Robust Learning of Optimal Auctions

Wenshuo Guo

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2022-265
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2022/EECS-2022-265.html

December 16, 2022

Robust Learning of Optimal Auctions

by Wenshuo Guo

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

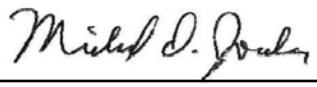

Professor Michael I. Jordan
Research Advisor

Dec 6, 2022

(Date)

* * * * * * *

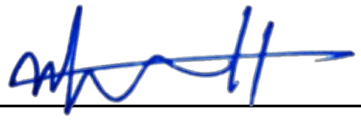

Professor Nika Haghtalab
Second Reader

Dec 14, 2022

(Date)

Robust Learning of Optimal Auctions

Wenshuo Guo[†], Michael I. Jordan[†,‡], Manolis Zampetakis[†]

[†]Department of Electrical Engineering and Computer Sciences,
[‡]Department of Statistics,
University of California, Berkeley

December 6, 2022

Abstract

We study the problem of learning revenue-optimal multi-bidder auctions from samples when the samples of bidders' valuations can be adversarially corrupted or drawn from distributions that are adversarially perturbed. First, we prove tight upper bounds on the revenue we can obtain with a corrupted distribution under a population model, for both regular valuation distributions and distributions with monotone hazard rate (MHR). We then propose new algorithms that, given only an "approximate distribution" for the bidder's valuation, can learn a mechanism whose revenue is nearly optimal simultaneously for all "true distributions" that are α-close to the original distribution in Kolmogorov-Smirnov distance. The proposed algorithms operate beyond the setting of bounded distributions that have been studied in prior works, and are guaranteed to obtain a fraction $1 - O(\alpha)$ of the optimal revenue under the true distribution when the distributions are MHR. Moreover, they are guaranteed to yield at least a fraction $1 - O(\sqrt{\alpha})$ of the optimal revenue when the distributions are regular. We prove that these upper bounds cannot be further improved, by providing matching lower bounds. Lastly, we derive sample complexity upper bounds for learning a near-optimal auction for both MHR and regular distributions.

1 Introduction

Optimal auctions play a crucial role in economic theory, with a wide range of applications across various industries, public sectors, and online platforms [4, 15, 16, 18, 21, 22]. In such auctions, pricing mechanisms need to be determined by the auction designer so as to satisfy various desired goals, such as revenue maximization and incentive compatibility. Often this determination is made based on information about the buyers that is assumed to be available a priori. For example, in a standard valuation model, each bidder has a valuation over the available items, and if the sellers knows the distribution of these valuations, they could design an optimal auction which maximizes the revenue.

Arguably the fundamental difficulty in the design of optimal auctions is that real valuations are private and unknown to the auction designer. Consider specifically the problem of selling one item to multiple buyers. Suppose that we model the buyers' valuations as arising as independent draws from buyer-specific prior distributions. In this scenario, what is the optimal mechanism in terms of the expected revenue? This problem was solved by Myerson [21] through a characterization of *virtual value functions*. In particular, we can define a virtual value function of each buyer based on their prior distributions. An optimal auction then lets the buyer with the largest non-negative virtual value win the item, and charges the winner a price that equals the threshold value above which she wins.[1]

Unfortunately, there is a further fundamental challenge in deploying these theoretical results in practice, which is that in real-world settings the auction designer may not even know the prior distributions on valuations. Instead, what the designer might hope for is that there is a stream of previous transactions, or some other relevant auxiliary data, that is helpful in inferring the buyers' private distributions. This perspective has motivated an active recent literature learning optimal auctions from samples [2, 7, 8, 9, 11, 12, 13, 14, 19, 20, 23, 24, 25]. In this line of work, the central question is: suppose we are only able to access the prior distributions in the form of independent samples, how many samples are sufficient and necessary for finding an approximately optimal auction?

While this merging of mechanism design and learning theory is appealing, a further concern arises. Given the potentially adversarial setting of auction design, do we really believe that the data that we observe are drawn in accord with our assumptions? More concretely, is the learning of optimal auctions robust to adversarial corruptions of the samples? This problem is arguably at the core of what it means to learn an optimal auction. It is a challenging problem; indeed, as we show in Counterexample 1 in Section 4, auction designs that are optimal in the absence of corruptions can become arbitrarily bad even if a small portion of the samples are corrupted. Building on earlier work by Cai and Daskalakis [6] and Brustle et al. [3], we tackle a key open problem—what is the best approximation to the optimal revenue for arbitrary levels of corruption for distributions with unbounded support? And what is the mechanism that achieves it?

In summary, in this work we explore the problem of the robust learning of optimal auctions, where the samples of bidders' valuations are subject to corruption and their support is unbounded. In particular, we consider having access to samples that are drawn from some distribution $\tilde{\mathcal{D}}$ which is within a Kolmogorov-Smirnov (KS) distance α of the true distribution $\mathcal{D}^*$. Denote OPT as the maximum revenue we can achieve under the true valuation distributions. Our goal is to design mechanisms that are guaranteed to achieve a revenue of at least $(1 - \rho(\alpha)) \cdot \text{OPT}$ for the smallest

[1] More generally, the optimal auction picks the winner based on the virtual value after an "ironing" procedure.

possible error $\rho(\alpha)$ and with the use of a minimal number of samples.

1.1 Our results

We study the problem of learning revenue-optimal multi-bidder auctions from samples when the samples of bidders' valuations can be adversarially corrupted or drawn from distributions that are adversarially perturbed. We summarize our main results as follows:

1. We derive tight upper bounds on the revenue we can obtain with a corrupted distribution under a population model. For distributions with monotone hazard rate (MHR), and with total corruption α, we obtain an approximation ratio of $1 - O(\alpha)$ compared to the optimal revenue under the true distribution (see Theorem 3.6). For regular valuation distributions, where for total corruption α, we get an approximation ratio of $1 - O(\sqrt{\alpha})$ (see Theorem 3.8).

2. To achieve these upper bounds, we propose a new *theoretical* algorithm for the population model (see Algorithm 1) that, given only an "approximate distribution" for the bidder's valuation, can learn a mechanism whose revenue is nearly optimal simultaneously for all "true distributions" that are α-close to the given distribution in Kolmogorov-Smirnov distance. The proposed algorithm operates beyond the setting of bounded distributions that have been studied in prior works; indeed, they apply to general unbounded MHR and regular distributions.

3. We further show that these upper bounds under the population model cannot be further improved (up to constant log factors), by providing matching lower bounds for both the MHR and regular distributions (see Theorem 3.7 and Theorem 3.9).

4. Lastly, we derive sample complexity upper bounds for learning a near-optimal auction for both MHR and regular distributions with multiple bidders (Theorem 4.3 and Theorem 4.4), and propose a *practical* algorithm (see Algorithm 2) which takes samples as input. We also provide accompanying sample complexity lower bounds (Theorem 4.5), and demonstrate a small gap relative to the corresponding upper bounds.

1.2 Related work

Designing revenue optimal auctions is a classic problem in economic theory that has attracted much research attention. We survey the most closely related work in two main areas.

Learning optimal auctions from samples. Recent work has explored settings of learning approximately optimal auction from samples, both for single-item auctions [7], and multi-item auctions [1, 2, 19, 25]. Most recently, Guo et al. [13] provide a complete set of sample complexity bounds for single-item auctions, by deriving matching upper and lower bounds up to a poly-logarithmic factor. While these approaches have obtained fruitful results on the sample complexity of learning optimal auctions, a key assumption that is commonly made in this work is that the samples are independently and identically drawn from the bidders' valuation distributions, with the goal of learning an auction which maximizes the expected revenue on the underlying, unknown distribution over bidder valuations. A major difference in our work is that we consider that the samples can suffer from potential corruptions, which is a significantly more challenging setting.

Robustness of learning optimal auctions. Our paradigm on the robust learning of optimal auctions is closely related to recent work that considers the learning of auctions from mismatched distributions or corrupted samples. Cai and Daskalakis [6] consider a multi-item auction setting, where there is a given "approximate distribution," and the goal is to compute an auction whose revenue is approximately optimal simultaneously for all "true distributions" that are close to the given one. They provide an algorithm that achieves a poly-α additive loss compared to the true optimal revenue. More recently, Brustle et al. [3] consider learning multi-item auctions where bidders' valuations are drawn from correlated distributions that can be captured by Markov random fields. However, they make a key simplifying assumption—that the bidders' valuation for the items lie in some bounded interval. Our results, by contrast, apply to the general setting of unbounded valuation distributions, a setting that requires new theoretical machinery. To the best of our knowledge, our work constitutes the first analysis of the learnability of single-item optimal auctions from corrupted samples for unbounded distributions.

Organization. In Section 2, we provide background on auction models and formally state our problem. Section 3 contains our main theoretical statements for the population model. We propose an algorithm that achieves optimal theoretical upper bounds, by providing matching lower bounds. Section 4 contains our main results on learning with finite samples. We provide a practical algorithm that takes samples from the corrupted distribution, and provides sample complexity upper and lower bounds for both the regular and MHR distributions cases. We conclude in Section 5.

2 Preliminaries

We begin by formally defining the setting we study for robust learning of optimal auctions, which includes the revenue objective and the general classes of valuation distributions that we consider.

2.1 Auction models

Single-bidder setting. Consider one item for sale to one bidder. The bidder has a private valuation $v \in \mathbb{R}_+$ for this item. We assume that v is a random variable distributed according to the distribution $\mathcal{D}^*$, with support $\mathbb{R}_+$, cumulative distribution function F, and probability density function f.

It is well known that the optimal auction in this setting is a reserve price auction, such that the task for the seller is to compute a reserve price p that optimizes revenue [21]. We assume that the bidder has a quasi-linear utility that is equal to $u(p) = v - p$ if she decides to buy the item and $u(p) = 0$ otherwise. The seller aims to set p such that her expected revenue—i.e., the received payment—is maximized. We consider the setting where both v and $\mathcal{D}^*$ are unknown to the seller. However, the seller can access i.i.d. samples that are drawn from a distribution $\tilde{\mathcal{D}}$, which is α-close to $\mathcal{D}$ with regard to the Kolmogorov distance:

Definition 2.1. (Kolmogorov-Smirnov distance) For probability measures μ and ν on $\mathbb{R}$, define

$$d_k(\mu, \nu) = \sup_{x \in \mathbb{R}} |\mu((-\infty, x)) - \nu((-\infty, x))|.$$

It is well known that $d_k(\mu,\nu) \leqslant d_{TV}(\mu,\nu)$, where d_{TV} denotes the total variation (TV) distance between μ and ν. The closeness of $\tilde{\mathcal{D}}$ to $\mathcal{D}^*$ is thus formalized as follows:

$$d_k(\mathcal{D}^*, \tilde{\mathcal{D}}) \leqslant \alpha,$$

for some $\alpha > 0$.

Multi-bidder setting. Consider one item for sale to n bidders. Each bidder has a private valuation, $v_i \in \mathbb{R}_+$, where v_i is independently drawn from the corresponding prior distribution $\mathcal{D}_i^*$. Thus, the valuations $\mathbf{v} = (v_1, v_2, \cdots, v_n)$ follow a product distribution $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$. Each bidder submits a bid $b_i \geqslant 0$. Denote all the bids as $\mathbf{b} = (b_1, \cdots, b_n)$. A mechanism in this setting consists of two rules: the allocation rule $\mathbf{x}(\mathbf{b})$ that takes the bids $\mathbf{b}$ and outputs the probability $x_i(\mathbf{b})$ that each bidder i will receive the item, and the payment rule $\mathbf{p}(\mathbf{b})$ that takes the bids $\mathbf{b}$ and outputs the payment of bidder i. Bidder i's utility is then $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$. The goal of the seller is to find a mechanism that maximizes the expected revenue $\mathbb{E}[\sum_{i\in[n]} p_i(\mathbf{b})]$, where the expectation is over $\mathbf{v} \sim \mathbf{D}^*$, under the following *Dominant Strategy Incentive Compatibility (DSIC)* and the *Individual Rationality (IR)* constraints:

$$u_i(v_i, \mathbf{b}_{-i}) \geqslant u_i(b_i, \mathbf{b}_{-i}) \qquad \text{for all } v_i, b_i \in \mathbb{R}_+ \text{ and all } \mathbf{b}_{-i} \in \mathbb{R}_+^{n-1} \tag{DSIC}$$

$$u_i(v_i, \mathbf{b}_{-i}) \geqslant 0 \qquad \text{for all } v_i \in \mathbb{R}_+ \text{ and all } \mathbf{b}_{-i} \in \mathbb{R}_+^{n-1}. \tag{IR}$$

We consider the setting in which the valuations and the prior distributions are unknown to the seller. Instead, the seller has access to a finite number of i.i.d. samples drawn from the product distribution $\tilde{\mathcal{D}} = \tilde{\mathcal{D}}_1 \times \cdots \times \tilde{\mathcal{D}}_n$, where each $\tilde{\mathcal{D}}_i$ satisfies

$$d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha_i,$$

for some $\alpha_i > 0, \forall i \in [n]$.

Revenue objective. Letting $\mathbf{D}$, $\mathbf{D}'$ be product or single bidder distributions as described above, we define $M_{\mathbf{D}}$ to be the mechanism that achieves the optimal revenue for the value distributions $\mathbf{D}$ and $\mathrm{OPT}(\mathbf{D})$ its expected revenue. Let also $\mathrm{Rev}(M_{\mathbf{D}}, \mathbf{D}')$ be the expected revenue of the mechanism $M_{\mathbf{D}}$ when applied to a setting where the values are drawn with respect to $\mathbf{D}'$.

2.2 Monotone hazard rate (MHR) and regular distributions

For any bidder i with a valuation $v_i \sim \mathcal{D}_i$, define the *virtual value function* for this bidder as $\phi_i(v) \stackrel{\text{def}}{=} v - \frac{1-F_i(v)}{f_i(v)}$, where F_i and f_i are the CDF and PDF of $\mathcal{D}_i$. The *hazard rate* of the distribution $\mathcal{D}_i$ is defined as the function $\frac{f_i(v)}{1-F_i(v)}$. Then, the distribution $\mathcal{D}_i$ is said to be *regular* if the virtual value $\phi_i(v)$ is monotonically non-decreasing in v. Further, distribution $\mathcal{D}_i$ has *monotone hazard rate* (MHR) if $\frac{f_i(v)}{1-F_i(v)}$ is monotone non-decreasing.

3 The Population Model

In this section, we study the problem of learning optimal auction assuming that we have the exact knowledge of the adversarially perturbed distributions $\tilde{\mathbf{D}}$. We relax this assumption in Section 4 where we show how to learn optimal auctions when we only have sample access to $\tilde{\mathbf{D}}$.

We begin in Section 3.1 with the description of our mechanism in the population model. Then, in Section 3.2, we present our analysis for the population mechanism for Monotone Hazard Rate distributions and we also present the sketch of our proof for the single-bidder case. Similarly, in Section 3.3 we state our analysis for the population mechanism for regular distributions and we present a proof sketch for the single-bidder case. Finally, we show that our proposed mechanism achieves optimal (up to constants) guarantees among any mechanism in the population model.

3.1 Robust Myerson auction in the population model

Our algorithm assumes as an input the exact knowledge of a product distribution, $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \times \tilde{\mathcal{D}}_n$, such that the $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha_i$ and its goal is to find a mechanism that achieves approximately optimal revenue for $\mathbf{D}^*$, where $\mathbf{D}^* = \Pi_i \mathcal{D}_i^*$. Without further assumptions, this is an impossible task, as we explain in Section 4 via an example. Thus we assume that the algorithm possesses some additional knowledge regarding $\mathcal{D}_i^*$, either that it is MHR or regular, and the mechanism needs to exploit this additional property.

To utilize the additional property of the distributions $\mathcal{D}_i^*$, our mechanism uses the important concept of the *link function* for MHR and regular distributions.

Definition 3.1 (Link Function)**.** The link function $h_M(x; F)$ for MHR distributions is defined as $h_M(x; F) = -\ln(1 - F(x))$ and the link function $h_r(x; F)$ for regular distributions is defined as $h_r(x; F) = 1/(1 - F(x))$. We also define the corresponding inverse link functions $h_M^{-1}(x; h) = 1 - \exp(-h(x))$ and $h_r^{-1}(x; h) = 1 - 1/h(x)$. Observe that $h_M^{-1}(x; h_M(x; F)) = F(x)$ and $h_r^{-1}(x; h_r(x; F)) = F(x)$. We may write $h_M(x)$ or $h_r(x)$ when F is clear from the context.

We provide some intuition on the link function. First, by construction, the link function of either an MHR distribution or a regular distribution is convex and non-decreasing. Second, the link function is monotone with regard to F. These two properties are important when we define the notion of a minimal MHR/regular distribution in a Kolmogorov ball, momentarily, which will be used as a necessary step in our algorithm.

Importantly, the link function provides a convenient characterization of the optimal reserve price and optimal revenue for a distribution F that is MHR or regular. To see this, first consider a single bidder with a valuation distribution F. Denote the optimal reserve price for selling one item to her as x^*, and the optimal expected revenue as $\text{OPT}(F)$. Then, when F is MHR, we show that x^* is also the unique minimizer of $(h_M(x) - \log(x))$. On the other hand, when F is regular, v^* is the point where $h_r(x)$ intersects with its tangent line kx, with $k = 1/\text{OPT}(F)$ (proof details in Appendix). Figure 1 illustrates such a useful property for h_M and h_r explicitly, for a single-item, single-bidder auction.

Next, we formally define stochastic dominance between two distributions, and state the property of strong revenue monotonicity.

Definition 3.2 (Stochastic dominance)**.** Given two distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ with CDFs as F_1 and F_2. Then, we say $\mathcal{D}_1$ (first-order) stochastically dominates $\mathcal{D}_2$ if for every $x \in \mathcal{X}$,

$$F_1(x) \leqslant F_2(x),$$

denoted as $\mathcal{D}_1 \succeq \mathcal{D}_2$. We say a product distribution $\mathbf{D} = \Pi_i \mathcal{D}_i$ (component-wise) stochastically dominates another product distribution $\mathbf{D}' = \Pi_i \mathcal{D}_i'$ if for every i, we have $\mathcal{D}_i \succeq \mathcal{D}_i'$.

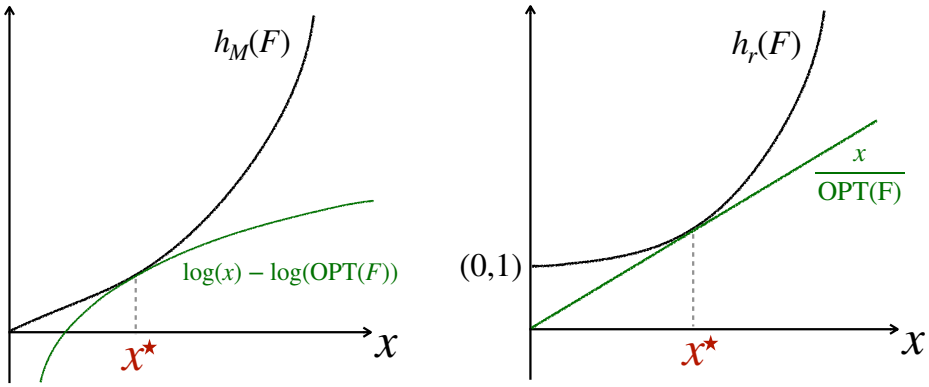

Figure 1: Optimal reserve price x^* with regard to the link function, for a single-item single-bidder auction with a valuation distribution F. *(left)* F is MHR; *(right)* F is regular.

Algorithm 1 Robust Myerson Auction in the Population Model

1: **Input:** $\alpha_1 \dots \alpha_n > 0$, link function $h(\cdot)$, possibly corrupted valuation distribution $\tilde{F} = \Pi_{i=1}^n \tilde{F}_i$.
2: **for** i = 1...n **do**
3: Compute a minimal regular / MHR distribution in $B_{d_k,\alpha_i}(\tilde{F}_i)$ according to Eq (1), denote as $\widehat{F}_i$.
4: **end for**
5: Set $\widehat{F} = \Pi_{i=1}^n \widehat{F}_i$.
6: Output Myerson's optimal auction $M_{\widehat{F}}$ w.r.t. the distribution $\widehat{F}$.

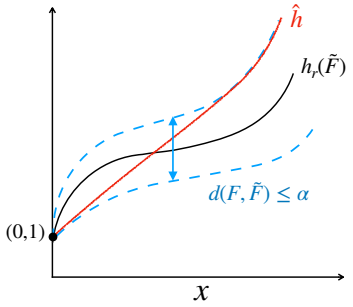

Figure 2: A minimal regular distribution in $B_{d_k,\alpha}$, in the space transformed by applying the link function.

Lemma 3.3 (Strong revenue monotonicity [13])**.** *Let* $\mathbf{D}$, $\mathbf{D}'$ *be two product distributions such that* $\mathbf{D}' \succeq \mathbf{D}$*, then, for* M *that is the optimal mechanism for* $\mathbf{D}$*, we have:*

$$\text{Rev}(M, \mathbf{D}) \leqslant \text{Rev}(M, \mathbf{D}').$$

The following lemma illustrates the importance of the link functions as well as their connection with first-order stochastic dominance. The proof of this lemma is given in Appendix A.

Lemma 3.4. *A distribution with CDF* F *is MHR if and only if* $h_M(x; F)$ *is a convex function of* x*. Similarly,* F *is regular if and only if* $h_r(x; F)$ *is a convex function of* x*. Moreover, for two MHR (resp. regular) distributions* F_1 *and* F_2*, such that* $F_1 \succeq F_2$*, we have that* $h_M(x; F_1) \leqslant h_M(x; F_2)$ *(resp.* $h_r(x; F_1) \leqslant h_r(x; F_2)$*) for all* x*.*

A key idea used in our algorithm is the minimal MHR/regular distribution within a Kolmogorov distance divergence ball. Formally,

Definition 3.5. For a given distribution with its cumulative distribution function as F, denote the set of all the distributions that are α-close to F in Kolmogorov distance as $B_{d_k,\alpha}(F)$:

$$B_{d_k,\alpha}(F) \stackrel{\text{def}}{=} \{F' : d_k(F', F) \leqslant \alpha\}.$$

Further, define a minimal MHR/regular distribution within $B_{d_k,\alpha}(F)$ as:

$$\widehat{F}(x) = h^{-1}(x; \widehat{h}), \quad \text{where} \quad \widehat{h}(x) \stackrel{\text{def}}{=} \max_{\substack{\tilde{F} \in B_{d_k,\alpha}(F) \\ \tilde{F} \text{ is MHR / regular}}} h\left(\tilde{F}(x)\right) \quad \forall x \in \mathbb{R}_+. \tag{1}$$

Figure 2 gives an illustration of a minimal regular distribution within $B_{d_k,\alpha}(F)$, in the space transformed by the link function of regular distributions.

3.2 Analysis for MHR distributions

In this section we state the results for the performance of Algorithm 1 for MHR distributions and we provide a proof sketch for the single-bidder case. The full proof of the following theorem can be found in Appendix B.

Theorem 3.6. *Let $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$ be a product distribution where every $\mathcal{D}_i^*$ is MHR. Let also $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \tilde{\mathcal{D}}_n$ be any product distribution such that for all $i \in [n]$ it holds that $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha_i$. If $\tilde{M}$ is the mechanism that Algorithm 1 outputs with input $\tilde{\mathbf{D}}$ then it holds that*

$$\mathrm{Rev}(\tilde{M}, \mathbf{D}^*) \geqslant \left(1 - \tilde{O}\left(\sum_{i=1}^{n} \alpha_i\right)\right) \cdot \mathrm{OPT}(\mathbf{D}^*).$$

In particular for $n = 1$, if $\alpha = \alpha_1$, then we have that $\mathrm{Rev}(\tilde{M}, \mathcal{D}^*) \geqslant (1 - O(\alpha)) \cdot \mathrm{OPT}(\mathcal{D}^*)$.

Proof sketch for $n = 1$. The first key step in our proof is the observation that, by construction, Algorithm 1 runs the Myerson optimal auction on an MHR distribution $\widehat{F}$, such that $\widehat{F}$ is stochastically dominated by any other MHR distribution that is within $B_{d_k,\alpha}(\tilde{F}')$. On the other hand we have $d_k(F^*(x), \tilde{F}(x)) \leqslant \alpha$. Applying the triangle inequality, we have $d_k(F^*(x), \widehat{F}(x)) \leqslant 2\alpha$. It is then sufficient for us to bound the ratio of the optimal revenue for any two MHR distributions F_1 and F_2, with $d_k(F_1, F_2) \leqslant 2\alpha$, and where F_1 is stochastically dominated by F_2.

The key part of our proof then considers such F_1, F_2, and due to the fact that the ratio of the revenues, $\mathrm{OPT}_{F_1}/\mathrm{OPT}_{F_2}$, is scale invariant, we assume without loss of generality that $\mathrm{OPT}_{F_1} = 1$. We then prove that this leads to $h(P_{F_1}^*) \leqslant 1$. The result then follows from two further key lemmas. First, for any reserve price $x < P_{F_1}^*$, $|h_1(x) - h_2(x)| = \left|\log\left(\frac{1-F_2(x)}{1-F_1(x)}\right)\right|$. Further applying the fact that by assumption $|F_1(x) - F_2(x)| \leqslant \alpha$ we show that $|h_1(x) - h_2(x)| = O(\alpha)$ for any reserve price $x < P_{F_1}^*$. Second, using the fact that F_1 is stochastically dominated by F_2, we derive that $P_{F_2}^* \leqslant P_{F_1}^*$. The conclusion then follows from bounding the ratio of $s_1(x) = h_1(x) - \log(x)$, and $s_2(x) = h_2(x) - \log(x)$, based on the definition of $P_{F_1}^*$ and $P_{F_2}^*$. ■

Next we show that the information-theoretic Algorithm 1 is optimal up to constants for MHR distributions. We provide the proof of the following theorem in Appendix C.

Theorem 3.7. *Let M be any DSIC and IR mechanism that takes as input a product distribution $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \times \tilde{\mathcal{D}}_n$. Then there exists a product distribution $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$ such that $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha$, $\mathcal{D}_i^*$ is MHR for every i, and*

$$\mathrm{Rev}(M, \mathbf{D}^*) \leqslant (1 - \tilde{\Omega}(n \cdot \alpha)) \cdot \mathrm{OPT}(\mathbf{D}^*).$$

3.3 Analysis for regular distributions

In this section we state the results for the performance of Algorithm 1 for regular distributions and we provide a proof sketch for the single-bidder case. The full proof of the following theorem can be found in Appendix B.

Theorem 3.8. *Let $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$ be a product distribution where every $\mathcal{D}_i^*$ is regular. Let also $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \tilde{\mathcal{D}}_n$ be any product distribution such that for all $i \in [n]$ it holds that $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha_i$. If $\tilde{M}$ is the mechanism that Algorithm 1 outputs with input $\tilde{\mathbf{D}}$ then it holds that*

$$\mathrm{Rev}(\tilde{M}, \mathbf{D}^*) \geqslant \left(1 - 5 \cdot \sqrt{\sum_{i=1}^{n} \alpha_i}\right) \cdot \mathrm{OPT}(\mathbf{D}^*).$$

Proof sketch for $n = 1$. We first prove a general result that for two regular distributions F and $\bar{F}$, such that $d_k(F, \bar{F}) \leqslant \alpha$, where $F(x)$ is stochastically dominated by $\bar{F}(x)$ for $x \in \mathbb{R}_+$. The optimal revenue of these two distributions is close, formally $\frac{\mathrm{OPT}(F)}{\mathrm{OPT}(\bar{F})} \geqslant 1 - O(\sqrt{\alpha})$. The first key step replies on using the link function $h_r(x) = \frac{1}{1-F(x)}$ for regular distributions. Since $h_r(x)$ preserves the same monotonicity property as $F(x)$, we first derive a lower bound on $\bar{h}_r(x, \bar{F})$ that is $\bar{h}_r(x, \bar{F}) \geqslant h_r(x, F) - \alpha h_r^2(x, F)$, using the fact that $d_k(F, \bar{F}) \leqslant \alpha$. This bound gives us useful constraints to discuss in different cases in the following part of the proof. Denote the corresponding optimal reserve prices for F and $\bar{F}$ as P and $\bar{P}$. We discuss separately two cases for $h(\bar{P})$, where, for case 1 we have $h(\bar{P}) \leqslant \frac{1}{\sqrt{\alpha}}$, and for case 2, we have $h(\bar{P}) > \frac{1}{\sqrt{\alpha}}$. Using the connection from the link function to the revenue (see Figure 1), case 1 directly leads to the conclusion that $\frac{\mathrm{OPT}(F)}{\mathrm{OPT}(\bar{F})} \geqslant 1 - \sqrt{\alpha}$. Case 2 is more subtle and requires a more careful argument. Lastly, by construction, Algorithm 1 runs the Myerson optimal auction on a regular distribution $\widehat{F}$, such that $\widehat{F} \geqslant \widehat{F}'(x)$ for all $x \in \mathbb{R}_+$, for any other regular distribution $F'(x)$ such that $d_k(F'(x), \tilde{F}(x)) \leqslant \alpha$. Applying the triangle inequality and combining with the conclusions obtained from the two cases concludes the proof. ■

Finally, we show that the information-theoretic Algorithm 1 is optimal up to constants for regular distributions. We provide the proof of the following theorem in Appendix C.

Theorem 3.9. *Let M be any DSIC and IR mechanism that takes as input a product distribution $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \times \tilde{\mathcal{D}}_n$. Then there exists a product distribution $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$ such that $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha$, $\mathcal{D}_i^*$ is regular for every i, and*

$$\mathrm{Rev}(M, \mathbf{D}^*) \leqslant (1 - \Omega(\sqrt{n \cdot \alpha})) \cdot \mathrm{OPT}(\mathbf{D}^*).$$

4 Finite Samples

We provide a practical algorithm that takes samples from the corrupted distribution $\tilde{\mathbf{D}}$ as an input. We show that this algorithm achieves almost optimal sample complexity for the MHR distribution case and the single-bidder regular distribution case, whereas for the multi-bidder regular distributions there is a small gap between our upper and lower bounds.

An important notion to explain our algorithm for the finite-sample case is the following notion of the convex envelope.

Definition 4.1 (Convex Envelope)**.** The convex envelope $Conv(f)$ of a function f is a function with the following property

$$Conv(f)(x) = \sup\{g(x) \mid g \text{ is convex and } g \leqslant f \text{ over } \mathbb{R}_+\}.$$

In words, $Conv(f)$ is the maximum convex function that is below f.

For our algorithm one important property of the convex envelope is expressed in the following lemma whose proof is presented in Appendix A.

Lemma 4.2. *Let f be a non-decreasing piecewise constant function with k pieces, then $Conv(f)$ can be computed in time* $\text{poly}(k)$ *and is a piecewise linear function with $O(k)$ pieces.*

Algorithm 2 Robust Empirical Myerson Auction

1: **Input:** m i.i.d. samples from (possibly corrupted) value distribution $\mathbf{D} = \Pi_{i=1}^n \mathcal{D}_i$, link function $h(\cdot)$.
2: Let $\mathbf{E} = \Pi_{i=1}^n E_i$ be the empirical distribution, i.e., the uniform distribution over the samples.
3: **for** $i = 1 \dots n$ **do**
4: Construct $\widehat{E_i}$ as following: let $q^{E_i}(v)$ be the quantile of E_i; the quantile of $\widehat{E_i}$ is as follows:

$$q^{\widehat{E}_i}(v) = \begin{cases} \max\left\{0, q^{E_i}(v) - \sqrt{\frac{2q^{E_i}(v)\left(1-q^{E_i}(v)\right)\ln(2mn\delta^{-1})}{m}} - \frac{4\ln(2mn\delta^{-1})}{m} - \alpha_i\right\} & \text{if } v > 0 \\ 1 & \text{if } v = 0 \end{cases}$$

5: Construct $\tilde{E}_i$ such that $h\left(\tilde{E}_i(\cdot)\right)$ is the convex envelope of $h\left(\widehat{E}(\cdot)\right)$, i.e.

$$\tilde{E}_i(\cdot) = h^{-1}\left(Conv\left(h(\widehat{E}_i(\cdot))\right)\right)$$

6: **end for**
7: Set $\tilde{\mathbf{E}} = \Pi_{i=1}^n \tilde{E}_i$
8: Output Myerson's optimal auction $M_{\tilde{\mathbf{E}}}$ w.r.t. $\tilde{\mathbf{E}}$.

The above algorithm resembles the main algorithm of [13] with the addition of step 5. We first show that step 5 is necessary if we wish to obtain any non-trivial result in the robust auction learning setting that we explore in this paper.

Counterexample 1. Imagine we have just one agent, i.e., $n = 1$, with true distribution $\mathcal{D}^*$ equal to an exponential distribution with parameter $\lambda = 1$. Also, to strengthen our counterexample imagine that we have available an infinite number of samples, i.e., $m \to \infty$. Now consider $\tilde{\mathcal{D}}$ to be the corrupted distribution where probability mass α is removed from the mass closer to 0 and it is placed as a point mass at the point c/α for some number c. In this case, running Algorithm 2 without step 5 will result is implementing an auction with reserve price that is very close to c/α. The probability though that the true agent with distribution $\mathcal{D}^*$ will buy this item goes to zero with a rate $\exp(-c/\alpha)$ as $c \to \infty$. Hence, the total revenue will be at most $(c/\alpha) \cdot \exp(-c/\alpha)$ and therefore we can make the total revenue to go to zero as we increase $c \to \infty$. Observe that this counterexample works even though we assumed that the initial distribution $\mathcal{D}^*$ is MHR.

We next provide the analysis of the performance of Algorithm 2 for MHR and regular distributions. The proof of the following result can be found in Appendix D.

Theorem 4.3 (Finite samples, Regular distribution)**.** *Let $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$ be a product distribution where every $\mathcal{D}_i^*$ is regular. Let also $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \tilde{\mathcal{D}}_n$ be any product distribution such that for all $i \in [n]$ it holds that $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha_i$. If $\tilde{M}$ is the mechanism that Algorithm 2 outputs with input m samples from $\tilde{\mathbf{D}}$ and assume that $m = \tilde{\Omega}\left(\max_{i\in[n]}\left\{\log(\frac{1}{\delta})/\alpha_i^2\right\}\right)$ then it holds that*

$$\mathbb{P}\mathrm{r}\left(\mathrm{Rev}(\tilde{M}, \mathbf{D}^*) \geqslant \left(1 - O\left(\sqrt{\sum_{i=1}^n \alpha_i}\right)\right) \cdot \mathrm{OPT}(\mathbf{D}^*)\right) \geqslant 1 - \delta.$$

Additionally, in the single-bidder case with $n = 1$ and $\alpha = \alpha_1$ the sample requirement becomes $m = \tilde{\Omega}\left(\log(\frac{1}{\delta})/\alpha^{3/2}\right)$.

The corresponding theorem for MHR distributions is the following, whose proof can be found in Appendix D.

Theorem 4.4 (Finite samples, MHR distribution)**.** *Let $\mathbf{D}^* = \mathcal{D}_1^* \times \cdots \times \mathcal{D}_n^*$ be a product distribution where every $\mathcal{D}_i^*$ is MHR. Let also $\tilde{\mathbf{D}} = \tilde{\mathcal{D}}_1 \times \cdots \tilde{\mathcal{D}}_n$ be any product distribution such that for all $i \in [n]$ it holds that $d_k(\mathcal{D}_i^*, \tilde{\mathcal{D}}_i) \leqslant \alpha_i$. If $\tilde{M}$ is the mechanism that Algorithm 2 outputs with input m samples from $\tilde{\mathbf{D}}$ and assume that $m = \tilde{\Omega}\left(\max_{i\in[n]}\left\{\log\left(\frac{1}{\delta}\right)/\alpha_i^2\right\}\right)$ then it holds that*

$$\mathbb{P}\mathrm{r}\left(\mathrm{Rev}(\tilde{M}, \mathbf{D}^*) \geqslant \left(1 - \tilde{O}\left(\sum_{i=1}^n \alpha_i\right)\right) \cdot \mathrm{OPT}(\mathbf{D}^*)\right) \geqslant 1 - \delta.$$

We make a few remarks about the sample complexity upper bounds in the sequel.

First, in both Theorem 4.3 and Theorem 4.4, the sample complexity upper bounds depend in a simple way on the sum of all the fractions of corruptions for each bidder; i.e., $\sum_{i=1}^n \alpha_i$, indicating the important effect of the *total* amount of corruption. Second, for regular distributions, in Theorem 4.3 we obtain a tight sample complexity bound for the single-bidder case, with $m = \tilde{\Omega}\left(\log(\frac{1}{\delta})/\alpha^{3/2}\right)$. For multi-bidder settings, our upper bound contains a small gap, with $m = \tilde{\Omega}\left(\max_{i\in[n]}\left\{\log(\frac{1}{\delta})/\alpha_i^2\right\}\right)$. Whether such a gap can be matched is an interesting open question for future work. Lastly, comparing Theorem 4.3 and Theorem 4.4, it appears that for the multi-bidder settings the sample complexity bounds are of the same order, but we emphasize the key difference that for regular distributions this sample size is needed to provide a much *weaker* guarantee on the revenue objective, which is a $\left(1 - O\left(\sqrt{\sum_{i=1}^n \alpha_i}\right)\right)$ fraction of the optimal revenue, while the guarantee for MHR distributions is a $\left(1 - O\left(\sum_{i=1}^n \alpha_i\right)\right)$ fraction of the optimal revenue.

We next provide an information-theoretic lower bound that establishes the tightness of our upper bounds for the single-bidder single-item case with regular and MHR distributions.

Theorem 4.5 (Sample complexity lower bounds)**.** *Let M be any DSIC and IR mechanism for a single-item single-buyer setting that takes as input m samples from a distribution $\tilde{\mathcal{D}}$. If*

$$\mathrm{Rev}(M, \mathcal{D}^*) \geqslant (1 - O(\sqrt{\alpha})) \cdot \mathrm{OPT}(\mathcal{D}^*),$$

for all distributions $\mathcal{D}^*$ such that $d_k(\mathcal{D}^*, \tilde{\mathcal{D}}) \leqslant \alpha$, where $\mathcal{D}^*$ is regular, then $m \geqslant \tilde{\Omega}\left(\log(\frac{2}{\delta})/\alpha^{3/2}\right)$. Additionally, if

$$\mathrm{Rev}(M, \mathcal{D}^*) \geqslant (1 - O(\alpha)) \cdot \mathrm{OPT}(\mathcal{D}^*),$$

for all distributions $\mathcal{D}^*$ such that $d_k(\mathcal{D}^*, \tilde{\mathcal{D}}) \leqslant \alpha$, where $\mathcal{D}^*$ is MHR, we have $m \geqslant \tilde{\Omega}\left(\log(\frac{2}{\delta})/\alpha^{3/2}\right)$.

Theorem 4.5 provides a general sample complexity lower bound on learning a near-optimal auction with at least a $(1 - O(\sqrt{n \cdot \alpha}))$ fraction of the optimal revenue under the true valuation distribution. In comparison to our upper bounds (see Theorem 4.3 and Theorem 4.4), there is a small gap and we leave the nature of this gap as an open question for future work.

5 Conclusions

We have studied the learning of revenue-optimal auctions for multiple bidders, in a setting in which the samples can be corrupted adversarially. We first consider the information-theoretic limit in a population model, assuming exact knowledge of the adversarially perturbed valuation distribution. We develop a theoretical algorithm which obtains a tight upper bound on the revenue for the MHR and regular distributions, obtaining the information-theoretic limit of the robustness guarantee. We then relax the population model and derive sample complexity bounds for learning optimal auctions from samples. We propose a practical algorithm which takes the corrupted samples as input, and provide the sample complexity upper bounds for the MHR distribution case and the single-bidder regular distribution case. We also provide accompanying sample complexity lower bounds, and demonstrate a small gap relative to the corresponding upper bounds.

Acknowledgments

This work was supported in part by the Mathematical Data Science program of the Office of Naval Research under grant number N00014-18-1-2764.

References

[1] M.-F. Balcan, T. Sandholm, and E. Vitercik. Sample complexity of automated mechanism design. *arXiv preprint arXiv:1606.04145*, 2016.

[2] M.-F. Balcan, T. Sandholm, and E. Vitercik. A general theory of sample complexity for multi-item profit maximization. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 173–174, 2018.

[3] J. Brustle, Y. Cai, and C. Daskalakis. Multi-item mechanisms without item-independence: Learnability via robustness. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 715–761, 2020.

[4] M. M. Bykowsky, R. J. Cull, and J. O. Ledyard. Mutually destructive bidding: The FCC auction design problem. *Journal of Regulatory Economics*, 17(3):205–228, 2000.

[5] Y. Cai and C. Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 522–531. IEEE, 2011.

[6] Y. Cai and C. Daskalakis. Learning multi-item auctions with (or without) samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 516–527. IEEE, 2017.

[7] R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, pages 243–252, 2014.

[8] N. R. Devanur, Z. Huang, and C.-A. Psomas. The sample complexity of auctions with side information. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pages 426–439, 2016.

[9] M. Dudík, N. Haghtalab, H. Luo, R. E. Schapire, V. Syrgkanis, and J. W. Vaughan. Oracle-efficient online learning and auction design. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, pages 528–539. IEEE, 2017.

[10] A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956.

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

[12] Y. A. Gonczarowski and S. M. Weinberg. The sample complexity of up-to-ε multi-dimensional revenue maximization. *Journal of the ACM*, 68(3):1–28, 2021.

[13] C. Guo, Z. Huang, and X. Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673, 2019.

[14] Z. Huang, Y. Mansour, and T. Roughgarden. Making the most of your samples. *SIAM Journal on Computing*, 47(3):651–674, 2018.

[15] P. Klemperer. What really matters in auction design. *Journal of Economic Perspectives*, 16 (1):169–189, 2002.

[16] S. Lahaie, D. M. Pennock, A. Saberi, and R. V. Vohra. Sponsored search auctions. *Algorithmic Game Theory*, 1:699–716, 2007.

[17] P. Massart. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, pages 1269–1283, 1990.

[18] P. Milgrom and P. R. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[19] J. Morgenstern and T. Roughgarden. The pseudo-dimension of near-optimal auctions. *arXiv preprint arXiv:1506.03684*, 2015.

[20] J. Morgenstern and T. Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pages 1298–1318. PMLR, 2016.

[21] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[22] A. E. Roth and A. Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *American Economic Review*, 92(4):1093–1103, 2002.

[23] T. Roughgarden and O. Schrijvers. Ironing in the dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 1–18, 2016.

[24] T. Roughgarden and J. R. Wang. Minimizing regret with multiple reserves. *ACM Transactions on Economics and Computation*, 7(3):1–18, 2019.

[25] V. Syrgkanis. A sample complexity measure with applications to learning optimal auctions. *arXiv preprint arXiv:1704.02598*, 2017.

Appendix

A Proofs of Technical Lemmas

Lemma 3.4. A distribution with CDF F is MHR if and only if $h_M(x;F)$ is a convex function of x. Similarly, F is regular if and only if $h_r(x;F)$ is a convex function of x. Moreover, for two MHR (resp. regular) distributions F_1 and F_2, such that $F_1 \succeq F_2$, then we have that $h_M(x;F_1) \leqslant h_M(x;F_2)$ (resp. $h_r(x;F_1) \leqslant h_r(x;F_2)$) for all x.

Proof. We first show that given the CDF of any MHR distribution $F(x) : \mathbb{R}_+ \to [0,1]$, $h_M(x) \stackrel{\text{def}}{=} -\log(1-F(x))$ is a convex, non-decreasing function with $h(0) = 0$. (Without loss of generality, we consider $x \in [0,\infty]$, i.e. $\arg\min_x h(x) = 0$.) We first present the analysis for the case when the distribution is continuous and smooth, and then generalize the same statement to discrete distributions.

MHR continuous distributions:

Denote the corresponding PDF of $F(x)$ as $f(x)$, and $g(x) \stackrel{\text{def}}{=} \frac{f(v)}{1-F(v)}$. By definition, $F(0) = 0$ implies $h_M(0) = 0$. Then, given that $F(x)$ is MHR, we have that $g(x)$ is monotone non-decreasing. By construction,

$$(h_M(x))'' = \left(\frac{f(v)}{1-F(v)}\right)' = g'(x) \geqslant 0.$$

Therefore, $h_M(x)$ is convex. Moreover, since $F(x)$ is a CDF thus non-decreasing, $h_M(x) = -\log(1-F(x))$ is also non-decreasing. We show that given any $h_M(x) : \mathbb{R}_+ \to \mathbb{R}_+$, such that $h_M(x)$ is convex, non-decreasing, $h_M(0) = 0$, and $\max_x h_M(x) = \infty$. Then, $F(x) \stackrel{\text{def}}{=} 1 - \exp(-h_M(x))$ is CDF of an MHR distribution.

By construction, $h_M(0) = 0$ implies $F(0) = 0$, and $\max_x h_M(x)$ implies $\max_x F(x) = 1$. Also given that $h_M(x)$ is convex, $g'(x) = \left(\frac{f(v)}{1-F(v)}\right)' = (h_M(x))'' \geqslant 0$, which by definition implies $F(x)$ is MHR.

MHR discrete distributions:

The lemma statement generalizes to the case when the valuation is discrete. We assume that the valuation can take a discrete set of values $\{x_i\}, i = 1, \cdots, n$. Without loss of generality, we will restrict these values to the set $\mathbb{N}_0$ with probability mass function $P(x = i) = p_i; i = 0 \cdots n$. We define the *discrete* hazard rate as:

$$g(x_i) = \frac{P(x = i)}{P(x \geqslant i)}.$$

Then, the valuation distribution is MHR iff the discrete hazard rate is non-decreasing:

$$g(x_{i+1}) \geqslant g(x_i), \tag{2}$$

for all $i = 0 \cdots n$.

In this case, our link function will also be discrete. Further, denote $s_i \stackrel{\text{def}}{=} P(x \geqslant i)$, then

$$h(x_i) = -\log(P(x \geqslant x_i)) = -\log(s_i).$$

Then $h(x)$ is convex if and only if for any $i \geqslant 0$,

$$h(x_{i+2}) - h(x_{i+1}) \geqslant h(x_{i+1} - h(x_i). \tag{3}$$

We show that Eq (2) and Eq (3) are equivalent. Notice that

$$\begin{aligned}
& h(x_{i+2}) - h(x_{i+1}) \geqslant h(x_{i+1} - h(x_i) \\
\iff & \frac{s_{i+1}}{s_{i+1} - p_{i+1}} \geqslant \frac{s_i}{s_i - p_i} \\
\iff & p_{i+1} s_i \geqslant p_i s_{i+1} \\
\iff & \frac{p_{i+1}}{s_{i+1}} \geqslant \frac{p_i}{s_i} \\
\iff & g(x_{i+1}) \geqslant g(x_i),
\end{aligned}$$

which completes the proof.

Regular continuous distributions:

We further prove a similar statement for regular continuous distributions. First, given a CDF of a regular distribution $F(x)$,

$$\left(\frac{1}{1 - F(x)}\right)'' = \frac{(1 - F(x))f(x)' + 2f(x)^2}{(1 - F(x))^3}.$$

By definition, the virtual value function is $\phi(x) \stackrel{\text{def}}{=} v - \frac{1-F(x)}{f(x)}$, and

$$\phi'(x) = \frac{(1 - F(x))f(x)' + 2f(x)^2}{f(x)^2}.$$

Therefore, $\left(\frac{1}{1-F(x)}\right)''$ and $\phi'(x)$ share the same sign. Moreover, the distribution with CDF as $F(x)$ is regular if and only if the virtual value $\phi(x)$ is monotonically non-decreasing, which is $\phi'(x) \geqslant 0$. Hence the regularity of $F(x)$ implies that $h_r(x) \stackrel{\text{def}}{=} \frac{1}{1-F(x)}$ is convex. Since $F(x)$ is a CDF thus non-decreasing, $h_r(x) = \frac{1}{1-F(x)}$ is also non-decreasing.

Regular discrete distributions:

Similar to the MHR distributions, the lemma statement generalizes to the case when the valuation is discrete for regular distributions. Assume that the valuation can take a discrete set of values $\{x_i\}, i = 1, \cdots, n$. Without loss of generality, we will restrict these values to the set $\mathbb{N}_0$ with probability mass function $P(x = i) = p_i; i = 0 \cdots n$. Further, consistent with the proof for MHR distributions, we denote $s_i \stackrel{\text{def}}{=} P(x \geqslant i)$.

The *discrete* virtual value function is defined as:

$$\phi(x_i) = x_i - \frac{s_i}{p_i},$$

and the valuation distribution is regular iff $\phi(x)$ is non-decreasing:

$$\phi(x_{i+1}) \geqslant \phi(x_i), \tag{4}$$

for all $i = 0 \cdots n$.

In this case, our link function will again be discrete:

$$h(x_i) = \frac{1}{P(x \geqslant x_i)} = \frac{1}{s_i}.$$

and $h(x)$ is convex if and only if for any $i \geqslant 0$,

$$h(x_{i+2}) - h(x_{i+1}) \geqslant h(x_{i+1}) - h(x_i). \tag{5}$$

We show that Eq (4) and Eq (5) are equivalent.

$$\begin{aligned}
&h(x_{i+2}) - h(x_{i+1}) \geqslant h(x_{i+1}) - h(x_i) \\
\iff &\frac{1}{s_{i+2}} + \frac{1}{s_i} \geqslant \frac{2}{s_{i+1}} \\
\iff &\frac{1}{s_{i+1} - p_{i+1}} + \frac{1}{s_i} \geqslant \frac{2}{s_{i+1}} \\
\iff &s_{i+1}^2 + p_i p_{i+1} \geqslant s_i s_{i+1} - s_i p_{i+1}. \\
\iff &p_i p_{i+1} + p_{i+1} s_i + s_{i+1}(s_{i+1} - s_i) \geqslant 0 \\
\iff &p_i p_{i+1} + p_{i+1} s_i - s_{i+1} p_i \geqslant 0
\end{aligned} \tag{6}$$

Moreover, from the regularity condition Eq (4), we have

$$\begin{aligned}
&\phi(x_{i+1}) \geqslant \phi(x_i) \\
\iff &i + 1 - \frac{s_{i+1}}{p_{i+1}} \geqslant i - \frac{s_i}{p_i} \\
\iff &1 - \frac{s_{i+1}}{p_{i+1}} + \frac{s_i}{p_i} \geqslant 0 \\
\iff &p_i p_{i+1} + p_{i+1} s_i - s_{i+1} p_i \geqslant 0.
\end{aligned} \tag{7}$$

Combining (6) and (7) together completes the proof.

Stochastic dominance:

Lastly, we show that for two MHR (resp. regular) distributions F_1 and F_2, such that $F_1 \succeq F_2$, then we have that $h_M(x; F_1) \leqslant h_M(x; F_2)$ (resp. $h_r(x; F_1) \leqslant h_r(x; F_2)$) for all x. This follows directly from the monotonicity of the link functions and the definition of stochastic dominance (see Definition 3.2).

Recall that the link function $h_M(x; F)$ for MHR distributions is defined as $h_M(x; F) = -\ln(1 - F(x))$, and the link function $h_r(x; F)$ for regular distributions is defined as $h_r(x; F) = 1/(1 - F(x))$. Therefore, for two MHR (resp. regular) distributions F_1 and F_2, $F_1(x) < F_2(x)$ implies $h_M(x, F_1) < h_M(x, F_2)$ (resp. $h_r(x, F_1) < h_r(x, F_2)$), which completes the proof.

■

Lemma 4.2. Let f be a non-decreasing piecewise constant function with k pieces, then $Conv(f)$ can be computed in time poly(k) and is a piecewise linear function with $O(k)$ pieces.

Proof. Given that $f(x)$ is a non-decreasing piecewise constant function with k pieces, we show that the following iterative procedure outputs its lower convex envelope $Conv(f)$, which can be computed in time poly(k) and is a piecewise linear function with $O(k)$ pieces. Figure 3 provides an illustration of the construction according to this procedure.

Procedure 1 Computing lower convex envelope for non-decreasing piecewise constant functions

1: **Input: a piecewise constant function** $f(x) : \mathbb{R} \to \mathbb{R}$ **with** k **pieces.** Denote the left starting point of each piece and the end point as $x_0, \ldots, x_k$.
2: **Initialize:** $i \leftarrow 0, i' \leftarrow 0$.
3: **while** $i \leqslant k-1$ **do**
4: $\quad \bar{x}_{i'} \leftarrow x_i, g(\bar{x}_{i'}) \leftarrow f(x_i)$.
5: $\quad i' \leftarrow i' + 1$.
6: $\quad$ Compute $i \leftarrow \arg\min_{i<j\leqslant k} \frac{f(x_j)-f(x_i)}{x_j-x_i}$.
7: **end while**
8: $\bar{x}_{i'} \leftarrow x_i, g(\bar{x}_{i'}) \leftarrow f(x_i)$; $k' \leftarrow i'$.
9: **Return: a piecewise linear function** $g(x) : \mathbb{R} \to \mathbb{R}$ **with** $k' < k$ **pieces.** The left starting points of each piece and the end points are $\bar{x}_0, \ldots, \bar{x}_{i'}$, with the corresponding function values as specified in the procedure.

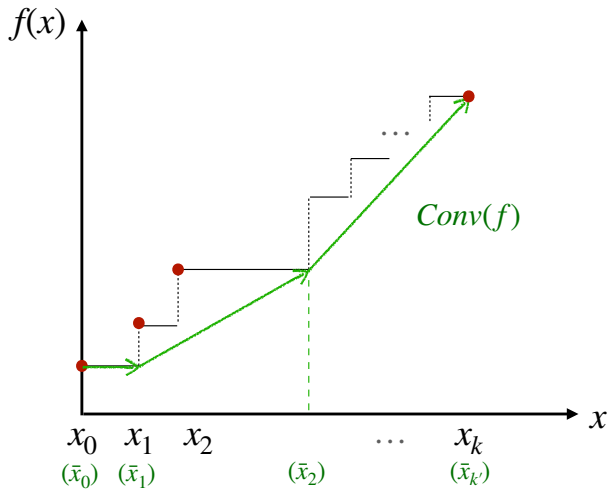

Figure 3: Lower convex envelope of a non-decreasing piecewise constant function $f(x)$.

First, the above procedure requires at most k^2 rounds. We show that its output, $g(x)$, is the lower convex envelope for $f(x)$. It is clear from construction that $g(x)$ is piecewise linear, with vertices at $\bar{x}_0, \ldots, \bar{x}_{k'}$. Moreover, $g(x) \leqslant f(x)$ for all x by construction.

Next we show that $g(x)$ is convex. Consider at a round t with $i = i_t, 1 < 1 < k$. Then, step (6) computes $i_{t+1} = \arg\min_{i_t<j\leqslant k} \frac{f(x_j)-f(x_{i_t})}{x_j-x_{i_t}}$. Further denote $\min_{i_t<j\leqslant k} \frac{f(x_j)-f(x_{i_t})}{x_j-x_{i_t}}$ as $s(i_t)$. We show that $s(i_{t+1}) \geqslant s_{i_t}$.

Suppose that $s(i_{t+1}) < s_{i_t}$. Then there exists $j^* > i_{t+1} > i_t$, such that

$$\frac{f(x_{j^*}) - f(x_{i_{t+1}})}{x_{j^*} - x_{i_{t+1}}} < \frac{f(x_{i_{t+1}}) - f(x_{i_t})}{x_{i_{t+1}} - x_{i_t}},$$

which further implies that

$$\frac{f(x_{j^*}) - f(x_{i_t})}{x_{j^*} - x_{i_t}} < \frac{f(x_{i_{t+1}}) - f(x_{i_t})}{x_{i_{t+1}} - x_{i_t}}.$$

Since $j^* > i_{t+1} > i_t$, this contradicts the fact that $i_{t+1} = \arg\min_{i_t<j\leqslant k} \frac{f(x_j)-f(x_{i_t})}{x_j-x_{i_t}}$. Therefore $s(i_{t+1}) \geqslant s_{i_t}$, which means that the slope of each piece for $g(x)$ is non-decreasing. Thus $g(x)$ is convex. Lastly, since $g(x)$ has all vertices with the same function values as $f(x)$, i.e. $g(x) = f(x)$ at all its vertices, and given that $g(x) \leqslant f(x)$ for all x, the values at these vertices are maximized and cannot be further improved. This completes the proof. ■

We further provide two lemmas which present useful properties of the link functions in connection to the revenue.

Lemma A.1. *Given an MHR distribution with the CDF as $F(x) : \mathbb{R}_+ \to [0, 1]$. Define $h(x) \stackrel{\text{def}}{=} -\log(1 - F(x))$. Then, at any reserve price x, the expected revenue $R(x) = \exp(-h(x) + \log(x))$. Moreover, the optimal reserve price P_F^* is the minimizer of $(h(x) - \log(x))$.*

Proof. First by construction, $h(x) - \log(x) = -\log(R(x))$. By definition, the optimal reserve price maximizes the revenue $R(x) = x(1 - F(x))$, thus

$$\begin{aligned} &\max \quad x(1 - F(x)) \\ \iff &\min \quad -\log(x(1 - F(x))) \\ \iff &\min \quad -\log(x) - \log(1 - F(x)) \\ \iff &\min \quad h(x) - \log(x), \end{aligned}$$

which completes the proof. ■

Lemma A.2. *Consider a valuation distribution $\mathcal{D}$ with CDF as $F(x)$. Denote the optimal reserve price as P_F^* and the optimal expected revenue at P_F^* as OPT_F. Then P_F^* should be $P_F^* \leqslant e$, assuming that $\mathrm{OPT}_F \leqslant 1$ and $F(x)$ is MHR.*

Proof. By Lemma A.1, $\mathrm{OPT}_F \leqslant 1$ implies that,

$$h(P_F^*) = \log(P_F^*) + b,$$

for some $b \geqslant 0$. Also by Lemma 3.4, h is convex. Combined with the fact that OPT_F is the optimal reserve price and the concavity of $\log(x)$, OPT_F is the only point where $h(P_F^*) = \log(P_F^*) + b$ holds.

Now consider a linear function $y = ax, a > 0$, which is a tangent line of the function $\log(x) + b$. Denote the tangent point as x^*. Solving the equation that $a = (\log(x))' = \frac{1}{x}$, and $ax = \log(x) + b$ give that:

$$x^* = e^{1-b} \leqslant e.$$

Suppose that $P_F^* > x^*$. Consider the linear function $g(x) = \frac{h(P_F^*)}{P_F^*}x$. Since x^* is the tangent point, there exists a point $\bar{x} < P_F^*$, such that $g(\bar{x}) = \log(\bar{x}) + b$. Further, since h is convex, for any point $0 < x < P_F^*$, we have $h(x) < g(x)$. By the continuity of $\log(x)$ and $h(x)$, there exists $\bar{x}' < P_F^*$, such that $h(\bar{x}') = \log(\bar{x}) + b$. This implies that $\bar{x}'$ achieves a larger revenue than P_F^*, and contradicts the fact that P_F^* is the optimal reserve price. Hence, $P_F^* < x^* \leqslant e$, which completes the proof. ■

B Proof of Upper Bounds for the Population Model

We first prove the following technical lemma that connects the coordinate Kolmogorov distance with the difference in expectation of increasing functions.

Definition B.1 (Increasing Functions and Sets)**.** Let $u : \mathbb{R}^n \to \mathbb{R}$, we say that u is increasing if for every $\mathbf{v} = (v_1, \ldots, v_n)$, $\mathbf{v}' = (v'_1, \ldots, v'_n)$ such that $v'_i \geqslant v_i$, it holds that $u(\mathbf{v}') \geqslant u(\mathbf{v})$. We say that the subset $A \subseteq \mathbb{R}^n$ is increasing if and only if its characteristic function $\mathbf{1}_A(\mathbf{x})$ is an increasing function of $\mathbf{x}$.

Lemma B.2. *Let $\mathbf{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$, $\mathbf{D}' = \mathcal{D}'_1 \times \cdots \times \mathcal{D}'_n$ be product n-dimensional distributions with $d_k(\mathcal{D}_i, \mathcal{D}'_i) \leqslant \alpha_i$. Then for every increasing function $u : \mathbb{R}^n \to [0, \bar{u}]$ it holds that*

$$\left| \mathop{\mathbb{E}}_{\mathbf{v}\sim\mathbf{D}}[u(\mathbf{v})] - \mathop{\mathbb{E}}_{\mathbf{v}'\sim\mathbf{D}'}[u(\mathbf{v}')] \right| \leqslant \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i \right).$$

Proof. Our first step is to prove that the lemma holds for any function u that is a characteristic function of an increasing set A and then we extend to all increasing functions.

Let $u = \mathbf{1}_A$ we have that $\mathbb{E}_{\mathbf{v}\sim\mathbf{D}}[u(\mathbf{v})] = \Pr_{\mathbf{v}\sim\mathbf{D}}(\mathbf{v} \in A)$. We define the sequence of distributions $\mathbf{D}_j = \mathcal{D}'_1 \times \cdots \times \mathcal{D}'_j \times \mathcal{D}_{j+1} \times \cdots \times \mathcal{D}_n$ for $j = 0, \ldots, n$, where obviously $\mathbf{D}_0 = \mathbf{D}$ and $\mathbf{D}_n = \mathbf{D}'$. Now via triangle inequality we have that

$$\left| \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}}(\mathbf{v} \in A) - \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}'}(\mathbf{v} \in A) \right| \leqslant \sum_{j=1}^{n} \left| \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_j}(\mathbf{v} \in A) - \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_{j-1}}(\mathbf{v} \in A) \right|. \tag{8}$$

Let $b_j(\mathbf{v}_{-j})$ be the threshold of the step function $\mathbf{1}_A(v_j, \mathbf{v}_{-j})$ when we fix $\mathbf{v}_{-j}$ and we view it as a function of v_j. Now we have that

$$\begin{aligned}
\mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_j}(\mathbf{v} \in A) &= \int_{\mathbb{R}^n} \mathbf{1}_A(x_j, \mathbf{x}_{-j}) \;\; d\mathcal{D}'_1(x_1) \cdots d\mathcal{D}'_j(x_j) \cdot d\mathcal{D}_{j+1}(x_{j+1}) \cdots d\mathcal{D}_n(x_n) \\
&= \int_{\mathbb{R}^{n-1}} (1 - \mathcal{D}'_j(b_j(\mathbf{x}_{-j}))) \;\; d\mathcal{D}'_1(x_1) \cdots d\mathcal{D}'_{j-1}(x_{j-1}) \cdot d\mathcal{D}_{j+1}(x_{j+1}) \cdots d\mathcal{D}_n(x_n)
\end{aligned}$$

similarly we have

$$\mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_{j-1}}(\mathbf{v} \in A) = \int_{\mathbb{R}^{n-1}} (1 - \mathcal{D}_j(b_j(\mathbf{x}_{-j}))) \;\; d\mathcal{D}'_1(x_1) \cdots d\mathcal{D}'_{j-1}(x_{j-1}) \cdot d\mathcal{D}_{j+1}(x_{j+1}) \cdots d\mathcal{D}_n(x_n).$$

Combining these we get that

$$\begin{aligned}
&\left| \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_j}(\mathbf{v} \in A) - \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_{j-1}}(\mathbf{v} \in A) \right| \leqslant \\
&\quad \leqslant \int_{\mathbb{R}^{n-1}} \left| \mathcal{D}'_j(b_j(\mathbf{x}_{-j})) - \mathcal{D}_j(b_j(\mathbf{x}_{-j})) \right| \;\; d\mathcal{D}'_1(x_1) \cdots d\mathcal{D}'_{j-1}(x_{j-1}) \cdot d\mathcal{D}_{j+1}(x_{j+1}) \cdots d\mathcal{D}_n(x_n).
\end{aligned}$$

from the latter we can use the fact that $d_k(\mathcal{D}_j, \mathcal{D}'_j) \leqslant \alpha_j$ and we get that

$$\left| \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_j}(\mathbf{v} \in A) - \mathop{\mathbb{P}\mathrm{r}}_{\mathbf{v}\sim\mathbf{D}_{j-1}}(\mathbf{v} \in A) \right| \leqslant \alpha_j.$$

Applying the above to (8) we get that

$$\left|\Pr_{\mathbf{v}\sim\mathbf{D}}(\mathbf{v}\in A)-\Pr_{\mathbf{v}\sim\mathbf{D}'}(\mathbf{v}\in A)\right|\leqslant\sum_{j=1}^{n}\alpha_j. \tag{9}$$

The last steps is to extend the above to arbitrary increasing functions. We are going to approximate the increasing function u via a sequence of functions u_k which uniformly converges to u. Then we will show the statement of the lemma for every function u_k which by uniform convergence implies the lemma for u as well. We set $A_{i,k}\triangleq\{\mathbf{x}\in\mathbb{R}^n\mid u(\mathbf{x})\geqslant\frac{i}{k}\bar{u}\}$ and we define

$$u_k(\mathbf{x})=\frac{\bar{u}}{k}\sum_{i=1}^{k}\mathbf{1}_{A_{i,k}}(\mathbf{x}).$$

Observe from the above definition that $u_k\to u$ uniformly and since u is increasing we also have that all the sets A_i are increasing. Also observe that

$$\mathop{\mathbb{E}}_{\mathbf{v}\sim\mathbf{D}}[u_k(\mathbf{v})]=\frac{\bar{u}}{k}\sum_{i=1}^{k}\Pr_{\mathbf{v}\sim\mathbf{D}}(\mathbf{v}\in A_{i,k})$$

therefore we get that

$$\left|\mathop{\mathbb{E}}_{\mathbf{v}\sim\mathbf{D}}[u_k(\mathbf{v})]-\mathop{\mathbb{E}}_{\mathbf{v}\sim\mathbf{D}'}[u_k(\mathbf{v})]\right|\leqslant\frac{\bar{u}}{k}\sum_{i=1}^{k}\left|\Pr_{\mathbf{v}\sim\mathbf{D}}(\mathbf{v}\in A_{i,k})-\Pr_{\mathbf{v}\sim\mathbf{D}'}(\mathbf{v}\in A_{i,k})\right|.$$

Now we can apply (9) and we get

$$\left|\mathop{\mathbb{E}}_{\mathbf{v}\sim\mathbf{D}}[u_k(\mathbf{v})]-\mathop{\mathbb{E}}_{\mathbf{v}\sim\mathbf{D}'}[u_k(\mathbf{v})]\right|\leqslant\bar{u}\cdot\left(\sum_{j=1}^{n}\alpha_j\right).$$

Finally, since this is true for every u_k and u converges uniformly to u the above should be true for u as well and hence the lemma follows. ■

We are going to use Lemma B.2 both for the regular distributions case and for the MHR distributions case.

B.1 Monotone Hazard Rate Distributions—Proof of Theorem 3.6

In this section we show the part of the Theorem 3.6 related to $n>1$. For the stronger result for the case $n=1$ we refer to Section B.3.

Let $\tilde{\mathbf{D}}$ be the corrupted product distribution that we observe, $\widehat{\mathbf{D}}$ be the output distribution of Algorithm 1, $\mathbf{D}^*$ be the original distribution that we are interested in. We know from the description of Algorithm 1 for $\widehat{\mathbf{D}}=\widehat{\mathcal{D}}_1\times\cdots\times\widehat{\mathcal{D}}_n$ that $\widehat{\mathcal{D}}_i$ is MHR, that $d_k(\widehat{\mathcal{D}}_i,\mathcal{D}_i^*)\leqslant\alpha_i$ and that $\widehat{\mathcal{D}}_i\preceq\mathcal{D}_i^*$. We also know that $\mathcal{D}_i^*$ is MHR. Finally, we know that the output M of Algorithm 1 is the Myerson optimal mechanism for the distribution $\widehat{\mathbf{D}}$ and hence $\mathrm{Rev}(M,\widehat{\mathbf{D}})=\mathrm{OPT}(\widehat{\mathbf{D}})$. So applying the strong revenue monotonicity lemma 3.3 we have that

$$\mathrm{OPT}(\widehat{\mathbf{D}})=\mathrm{Rev}(M,\widehat{\mathbf{D}})\leqslant\mathrm{Rev}(M,\mathbf{D}^*). \tag{10}$$

Therefore to show Theorem 3.6, it suffices to show that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) \geqslant \left(1 - \tilde{O}\left(\sum_{i=1}^{n} \alpha_i\right)\right) \cdot \mathrm{OPT}(\mathbf{D}^*). \tag{11}$$

We are going to use the following result from [5] but with the formulation obtained in Lemma 17 of [13], combined with the weak revenue monotonicity (Lemma 3 of [13]).

Theorem B.3 ([5]). *For any product MHR distribution* $\mathbf{D}$*, and any* $\frac{1}{4} \geqslant \varepsilon \geqslant 0$ *and* $u \geqslant c \cdot \log\left(\frac{1}{\varepsilon}\right) \mathrm{OPT}(\mathbf{D})$*. Let* $t_u(\mathcal{D}_1)$*, …,* $t_u(\mathcal{D}_n)$ *be the distributions obtained by truncating* $\mathcal{D}_1$*, …,* $\mathcal{D}_n$ *at the value* $\bar{u}$ *and let* $t_u(\mathbf{D})$ *be their product distribution, where* c *is an absolute constant. Then, we have that*

$$\mathrm{OPT}(\mathbf{D}) \geqslant \mathrm{OPT}(t_u(\mathbf{D})) \geqslant (1 - \varepsilon) \cdot \mathrm{OPT}(\mathbf{D}).$$

Now let $\bar{u} = c \cdot \log\left(\frac{1}{\varepsilon}\right) \mathrm{OPT}(\mathbf{D}^*)$, then we also have that $\bar{u} \geqslant c \cdot \log\left(\frac{1}{\varepsilon}\right) \mathrm{OPT}(\widehat{\mathbf{D}})$ due to weak revenue monotonicity (Lemma 3 of [13]). Hence, applying Theorem B.3 we have that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) \geqslant \mathrm{OPT}(t_{\bar{u}}(\widehat{\mathbf{D}})) \qquad \text{and} \qquad \mathrm{OPT}(t_{\bar{u}}(\mathbf{D}^*)) \geqslant (1 - \varepsilon) \cdot \mathrm{OPT}(\mathbf{D}^*). \tag{12}$$

Since we know that $d_k(\widehat{\mathcal{D}}_i, \mathcal{D}_i^*) \leqslant \alpha_i$ we also have that $d_k(t_{\bar{u}}(\widehat{\mathcal{D}}_i), t_{\bar{u}}(\mathcal{D}_i^*)) \leqslant \alpha_i$. Let now $M_{\bar{u}}^*$ be the optimal mechanism for the distribution $t_{\bar{u}}(\mathbf{D}^*)$. It is easy to see that the ex-post revenue obtained from the mechanism $M_{\bar{u}}^*$ is an increasing function of the observed bids. Hence, we can apply Lemma B.2 to the $[0, \bar{u}]$ bounded distributions $t_{\bar{u}}(\widehat{\mathbf{D}})$ and $t_{\bar{u}}(\mathbf{D}^*)$ and we get that

$$\begin{aligned} \mathrm{OPT}(t_{\bar{u}}(\widehat{\mathbf{D}})) &\geqslant \mathrm{Rev}(M_{\bar{u}}^*, t_{\bar{u}}(\widehat{\mathbf{D}})) \geqslant \mathrm{Rev}(M_{\bar{u}}^*, t_{\bar{u}}(\mathbf{D}^*)) - \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i\right) \\ &= \mathrm{OPT}(t_{\bar{u}}(\mathbf{D}^*)) - \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i\right). \end{aligned} \tag{13}$$

If we combine (12) and (13) then we have that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) \geqslant (1 - \varepsilon) \cdot \mathrm{OPT}(\mathbf{D}^*) - \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i\right). \tag{14}$$

Now we can substitute the value of $\bar{u}$ to the above inequality and we get that

$$\mathrm{OPT}(\tilde{\mathbf{D}}) \geqslant \left(1 - c \cdot \log\left(\frac{1}{\varepsilon}\right) \cdot \left(\sum_{i=1}^{n} \alpha_i\right) - \varepsilon\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

Finally, setting $\varepsilon = \sum_{i=1}^{n} \alpha_i$ we get

$$\mathrm{OPT}(\tilde{\mathbf{D}}) \geqslant \left(1 - (c+1) \cdot \left(\sum_{i=1}^{n} \alpha_i\right) \cdot \log\left(\frac{1}{\sum_{i=1}^{n} \alpha_i}\right)\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

Hence, (11) follows and as we explained this proves Theorem 3.6.

B.2 Regular Distributions—Proof of Theorem 3.8

Let $\tilde{\mathbf{D}}$ be the corrupted product distribution that we observe, $\widehat{\mathbf{D}}$ be the output distribution of Algorithm 1, $\mathbf{D}^*$ be the original distribution that we are interested in. We know from the description of Algorithm 1 for $\widehat{\mathbf{D}} = \widehat{\mathcal{D}}_1 \times \cdots \times \widehat{\mathcal{D}}_n$ that $\widehat{\mathcal{D}}_i$ is a regular distribution, that $d_k(\widehat{\mathcal{D}}_i, \mathcal{D}_i^*) \leqslant \alpha_i$ and that $\widehat{\mathcal{D}}_i \preceq \mathcal{D}_i^*$. We also know that $\mathcal{D}_i^*$ is regular. Finally, we know that the output M of Algorithm 1 is the Myerson optimal mechanism for the distribution $\widehat{\mathbf{D}}$ and hence $\mathrm{Rev}(M, \widehat{\mathbf{D}}) = \mathrm{OPT}(\widehat{\mathbf{D}})$. So applying the strong revenue monotonicity lemma 3.3 we have that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) = \mathrm{Rev}(M, \widehat{\mathbf{D}}) \leqslant \mathrm{Rev}(M, \mathbf{D}^*). \tag{15}$$

Therefore to show Theorem 3.8, it suffices to show that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) \geqslant \left(1 - \tilde{O}\left(\sum_{i=1}^{n} \alpha_i\right)\right) \cdot \mathrm{OPT}(\mathbf{D}^*). \tag{16}$$

We are going to use the following theorem from [8], combined with the weak revenue monotonicity (Lemma 3 of [13]).

Theorem B.4 (Lemma 2 of [8])**.** *Let $\mathbf{D}$ be a product of n regular distributions and $\mathrm{OPT}(\mathbf{D})$ be the optimal revenue of $\mathbf{D}$. Suppose $\frac{1}{4} \geqslant \varepsilon \geqslant 0$ and $u \geqslant \frac{1}{\varepsilon}\mathrm{OPT}(\mathbf{D})$. Let $t_u(\mathcal{D}_1), \ldots, t_u(\mathcal{D}_n)$ be the distributions obtained by truncating $\mathcal{D}_1, \ldots, \mathcal{D}_n$ at the value u and let $t_u(\mathbf{D})$ be their product distribution. Then, we have that*

$$\mathrm{OPT}(\mathbf{D}) \geqslant \mathrm{OPT}(t_u(\mathbf{D})) \geqslant (1 - 4\varepsilon) \cdot \mathrm{OPT}(\mathbf{D}).$$

Now let $\bar{u} = \frac{1}{\varepsilon}\mathrm{OPT}(\mathbf{D}^*)$, then we also have that $\bar{u} \geqslant \frac{1}{\varepsilon}\mathrm{OPT}(\widehat{\mathbf{D}})$ due to weak revenue monotonicity (Lemma 3 of [13]). Hence, applying Theorem B.4 we have that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) \geqslant \mathrm{OPT}(t_{\bar{u}}(\widehat{\mathbf{D}})) \qquad \text{and} \qquad \mathrm{OPT}(t_{\bar{u}}(\mathbf{D}^*)) \geqslant (1 - \varepsilon) \cdot \mathrm{OPT}(\mathbf{D}^*). \tag{17}$$

Since we know that $d_k(\widehat{\mathcal{D}}_i, \mathcal{D}_i^*) \leqslant \alpha_i$ we also have that $d_k(t_{\bar{u}}(\widehat{\mathcal{D}}_i), t_{\bar{u}}(\mathcal{D}_i^*)) \leqslant \alpha_i$. Let now $M_{\bar{u}}^*$ be the optimal mechanism for the distribution $t_{\bar{u}}(\mathbf{D}^*)$. It is easy to see that the ex-post revenue obtained from the mechanism $M_{\bar{u}}^*$ is an increasing function of the observed bids. Hence, we can apply Lemma B.2 to the $[0, \bar{u}]$ bounded distributions $t_{\bar{u}}(\widehat{\mathbf{D}})$ and $t_{\bar{u}}(\mathbf{D}^*)$ and we get that

$$\begin{aligned}
\mathrm{OPT}(t_{\bar{u}}(\widehat{\mathbf{D}})) \geqslant \mathrm{Rev}(M_{\bar{u}}^*, t_{\bar{u}}(\widehat{\mathbf{D}})) &\geqslant \mathrm{Rev}(M_{\bar{u}}^*, t_{\bar{u}}(\mathbf{D}^*)) - \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i\right) \\
&= \mathrm{OPT}(t_{\bar{u}}(\mathbf{D}^*)) - \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i\right).
\end{aligned} \tag{18}$$

If we combine (17) and (18) then we have that

$$\mathrm{OPT}(\widehat{\mathbf{D}}) \geqslant (1 - \varepsilon) \cdot \mathrm{OPT}(\mathbf{D}^*) - \bar{u} \cdot \left(\sum_{i=1}^{n} \alpha_i\right). \tag{19}$$

Now we can substitute the value of $\bar{u}$ to the above inequality and we get that

$$\mathrm{OPT}(\tilde{\mathbf{D}}) \geqslant \left(1 - \frac{1}{\varepsilon} \cdot \left(\sum_{i=1}^{n} \alpha_i\right) - 4\varepsilon\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

Finally, setting $\varepsilon = \sqrt{\sum_{i=1}^{n} \alpha_i}$ we get

$$\mathrm{OPT}(\tilde{\mathbf{D}}) \geqslant \left(1 - 5 \cdot \sqrt{\sum_{i=1}^{n} \alpha_i}\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

Hence, (16) follows and as we explained this proves Theorem 3.8.

B.3 MHR Distributions – Proof of Theorem 3.6, $n = 1$ Case

In this subsection we show the part of the Theorem 3.6 related to $n = 1$, for which we obtain a stronger result compared to the case $n > 1$. We first show a useful proposition:

Proposition B.5. *Consider two MHR distributions $\mathcal{D}_1$, $\mathcal{D}_2$ with CDFs as F_1 and F_2, such that $d_k(\mathcal{D}_1, \mathcal{D}_1) \leqslant \alpha$, and $F_1(x) \geqslant F_2(x)$ for all $x \in \mathbb{R}_+$. Denote the optimal expected revenue under $\mathcal{D}_1$ and $\mathcal{D}_2$ as OPT_{F_1} and OPT_{F_2}, and the corresponding optimal reserve prices as $P^*_{F_1}$ and $P^*_{F_2}$. Then,*

$$(1 + \alpha e)^{-1} \leqslant \frac{\mathrm{OPT}_{F_1}}{\mathrm{OPT}_{F_2}} \leqslant 1 + \alpha e.$$

Proof. Consider two MHR distributions $\mathcal{D}_1$, $\mathcal{D}_2$ with CDFs as F_1 and F_2, such that $d_k(\mathcal{D}_1, \mathcal{D}_1) \leqslant \alpha$, and $F_1(x) \geqslant F_2(x)$ for all $x \in \mathbb{R}_+$. Denote the optimal expected revenue under $\mathcal{D}_1$ and $\mathcal{D}_2$ as OPT_{F_1} and OPT_{F_2}, and the corresponding optimal reserve prices as $P^*_{F_1}$ and $P^*_{F_2}$. Without loss of generality, we consider $\mathrm{OPT}_{F_1} \geqslant \mathrm{OPT}_{F_2}$. Further, since the ratio of the revenues, e.g. $\frac{\mathrm{OPT}_{F_1}}{\mathrm{OPT}_{F_2}}$ is scale invariant, we assume without loss of generality that $\mathrm{OPT}_{F_1} = 1$.

By Lemma A.2, we have $P^*_{F_1} \leqslant e$. By Lemma A.1, $\mathrm{OPT}_{F_1} = 1$ implies that $h_1(P^*_{F_1}) = \log(P^*_{F_1})$. Since $P^*_{F_1} \leqslant e$, we have

$$\begin{aligned}
& h_1(P^*_{F_1}) \leqslant 1 \\
\iff & -\log(1 - F_1(P^*_{F_1})) \leqslant 1 \\
\iff & F_1(P^*_{F_1}) \leqslant 1 - \frac{1}{e} \\
\iff & 1 - F_1(P^*_{F_1}) \geqslant \frac{1}{e}.
\end{aligned}$$

Therefore, since F_1 is non-decreasing, for any $x < P^*_{F_1}$, $1 - F_1(x) \geqslant \frac{1}{e}$. So for any $x < P^*_{F_1}$, we have

$$\begin{aligned}
|h_1(x) - h_2(x)| &= \left|\log\left(\frac{1 - F_2(x)}{1 - F_1(x)}\right)\right| \\
&= \left|\log\left(1 + \frac{F_1(x) - F_2(x)}{1 - F_1(x)}\right)\right| \\
&\leqslant \log(1 + \alpha e) \\
&= O(\alpha),
\end{aligned}$$

where the at the second last step, the inequality follows from the fact that $d_k(\mathcal{D}_1, \mathcal{D}_1) \leqslant \alpha$, and $x < P^*_{F_1}$.

Further, $F_1(x) \geqslant F_2(x)$ for all $x \in \mathbb{R}_+$ implies that $h_1(x) \geqslant h_2(x)$ for all $x \in \mathbb{R}_+$. Therefore, $h_1(P^*_{F_1}) = \log(P^*_{F_1}) \geqslant h_2(P^*_{F_1})$. Therefore, we have $P^*_{F_2} \leqslant P^*_{F_1}$, and

$$|h_1(P^*_{F_2}) - h_2(P^*_{F_2})| \leqslant \log(1 + \alpha e).$$

Now define functions $s_1(x) = h_1(x) - \log(x)$, and $s_2(x) = h_2(x) - \log(x)$. Then by the definition of $P^*_{F_1}$, $P^*_{F_2}$ and Lemma A.1,

$$\begin{aligned}\min_{x \leqslant P^*_{F_1}} s_1(x) = s_1(P^*_{F_1}) &\leqslant s_1(P^*_{F_2}) \\ &\leqslant s_2(P^*_{F_2}) + \log(1 + \alpha e) \\ &= \min_{x \leqslant P^*_{F_2}} s_2(x) + \log(1 + \alpha e).\end{aligned}$$

Therefore, by the definitions of s_1 and s_2,

$$\begin{aligned}&\left|\min_{x \leqslant P^*_{F_1}} s_1(x) - \min_{x \leqslant P^*_{F_2}} s_2(x)\right| \leqslant \log(1 + \alpha e) \\ \iff &|\log(\mathrm{OPT}_{F_2}) - \log(\mathrm{OPT}_{F_1})| \leqslant \log(1 + \alpha e) \\ \iff &-\log(1 + \alpha e) \leqslant \log(\mathrm{OPT}_{F_2}) \leqslant \log(1 + \alpha e) \\ \iff &(1 + \alpha e)^{-1} \leqslant \mathrm{OPT}_{F_2} \leqslant 1 + \alpha e.\end{aligned}$$

The above directly implies:

$$(1 + \alpha e)^{-1} \leqslant \frac{\mathrm{OPT}_{F_1}}{\mathrm{OPT}_{F_2}} \leqslant 1 + \alpha e.$$

which completes the proof. ■

Now we are ready to prove Theorem 3.6 for the $n = 1$ case.

Proof. First, by construction, Algorithm 1 runs the Myerson optimal auction on an MHR distribution $\widehat{F}$, such that $\widehat{F} \geqslant \widehat{F}'(x)$ for all $x \in \mathbb{R}_+$, for any MHR distribution $F'(x)$ such that $d_k(F'(x), \tilde{F}(x)) \leqslant \alpha$. Also by assumption, $d_k(F^*(x), \tilde{F}(x)) \leqslant \alpha$. Therefore by triangle inequality, $d_k(F^*(x), \widehat{F}(x)) \leqslant d_k(F^*(x), \tilde{F}(x)) + d_k(\tilde{F}(x), \widehat{F}(x)) \leqslant 2\alpha$.

Denote $\alpha' = 2\alpha$. By Proposition B.5,

$$\left(1 + \alpha' e\right)^{-1} \leqslant \frac{\mathrm{OPT}_{F_1}}{\mathrm{OPT}_{F_2}} \leqslant 1 + \alpha' e.$$

Note that $(1 + \alpha' e)^{-1} = (1 + 2\alpha e)^{-1} = 1 - O(\alpha)$, which completes the proof. ■

C Proof of Optimality for the Upper Bounds

For these lower bounds we follow the idea of the lower bounds from [13] adapted to the corrupted case that we consider in this paper. The lower bound constructions of [13] are based on a family of distributions

$$\mathcal{H} = \{\mathbf{D} \mid \mathcal{D}_1 = \mathcal{D}^b, \mathcal{D}_i = \mathcal{D}^h \quad \text{or} \quad \mathcal{D}_i = \mathcal{D}^\ell \text{ for all } 2 \leqslant i \leqslant n\}.$$

Observe that this family is characterized by the triplet of distributions $\mathcal{D}^b$, $\mathcal{D}^h$, and $\mathcal{D}^\ell$ for which we ask for the following conditions.

a) $\mathcal{D}^b$ is a point mass at v_0.

b) The propability of $v \geqslant v_2$ is at most $1/n$ both when $v \sim \mathcal{D}^h$ and when $v \sim \mathcal{D}^\ell$.

c) The probability of $v_1 > v \geqslant v_2$ is at least p both when $v \sim \mathcal{D}^h$ and when $v \sim \mathcal{D}^\ell$.

d) For any value v such that $v_1 > v \geqslant v_2$, we have $\phi^\ell(v) + \Delta \leqslant v_0 \leqslant \phi^h(v) - \Delta$, where ϕ^ℓ is the virtual value function of $\mathcal{D}^\ell$ and correspondingly for ϕ^h.

e) For any value v such that $v < v_2$, we have that $\phi^h(v), \phi^\ell(v) \leqslant v_0$.

f) For any value $v_1 > v \geqslant v_2$ we have that the ratio $\frac{d\mathcal{D}^h}{d\mathcal{D}^\ell}(v)$ is upper and lower bounded by a constant, where $\frac{d\mathcal{D}^h}{d\mathcal{D}^\ell}$ is the Radon–Nikodym derivative between $\mathcal{D}^h$ and $\mathcal{D}^\ell$.

g) $\mathcal{D}^h$ is regular.

h) The point v_1 is either $+\infty$ or is a point mass and an upper bound on the support in both $\mathcal{D}^\ell$ and $\mathcal{D}^h$.

Under these conditions and using the exact same proof as the Lemma 18 from [13] we can show the following.

Lemma C.1. *Let $\mathcal{H}$ be a class of distributions that satisfies the conditions a) - h) and additionally satisfies the following.*

i) We have that $d_k(\mathcal{D}^\ell, \mathcal{D}^h) \leqslant \alpha/n$.

Then any algorithm that is robust to a total corruption α in Kolmogorov distance across all bidders achieves revenue of at most

$$\mathrm{OPT}(\mathbf{D}) - \Omega(n \cdot p \cdot \Delta)$$

for any distribution $\mathbf{D} \in \mathcal{H}$.

C.1 MHR Distributions – Proof of Theorem 3.7

Let $a = \ln(n) - \ln(1-\beta)$, $b = \ln(n)$, $v_0 = a - 1$, $v_1 = \ln(n) - 2 \cdot \ln(1-\beta)$, $v_2 = a$, $p = \beta \cdot (1-\beta)/n$, $\Delta = 1/2$. Then we define $\mathcal{D}^\ell$ and $\mathcal{D}^h$ according to their CDFs F^ℓ and F^h which are the following:

$$F^\ell(v) = \begin{cases} 1 - \exp(-v) & v < v_1 \\ 0 & v \geqslant v_1 \end{cases},$$

$$F^h(v) = \begin{cases} 1 - \exp\left(-\frac{b}{a} \cdot v\right) & v < v_2 \\ 1 - \exp\left(-\frac{v_1 - b}{v_1 - a} \cdot (v - a) + b\right) & v_2 \leqslant v < v_1 \\ 0 & v \geqslant v_1 \end{cases}.$$

Observe also that for this choice of distributions it holds that

$$\phi^\ell(v) = \begin{cases} v - 1 & v < v_1 \\ v_1 & v \geqslant v_1 \end{cases},$$

$$\phi^h(v) = \begin{cases} v - \frac{a}{b} & v < v_2 \\ v - \frac{v_1 - a}{v_1 - b} & v_2 \leqslant v < v_1 \\ v_1 & v \geqslant v_1 \end{cases}.$$

Now the conditions a) - h) are easy to verify. For the condition i) we observe that the maximum difference between the two CDFs is at $v = v_2$ for which we have that $\left|F^\ell(v_2) - F^h(v_2)\right| \leqslant \beta/n$. Hence, Lemma C.1 implies that the maximum revenue achievable by any robust mechanism is

$$\mathrm{OPT}(\mathbf{D}) - \Omega(n \cdot p \cdot \Delta) = \mathrm{OPT}(\mathbf{D}) - \Omega(\beta).$$

Observe that since the maximum value of any bidder is at most $\ln(n)$ we have that the maximum revenue is

$$\left(1 - \frac{\beta}{\ln(n)}\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

If we write this expression with respect to the amount of corruption per bidder, then we have that the maximum possible revenue is

$$\left(1 - \frac{n \cdot \alpha}{\ln(n)}\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

Finally, we observe that all of $\mathcal{D}^b$, $\mathcal{D}^\ell$, and $\mathcal{D}^h$ are MHR and hence Theorem 3.7 follows.

C.2 Regular Distributions – Proof of Theorem 3.9

For the case of regular distributions we will use the same distributions used by [13] in their proof of their Theorem 2. In particular, let $v_0 = 3/2$, $v_1 = +\infty$, $v_2 = 1 + \frac{1}{\beta}$, $p = \frac{\beta}{n}$, and $\Delta = 1/2$. We define $\mathcal{D}^\ell$ and $\mathcal{D}^h$ through their CDFs as follows

$$F^\ell(v) = 1 - \frac{1}{n \cdot (v - 1)},$$

$$F^h(v) = \begin{cases} 0 & v < 1 + \frac{1}{n} \\ 1 - \frac{1}{n \cdot (v-1)} & 1 + \frac{1}{n} \geqslant v < v_2 \\ 1 - \frac{1-\beta}{n \cdot (v-2)} & v \geqslant v_2 \end{cases}.$$

The fact that these distributions satisfy a) - h) can be found in [13]. We will focus on proving i). It is not hard to see that the two CDFs appears when $v = \bar{v} = 1 + \frac{1}{\sqrt{1-\beta}}$. For this value we have

$$\left|F^\ell(\bar{v}) - F^h(\bar{v})\right| = \frac{1}{n}\left(2 - \beta - 2\sqrt{1-\beta}\right) \leqslant \frac{\beta^2}{n},$$

where the last inequality can be easily verifies for $\beta \leqslant 1$. Now setting $\alpha = \frac{\beta^2}{n}$, observing that $n \cdot p \cdot \Delta = \Omega(\beta)$, and observing that $\mathrm{OPT}(\mathbf{D}) \leqslant O(1)$ we can apply Lemma C.1 and we get that the maximum possible revenue is

$$\left(1 - \Omega\left(\sqrt{n \cdot \alpha}\right)\right) \cdot \mathrm{OPT}(\mathbf{D}).$$

Finally by observing that all of $\mathcal{D}^b$, $\mathcal{D}^\ell$, and $\mathcal{D}^h$ are regular Theorem 3.9 follows.

D Proofs of Sample Complexity Bounds

D.1 Proof of Theorem 4.3, $n > 1$ Case

This follows easily from Theorem 3.8 and the DKW inequality Dvoretzky et al. [10], Massart [17] that states that the empirical CDF with m samples is close to the population CDF with an error of at most

$$O\left(\sqrt{\frac{\log(1/\delta)}{m}}\right)$$

with probability at least $1 - \delta$. ■

D.2 Proof of Theorem 4.3, $n = 1$ Case

We present in this section a proof of Theorem 4.3 for the case with $n = 1$ and regular distributions. In this case, we show that Algorithm 2 achieves the optimal sample complexity, up to a poly-logarithmic factor.

First, by [Lemma 5, Guo et al. [13]], we have that with probability at least $1 - \delta$, for any value $v \geqslant 0$, the quantiles of $\tilde{\mathcal{D}}$ and its empirical counterpart E satisfy that:

$$|q^E(v) - q^{\tilde{\mathcal{D}}}(v)| \leqslant \sqrt{\frac{2q^{\tilde{\mathcal{D}}}(v)(1 - q^{\tilde{\mathcal{D}}(v)})\ln(2m\delta^{-1})}{m}} + \frac{\ln(2m\delta^{-1})}{m}. \tag{20}$$

Further note that by construction, we have

$$q^E - q^{\widehat{E}} \leqslant \sqrt{\frac{2q^E(v)\,(1 - q^E(v))\ln(2m\delta^{-1})}{m}} + \frac{4\ln(2m\delta^{-1})}{m} + \alpha.$$

Given that Algorithm 2 runs the Myerson optimal auction on $\tilde{E}$, which is a minimal regular distribution that dominates $\tilde{E}$. Further, $\widehat{E} \succeq D^*$ by construction, assuming Eq (20) holds. Therefore, we have $D^* \succeq \tilde{E}$ assuming Eq (20) holds. Applying Lemma 3.3 yields:

$$\mathrm{Rev}(M_{\tilde{E}}, \mathcal{D}^*) \geqslant \mathrm{Rev}(M_{\tilde{E}}, \tilde{E}) = \mathrm{OPT}(\tilde{E}).$$

Therefore, the remaining task is to ensure that m is sufficiently large such that

$$\mathrm{OPT}(\tilde{E}) \geqslant (1 - \sqrt{\alpha})\mathrm{OPT}(\mathcal{D}^*).$$

We will use a useful lemma below which connects the ratio of revenues that we are interested in with the value of link function at an optimal reserve price.

Lemma D.1. *Given two regular distributions $\mathcal{D}, \bar{\mathcal{D}}$ with CDFs $F, \bar{F}$, such that $\bar{F} \succeq F$ and $d_k(\mathcal{D}, \bar{\mathcal{D}}) \leqslant \beta$. Denote the optimal reserve price for $\bar{F}$ as $\bar{P}$, and the optimal expected revenue for $F, \bar{F}$ as $\mathrm{OPT}_F, \mathrm{OPT}_{\bar{F}}$. Then we have*

$$\frac{\mathrm{OPT}_F}{\mathrm{OPT}_{\bar{F}}} \geqslant 1 - \beta h_r(\bar{P})$$

Proof. Recall that $h_r(x) = \frac{1}{1-F(x)}$, and $\bar{h}_r(x) = \frac{1}{1-\bar{F}(x)}$. Then, $F(x) \geqslant \bar{F}(x)$ implies $h_r(x) \geqslant \bar{h}_r(x)$.

By definition, $d_k(\mathcal{D}, \bar{\mathcal{D}}) \leqslant \beta$ implies that $\max_x F(x) - \bar{F}(x) \leqslant \beta$. So we have:

$$h_r(x) - \bar{h}_r(x) = \frac{F(x) - \bar{F}(x)}{(1 - F(x))(1 - \bar{F}(x)} = (F(x) - \bar{F}(x))h_r(x)\bar{h}_r(x) \leqslant \beta h_r^2(x),$$

where the last inequality follows from the fact that $\max_x F(x) - \bar{F}(x) \leqslant \beta$, and $h_r(x) \geqslant \bar{h}_r(x)$. Thus, for all x,

$$\bar{h}_R(x) \geqslant h_r(x) - \beta h_r^2(x). \tag{21}$$

Note that the expected revenue, $R(x) = x(1 - F(x))$, at any x, equals to $\frac{x}{h_r(x)}$, which is the reciprocal of the slope for the linear function $g(a) = h_r(x) \cdot a$. Hence, the revenue is maximized when the slope for the linear function $g(a) = h_r(x) \cdot a$ is minimized.

Denote the corresponding optimal reserve prices for F and $\bar{F}$ as P and $\bar{P}$. Then at $\bar{P}$,

$$\bar{h}_r(\bar{P}) = \frac{1}{1 - \bar{F}(\bar{P})} = \frac{1}{\mathrm{OPT}_{\bar{F}}} \cdot \bar{P}.$$

Denote $\mathrm{Rev}(F, x)$ as the expected revenue with a reserve price at x for a valuation distribution with CDF as F. Then,

$$\frac{\mathrm{OPT}_F}{\mathrm{OPT}_{\bar{F}}} \geqslant \frac{\mathrm{Rev}(F, \bar{P})}{\mathrm{OPT}_{\bar{F}}} = \frac{\bar{h}_r(\bar{P})}{h(\bar{P})} \geqslant \frac{h_r(\bar{P}) - \beta h_r^2(\bar{P})}{h_r(\bar{P})} = 1 - \beta h_r(\bar{P}),$$

where the first inequality follows directly from the definition of the optimal revenue, and the second inequality is from Eq (21). ■

Now we will use Lemma D.1 to proceed. Denote the optimal reserve price for $\mathcal{D}^*$ as P^*. Denote the link function applied to $\tilde{E}$ and $\mathcal{D}^*$ as $\tilde{h}$, h^*, respectively. Then, we will discuss two cases for $\tilde{h}(P^*)$.

Case 1: $\tilde{h}(P^*) > \frac{1}{\sqrt{\alpha}}$. For this case, $\tilde{h}(P^*) > \frac{1}{\sqrt{\alpha}}$ implies that $q^{\tilde{E}}(P^*) < \sqrt{\alpha}$. Applying [Lemma 5, Guo et al. [13]] and triangle inequalities, we have

$$|q^{\tilde{E}} - q^{\mathcal{D}^*}| \leqslant \sqrt{\frac{2q^{\tilde{E}}(v)\left(1 - q^{\tilde{E}}(v)\right)\ln(2m\delta^{-1})}{m}} + \frac{4\ln(2m\delta^{-1})}{m} + \alpha.$$

Given that $q^{\tilde{E}}(P^*) < \sqrt{\alpha}$, we have $q^{\tilde{E}}(1 - q^{\tilde{E}}) \leqslant q^{\tilde{E}} \leqslant \sqrt{\alpha}$. Therefore, it suffices to have

$$\sqrt{\frac{\sqrt{\alpha}}{m}} \leqslant C_1\alpha,$$

for some universal constant C_1 to ensure that $|q^{\tilde{E}} - q^{\mathcal{D}^*}| = O(\alpha)$, which implies $m \geqslant 1/\{C_1^2\alpha^{3/2}\}$ for some universal constant C_1.

Case 2: $\tilde{h}(P^*) \leqslant \frac{1}{\sqrt{\alpha}}$. For this case, $\tilde{h}(P^*) \leqslant \frac{1}{\sqrt{\alpha}}$ implies that $q^{\tilde{E}}(P^*) \geqslant \sqrt{\alpha}$.

By lemma D.1, we have that

$$\frac{\mathrm{OPT}_{\tilde{E}}}{\mathrm{OPT}_{\mathcal{D}^*}} \geqslant 1 - \beta\tilde{h}_r(P^*),$$

therefore it suffice to ensure that $1 - \beta\tilde{h}_r(P^*) \geqslant 1 - C_2\sqrt{\alpha}$ for some universal constant C_2, which implies that $\beta \leqslant q^{\tilde{E}}(P^*) \cdot C_2\sqrt{\alpha}$. Applying [Lemma 5, Guo et al. [13]], it suffices to have that $\sqrt{\frac{q^{\tilde{E}}(P^*)}{m}} \leqslant \beta \leqslant q^{\tilde{E}}(P^*) \cdot C_2\sqrt{\alpha}$, which yields that $m > \frac{1}{C_2^2\alpha q^{\tilde{E}}}$. Lastly, applying the fact that we are in the case where $q^{\tilde{E}}(P^*) \geqslant \sqrt{\alpha}$ we get that is suffices to have $m > \frac{1}{C_2^2\alpha^{3/2}}$ for some universal constant C_2. This completes the proof.

■

D.3 Proof of Theorem 4.4

This follows easily from Theorem 3.6 and the DKW inequality [10, 17] that states that the empirical CDF with m samples is close to the population CDF with an error of at most

$$O\left(\sqrt{\frac{\log(1/\delta)}{m}}\right)$$

with probability at least $1 - \delta$. ■

D.4 Proof of Theorem 4.5

We omit the details of this proof since it follows from Theorem 2 and Appendix E of [13] applied for the case $n = 1$. The reason is that if we could get a better bound in our corrupted case then this algorithm could be used to improve our sample complexity result in the non-corrupted case.