m)} := \frac{h + C_{k-1}^{(m)}}{1 + hx_k}, \qquad k \in [m].$$

Then

$$C_m^{(m)} \longrightarrow c(X) \qquad \text{as } m \to \infty,$$

where c is the continuous envelope.

Proof of Lemma EC.4 Let $S := A(X)$ and $s_k := kh$ for $k \in [m]$. We show that the discrete continuation recursion is a first-order approximation to the continuous envelope after the change of variables $s = A(x)$.

Taking logarithms in Eq. (EC.31) and dividing both sides by x_k gives

$$A(x_k) = k \frac{\log(1 + hx_k)}{x_k}. \tag{EC.32}$$

Using the fact that $u - \frac{u^2}{2} \leq \log(1+u) \leq u$ in Eq. (EC.32), we obtain

$$s_k - A(x_k) = kh - k\frac{\log(1+hx_k)}{x_k} \geq 0, \quad \text{and} \quad s_k - A(x_k) = \frac{k(h - \log(1+hx_k))}{x_k} \leq \frac{kh^2 x_k}{2} \leq \frac{SX}{2}h.$$

Hence, we have

$$\max_{k \in [m]} |A(x_k) - s_k| = O(h). \tag{EC.33}$$

Because $A'(x)$ is continuous and strictly positive on $[0, X]$, $A^{-1}(x)$ is Lipschitz on $[0, S]$. Therefore,

$$x_k = A^{-1}(s_k) + O(h), \tag{EC.34}$$

uniformly over $k \in [m]$.

We now rewrite the continuous envelope. Define $y(s) := c(A^{-1}(s))$, $0 \leq s \leq S$. By the chain rule and the defining ODE for $c(x)$,

$$y'(s) = 1 - A^{-1}(s)y(s), \qquad y(0) = 1. \tag{EC.35}$$

Next, rearrange the discrete continuation recursion as

$$C_k^{(m)} - C_{k-1}^{(m)} = h\frac{1 - x_k C_{k-1}^{(m)}}{1 + hx_k}. \tag{EC.36}$$

The continuation sequence is uniformly bounded. Indeed, $C_k^{(m)} \le C_{k-1}^{(m)} + h$, and hence $0 < C_k^{(m)} \le 1 + kh \le 1 + S$. Combining this bound with Eq. (EC.34) and $x_k \le X$, we obtain $\frac{1-x_k C_{k-1}^{(m)}}{1+hx_k} = 1 - A^{-1}(s_k)C_{k-1}^{(m)} + O(h)$, uniformly in k. Thus,

$$C_k^{(m)} = C_{k-1}^{(m)} + h\left(1 - A^{-1}(s_k)C_{k-1}^{(m)}\right) + O(h^2). \tag{EC.37}$$

Likewise, integrating Eq. (EC.35) over $[s_{k-1}, s_k]$ gives

$$y(s_k) = y(s_{k-1}) + h\left(1 - A^{-1}(s_k)y(s_{k-1})\right) + O(h^2), \tag{EC.38}$$

where the error term is uniform in k. Let $e_k := C_k^{(m)} - y(s_k)$. Subtracting Eq. (EC.38) from Eq. (EC.37), and using $A^{-1}(s_k) \le X$, yields

$$|e_k| \le (1 + Xh)|e_{k-1}| + Kh^2$$

for some constant $K > 0$ independent of m and k. Since $e_0 = 0$, iteration gives

$$|e_m| \le Kh^2 \sum_{j=0}^{m-1} (1 + Xh)^j \le KSh\, e^{XS} \longrightarrow 0,$$

because $h = S/m \to 0$ as m increases.

Finally, $s_m = mh = S = A(X)$, and therefore $y(s_m) = c(A^{-1}(A(X))) = c(X)$. Hence,

$$C_m^{(m)} = y(s_m) + e_m \longrightarrow c(X),$$

as claimed. □

LEMMA EC.5 **(Continuous Approximation of Triangular Instances)**. *For every finite triangular-revenue instance $\mathcal{I}$, let $\mathbf{v}$ denote its cap-price vector. If $\mathcal{R}_{UP}(\mathcal{I}) \le 1$, then for sufficiently small $\delta > 0$, there exists an independent continuous regular instance $\mathcal{I}^\delta$ such that*

$$\mathcal{R}_{DP}(\mathbf{v}; \mathcal{I}^\delta) \ge \mathcal{R}_{DP}(\mathbf{v}; \mathcal{I}) \qquad \text{and} \qquad \mathcal{R}_{UP}(\mathcal{I}^\delta) \le \frac{1}{1-\delta}.$$

Proof of Lemma EC.5 Consider an item with $V_i \sim \mathrm{Tri}(a_i, v_i)$, and let $q_i = \frac{a_i}{a_i + v_i}$ be its acceptance probability at the cap. By the definition of $\mathrm{Tri}(a_i, v_i)$, its revenue-quantile curve is

$$\mathcal{R}_{\mathrm{Tri}(a_i,v_i)}(q) = \begin{cases} v_i q, & 0 \le q \le q_i, \\ a_i(1-q), & q_i \le q \le 1. \end{cases}$$

Now, for $\delta > 0$ small enough that $\delta \le \min\{a_i + v_i, 1\}$, we consider a similar revenue function

$$\mathcal{R}_i^\delta(q) = \begin{cases} v_i q + \delta q\left(1 - \frac{q}{q_i}\right), & 0 \le q \le q_i, \\ a_i(1-q), & q_i \le q \le 1. \end{cases}$$

Note that the left derivative of $\mathcal{R}_i^\delta(q)$ at q_i equals $v_i - \delta$, whereas the right derivative $\mathcal{R}_i^\delta(q)$ at q_i equals $-a_i$. The first branch $\mathcal{R}_i^\delta(q)$ $(q \le q_i)$ is strictly concave and the second is linear, thus $\mathcal{R}_i^\delta(q)$ is concave. Hence, the induced distribution F_i^δ is regular. Its upper quantile function is

$$Q_i^\delta(q) = \frac{\mathcal{R}_i^\delta(q)}{q} = \begin{cases} v_i + \delta\left(1 - \frac{q}{q_i}\right), & 0 < q \le q_i, \\ \frac{a_i(1-q)}{q}, & q_i \le q \le 1. \end{cases}$$

Inverting $Q_i^\delta(q)$ gives us,

$$F_i^\delta(x) = \begin{cases} \frac{x}{a_i+x}, & 0 \le x \le v_i, \\ 1 - \frac{q_i}{\delta}(v_i + \delta - x), & v_i \le x \le v_i + \delta \\ 1, & x \ge v_i + \delta. \end{cases}$$

Therefore, F_i^δ is continuous and regular.

Moreover, $Q_i^\delta(q_i) = v_i$, so $\Pr(V_i^\delta \ge v_i) = q_i$ when $V_i^\delta \sim F_i^\delta(x)$. Therefore, pricing every item at its original cap preserves every itemwise acceptance probability. Since independence and the ordering of the caps are unchanged, the revenue of the cap-pricing is preserved exactly. Hence,

$$\mathcal{R}_{DP}(\mathbf{v}; \mathcal{I}^\delta) \ge \mathcal{R}_{DP}(\mathbf{v}; \mathcal{I}).$$

Now, the only thing remaining is to bound the revenue of uniform pricing. By the construction of $F_i^\delta(x)$, we have

$$F_i^\delta(p) \ge F_{\mathrm{Tri}(a_i, v_i)}(p - \delta), \qquad p \ge \delta.$$

This follows directly from the explicit form of $F_i^\delta(x)$: smoothing shifts each valuation upward by at most δ. Hence, by independence, we can easily obtain

$$\prod_i F_i^\delta(p) \ge \prod_i F_{\mathrm{Tri}(a_i, v_i)}(p - \delta).$$

Therefore, for any $p > 1$,

$$\begin{aligned} \mathcal{R}_{UP}(p; \mathcal{I}^\delta) &= p\left(1 - \prod_i F_i^\delta(p)\right) \\ &\le p\left(1 - \prod_i F_{\mathrm{Tri}(a_i, v_i)}(p - \delta)\right) \\ &= \frac{p}{p - \delta} \mathcal{R}_{UP}(p - \delta; \mathcal{I}) \\ &\le \frac{1}{1 - \delta}, \end{aligned}$$

where the last inequality follows since $\mathcal{R}_{UP}(p; \mathcal{I}) \le 1$ and $\frac{p}{p-\delta} \le \frac{1}{1-\delta}$ for any $p > 1$ and δ small enough. For the case $p \le 1$, $\mathcal{R}_{UP}(p; \mathcal{I}^\delta) \le p \le 1$ is automatically true. □