Algorithmic Price Discrimination

Rachel Cummings * Nikhil R. Devanur † Zhiyi Huang ‡ Xiangning Wang §

Abstract

We consider a generalization of the third degree price discrimination problem studied in [4](Bergemann et al., 2015), where an intermediary between the buyer and the seller can design market segments to maximize any linear combination of consumer surplus and seller revenue. Unlike in [4], we assume that the intermediary only has partial information about the buyer's value. We consider three different models of information, with increasing order of difficulty. In the first model, we assume that the intermediary's information allows him to construct a probability distribution of the buyer's value. Next we consider the sample complexity model, where we assume that the intermediary only sees samples from this distribution. Finally, we consider a bandit online learning model, where the intermediary can only observe past purchasing decisions of the buyer, rather than her exact value. For each of these models, we present algorithms to compute optimal or near optimal market segmentation.

1 Introduction

Third degree price discrimination occurs when a seller uses auxiliary information about buyers to offer different prices to different populations, e.g., student and senior discounts for movie tickets. A modern version of this arises in the context of online platforms that match sellers and buyers. Here an intermediary observes information about buyers and may pass on some of this information to the seller to help him price discriminate. One natural example where price discrimination could be (and often is) used in practice is an *ad exchange*, which matches buyers and sellers of online ad impressions. A buyer is an advertiser, and a seller is a publisher, and the impression is sold via an auction where the seller sets a reserve price. The ad exchange commonly has additional data about the user viewing the impression or about the buyers. It could share some of this data with the seller before he sets the reserve price.

The seminal work of [4] shows the following surprising result in such a setting. Usually, there is a tradeoff between *social welfare* which is the value generated by the sale, and seller *revenue*. Seller revenue is maximized by setting an appropriate price. Social welfare is maximized by selling the item to the buyer as long as his value for the item is ≥ 0, but this generates 0 revenue for the seller. Almost magically, [4] show that an intermediary can *segment the market* such that it not only maximizes social welfare, but also guarantees that the seller revenue doesn't change in the process. This shows that price discrimination can be used to benefit the customer, contrary to the belief that it exploits the customer, thus making it palatable.

While this is a strong result, it requires that the intermediary knows the buyer's exact value, which is a very strong assumption, and is often not satisfied in practice. What is more reasonable is that the intermediary can estimate a personalized probability distribution once the buyer is seen. For instance, if the intermediary observes that the buyer is a student, it may estimate a lower willingness-to-pay, but is unlikely to know the buyer's exact value. Realistically, the intermediary may wish to use machine learning techniques to estimate the personalized probability distribution for a new buyer based upon their observed characteristics and past market data. *In this paper, we analyze the power of third degree price discrimination in this setting where the intermediary has only a noisy signal of a buyer's value.*

1.1 Model and Results

The seller sells a single item, and there is a single buyer. We consider value distributions with a finite support. We assume that the intermediary observes finitely many types of buyers; each type is associated with a different distribution over the values. We denote the set of values by $[V] = \{1, 2, \ldots, V\}$, the set of types by $[T] = \{1, 2, \ldots, T\}$, and the distribution over values given a type t by $\mathcal{F}(t)$. We denote the distribution over types by $\mathcal{T}$. Given this, the mechanism proceeds as follows. This is illustrated in Figure 1.

1. A segmentation is a pair of a segment set Σ, and a segment map $\mathcal{G} : [T] \mapsto \Delta(\Sigma)$, where $\Delta(\cdot)$ denotes the set of all probability distributions over

*Georgia Institute of Technology. Email: rachelc@gatech.edu. R.C. supported in part by a Mozilla Research Grant, a Google Research Fellowship, and NSF grant CNS-1850187. Part of this work was completed while R.C. was visiting the Simons Institute for the Theory of Computing.

†Microsoft Research. Email: nikdev@microsoft.com

‡The University of Hong Kong. Email: zhiyi@cs.hku.hk. This work is supported in part by an RGC grant HKU17203717E.

§The University of Hong Kong. Email: xnwang@cs.hku.hk

a given domain. Once the intermediary decides on a segmentation, it is revealed to the seller.

2. When a buyer arrives, her type t and value v are drawn from the prior distribution. The intermediary observes only her type t but not the value v.

3. Intermediary draws a segment σ from the distribution $\mathcal{G}(t)$ and reveals it to the seller.

4. On observing a segment σ, the seller posts the monopoly price p for the value distribution conditioned on observing σ.

5. Buyer buys the item if and only if her value $v \geq p$.

The model in [4] is the special case where the type set is identical to the value set, and the distributions $\mathcal{F}(t)$ are point masses.

A price is said to be a *monopoly price* if it maximizes seller revenue for a given distribution, and this revenue is called the *monopoly revenue*. We call the marginal distribution over values $[V]$ as the *prior* distribution. The seller is always guaranteed at least the monopoly revenue for the prior distribution, since he can ignore the segment information and set a monopoly price.

The intermediary's objective is some given positive linear combination of seller revenue and *consumer surplus*. Consumer surplus is the expectation of the buyer's utility, which is $v - p$ if the buyer with value v buys the item at price p, and is 0 otherwise. Of particular interest is the special case of maximizing consumer surplus alone. We consider three informational models of increasing difficulty for the intermediary, and show the following results.

Bayesian: The intermediary and the seller know the value-type distributions: $\mathcal{F}(t)$ for all t, and $\mathcal{T}$. We show that the optimal segmentation can be computed using a linear program (LP). The range of achievable values for consumer surplus and revenue depend on the distribution, and one may not always be able to achieve the full consumer surplus as in [4]. Some other nice properties may not hold as well, see Appendix A for examples.

Sample Complexity: The intermediary and the seller observe a batch of signal-value pairs sampled from the underlying distribution. We are interested in the number of samples required to get an ϵ approximation. We first cosntruct a distribution for which no bounded function of ϵ is sufficient. The $\mathcal{F}(t)$'s in this example satisfy both boundedness and regularity, which are standard assumptions in the sample complexity of mechanism design. This points to further limitations on what such an intermediary can do: in case of noisy signals, the distribution cannot be arbitrary. Motivated by this, we identify a property about the distributions, which we call MHR-like, and show (via an algorithmic construction) that a polynomial number of samples are sufficient. This is the technically most challenging part of the paper and most of the focus in the main body of the paper is on this part.

Online Learning: The intermediary must learn the segmentation online using only bandit feedback from the buyer's decision to purchase or not at the seller's chosen price. The last step of the timeline depicted in Figure 1 is modified in this setting so that the intermediary and seller only observe the buyer's purchase decision, not her value. Here we give *no-regret* learning algorithms. Clearly, we need certain assumptions on the seller's behavior for any nontrivial result; there is not much we can do if the seller picks prices randomly all the time. Our assumptions can accommodate natural no regret learning algorithms on the seller side, including the Upper-Confidence-Bound (UCB) algorithm and the Explore-then-Commit (ETC) algorithm.

1.2 Contributions to the Sample Complexity of Mechanism Design

Pioneered by [1], [18], and [15], and formalized by [12], the sample complexity of mechanism design, in particular, the revenue maximization problem, has been a focal point in algorithmic game theory in the last few years [33, 2, 14, 32, 24, 10, 20, 21, 26, 22].

This paper adds to the literature of sample complexity of mechanism design in two-folds. The first one is conceptual: we formulate the first sample complexity problem from the viewpoint of an *intermediary* rather than the seller, and for the task of designing *information dispersion* rather than allocations and payments. We show impossibility results for the general case and, more importantly, identify sufficient conditions under which we derive positive algorithmic results.

Conceptually new models often lead to new technical challenges. Our second contribution is an algorithmic ingredient that tackles such a new challenge. Let us start with a thought experiment: consider a more powerful intermediary who knows the true distributions; the seller, however, still acts according to some beliefs formed from the observed samples. *Does the problem become trivial? Can the intermediary simply run the optimal segmentation w.r.t. the true distributions and expect near optimal outcomes?*

The answers turn out to be negative. Consider a segment for which there are two prices p^* and p, such that p^* is the monopoly price with a sale probability close to 1, while p gets near optimal revenue with a sale probability close to 0. If the intermediary includes this

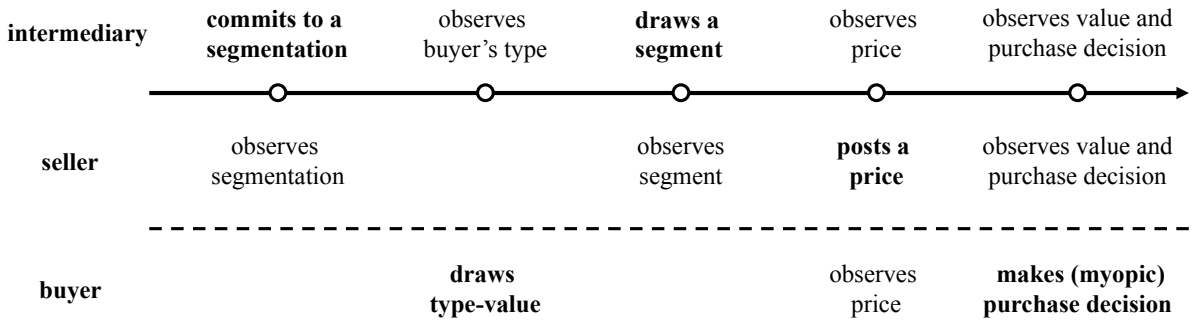

Figure 1: Time line of a single round.

segment, however, the seller's beliefs may overestimate the revenue of p and/or underestimate that of p^* and, thus, deduce that p is the monopoly price instead of p^*. As a result, the resulting social welfare may be much smaller than what the intermediary expects from the true distributions.

This example shows that, unlike existing works on the sample complexity of mechanism design, where the difficulties arise purely from the learning perspective of the problem, our problem presents an extra challenge from the uncertainties in the seller's behavior due to his inaccurate beliefs.

Intuitively, the intermediary would like to convert the optimal segmentation w.r.t. the true distributions into a more robust version, such that for any approximately accurate beliefs that the seller may have, the resulting objective is always close to optimal. We will refer to this procedure as *robustification*, and the result as the *robustified segmentation*.

2 Bayesian Model

We start with an example through which we will illustrate the main ideas in this section.

Example. The example is parameterized by a noise level, $1 - z \in [0, 1]$. The value set $[V] = \{1, 2, 3\}$ is identical to the type set $[T]$. Each type $t \in [T]$ corresponds to the distribution $\mathcal{F}(t)$:

$$\mathbf{Pr}_{v\sim\mathcal{F}(t)}[v] = \begin{cases} z & \text{if } v = t, \\ \frac{1-z}{2} & \text{otherwise.} \end{cases}$$

When $z = 1/3$, all $\mathcal{F}(t)$'s equal the uniform prior, and no non-trivial segmentation is possible. At the other extreme, when $z = 1$ each $\mathcal{F}(t)$ is a pointmass at t, which is the [4] model. □

Simplex View. As observed by [4], the key idea is to identify segments with probability distributions over $[V]$. The only thing that matters given a segment σ is the posterior distribution on $[V]$ conditioned on the intermediary choosing σ. Since $[V]$ is finite, it is easier to think of $\Delta([V])$ as the unit simplex in the appropriate dimensions. Then, a segment σ is simply a point in this simplex. Further, all that matters for a segmentation is the distribution over σ's as observed by the seller, i.e., it is sufficient to specify a distribution over the simplex $\Delta([V])$. The only constraint on this distribution is that its expectation must equal the prior distribution over values, which is another point on the simplex, denoted by $\boldsymbol{\tau}$.

Going further, it is sufficient to only consider some special points on the simplex. We denote these special points by x^S, for a subset $S \subseteq [V]$: this is the equal revenue distribution with support equal to the set S, i.e., these are distributions supported on S such that $p \cdot \mathbf{Pr}[v \geq p]$ is the same for all $p \in S$. These special points partition the simplex into regions X_v for each $v \in [V]$: each distribution in X_v is such that v is a monopoly price for it.

We now describe this through Example 1. Figure 2a shows the unit simplex, and the points x^S for all $S \subseteq [V]$. The light gray region is X_1, white is X_2 and dark gray is X_3. The uniform distribution over $[V]$ is represented by $\boldsymbol{\tau}$. An optimal segmentation with no noise (when $z = 1$) corresponds to representing $\boldsymbol{\tau}$ as the following convex combination of the vertices of the white polytope:

$$\boldsymbol{\tau} = \tfrac{1}{6}x^{\{2\}} + \tfrac{1}{6}x^{\{2,3\}} + \tfrac{2}{3}x^{[V]}.$$

This corresponds to the following segmentation with $\Sigma = \{\sigma_1, \sigma_2, \sigma_3\}$ corresponding to $x^{\{2\}}, x^{\{2,3\}}$ and $x^{[V]}$ resp. and $\mathcal{G}(1) = (0, 0, 1)$; $\mathcal{G}(2) = (1/2, 1/6, 1/3)$; and $\mathcal{G}(3) = (0, 1/3, 2/3)$.

[4] consider segmentations that only consist of vertices of X_{p^*} where p^* is a monopoly price of the prior distribution (which is the white region X_2 in Figure 2a). They assume that ties for monopoly price are broken in favor of the lowest price, which is the lowest value in

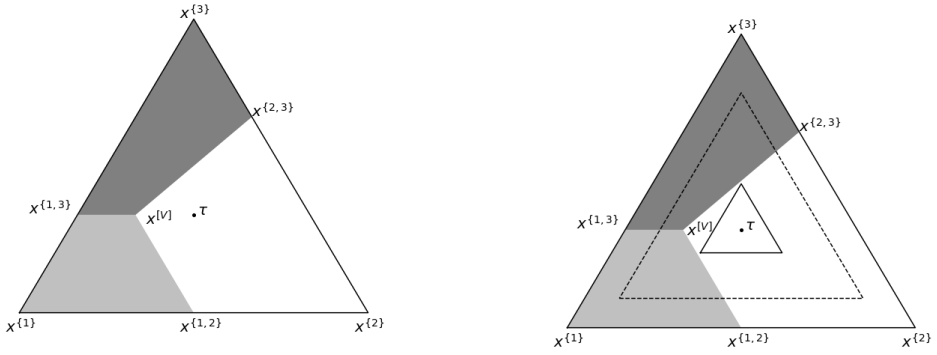

(a) Partition of the simplex into X_v's and τ = the uniform distribution over $[V]$.

(b) Convex hull of types for $z < 1$. Solid inner triangle is for $z = 0.49$; dashed one is for $z = 0.8$.

Figure 2: Simplex view for Example 1

the support. This implies that the item is always sold thus maximizing social welfare.

Generalization of the Simplex View. We extend this simplex view to our model. When the number of types T is at most the number of values V, and the type distributions are non-degenerate, we can continue to consider the simplex on the set of values, as we have done so far. This will be the case for Example 1. A more general view is to consider the simplex on the set of types. Most of the intuition extends to this view, although geometrically the picture is somewhat different. (This is the view we use in the proofs; the simplex on values is used just for illustration.)

The case when $z < 1$ is depicted in Figure 2b. The main difference from the previous picture is that we are not allowed to choose any point on the simplex for our segmentation. Instead, we are restricted to only choose the points in the convex hull of the $\mathcal{F}(t)$s, for all $t \in [T]$. We denote this convex hull by $\Delta([T])$, by abuse of notation. For $z = 0.49$, the figure shows that $\Delta([T])$ is contained entirely inside the white region. Thus no matter what segmentation is used, the seller always sets the monopoly price of 2; segmentation is therefore useless. For $z = 0.8$, the figure shows that segmentation is possible because $\Delta([T])$ intersects with all three regions, X_1, X_2, X_3.

We introduce some notation now. Given any segmentation $(\Sigma, \mathcal{G})$, this induces a distribution over segments, denoted $\mathcal{S}$, and a posterior distribution on the values $[V]$ for each segment $\sigma \in \Sigma$, which we abuse notation and denote by $\mathcal{F}(\sigma)$. For any distribution $\mathcal{F}$, let $\text{Rev}(\mathcal{F})$ denote its monopoly revenue. Let $\text{CS}(\mathcal{F})$ denote the consumer surplus when the seller sets the monopoly price for distribution $\mathcal{F}$. Our goal is to find a segmentation to maximize a linear combination of revenue and consumer surplus, i.e., for some parameter $\lambda \in [0, 1]$, maximize:

$$\mathbf{E}_{\sigma \sim \mathcal{S}} \Big[\lambda \cdot \text{Rev}\big(\mathcal{F}(\sigma)\big) + (1 - \lambda) \cdot \text{CS}\big(\mathcal{F}(\sigma)\big) \Big] .$$

From now on, we let $\Delta([T]) = \{\boldsymbol{x} \in \mathbb{R}_+^T : \|\boldsymbol{x}\|_1 = 1\}$ denote probability distributions over the types. We first formalize the claim that segmentation schemes correspond to probability distributions over $\Delta([T])$ with a given expectation. The proofs in the section are deferred to Appendix B.

LEMMA 2.1. *Let $\boldsymbol{\tau}$ denote the point in the simplex $\Delta([T])$ corresponding to the distribution $\mathcal{T}$: $\boldsymbol{\tau} = (\mathbf{Pr}_{\mathcal{T}}[1], \mathbf{Pr}_{\mathcal{T}}[2], \ldots, \mathbf{Pr}_{\mathcal{T}}[T])$. There is a 1:1 correspondence between segmentations $(\Sigma, \mathcal{G})$ and probability distributions μ over $\Delta([T])$ such that the expectation is $\boldsymbol{\tau}$, i.e.,*

$$\int \boldsymbol{x} d\mu = \boldsymbol{\tau}. \tag{2.1}$$

Using this lemma, we switch our design space to probability distributions μ that satisfy (2.1). We use $\sigma \in \Sigma$ and $\boldsymbol{x} \in \Delta([T])$ interchangably.

We now partition the simplex $\Delta([T])$ into V areas, $X_1, X_2, \ldots, X_V$, one for each value/price in $[V]$, such that the price p is a monopoly price for any segment in X_p. For any distribution $\mathcal{F}$ and any price $p \in [V]$, let $\mathrm{Rev}(\mathcal{F}, p)$ denote the revenue of price p on distribution $\mathcal{F}$, and $\mathrm{CS}(\mathcal{F}, p)$ denote the consumer surplus. For any $p \in [V]$, define:

$$X_p = \{\boldsymbol{x} : \mathrm{Rev}(\mathcal{F}(\boldsymbol{x}), p) \geq \mathrm{Rev}(\mathcal{F}(\boldsymbol{x}), p'), \forall p' \neq p\} \ .$$

Since the revenue function $\mathrm{Rev}(\mathcal{F}(\boldsymbol{x}), p) = \sum_t x_t \cdot \mathrm{Rev}(\mathcal{F}(t), p))$ is linear in $\boldsymbol{x}$, the set X_p is the intersection of the simplex $\Delta([T])$ and a polytope defined by $V-1$ linear constraints. Further, if we restrict our domain to points $\boldsymbol{x} \in X_p$, we have that these are linear functions in $\boldsymbol{x}$.

$$\mathrm{Rev}(\mathcal{F}(\boldsymbol{x})) = \mathrm{Rev}(\mathcal{F}(\boldsymbol{x}), p) = \sum_t x_t \cdot \mathrm{Rev}(\mathcal{F}(t), p) \ , \text{ and}$$

$$\mathrm{CS}(\mathcal{F}(\boldsymbol{x})) = \mathrm{CS}(\mathcal{F}(\boldsymbol{x}), p) = \sum_t x_t \cdot \mathrm{CS}(\mathcal{F}(t), p) \ .$$

We next observe that this implies that it is sufficient to choose at most one point from each X_p. The idea is that we can replace the distribution conditioned on $\boldsymbol{x} \in X_p$ by its expectation.

LEMMA 2.2. *There is an optimal segmentation such that the distribution μ is supported on at most one point from each X_p, i.e., a finite set of the form $\{\boldsymbol{x}_p \in X_p, \forall\ p \in [V]\}$.*

Using this lemma, we now show that the following linear program (LP) captures the optimal segmentation. The variables are $\boldsymbol{z}_p = \boldsymbol{x}_p \cdot \mu(\boldsymbol{x}_p)$. We denote by Z_p the region that is the convex hull of X_p and the origin. Rev and CS extend naturally to Z_ps.

(2.2)
$$\begin{aligned} \max \quad & \sum_{p \in [V]} \lambda \cdot \mathrm{Rev}\big(\mathcal{F}(\boldsymbol{z}_p)\big) + (1-\lambda) \cdot \mathrm{CS}\big(\mathcal{F}(\boldsymbol{z}_p)\big) \\ \text{s.t.} \quad & \forall\ p \in [V], \boldsymbol{z}_p \in Z_p \text{ and } \sum_{p \in [V]} \boldsymbol{z}_p = \boldsymbol{\tau} \ . \end{aligned}$$

THEOREM 2.1. *We can find an optimal segmentation in polynomial time by solving LP* (2.2).

3 Sample Complexity Model

We scale the values to be in $(0, 1]$, i.e., $[V] = \left\{\frac{1}{V}, \frac{2}{V}, \ldots, 1\right\}$. This treatment simplifies the notations in the proofs, and separates the two roles of V: the scale of the values (less interesting, always has the same degree as ϵ), and the number of possible values. To translate the bounds into the original scaling, replace ϵ with $\frac{\epsilon}{V}$ everywhere. We further assume the type distribution to be uniform to simplify discussions. This is w.l.o.g. up to duplication of types.

Following standard notations in algorithmic mechanism design, we refer to the sale probability of a price as its quantile. We will consider the revenue curve in the quantile space where the x and y coordinates are the quantile of a price and its revenue, respectively.

3.1 Model and Results

Intermediary: The intermediary has access to the value distributions of different types only in the form of m i.i.d. samples per type. She chooses a segmentation based on these samples, and then the chosen segmentation is evaluated on a freshly drawn type-value pair, i.e., the test sample. The expectation of the objective is taken over the random realization of the m samples per type as well as the test sample, and potentially the randomness in the choice of the segmentation.

Buyer: The buyer bids truthfully since the seller effectively posts a take-it-or-leave-it price.

Seller: We need to further define how the seller acts. Consider the following candidate models:

1. The seller knows the value distributions exactly. Hence, given the segmentation and the realized segment, which induces a mixture of the value distributions of different types, the seller posts the monopoly price of the mixture.

2. The seller can access the same set of m samples per type, and believes that the value distributions are the empirical distributions, i.e., the uniform distributions over the corresponding samples. Hence, she posts the monopoly price of the mixture of empirical distributions.

3. The seller further has access to other sources of samples.

4. The seller further has access to other sources of prior knowledge.

This is only a nonexclusive list of many potential models that are equally well-motivated in our opinion, depending on the actual applications. *Is there a unifying model that allows us to study all these settings in one shot and get non-trivial positive results?*

To this end, this paper considers the following overarching model (the subscript S indicates that these variables are associated with the *seller*):

> For $\epsilon_S = O\big(m^{-1/2} \log(mV)\big)$, the seller forms beliefs $\mathcal{F}_S(t)$'s, $t \in [T]$, such that for any type t the Kolmogorov-Smirnov distance between

$\mathcal{F}_S(t)$ and $\mathcal{F}(t)$ is at most ϵ_S, i.e., for any value $v \in [V]$, v's quantiles w.r.t. $\mathcal{F}(t)$ and $\mathcal{F}_S(t)$ differ by at most ϵ_S. Then, she posts the monopoly price of the mixture of the beliefs.

The choice of ϵ_S is based on a standard concentration plus union bound combination on the empirical distributions over the m samples that the intermediary can access. In other words, we assume that the seller's beliefs are at least as good as what could have been estimated using the intermediary's samples. All aforementioned candidate models are special cases of ours.

We start with an impossibility result for general value distributions. See Section A for details.

THEOREM 3.1. *If the value distributions are allowed to have multiple monopoly prices whose social welfare differ by at least $\Omega(1)$, e.g., the uniform distribution over $\{\frac{1}{2}, 1\}$, no algorithm can obtain any $o(1)$-approximation using a bounded number of samples.*

(Discrete) MHR-like Distributions. Given the above impossibility result that relies on value distributions that have multiple monopoly prices whose respective values of social welfare are vastly different, we intuitively need the value distributions to be unimodal and far from having a plateau. The family of continuous monotone hazard rate (MHR) distributions, a standard family of distributions in the literature, has all the nice properties that we need, except that they are continuous. They are unimodal since they have concave revenue curves in the quantile space (folklore). In fact, their revenue curves in the quantile space are strongly concave near the monopoly price [26, Lemma 3.3]. They also admit other useful properties: the optimal revenue is at least a constant fraction of the social welfare [15, Lemma 3.10]; and the monopoly price has a sale probability lower bounded by some constant [23, Lemma 4.1].

There is an existing notion of discrete MHR distributions by [3] that mimics the functional form of the continuous version. However, it loses some useful properties. In particular, it contains some distributions that have two monopoly prices, e.g., the uniform distribution over $\{\frac{1}{2}, 1\}$, and as a result still suffers from the impossibility result.

Instead, we define a family of (discrete) MHR-like distributions directly from the aforementioned benign properties of continuous MHR distributions. Hence, unlike the existing notion of discrete MHR distributions, our definition truly inherits the main features of continuous MHR distributions. We remark that the constants $\frac{1}{4}$ and $\frac{1}{e}$ in the following definition are merely copied from the continuous counterparts; our results still hold asymptotically if they are replaced by other constants.

DEFINITION 3.1. (MHR-LIKE DISTRIBUTIONS) *A discrete distribution $\mathcal{F}$ is* MHR-like *if it satisfies:*

1. *(Concavity) Its revenue curve is concave in the quantile space.*
2. *(Strong concavity near monopoly price) For its monopoly price p^* and any other price p', suppose their quantiles are q^* and q' respectively; then, we have:*
$$\mathrm{Rev}\big(p', \mathcal{F}\big) \le \big(1 - \tfrac{1}{4}(q^* - q')^2\big) \cdot \mathrm{Rev}\big(p^*, \mathcal{F}\big) \ .$$
3. *(Large monopoly sale probability) Its monopoly price's sale probability is at least $\frac{1}{e}$.*
4. *(Small revenue and welfare gap) Its monopoly revenue is at least $\frac{1}{e}\mathbf{E}_{v\sim\mathcal{F}}[v]$.*

The main difference of our MHR-like distribution and the notion in [3] is property 2 in definition 3.1. An MHR-like distribution can be made by discretizing an continuous MHR distribution, meanwhile ensuring that there is a gap between the optimal revenue and any sub-optimal ones.

We show that polynomially many samples are sufficient for learning an ϵ-optimal segmentation, with only the mild assumption on seller's behavior discussed earlier in the section.

THEOREM 3.2. *With $m = \mathrm{poly}\big(\epsilon^{-1}, T, \log V\big)$ i.i.d. samples, we can learn a segmentation that is optimal up to an ϵ additive factor in $\mathrm{poly}\big(\epsilon^{-1}, T, V\big)$ time.*

3.2 Robustification: Motivation and Definition

Recall the thought experiment in Section 1. Consider a more powerful intermediary who has exact knowledge of the true distributions; the seller, however, still acts according to her approximately accurate beliefs. Further, recall the example where the problem remains nontrivial even when the intermediary has more power; we give more details below. Consider a segment for which there are two prices p^* and p, such that p^* is the monopoly price with a quantile close to 1, while p gets close-to-optimal revenue with a quantile close to 0. Even though a single MHR-like distribution cannot have such a plateau, a mixture of MHR-like distributions can.[1] See Figure 3a for an illustrative example. If the intermediary includes this segment, the seller's belief may overestimate the revenue of p and/or underestimate that of p^* and thus deduce that p is the monopoly price

[1] In fact, every distribution on $[V]$ is a mixture of MHR-like distributions, because point masses are MHR-like.

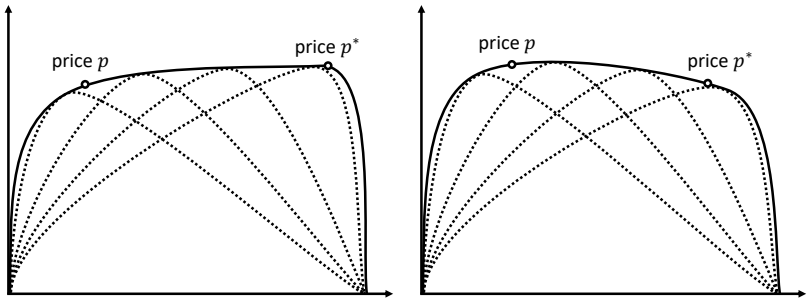

(a) True distribution (b) Seller's belief

Figure 3: Plateau Example. In a revenue curve in the quantile space, the x and y coordinates are the quantile of a price and its revenue respectively. On the left, the solid curve is the revenue curve of a segment w.r.t. the true distributions; the dotted curves are those of the type distributions mixed in the segment. On the right are the counterparts w.r.t. the seller's beliefs. Prices p^* and p are monopoly prices of the segment w.r.t. the true distributions and the seller's beliefs respectively.

instead of p^*. See Figure 3b. The resulting social welfare may be much smaller than what the intermediary expects from the true distributions.

Intuitively, we would like to convert the optimal segmentation w.r.t. the true distributions into a more robust version, such that for any approximately accurate beliefs that the seller may have, the resulting social welfare and revenue are both close to optimal. As mentioned in Section 1, we will refer to this procedure as *robustification*, and the result as the *robustified segmentation*.

In the following definition, the subscripts of ϵ_S and ϵ_I indicate they are the additive errors that the seller and intermediary are aiming for, respectively. Further, it states that the robustified segmentation must keep all segments in the original version ($\Sigma^r \supseteq \Sigma^*$). We ignore insignificant segments due to technical difficulties in achieving the stated properties for them,[2] and that their roles in the revenue and social welfare are negligible. For any significant segment, the first two conditions state that its weight and mixture of types are preserved approximately; the third condition gives the desirable robustness against the uncertainties in the seller's behavior.

Definition 3.2. (Robustified Segmentation) *Suppose that (1) $(\Sigma^*, \mathcal{G}^*)$ is a segmentation, represented by $\boldsymbol{x}^*_\sigma$ and weight $w^*_\sigma = \mu(\boldsymbol{x}^*_\sigma)$, $\forall\ \sigma \in \Sigma^*$; and (2) p^*_σ is an optimal price w.r.t. $\mathcal{F}(\boldsymbol{x}^*_\sigma)$, $\forall\ \sigma \in \Sigma^*$. For any $\epsilon_I \geq \epsilon_S > 0$, $(\Sigma^r, \mathcal{G}^r)$ is an (ϵ_I, ϵ_S)-robustified segmentation, represented by $\boldsymbol{x}^r_\sigma$ and weight w^r_σ with $\Sigma^r \supseteq \Sigma^*$, if for any $\sigma \in \Sigma^*$, either σ is insignificant in that $\mathbf{E}_{v\sim\mathcal{F}(\boldsymbol{x}^*_\sigma)}[v] < \epsilon_I$, or:*

1. *(Weight preservation) $w^r_\sigma \geq (1-\epsilon_I)\cdot w^*_\sigma$;*
2. *(Mixture preservation) $\|\boldsymbol{x}^*_\sigma - \boldsymbol{x}^r_\sigma\|_1 \leq \epsilon_I$; and*
3. *(Robustness) no ϵ_S-optimal price w.r.t. $\mathcal{F}(\boldsymbol{x}^r_\sigma)$ has a quantile smaller than that of p^*_σ by ϵ_I.*

The next lemma shows that the technical conditions in the definition of robustified segmentation indeed lead to robust bounds in terms of both social welfare and revenue and, by induction, their linear combinations. The proof follows from straightforward calculations and therefore is deferred to full version paper.

Lemma 3.1. *For any prices p^r_σ's that are ϵ_S-optimal w.r.t. segments $\sigma \in \Sigma^r$ in the robustified segmentation, we have the following in terms of social welfare and revenue:*

$$\sum_{\sigma\in\Sigma^r} w^r_\sigma \cdot \mathrm{SW}\big(p^r_\sigma, \mathcal{F}(\boldsymbol{x}^r_\sigma)\big) \geq \sum_{\sigma\in\Sigma^*} w^*_\sigma \cdot \mathrm{SW}\big(p^*_\sigma, \mathcal{F}(\boldsymbol{x}^*_\sigma)\big) - O\big(\epsilon_I\big)\ ,$$

$$\Big|\sum_{\sigma\in\Sigma^r} w^r_\sigma \cdot \mathrm{Rev}\big(p^r_\sigma, \mathcal{F}(\boldsymbol{x}^r_\sigma)\big) - \sum_{\sigma\in\Sigma^*} w^*_\sigma \cdot \mathrm{Rev}\big(p^*_\sigma, \mathcal{F}(\boldsymbol{x}^*_\sigma)\big)\Big| \leq O\big(\epsilon_I\big)\ .$$

3.3 Robustification: Algorithm This subsection introduces an algorithm that finds such an (ϵ_I, ϵ_S)-robustified segmentation in polynomial time for any sufficiently large ϵ_I, i.e.:

$$\epsilon_I \geq \epsilon_S^{\frac{1}{6}} T^{\frac{2}{3}} \log^{\frac{1}{6}} V = \tilde{O}\left(m^{-\frac{1}{12}} T^{\frac{2}{3}}\right)\ , \tag{3.3}$$

Lemma 3.2. *There is an algorithm that computes in polynomial time an (ϵ_I, ϵ_S)-robustified segmentation, for any ϵ_I and ϵ_S that satisfy Eqn. (3.3).*

[2] If the expected value is tiny in the first place, all prices are ϵ_S-optimal. Hence, we cannot achieve robustness.

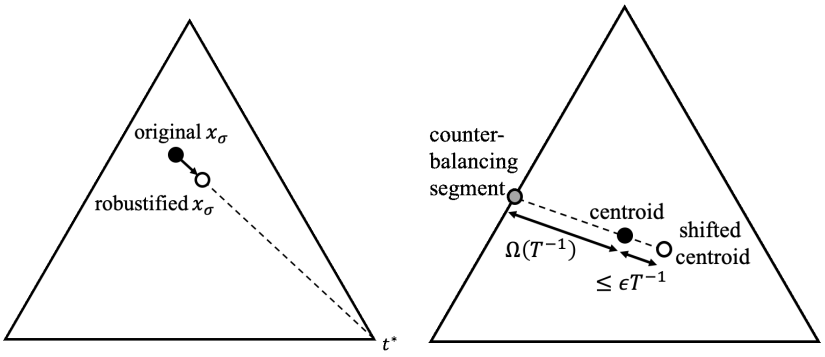

(a) Robustify a significant segment (b) Counterbalance to restore the centroid

Figure 4: Robustification.

3.3.1 Proof Sketch of Lemma 3.2

Step 1: Robustify the significant segments one by one, ignoring the centroid constraint. Any significant segment σ is represented by a point $\boldsymbol{x}_\sigma$ and weight $w_\sigma = \mu(\boldsymbol{x}_\sigma)$, whose intended price is p^*. In the simplex view, we want to move slightly away from $\boldsymbol{x}_\sigma$ such that we end up far from all regions X_p where price p gives a small consumer surplus. How do we find which direction to move towards? (Since we are in high dimensions, we cannot rely on geometric intuition.)

The choice of direction to move towards relies on a structural result about the mixtures of MHR-like distributions stated as Lemma 3.3. The lemma promises that there exists a type t^* such that for the distribution $\mathcal{F}(t^*)$, for prices whose quantiles are less than that of p^* by at least ϵ_I, there is a revenue gap of $\frac{\epsilon_S T}{\epsilon_I}$. Once we prove existence, we can find such a type by enumerating over all types and checking if the property holds.

In the simplex view, see Figure 4a. We want to move $\boldsymbol{x}_\sigma$(black point) towards the vertex that corresponds to type t^*(to the white point); we want to be at $(1-\frac{\epsilon_I}{T})\boldsymbol{x}_\sigma + \frac{\epsilon_I}{T}\boldsymbol{x}_{t^*}$. To do this, decrease the probability of mapping each type to σ by an $1-\frac{\epsilon_I}{T}$ factor; then, increase the probability of mapping t^* to σ additively by $\frac{w_\sigma \epsilon_I}{T}$ to restore the original weight. Clearly, this satisfies the *mixture preservation* condition.

Let p be any price whose quantile is smaller than that of p^* by at least ϵ_I. The revenue gap between p and p^* for $\mathcal{F}(t^*)$ is $\frac{\epsilon_S T}{\epsilon_I}$, and we moved towards $\boldsymbol{x}_{t^*}$ by $\frac{\epsilon_I}{T}$, therefore the revenue gap between p and p^* for the mixture is at least ϵ_S. As a result, p cannot be an ϵ_S-optimal price in the resulting segment. Thus, we have the *robustness* condition.

Step 2: See Figure 4b. We will add a counterbalancing segment(gray) to restore the centroid. After the first step, the centroid may be shifted from its intended location, i.e., the middle of the simplex, by up to $\frac{\epsilon_I}{T}$. Consider the line that crosses the intended centroid and the shifted one. Add a counterbalancing segment at its intersection with the boundary of the simplex on the opposite side of shifted centroid, with an appropriate weight that restores the centroid. The weight is only $O(\epsilon_I)$ because the distance between the intended centroid and the counterbalancing segment, in fact, any point on the boundary of the simplex in general, is at least $\Omega(\frac{1}{T})$ by basic geometry.

Finally, the total weight may now exceed 1 by up to $O(\epsilon_I)$. Normalize the weights of all segments to restore a total weight of 1. It decreases the weights of the segments by at most $1 - O(\epsilon_I)$ and therefore satisfies the *weight preservation* condition.

3.3.2 Structural Lemma and the Proof Sketch

LEMMA 3.3. *For any segment $\boldsymbol{x} \in \Delta$ that is significant in the sense that $\mathbf{E}_{v\sim\mathcal{F}(\boldsymbol{x})}[v] \geq \epsilon_I$, and its corresponding monopoly price p^*, there is a type $t^* \in [T]$ such that for any price p whose quantile w.r.t. $\mathcal{F}(\boldsymbol{x})$ is smaller than that of p^* by ϵ_I, we have:*

$$\mathrm{Rev}\big(p, \mathcal{F}(t^*)\big) < \mathrm{Rev}\big(p^*, \mathcal{F}(t^*)\big) - \tfrac{\epsilon_S T}{\epsilon_I} \ .$$

This is technically the most challenging part of the proof. It is long and we show proof sketch here. For formal proof please refer to the full version. By concavity of the revenue curves of MHR-like distributions, it suffices to consider the inequality when p is the smallest price whose quantile w.r.t. $\mathcal{F}(\boldsymbol{x})$ is smaller than that of p^* by ϵ_I. Let this price be $\bar{p}$.

Recall the plateau example in Figure 3. From the picture, it is tempting to pick the type that corresponds to the "right-most" dotted revenue curve, as it has the desirable shape that the revenue rapidly decreases when the price increases from p^*. There are several problems with this approach. First, the concept of "right-most"

revenue curve is underdefined. Is it the one with the smallest monopoly price? Or the one with the largest monopoly sale probability? Second, even if we find a type whose revenue curve has the desirable shape, it still may not prove Lemma 3.3. For example, it may not have a large enough optimal revenue in the first place and, thus, the RHS of the inequality in the lemma is negative.

Instead, we will prove the lemma by contradiction. Intuitively, the contradiction will be that there is a type t such that the revenue curve of $\mathcal{F}(t)$ has a large plateau; this is not possible for MHR-like distributions. The assumption to the contrary guarantees that the revenue of p^* is not much above that of any price between $\bar{p}$ and p^*. The following additional conditions formalize the 'large plateau' notion:

1. Revenue between $\bar{p}$ and p^* is not much higher: $\forall\ p' \in [p^*, \bar{p}]$,
$$\mathrm{Rev}\big(p', \mathcal{F}(t)\big) - \mathrm{Rev}\big(p^*, \mathcal{F}(t)\big) \leq \tilde{O}\big(\tfrac{\epsilon_S T^2}{\epsilon_I^2}\big).$$
2. The plateau is high, i.e., revenue of p^* is large: $\mathrm{Rev}\big(p^*, \mathcal{F}(t)\big) \geq \Omega\big(\tfrac{\epsilon_I}{T}\big)$.
3. The plateau is wide, i.e., the quantiles of $\bar{p}$ and p^* differ by at least $\Omega\big(\tfrac{\epsilon_I^2}{T}\big)$.

The rest of the subsection assumes to the contrary that the inequality in the lemma fails to hold for all types. Then, we use a probabilistic argument to show that there must be a type t such that the distribution $\mathcal{F}(t)$ satisfies the conditions mentioned above, and argue that these lead to a contradiction.

Probabilistic Argument. Consider sampling a type t according to the mixture induced by the segment. We show that under the assumption to the contrary, the probability that condition 1 is violated is less than $O\big(\tfrac{\epsilon_I}{T}\big)$. Further, we prove that the revenue of p^* w.r.t. $\mathcal{F}(\boldsymbol{x})$, which is optimal for this distribution, is at least an $\Omega\big(\tfrac{1}{T}\big)$ fraction of the social welfare w.r.t.$\mathcal{F}(\boldsymbol{x})$. Then by the assumption that this segment is significant, this is at least $\Omega\big(\tfrac{\epsilon_I}{T}\big)$. Then, by a Markov inequality type argument, there is at least an $\Omega\big(\tfrac{\epsilon_I}{T}\big)$ probability that Condition 2 is satisfied. Putting together, there is a positive chance that we sample a type t that satisfies the first two conditions. Finally, we finish the argument by showing that the first two conditions actually imply the third one.

Contradiction. The proof is a case by case analysis, so we present the bottleneck case which forces the choice we made in Eqn. (3.3). This is when the monopoly price $p(t)$ of type t is smaller than both $\bar{p}$ and p^*. A complete proof that includes the other cases are deferred to full version of the paper.

The concavity, and strong concavity near monopoly price, of MHR-like distributions, along with the fact that both $\bar{p}$ and p^* are larger than the monopoly price, imply that the revenue gap is at least the revenue of p^* times the square of the quantile gap between the prices. Further by the second and third conditions above, of having large revenue and large quantile gap, the revenue gap between prices $\bar{p}$ and p^* is at least:

$$\Omega\left(\frac{\epsilon_I}{T}\right) \cdot \Omega\left(\frac{\epsilon_I^2}{T}\right)^2 = \Omega\left(\frac{\epsilon_I^5}{T^3}\right)$$

This is greater than $\frac{\epsilon_S T}{\epsilon_I}$ by our choice of $\epsilon_I \geq \epsilon_S^{\frac{1}{6}} T^{\frac{2}{3}} \log^{\frac{1}{6}} V$ in Eqn. (3.3).

3.4 Proof of Theorem 3.2: Project, Optimize, and Robustify Finally, we show how to use the robustification technique to design an algorithm, presented as Algorithm 1, that learns a (robust) $O(\epsilon_I)$ segmentation in the sample complexity model.

Algorithm. Similar to the existing works on the sample complexity of mechanism design, the algorithm starts by constructing the empirical distributions. Then, we *project* them back to the space of MHR-like distributions w.r.t. the Kolmogorov-Smirnov distance d_{KS}, i.e., the maximum difference in the quantile of any value. The feasibility of this step comes from the fact that the true distributions are MHR-like and satisfy the inequality. We explain in Appendix C how to compute in polynomial time approximate projections that relax the RHS of the inequality by a constant factor; other constants in our analysis need to be changed accordingly but the bounds stay the same asymptotically. Further, we *optimize* the segmentation according to the MHR-like empirical distributions. Finally, we *robustify* the resulting segmentation using Lemma 3.2.

Analysis. Note that for any type t, the distances between the true distribution $\mathcal{F}(t)$ and the seller's belief $\mathcal{F}_S(t)$, between $\mathcal{F}(t)$ and the empirical distribution $\mathcal{E}(t)$, and between $\mathcal{E}(t)$ and the MHR-like empirical distribution $\widetilde{\mathcal{E}}(t)$ are bounded by ϵ_S. Hence, the distance between the seller's belief $\mathcal{F}_S(t)$ and the MHR-like distribution $\widetilde{\mathcal{E}}(t)$ is at most $3\epsilon_S$ by the triangle inequality and, thus, the same conclusion holds replacing types with mixtures induced from the segments. Therefore, for any segment in the segmentation chosen by the algorithm, the seller's monopoly price w.r.t. her beliefs is a $6\epsilon_S$-optimal price w.r.t. the MHR-like empirical distributions. By Lemma 3.1, the performance of the algorithm is an $O(\epsilon_I)$-approximation comparing with the optimal w.r.t. the MHR-like empirical distributions.

It remains to show that the optimal w.r.t. the MHR-like empirical distributions is an $O(\epsilon_I)$-approximation to

Algorithm 1 Learn a (Robust) Segmentation from Samples

1. Construct empirical distributions $\mathcal{E}(t)$'s, $t \in [T]$, from samples.
2. Find MHR-like empirical distributions $\widetilde{\mathcal{E}}(t)$'s, $t \in [T]$, such that the Kolmogorov-Smirnov distances between them and the corresponding empirical distributions are small: $d_{KS}\big(\widetilde{\mathcal{E}}(t), \mathcal{E}(t)\big) \le \epsilon_S$.
3. Find optimal segmentation $(\Sigma^*, \mathcal{G}^*)$ w.r.t. MHR-like empirical distributions $\widetilde{\boldsymbol{\mathcal{E}}}$.
4. Construct the robustified segmentation $(\Sigma^r, \mathcal{G}^r)$ and return it.

Algorithm 2 Robustify a segmentation

Input: MHR-like distributions $\mathcal{F}(t)$, $t \in [T]$;
segmentation $(\Sigma^*, \mathcal{G}^*)$, represented by point-weight pairs $(\boldsymbol{x}^*_\sigma, w^*_\sigma)$, $\sigma \in \Sigma^*$;
prices p^*_σ, $\sigma \in \Sigma^*$, which are optimal w.r.t. the corresponding $\mathcal{F}(\boldsymbol{x}^*_\sigma)$'s.

Output: Robustified segmentation $(\Sigma^r, \mathcal{G}^r)$, represented by $\boldsymbol{x}^r_\sigma$'s and weights w^r_σ's.

1. For every segment $\sigma \in \Sigma^*$:
   (a) If segment σ is insignificant, i.e., $\mathbf{E}_{v\sim\mathcal{F}(\boldsymbol{x}^*_\sigma)}[v] < \epsilon_I$, let $\boldsymbol{x}^r_\sigma = \boldsymbol{x}^*_\sigma$.
   (b) Otherwise, construct $\boldsymbol{x}^r_\sigma$ using Algorithm 3;
   (c) Decrease the weight by a $1-\epsilon_I$ multiplicative factor, i.e., $w^r_\sigma = (1-\epsilon_I)w^*_\sigma$ in both cases.
2. For every type $t \in [T]$:
   (a) Add a new segment $\sigma(t) \notin \Sigma^*$ to Σ^r such that $\boldsymbol{x}^r_{\sigma(t)}$ is the vertex of simplex Δ that corresponds to type t.
   (b) Let its weight be $w^r_{\sigma(t)} = \frac{1}{T} - \sum_{\sigma\in\Sigma^*} w^r_\sigma \cdot x^r_{\sigma,t}$.

the optimal w.r.t. the true distributions. To do that, it suffices to find a good enough segmentation achieving this approximation. For this we once again resort to Lemma 3.2, in particular, the existence of an (ϵ_I, ϵ_S)-robustified segmentation $(\Sigma^r, \mathcal{G}^r)$ for the optimal segmentation w.r.t. the true distributions $\mathcal{F}(t)$'s. Note that the MHR-like empirical distributions $\widetilde{\mathcal{E}}(t)$'s are at most $O(\epsilon_S)$ away from the corresponding true distributions $\mathcal{F}(t)$'s, by triangle inequality. Therefore, running $(\Sigma^r, \mathcal{G}^r)$ on the MHR-like distributions, with a seller who posts the monopoly price w.r.t. the MHR-like distributions, gives an $O(\epsilon_I)$-approximation by Lemma 3.1.

4 Bandit Model

In the bandit model, the intermediary interacts with the seller and the buyer repeatedly for m rounds for some positive integer m, with the buyer's type-value pair freshly sampled in each round. The goal is to maximize the cumulative objective during all m rounds. There are a large variations of models depending on the modeling assumptions. Next, we explain our choice.

Intermediary's Information: The intermediary does not know the value distributions at the beginning and, therefore, must learn such information through the interactions in order to find a good enough segmentation. Further, the intermediary observes in each round only the purchase decision of the buyer, but not her value. This is similar to the bandit feedback in online learning and hence the name of our model. We remark that the alternative model where the intermediary can observe the values, which corresponds to full-information feedback in online learning, easily reduces to the sample complexity model as the intermediary may simply run the algorithm in the sample complexity model using the bids in previous rounds as the samples.

Since the intermediary can observe the buyer's type in each round, she can easily learn the type distribution through repeated interactions. To simplify the discussions, we will omit this less interesting aspect of the problem and will assume that the type distribution is

Algorithm 3 Robustify a (significant) segment

Input: MHR-like distributions $\mathcal{F}(t)$, $t \in [T]$;
segment σ, represented by $\boldsymbol{x}^*_\sigma$, such that $\mathbf{E}_{v\sim\mathcal{F}(\boldsymbol{x}^*_\sigma)}[v] \geq \epsilon_I$;
price p^*, which is optimal w.r.t. $\mathcal{F}(\boldsymbol{x}^*_\sigma)$.

Output: Robustified version of σ, represented by $\boldsymbol{x}^r_\sigma$.

1. Find type t_σ such that for any price p whose quantile w.r.t. $\mathcal{F}(\boldsymbol{x}^*_\sigma)$ is at most that of p^* less ϵ_I, we have $\mathrm{Rev}\big(p, \mathcal{F}(t_\sigma)\big) < \mathrm{Rev}\big(p^*, \mathcal{F}(t_\sigma)\big) - \frac{\epsilon_S T}{\epsilon_I}$.
2. Construct the robustified version as $\boldsymbol{x}^r_\sigma = \big(1 - \frac{\epsilon_I}{T}\big) \cdot \boldsymbol{x}^*_\sigma + \frac{\epsilon_I}{T} \cdot \boldsymbol{x}_{t_\sigma}$, where $\boldsymbol{x}_{t_\sigma} \in \Delta$ has its t_σ-th entry being one and all others being zeros.

publicly known. Following the treatment in previous models, we further assume that it is a uniform distribution.

Buyer's Behavior: We assume the buyer is myopic and, therefore, buys the item in each round if and only if the price posted is at most her value. In other words, the buyer does not take into account that her behavior in the current round may influence how the intermediary and the seller acts in the future. This challenge of non-myopic buyers was partly addressed in the online auction problem by [25]. Their techniques, however, do not directly apply to our problem.

Seller's Information: We assume that, like the intermediary, the seller does not have any information about the value distributions of the buyer at the beginning, and must learn such information through bandit feedback. With this assumption, we will investigate how to encourage the seller to explore on the intermediary's behalf. What makes it challenging is that the seller's objective (revenue) and the intermediary's objective (e.g., social welfare) may not be aligned.

Seller's Behavior: Any algorithm by the intermediary must rely on some assumptions on the seller's behavior to get a non-trivial performance guarantee. Informally, we need the seller to pick an (approximately) optimal price in terms of revenue when there is enough information for finding one; there is not much we can do if the seller simply ignores any information and picks prices randomly. On the other hand, we also need the seller to explore at a reasonable rate in order to learn the value distributions. If the seller could have other sources of information which allow him to estimate the distribution accurately, he may severely limit his exploration on prices whose confidence intervals suggest high-potential (and high uncertainty). Our assumption must disallow such strategies and ensure that the seller learns the distributions only via observing the buyer's actions. *What are the mildest behavioral assumptions (on the seller) that allow the intermediary to have a non-trivial guarantee in bandit model?*

Note that the seller herself faces an online learning problem with bandit feedback. Our model is driven by the exploration-exploitation dilemma in her viewpoint. First, we introduce the upper confidence bound (UCB) and the lower confidence bound (LCB) of the quantile of any value v and any type t given past observations in the form of (type, price, purchase decision)-tuples.

1. For any value $v \in [V]$ and any type $t \in [T]$, suppose there are $m(v,t)$ past observations with type t and price v, among which the buyer purchases the item in $m^+(v,t)$ observations. Then, for some constant $C > 0$ that depends on the desired confidence level, let:
$$\widetilde{U}(v,t) = \tfrac{m^+(v,t)}{m(v,t)} + \sqrt{\tfrac{C}{m(v,t)}}$$
$$\widetilde{L}(v,t) = \tfrac{m^+(v,t)}{m(v,t)} - \sqrt{\tfrac{C}{m(v,t)}}$$
2. Noting that quantiles are monotone, we define the UCB and LCB as follows:
$$U(v,t) = \min_{v' \leq v} \widetilde{U}(v',t) \quad L(v,t) = \max_{v' \geq v} \widetilde{L}(v',t)$$
3. This further induces the UCB and LCB of the quantile of value v w.r.t. each segment $\boldsymbol{x} \in \Delta$:
$$U(v,\boldsymbol{x}) = \sum_{t\in[T]} x_t \cdot U(v,t)$$
$$L(v,\boldsymbol{x}) = \sum_{t\in[T]} x_t \cdot L(v,t)$$

For some target average regret $0 < \epsilon_S < 1$ of the seller, we say that she **exploits** in a round if the segment $\boldsymbol{x}$ and her price p satisfy that:
$$U(p,\boldsymbol{x}) \geq \max_{p'\in[V]} U(p',\boldsymbol{x}) - \epsilon_S \ .$$

Otherwise, we say that she **explores**. We assume that the seller is an ϵ_S-canonical learner in the sense that she exploits in all but at most an ϵ_S fraction of the rounds.

Among others, we give two example algorithms that satisfy this definition. First, the Upper-Confidence-Bound (UCB) algorithm satisfies this with $\epsilon_S = 0$. Further, consider the following simple Explore-then-Commit (ETC) algorithm, with $\epsilon_S = \tilde{O}\big(m^{-\frac{1}{3}}TV\big)$. The seller explores in the first $\epsilon_S m = \tilde{\Omega}\big(TV\epsilon_S^{-2}\big)$ rounds by posting random prices. Since a price-type pair shows up with probability $\frac{1}{TV}$ in each of these rounds, she learns the quantile of every price $p \in [V]$ w.r.t. the value distribution $\mathcal{F}(t)$ of every $t \in [T]$ up to an additive error of ϵ_S . Then, in the remaining rounds, she can pick any price that is not obviously suboptimal in the sense that its UCB is smaller than the LCB of some other price.

4.1 Algorithm In each round, the algorithm seeks to place the intermediary in a win-win situation by maintaining a set of optimistic hypothetical value distributions for the types, together with a robustified version of the optimal segmentation w.r.t. the hypothetical distributions. If the seller indeed posts a price that is consistent with our optimistic hypothesis, we use the analysis from the previous section (Thm. 3.2) to show that the objective in this round is close to optimal. Otherwise, if the seller posts a price that is inconsistent with our optimistic hypothesis, we argue that there must be a sufficiently large gap between its UCB and LCB and, thus, the intermediary gets some useful new information.

There is a caveat, however, when the algorithm constructs the robustified segmentation: it needs to replace ϵ_S with some slightly larger parameter ϵ_M. In particular, let ϵ_M be such that if we define ϵ_I using Eqn. (3.3), replacing ϵ_S with ϵ_M, we have: $\epsilon_I m = TV\epsilon_M^{-3}$. Solving it gives that $\epsilon_M = \Theta\big(m^{-\frac{6}{19}}\mathrm{poly}(T,V)\big)$ and $\epsilon_I = \Theta\big(m^{-\frac{1}{19}}\mathrm{poly}(T,V)\big)$.

We show that for ϵ_S-canonical learners with a sufficiently small ϵ_S, which is satisfied by both aforementioned examples, we can get sublinear regret.

THEOREM 4.1. *Algorithm 4 gets at least* $\mathrm{Opt} - O\big(m^{-\frac{1}{19}} \cdot \mathrm{poly}(V,T)\big)$ *per round on average, provided that the seller is an* ϵ_S*-canonical learner with* $\epsilon_S \le O(\epsilon_M) = O\big(m^{-\frac{6}{19}} \cdot \mathrm{poly}(V,T)\big)$.

If the seller explores in a round such that the corresponding UCB and LCB differs by not only ϵ_S but by at least ϵ_M, we call it a *major exploration*. We first upper bound the number of rounds that involve such major explorations in the following lemma.

LEMMA 4.1. *The expected number of rounds that are major explorations is at most* $\tilde{O}(TV\epsilon_M^{-3})$.

Proof. Every time that the seller makes a major exploration on some price p in a round, say, in response to a segment represented by a point $\boldsymbol{x} \in \Delta$, the gap between the UCB and the LCB is at least ϵ_M. Then, the expected gap between $\widetilde{U}(p,t)$ and $\widetilde{L}(p,t)$ is also at least ϵ_M when type t is sampled according to $\boldsymbol{x}$. This implies that, with probability at least $\frac{\epsilon_M}{2}$, the realized type t actually has a gap of at least $\frac{\epsilon_M}{2}$ between $\widetilde{U}(p,t)$ and $\widetilde{L}(p,t)$. Note that for any type t, and any price p, this cannot happen by more than $O\big(\epsilon_M^{-2}\big)$ times by the definitions of $\widetilde{U}(p,t)$ and $\widetilde{L}(p,t)$. Hence, the expected number of times that the seller explores cannot exceed $O\big(TV\epsilon_M^{-3}\big)$ times. ∎

We now prove Theorem 4.1.

Case 1: Seller picks an undesirable price. Suppose the seller fails to pick a price that is at least $14\epsilon_M$-optimal w.r.t. the optimistically chosen distributions $\boldsymbol{\mathcal{F}}^*$. Instead, she chooses a price p. Then, either she is not exploiting in the sense of the definition in Section 4, which cannot happen in more than an ϵ_S fraction of the rounds, or the UCB of p is at least the maximum UCB among all prices less ϵ_S. This is larger than the expected revenue induced from the optimistically chosen distribution $\boldsymbol{\mathcal{F}}^*$ by at least $14\epsilon_M - \epsilon_S > \epsilon_M$, by the assumption that p is not $14\epsilon_M$-optimal. Note that the latter is weakly larger than the LCB. Hence, we conclude that the UCB and LCB differs by at least ϵ_M, which means that this is a major exploration that cannot happen in more than an $m^{-1} \cdot O\big(TV\epsilon_M^{-3}\big) = O(\epsilon_I)$ fraction of the rounds (Lemma 4.1).

Case 2: Sale probability is lower than expected. Next, consider the case when the seller picks a price that is indeed $14\epsilon_M$-optimal w.r.t. the optimistically chosen distributions $\boldsymbol{\mathcal{F}}^*$, but the sale probability of the price given by the true distributions is smaller than by $\boldsymbol{\mathcal{F}}^*$ by more than ϵ_I. In this case, note that both sale probabilities are bounded between the UCB and LCB; we again conclude that there is a gap of at least $\epsilon_I > \epsilon_M$ between them. Hence, this is a major exploration which cannot happen in more than $O(\epsilon_I)$ fraction of the rounds.

Case 3: Everything goes as expected. Finally, consider the good case, when the seller indeed picks a price that is at least $14\epsilon_M$-optimal least 4 w.r.t. the optimistically chosen distributions $\boldsymbol{\mathcal{F}}^*$, and that the sale probability of the price given by the true distributions is at least that by $\boldsymbol{\mathcal{F}}^*$ less ϵ_I. Then, by Lemma 3.1, we get that the expected objective in this round is at least $\mathrm{Opt} - O(\epsilon_I)$.

Algorithm 4 Segmentation algorithm in the bandit model (in each round)

1. Let $\mathcal{M}(t)$ be the set of MHR-like distributions with support $[V]$ such that the quantile of each value $v \in [V]$ is between the corresponding UCB and LCB.
2. Let $\mathcal{F}^*(t) \in \mathcal{M}(t)$, $t \in [T]$, be such that $\text{Opt}(\boldsymbol{\mathcal{F}}^*)$ is maximized.
3. Find $(\Sigma^*, \mathcal{G}^*)$, the optimal segmentation w.r.t. $\boldsymbol{\mathcal{F}}^*$, via the algorithm in Section 2.
4. Construct $(\mathcal{G}^r, \Sigma^r)$, a robustified version of $(\Sigma^*, \mathcal{G}^*)$, via Algorithm 2, using ϵ_M in place of ϵ_S.

Since the first two cases cannot happen in more than an $O(\epsilon_I)$ fraction of the rounds, the bound stated in Theorem 4.1 follows.

5 Further Related Work

The problem of price discrimination is highly related to *screening* in games of asymmetric information, pioneered by [37, 13], where the less informed player moves first in hopes of combating adverse selection. In our setting, the seller wishes to screen buyers by charging different prices depending on the buyer's value. The intermediary's segmentation allows the seller to screen more effectively. The intermediary themselves face a signaling problem, as their choice of segmentation is effectively a signaling scheme to the seller.

As such, our work is related to the broad literature on signaling and information design, where a mediator designs the information structures available to the players in a game [5]. A special case of this is known as *Bayesian persuasion* [28, 17]: an informed sender (here the intermediary) sends a signal about the state of the world to a receiver (here the seller), who must take an action that determines the payoff of both parties. The goals of the sender and receiver may not be aligned, so the sender must choose a signaling scheme such that the receiver's best response still yields high payoff for the sender. See [16, 6] for surveys on these topics.

Our results for online learning are also related to work on iteratively learning prices [8, 31, 11, 34, 9, 25]. Both lines of work consider the seller's problem of incentive design or learning, but do not have an intermediary or market segmentation component. Our model is also somewhat related to the literature on *dynamic mechanism design*, which considers the incentive guarantees of multi-round auctions where the same bidders may participate in multiple rounds. [7, 27, 35] gave truthful dynamic mechanisms for maximizing social welfare and revenue.

Our results are related to recent work on incentivizing exploration in a bandit model [19, 29, 30]. These papers typically model a myopic decision-maker in each round, and an informed non-myopic principle who can influence the decision-maker to explore rather than exploit. In our setting, the seller is myopic decision-maker who sets prices, and the intermediary can influence that decision by changing the segmentation. The previous results do not directly apply to our setting, as an action corresponds to setting a price in the observed segment. Hence there are exponentially many actions, so one should not hope for polynomial run time or good regret guarantees by directly applying those results. Additionally, the intermediary chooses the segmentation but the observed segment is chosen randomly, so the intermediary cannot force the seller to play any particular action.

6 Future Work

We view our results as initiating a new line of work on algorithmic price discrimination under partial information. We believe there are many promising open problems left to be explored in this direction, and hope this paper inspires future work under other informational models and market environments. We now present some of the most interesting directions for future work.

Competitive Markets. This paper and [4] consider a monopolist seller, which is a good fit for something like an ad exchange. In many online marketplaces the sellers are in a competitive rather than a monopolistic setting. The products are differentiated so sellers can exert some pricing power, which still incurs deadweight loss. It would be very interesting to extend this theory of price discrimination to such competitive markets.

Strategic Buyers. When a seller uses past buyer behavior in the form of auction bids or purchase decisions to decide future prices, and a buyer has repeated interactions with such a seller, the buyer may be incentivized to strategize. Even if each interaction in isolation is strategyproof, the buyer may forgo winning an earlier auction in order to get a lower price in the future. When each buyer represents an insignificant fraction of the entire market, techniques from differential privacy can address this issue [25]. In this paper we ignore this

strategic aspect and assume that buyers are myopic. It would be very interesting to get results analogous to [25], since their techniques do not readily extend to our model.

Worst Case Model. In online learning, even for arbitrary sequence of inputs, we can often get a regret guarantee matching that for an i.i.d. input sequence. In particular this is true for a monopolistic seller learning an optimal price or an optimal auction [9, 8]. It is tempting to conjecture that the same holds for our setting as well, but we run into difficulties even modeling the problem. How does the seller behave in such a scenario? In this paper, we modeled the seller behavior based on the underlying distribution. Defining a seller behavior model in the absence of such a distribution that is both reasonably broad and allows regret guarantees in the worst case setting is an interesting challenge.

References

[1] Maria-Florina Balcan, Avrim Blum, Jason D Hartline, and Yishay Mansour. Mechanism design via machine learning. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, pages 605–614. IEEE, 2005.

[2] Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*, pages 2083–2091, 2016.

[3] Richard E Barlow, Albert W Marshall, Frank Proschan, et al. Properties of probability distributions with monotone hazard rate. *The Annals of Mathematical Statistics*, 34(2):375–389, 1963.

[4] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.

[5] Dirk Bergemann and Stephen Morris. Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review*, 106(5):586–591, 2016.

[6] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. Cowles Foundation Discussion Paper No. 2075R3, 2018.

[7] Dirk Bergemann and Juuso Valimaki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

[8] Avrim Blum and Jason D Hartline. Near-optimal online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1156–1163. Society for Industrial and Applied Mathematics, 2005.

[9] Sebastien Bubeck, Nikhil R Devanur, Zhiyi Huang, and Rad Niazadeh. Online auctions and multi-scale online learning. In *Proceedings of the 18th ACM Conference on Economics and Computation*, pages 497–514. ACM, 2017.

[10] Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science*, pages 516–527. IEEE, 2017.

[11] Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.

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

[13] Pascal Courty and Li Hao. Sequential screening. *The Review of Economic Studies*, 67(4):697–717, 2000.

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

[15] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[16] Shaddin Dughmi. Algorithmic information structure design: A survey. *SIGecom Exchanges*, 15(2):2–24, 2017.

[17] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 412–425, 2016.

[18] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 736–745, 2007.

[19] Peter Frazier, David Kempe, Jon Kleinberg, and Robert Kleinberg. Incentivizing exploration. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 5–22, 2014.

[20] Yannai A. Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing*, pages 856–868, 2017.

[21] Yannai A Gonczarowski and S Matthew Weinberg. The sample complexity of up-to-ε multi-dimensional revenue maximization. In *Proceedings*

of the 59th IEEE Annual Symposium on Foundations of Computer Science, pages 416–426. IEEE, 2018.

[22] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st annual ACM Symposium on Theory of Computing*, 2019.

[23] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th International Conference on World Wide Web*, pages 189–198. ACM, 2008.

[24] Jason Hartline and Samuel Taggart. Sample complexity for non-truthful mechanisms. In *Proceedings of the 20th ACM Conference on Economics and Computation*, 2019.

[25] Zhiyi Huang, Jinyan Liu, and Xiangning Wang. Learning optimal reserve price against non-myopic bidders. In *Advances in Neural Information Processing Systems*, pages 2038–2048, 2018.

[26] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. *SIAM Journal on Computing*, 47(3):651–674, 2018.

[27] Sham Kakade, Ilan Lobel, and Hamid Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 64(4):837–854, 2013.

[28] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[29] Yishay Mansour, Aleksandrs Slivkins, and Vasilis Syrgkanis. Bayesian incentive-compatible bandit exploration. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 565–582, 2015.

[30] Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. Bayesian exploration: Incentivizing exploration in bayesian games. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 661–661, 2016.

[31] Andres M Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning*, pages 262–270, 2014.

[32] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pages 1298–1318, 2016.

[33] Jamie H Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems*, pages 136–144, 2015.

[34] Renato Paes Leme, Martin Pal, and Sergei Vassilvitskii. A field guide to personalized reserve prices. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1093–1102, 2016.

[35] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A Myersonian approach. *Econometrica*, 82(2):601–653, 2014.

[36] Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 1–18. ACM, 2016.

[37] Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–0374, 1973.

A Further Examples

A.1 Benefit of Segmentation We now reproduce an example from [4] that shows how a segmentation can eliminate deadweight loss.

Example. The value set $[V] = \{1, 2, 3\}$ is identical to the type set $[T]$, and each distribution $\mathcal{F}(t)$ is a pointmass at t. The prior distribution over values/types is uniform. The monopoly reserve of the uniform prior distribution is 2. When the seller does not segment the market and posts the monopoly reserve price, revenue is 4/3, consumer surplus is 1/3, and the deadweight loss is 1/3. □

Consider instead the following segmentation with $\Sigma = \{\sigma_1, \sigma_2, \sigma_3\}$ where $\mathcal{G}(1) = (1, 0, 0)$; $\mathcal{G}(2) = (1/3, 1/6, 1/2)$; and $\mathcal{G}(3) = (2/3, 1/3, 0)$. Recall that $\mathcal{G}(t)$ is the distribution of signals sent by the intermediary upon observing type t. This signaling scheme generates three market segments, one corresponding to each of $\sigma_1, \sigma_2, \sigma_3$. The seller can compute the conditional distribution of values within each segment, and will post the monopoly price for that distribution. Within segment σ_1, the distribution of values is: x_1 with probability 1/2, x_2 with probability 1/6, and x_3 with probability 1/3. This happens to be the equal revenue distribution on values $\{1, 2, 3\}$. Within segment σ_2, only buyers of types 2 or 3 will be present, since a buyer of type 1 will never be mapped into this segment. The conditional distribution of values in this segment is: x_2 with probability 1/3 and x_3 with probability 2/3, which also happens to be the equal revenue distribution on values $\{2, 3\}$. Type 2 is the only type with positive probability of being mapped to segment σ_3, so the value distribution in segment σ_3 is a point mass on value 2.

Since each market segment has an equal revenue distribution of values, the seller can maximize his profit by posting any price in the support of that distribution.

For the sake of this example, we will assume the seller breaks ties by posting the lowest optimal price.[3] In market segment σ_1, the seller will post price $p = 1$, which will generate revenue of 1 and consumer surplus of $(1/6) \times 1 + (1/3) \times 2 = 5/6$. In σ_2, the seller will post price $p = 2$, which will generate revenue of 2 and consumer surplus of $(2/3) \times 1 = 2/3$. Finally, in market segment σ_3, the seller will post price $p = 2$, which will generate revenue of 2 and consumer surplus of 0.

Given the prior distribution of values, one can also compute the probability of the intermediary generating each market segment. In this case, segment σ_1 is drawn with probability 2/3, segment σ_2 is drawn with probability 1/6, and segment σ_3 is drawn with probability 1/6. Combining these, we see that the expected revenue of this segmentation is $(2/3) \times 1 + (1/6) \times 2 + (1/6) \times 2 = 4/3$, and the consumer surplus is $(2/3) \times (5/6) + (1/6) \times (2/3) + (1/6) \times 0 = 2/3$. This segmentation has zero deadweight loss, which means that full economic efficiency is achieved.

A.2 Noisy Types Now consider the above example with noise. This example has been briefly discussed in Section 2. This subsection includes more details.

Example. Given a noise level of $1 - z$, each type $t \in [T]$ corresponds to the distribution $\mathcal{F}(t)$ given by

$$\mathbf{Pr}_{v \sim \mathcal{F}(t)}[v] = \begin{cases} z & \text{if } v = t, \\ \frac{1-z}{2} & \text{otherwise.} \end{cases}$$

□

Note that when $z = 1/3$, all $\mathcal{F}(t)$ equal the uniform prior, and no further (non-trivial) segmentation is possible. At the other extreme, when $z = 1$ each $\mathcal{F}(t)$ is a pointmass at t, which corresponds to Example 1. For $z = 0.49$, it turns out that market segmentation cannot help at all. Any segmentation will result in the seller always setting the monopoly reserve price of 2, and the result is the same as no segmentation! As in Example 1, revenue from no segmentation is 4/3, consumer surplus is 1/3, and the deadweight loss is 1/3.

On the other hand, when $z = 0.8$ segmentation is possible, but it is no longer possible to achieve full economic surplus due to the noisy types. If the intermediary implemented the same segment map as in Example 1, the resulting conditional value distributions in each segment would no longer be equal revenue distributions because types are no longer perfectly correlated with values. If the intermediary perfectly segments the market by types using the deterministic segment map $\mathcal{G}(t) = \sigma_t$, this will result in revenue $2 - \frac{4(1-z)}{3} = 1.7333$, consumer surplus $\frac{2(1-z)}{3} = 0.1333$, and deadweight loss $\frac{2(1-z)}{3} = 0.1333$. Note that under the noiseless setting of Example 1, this segmentation would have allowed the seller to perfectly price discriminate, resulting in revenue 2, consumer surplus 0, and deadweight loss 0. These changes in economic outcome are a direct result of the fact that types are only a noisy signal of the buyer's value. For example in market segment σ_3, the seller will still set monopoly price $p = 3$, but only a $(1-z)$-fraction of the segment will have value 3 and purchase the item.

A.3 Impossibility of [4] Style Characterization Recall the simplex view and Figure 2. In the setting of [4], the segmentation only consists of points in X_{p^*} where p^* is the monopoly price of the prior distribution (which is the white region X_2 in Figure 2a). As mentioned earlier, [4] assume that ties for monopoly price are broken in favor of the lowest price. This implies that even though the market is now segmented, the price in each segment is at most p^*. In other words, segmentation can only lower prices!

One could hope for a similar characterization even in the noisy signal case: after all we can still write $\boldsymbol{\tau}$ as a convex combination of points in the intersection of X_{p^*} (the white region) and $\Delta([T])$ (the convex hull of the type distributions, depicted as the triangles in the interior of the simplex in Figure 2a), as can be seen in Figure 2b. Unfortunately we show that this is not without loss of generality in the following example. In particular, the example shows that restricting to such segments may lead to a strictly smaller social welfare than otherwise.

Example. In this example, there are two types, corresponding to the points $x^{\{1,2\}}$ and $x^{\{3\}}$, i.e., type 1 is the uniform (equal revenue) distribution over $\{1, 2\}$, and type 2 is the point mass on 3. The distribution over types is still uniform, which once again gives the same prior distribution $\boldsymbol{\tau}$ as before. □

The difference now is that we can only pick points from the line joining these two points, as depicted in Figure 5a. Figure 5b shows the social welfare obtained as we move along this line. The point mass on the left corresponds to $x^{\{1,2\}}$, for which we assume that the seller picks a price of 1. The left segment corresponds to the white region X_2, and the right to the dark gray region X_3. Using only the white region means using either the point mass on the left or the left segment. The dotted line in Figure 5b shows the social welfare obtained from taking a convex combination of the end points. This corresponds to the segmentation where the intermediary simply reveals the type that he observes.

[3] This can be strictly enforced using arbitrarily small perturbations.

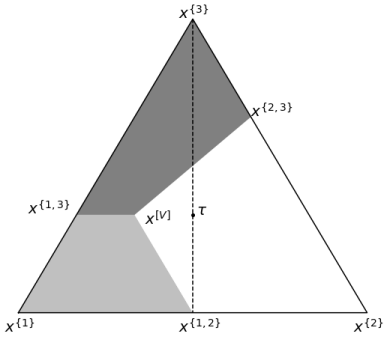

(a) The two types correspond to the points $x^{\{1,2\}}$ and $x^{\{3\}}$. The line joining them is the convex hull $\Delta([T])$ from which we can pick our segments.

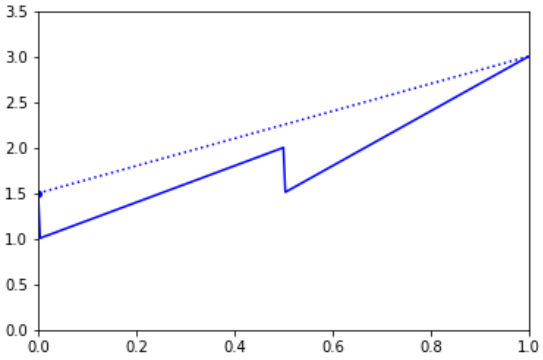

(b) The solid line represents the social welfare as you move along the line from $x^{\{1,2\}}$ to $x^{\{3\}}$. The dotted line corresponds to the social welfare of the corresponding convex combination of the two end points

Figure 5: Example in Appendix A.3

As can be seen from the figure, this is strictly better than restricting to points in X_2.

Nonetheless, we can add this as an additional constraint if so desired (at some loss in the objective). Our algorithm extends to handle this easily: just skip iterating over X_p for prices $p > p^*$.

A.4 Unbounded Sample Complexity in the General Case So far we have assumed that the prior distribution is given to us as input and is common knowledge to all players. How does this happen? What if you only have samples from the distribution? How many samples do you need in order to get within an ϵ of the optimum? These questions have been studied under 'sample complexity of auction design' quite intensively in the last few years. (See Section 5 for details on this line of work.) Only recently, the optimal sample complexity of single item auctions has been resolved [22]. In this paper, we consider the sample complexity of market segmentation.

Unlike auction design, here the seller sets the price to maximize revenue but the intermediary's objective may be something different, such as consumer surplus. For an equal revenue distribution, statistically the samples will indicate that the high price is revenue-optimal with a significant probability. This is still true even for distributions that are "close to" equal revenue where the low price is strictly revenue-optimal. Higher prices correspond to lower social welfare, since the buyer is less likely to purchase the good; thus the segmentation based on the samples could have a much smaller consumer surplus compared to the optimal segmentation for the distribution. This is particularly problematic because as saw earlier, the optimal segmentation only picks segments with equal revenue distributions. (Recall that the vertices of the colored regions correspond to equal revenue distributions.) We demonstrate this via the following example.

Example. Consider the distribution on $V = \{1, 2\}$ where the probability of seeing 1 is $0.5 + \delta$, for some very small $\delta > 0$. We will make the segmentation problem trivial: there is only one type, i.e., the intermediary observes no signal. The monopoly price is 1, and consumer surplus is therefore $0.5 - \delta$; this is trivially the optimal consumer surplus we can obtain via segmentation. □

When drawing multiple samples from this distribution, there is a constant probability of seeing see more 2s than 1s. In this case, the seller sets a monopoly price of 2, leading to a consumer surplus of 0. For any bounded function f of ϵ, we can set δ small enough such that with $f(\epsilon)$ samples, this happens with probability $> \epsilon$; we cannot therefore hope to get to within ϵ of the optimum. This example shows that we need a stronger assumption on the input distributions, as compared to those for single item auctions. The standard assumptions there are regularity and boundedness, both of which are satisfied by the distribution in the example above.

In the above example, we assumed that the seller

sets the monopoly price on the empirical distribution on the samples. This is not necessarily an accurate assumption. If the seller follows the literature on the sample complexity of single item auctions, he should consider a robust or a regularized version of the empirical distribution. The seller might also have some additional sources of information that allow him get an even more accurate estimate. We make a mild assumption on the seller behavior: that his beliefs are at least as accurate as the intermediary's estimate from the samples. For a formal definition and more detailed discussion, see Section 3.

Given the discussion so far, it may be tempting to assume that the type distributions are far from the boundaries of X_vs. This is too strong when you consider larger value ranges: two prices, say p and $p+1$, may have almost identical quantiles and, thus, revenues (which means the distribution close to the boundary between X_p and X_{p+1}), but this is not a problem if they both give similar consumer surpluses. We make a milder assumption, which is a discrete version of the monotone hazard rate (MHR) assumption.[4] Here, the naive way to generalize MHR to discrete distributions is to use the functional form. This doesn't seem to work; we instead require the distribution to have some of the properties that continuous MHR distributions are known to satisfy. Once again, a detailed discussion on this with formal statements are in Section 3.

B Missing Proofs in Section 2

Proof of Lemma 2.1 *From Segmentations to Distributions.* Given any segmentation $(\Sigma, \mathcal{G})$, consider the set X that is the union of the following points. For any segment $\sigma \in \Sigma$, let there be a point $\boldsymbol{x}(\sigma) \in \Delta([T])$ such that:

$$x_t(\sigma) = \frac{\mathbf{Pr}_{\mathcal{T}}[t] \cdot \mathbf{Pr}[\mathcal{G}(t) = \sigma]}{\mathbf{Pr}_{\mathcal{S}}[\sigma]} \ .$$

The probability distribution over X is defined as $\mu(\boldsymbol{x}(\sigma)) := \mathbf{Pr}_{\mathcal{S}}[\sigma]$. It is easy to verify that this is indeed a probability distribution and that it satisfies the expectation constraint (2.1).

From Distributions to Segmentations. Given any (X, μ) that satisfies (2.1), consider the following segmentation $(\Sigma, \mathcal{G})$. Let Σ be such that there is some bijection between Σ and X, where $\sigma \in \Sigma$ is uniquely mapped to $\boldsymbol{x}(\sigma) \in X$. For any type t, let $\mathcal{G}(t)$ follow a distribution given by

$$\mathbf{Pr}[\mathcal{G}(t) = \sigma] = \frac{x_t(\sigma)\mu(\boldsymbol{x})}{\mathbf{Pr}_{\mathcal{T}}[t]} \ .$$

[4] MHR distributions have non-decreasing hazard rate $\frac{f(x)}{1-F(x)}$, where f and F are the pdf and the cdf of the distribution respectively.

The pair (X, μ) satisfies (2.1) implies that $\mathcal{G}(t)$ are valid probability distributions:

$$\forall\ t \in [T], \sum_{\sigma \in \Sigma} \mathbf{Pr}[\mathcal{G}(t) = \sigma] = \sum_{\sigma \in \Sigma} \frac{x_t(\sigma)\mu(\boldsymbol{x})}{\mathbf{Pr}_{\mathcal{T}}[t]} = 1 \ .$$

Proof of Lemma 2.2 Consider any optimal segmentation with an arbitrary, or even unbounded number of segments. We show how to transform it into a segmentation with the same objective and yet having at most one segment within each region X_p, replacing the distribution conditioned on $\boldsymbol{x} \in X_p$ by its expectation.

Consider a specific area X_p. Let $\mu_p = \mathbf{Pr}_{\mu}\big[\boldsymbol{x} \in X_p\big]$ be the probability of realizing a segment in X_p. Let $\boldsymbol{x}_p = \mathbf{E}_{\mu}\big[\boldsymbol{x} \,|\, \boldsymbol{x} \in X_p\big]$. By the convexity of X_p, $\boldsymbol{x}_p$ is also in X_p.

Further, recall that when we restrict our domain to points $\boldsymbol{x} \in X_p$, we have that the revenue and consumer surplus are both linear function because the seller chooses a fix price p within X_p. Therefore, removing all segments in X_p and adding a new segment at $\boldsymbol{x}_p$ with probability mass μ_x do not change the objective. Repeating this process for all areas proves the lemma.

Proof of Theorem 2.1 Consider the following more direct mathematical program that uses $\boldsymbol{x}_p$'s and $\mu(\boldsymbol{x}_p)$'s as the variables:

(B.1)

$$\max \sum_{p \in [V]} \mu(\boldsymbol{x}_p)\Big(\lambda \mathrm{Rev}\big(\mathcal{F}(\boldsymbol{x}_p)\big) + (1-\lambda)\mathrm{CS}\big(\mathcal{F}(\boldsymbol{x}_p)\big)\Big)$$

$$\text{s.t.} \quad \forall\ p \in [V], \boldsymbol{x}_p \in X_p \text{ and } \sum_{p \in [V]} \mu(\boldsymbol{x}_p)\boldsymbol{x}_p = \boldsymbol{\tau} \ .$$

We need to show that (i) there is a one-to-one mapping between the variable space of program (B.1) and that of LP (2.2), and (ii) under this mapping, the above program becomes LP (2.2).

The mapping from $\boldsymbol{x}_p$ and $\mu(\boldsymbol{x}_p)$ to $\boldsymbol{z}_p$ is already given, i.e., $\boldsymbol{z}_p = \mu(\boldsymbol{x}_p)\boldsymbol{x}_p$. The other direction goes as follows. Given any $\boldsymbol{z}_p$, let $\mu(\boldsymbol{x}_p) = \|\boldsymbol{z}_p\|_1$ and $\boldsymbol{x}_p = \boldsymbol{z}_p/\|\boldsymbol{z}_p\|_1$. Here, we use the fact that $\boldsymbol{x}_p$ lies on the probability simplex.

Under this mapping, the objective of (B.1) becomes that of LP (2.2) due to linearity of $\mathrm{Rev}\big(\mathcal{F}(\cdot)\big)$ and $\mathrm{CS}\big(\mathcal{F}(\cdot)\big)$ for any fixed p. Finally, the equivalence of the constraints follow by the definition of the mapping.

C Poly-time Approximate Projection to MHR-like Distributions

This subsection explains how we can find an MHR-like distribution that is $6\epsilon_S$-close to the true distribution in polynomial time, given an empirical distribution that is ϵ_S-close. Our algorithm will guess which price is the

Algorithm 5 Find a Nearby MHR-like Distribution

Input: A distribution $\mathcal{E}$ that is ϵ_S-close to some MHR-like distribution $\mathcal{F}$ in terms of d_q; the optimal price p^* of the underlying MHR-like value distribution.

Output: An MHR-like distribution $\widetilde{\mathcal{E}}$ that is $6\epsilon_S$-close to $\mathcal{F}$ in terms of norm d_q.

1. Construct a sequence of distributions as follows:
   (a) $\mathcal{E}_1$: Subtract ϵ_S from the quantile of every value; round it to 0 if it becomes negative.
   (b) $\mathcal{E}_2$: Increase the quantile of p^* to that in $\mathcal{E}$ plus ϵ_S.
   (c) $\mathcal{E}_3$: Increase the quantile of each value until it meets the convex hull of the revenue curve.
2. Return $\widetilde{\mathcal{E}} = \mathcal{E}_3$.

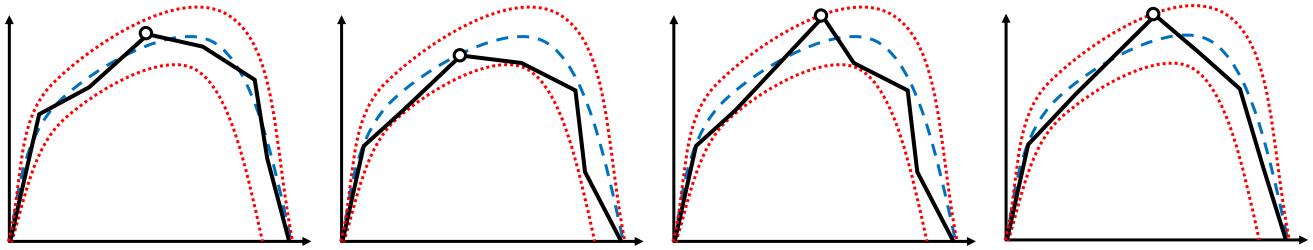

(a) $\mathcal{E}$ (empirical) (b) $\mathcal{E}_1$ (dominated) (c) $\mathcal{E}_2$ (strengthening p^*) (d) $\widetilde{\mathcal{E}} = \mathcal{E}_3$ (ironing)

Figure 6: Illustrative figures (in terms of the revenue curves in the quantile space) of how Algorithm 5 finds a nearby MHR-like distribution. The dashed curve is the revenue curve of the true distribution. The dotted curves are those obtain by subtracting/adding $O(\epsilon_S)$ to the quantiles of every value (rounded to 0 or 1 if necessary). The bold black curve corresponds to the distributions maintained by the algorithm.

optimal one w.r.t. the true distribution by brute-force. Given each guess p^*, we try to find a nearby MHR-like distribution conditioned on p^* being optimal using Algorithm 5. It constructs a sequence of distributions with the last one being the desired output, provided that the guess of p^* is correct.

The first distribution $\mathcal{E}_1$ is obtained by subtracting ϵ_S from the quantiles of all values, rounding up to 0 if necessary. The quantile of each value v w.r.t. the empirical distribution $\mathcal{E}$ is within a $[-\epsilon_S, \epsilon_S]$ window near that w.r.t. the true distribution; in contrast, the quantile w.r.t. $\mathcal{E}_1$ is within a $[-2\epsilon_S, 0]$ window. Distribution $\mathcal{E}_1$ is dominated by the true distribution in the sense of first-order stochastic dominance and, hence, is called the dominated empirical distribution (e.g., [22, 36]).

Then, we construct the second distribution $\mathcal{E}_2$ by increasing the quantile of the conjectured monopoly price p^* to ensure that it is at least as large as in the true distribution. The purpose of this step is to ensure the strong concavity property near the monopoly price.

Finally, we run an ironing step to restore concavity. The result, i.e., $\mathcal{E}_3$, is the final output. See Figure 6 for an illustrative picture of relations between the revenue curves of the sequence of distributions constructed by the algorithm.

LEMMA C.1. *Given any distribution $\mathcal{E}$ that is ϵ_S-close to an MHR-like distribution $\mathcal{F}$, and the monopoly price $p*$ of $\mathcal{F}$, Algorithm 5 computes in polynomial time an MHR-like distribution $\widetilde{\mathcal{E}}$ that is $6\epsilon_S$-close to $\mathcal{F}$ in terms of norm $d_q(\cdot,\cdot)$.*

Proof. The running time is trivial, noting that convex hulls can be computed in polynomial time. It remains to show that the output is an MHR-like distribution, and is $6\epsilon_S$-close to $\mathcal{E}$.

MHR-like – concavity: This part follows by the definition of the algorithm 5 (step 1c).

MHR-like – strong concavity near monopoly price: We will show that (1) p^* is the monopoly price of the final distribution $\widetilde{\mathcal{E}}$, and that (2) for any price p whose quantile w.r.t. $\widetilde{\mathcal{E}}$ is q, we have:

$$\left(1 - \tfrac{1}{4}(q^* - q)^2\right) \cdot \mathrm{Rev}\left(p^*, \widetilde{\mathcal{E}}\right) - \mathrm{Rev}\left(p, \widetilde{\mathcal{E}}\right) \geq 0 \ ,$$

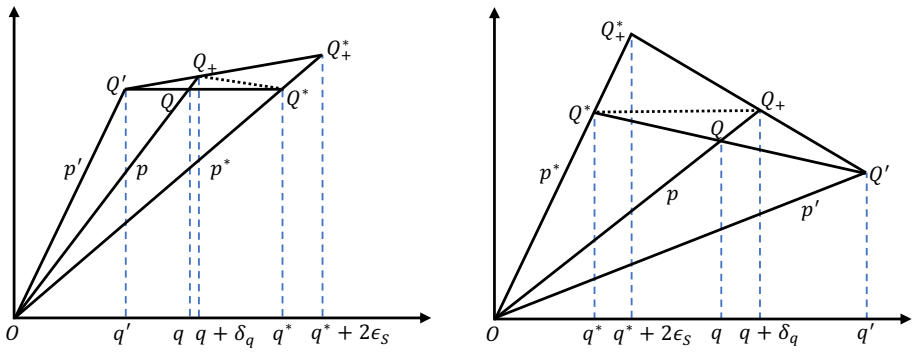

Figure 7: Illustrative pictures for the proof of Lemma C.1

$$\left(1-\tfrac{1}{4}(q^*-q)^2\right)\cdot p^*q^*-pq\geq 0\ .$$

We first prove them for distribution $\mathcal{E}_2$. Note that the quantiles of all values other than p^* are at most their counterparts in the true distribution $\mathcal{F}$ (step 1a), while the quantile of p^* is at least that in $\mathcal{F}$ (step 1b). Hence, p^* is also the monopoly price for $\mathcal{E}_2$. Further, the LHS of the above inequality is increasing in q^* (fixing any p, p^*, and q), and decreasing in q (fixing any p, p^*, and q^*). Since that the inequality holds for the true distribution, and that q^* weakly increases and q weakly decreases compared to the true distribution, it also holds for $\mathcal{E}_2$. Finally, we argue that, moving from $\mathcal{E}_2$ to $\widetilde{\mathcal{E}}=\mathcal{E}_3$, i.e., the ironing step, will not make p^* suboptimal, as the highest point of the revenue curve will not be ironed. The inequality will also continue to hold, because the first term is a concave function of q that stays the same while the second term becomes the convex hull of the counter part in $\mathcal{F}_2$. If a concave function dominates another one, it dominates its convex hull as well.

MHR-like – large monopoly sale probability: It follows from the fact that the quantile, i.e., the sale probability, of p^* weakly increases by our construction.

MHR-like – small revenue and welfare gap: First, imagine that the quantile of p^* is only increase to that in the true distribution $\mathcal{F}$, while the quantiles of other values weakly decrease by our construction. Then, the optimal revenue stays the same and the social welfare weakly decreases and therefore the small gap property continues to hold. Then, we further increase the quantile of p^* to that in our final output distribution $\widetilde{\mathcal{E}}$. As a result, the social welfare as well as the optimal revenue increase by the same amount. This will not change the fact that their gap is small.

$6\epsilon_S$-close: This is true by definition up to distribution $\mathcal{E}_2$. It remains to show that the ironing step will not make the quantile of some value v exceeds its quantile w.r.t. the true distribution $\mathcal{F}$ by more than $2\epsilon_S$. For any ironing that does not involved p^*, this is trivial: since both endpoints of the ironed interval are below the revenue curve of the true distribution, which is concave, the entire ironed interval is below that as well.

Next, consider an ironed interval with p^* being one of the endpoints. Suppose p^* is the right end point, i.e., there is another price $p'>p^*$. See Figure 7a for an illustrative picture of the argument below. The quantile of p' w.r.t. $\widetilde{\mathcal{E}}$ is upper bounded by that w.r.t. the true distribution $\mathcal{F}$, which we denote as q'; The quantile of p^*, on the other hand, is upper bounded by that w.r.t. the true distribution, denoted as q^*, plus $2\epsilon_S$ by our construction. Then, consider any price p between p^* and p'. We can lower bound the quantile of p as a function of p', p^*, q', and q^*, by the concavity of the revenue curve of $\mathcal{F}$. We denote this quantile by q: $Q'=(q',p'q')$, $Q=(q,pq)$, and $Q^*=(q^*,p^*q^*)$ lie on the same line. Finally, p's quantile in the final distribution, by definition, is determined by the ironing step. We denote this quantile by $q+\delta_q$: $Q'=(q',p'q')$, $Q_+=(q+\delta_q,p(q+\delta_q))$, and $Q^*_+=(q^*+2\epsilon_S,p^*(q^*+2\epsilon_S))$ lie on the same line. It suffices to bound δ_q. To do so, further let $O=(0,0)$ be the origin. We have:

$$\text{(C.2)}\quad \frac{\delta_q}{q}=\frac{\text{area}(Q'Q_+Q^*)}{\text{area}(OQ'Q^*)}\leq\frac{\text{area}(Q'Q^*_+Q^*)}{\text{area}(OQ'Q^*)}=\frac{2\epsilon_S}{q^*}\ .$$

The other case, when p^* is the left endpoint is almost verbatim up to the point of Eqn. (C.2). See Figure 7b for an illustrative picture. The only catch is that we now have $q'\geq q^*$. Fortunately, q', which is at most 1, cannot much bigger since $q^*\geq\frac{1}{e}$ (Condition 3 of Definition 3.1). As a result, we get a weaker bound that $\delta_q\leq e\cdot 2\epsilon_S\leq 6\epsilon_S$. ∎