Sample Complexity for Non-truthful Mechanisms

Jason Hartline* and Samuel Taggart†

June 24, 2019

Abstract

This paper considers the design of non-truthful mechanisms from samples. We identify a parameterized family of mechanisms with strategically simple winner-pays-bid, all-pay, and truthful payment formats. In general (not necessarily downward-closed) single-parameter feasibility environments we prove that the family has low representation and generalization error. Specifically, polynomially many bid samples suffice to identify and run a mechanism that is ϵ-close in Bayes-Nash equilibrium revenue or welfare to that of the optimal truthful mechanism with high probability.

*Northwestern University, hartline@eecs.northwestern.edu

†Oberlin College, Sam.Taggart@oberlin.edu

1 Introduction

The classical theory of revenue-maximizing mechanism design requires knowledge of agents' value distributions. As a result, the sample complexity of revenue maximization has received significant attention in recent years. This work has placed Bayesian mechanism design on more practical footing by analyzing the amount of sampled data necessary to produce a nearly-optimal mechanism. A key assumption is that samples are agents' values; e.g. they are obtained from past runs of a truthful mechanism. Moreover, the mechanisms produced are themselves truthful, and hence generate data suitable for future inference.

For many applications, however, truthful data may not be available. Samples are often equilibrium bids, produced by mechanisms where agents are not incentivized to report their values. Moreover, practical constraints often require a designer to employ non-truthful payment formats like those of the first-price (i.e. winner-pays-bid) or all-pay auction. This paper develops a theory of non-truthful mechanism design from samples. We identify a family of non-truthful mechanisms that have near-optimal revenue or welfare and require only polynomially many samples to design and implement. Our mechanisms' performance guarantees hold in equilibrium under standard non-truthful (or truthful) payment formats. We assume sampled data come from mechanisms within our family, and therefore need not be truthful. Our mechanisms may thus be redesigned as necessary when the environment changes.

Both practical and theoretical considerations necessitate the study of non-truthful mechanism design. In some common applications, outcomes are contracts, e.g., government procurement auctions, variable commission mechanisms of third-party listing agencies like Booking.com, and ad exchanges. For these applications, the theory of winner-pays-bid mechanisms is most appropriate. For games of effort – like crowdsourcing contests (e.g., Chawla et al., 2015), forecasting (e.g., Osband, 1989), and peer prediction (e.g., Dasgupta and Ghosh, 2013) – the theory of all-pay mechanisms is most appropriate. Ausubel and Milgrom (2006) discuss a number of other pragmatic concerns and describe why truthful mechanisms are rarely seen in practice. These observations justify revisiting standard theoretical questions in mechanism design (such as sample complexity) with non-truthfulness imposed as an exogenous constraint.

Non-truthful mechanisms can also exhibit surprising theoretical advantages. For example, Feng and Hartline (2018) proved that non-truthful mechanisms can be strictly better than truthful mechanisms when robustness to distributional assumptions is desired. In other words, the so-called "revelation principle" fails when looking for simple and robust mechanisms. While understanding the theoretical limits of non-truthful mechanisms is a potentially fruitful avenue of future research, such studies are hindered by the sparse literature on designing such mechanisms.

Non-truthful mechanisms add two significant complications to the study of sample complexity. First, for non-truthful mechanisms it should be assumed that there is sample access to equilibrium bids rather than values of the agents. Second, and more fundamentally, the literature has not yet shown how to design non-truthful mechanisms that have near-optimal equilibria with only estimates of the value distribution. (This task is straightforward with full distributional knowledge; see Appendix A.) To circumvent these issues, we give a black-box reduction from non-truthful mechanism design in arbitrary single-parameter settings to the design of rank-based position auctions for i.i.d. agents. In the latter setting, bid distributions are well-suited for estimation and it suffices to infer a limited set of parameters to design a near-optimal mechanism. The resulting approach is compatible with winner-pays-bid, all-pay, or truthful payment formats.

A theory of sample complexity for non-truthful mechanisms further requires careful delineation of two aspects of sample complexity: where samples come from and how samples are used. We assume samples are from the bid distributions of non-truthful mechanisms. A non-issue for truthful sample complexity, it is important for non-truthful formats which mechanism the bids are from; the Bayes-Nash equilibrium bid distributions of distinct non-truthful mechanisms are generally distinct. Our mechanisms use samples in two ways. First, samples can be used at design time to select the mechanism to be run, in which case they are usually data from past mechanisms, and can be assumed to be observed by the agents before bids are placed. Second, samples can be used at run time when the mechanism executes, in which case they are viewed as a source of randomness as agents are bidding. Run-time samples are most plausible when they are from concurrent runs of the same mechanism, e.g., in high-frequency settings such as advertising auctions. In analyses of mechanisms from samples, incentives and performance guarantees hold with high probability in design-time samples and in expectation in run-time samples.

Both uses have analogs in the literature on truthful mechanism design from samples. Most papers on sample complexity use design-time samples and elide computational considerations around producing desirable allocations and truthful payments. The literature on black-box reductions from Bayesian incentive compatible mechanism design to Bayesian algorithm design, meanwhile, relies heavily on run-time samples to achieve exact incentive compatibility. Even for the simplest setting of single-parameter agents (Hartline and Lucier, 2015) designing an exactly incentive compatible black-box reduction from design-time samples alone remains an open problem. A more detailed discussion of the literatures on sample complexity and black-box reductions, as well as the way they use samples, can be found in the related work section below.

Problem Statement We consider the problem of designing good mechanisms from samples in general single-parameter environments with independently distributed values in Bayes-Nash equilibrium where a general set system governs the subsets of agents that can be simultaneously served. Truthful mechanisms for this environment are well understood. The Vickrey-Clarke-Groves mechanism maximizes welfare, and a straightforward generalization of Myerson (1981) gives the truthful mechanism that maximizes expected revenue (e.g., Hartline, 2013). The sample complexity of truthful mechanisms of this setting was largely resolved by Devanur et al. (2016), Gonczarowski and Nisan (2017), and Guo et al. (2019).

We generalize truthful sample complexity to non-truthful mechanisms in the following way. The problem of *non-truthful sample complexity* is to identify in a parameterized family of winner-pays-bid (or all-pay) mechanisms and polynomials p_{design} and p_{run} such that with n-agent environments and desired loss ϵ:

C1. With $m_{\text{design}} = p_{\text{design}}(n, \epsilon^{-1}, \delta^{-1})$ design-time samples of profiles of Bayes-Nash equilibrium bids from any mechanism in the family, parameters of the designed mechanism can be selected.

C2. With $m_{\text{run}} = p_{\text{design}}(n, \epsilon^{-1})$ run-time samples of profiles of Bayes-Nash equilibrium bids in the selected mechanism, the selected mechanism can be run.

C3. With probability in the m_{design} design-time samples of at least $1 - \delta$, the expected performance, in agents' values and the m_{run} run-time samples of the selected mechanism, is at most ϵ less than that of the Bayesian optimal mechanism.[1]

The following story fits the above problem and is implicit in previous papers on mechanism design from samples. Per a standard interpretation of the Bayesian model for auctions, a designer aims to run a mechanism on agents drawn from one or several large populations. A mechanism is sought that performs well on a fresh draw of agents from each population. Our non-truthful mechanism designer fixes a large parameterized family of mechanisms and has independently drawn profiles of historical bids in one mechanism in the family. The designer uses these historical bid profiles as design-time samples to select new parameters of the mechanism. The agents adapt to the new equilibrium in the new mechanism. The designer collects historical samples in the new mechanism and uses them as run-time samples in its execution.

As is common in the literature on Bayes-Nash mechanism design, we consider runtime samples and agent strategies which follow a steady state equilibrium. While it is beyond the scope of this paper to model the adaptive process by which the agents might learn this equilibrium, we shall see that the mechanisms we produce have straightforward bidding problems for each agent. Similarly, we consider it beyond the scope of the paper to explicitly model the process by which runtime samples are obtained. For motivation, however, we note here several scenarios which justify their use. In practical applications such as ad auctions, agents bid in advance of the auction, and it is possible to batch the bid collection for many individual executions of the mechanism together. For these batched executions, the run-time samples can be from the other bid profiles that are collected within the same batch. If batching is infeasible, it may still be possible to produce run-time samples in an online manner by taking bid data from the most recent iterations of the mechanism.

Approach and Results. We solve the stated problem for general single-parameter environments and independent but non-identically distributed agents. With n agents, polynomial in n, ϵ^{-1}, and δ^{-1} design-time samples are sufficient to identify a mechanism that, with polynomially many run-time samples and probability at least $1 - \delta$, approximates the performance of the optimal mechanism to within precision ϵ in the following environments:

- (non-truthful) winner-pays-bid and all-pay mechanisms, additive welfare approximation, and bounded value distributions;[2]
- (non-truthful) winner-pays-bid and all-pay mechanisms, additive revenue approximation, and bounded and regular value distributions; and
- truthful mechanisms, multiplicative revenue approximation, and (unbounded) regular value distributions.[3]

Regular distributions are ones that satisfy a natural convexity property; details are given in Section 2.

A key primitive in our results is the *i.i.d. rank-based position auction*, a model popularized in the study of ad auctions on search engines, cf. Jansen and Mullen (2008). In such an auction agents are assigned to positions, with higher positions having higher allocation probabilities. In

[1] Multiplicative versions of C3 are also interesting. In the multiplicative version a mechanism with expected welfare or revenue at least $(1 - \epsilon)$ times optimal is required.

[2] Because of asymmetries in the value distributions of agents and the feasibility environment, welfare maximization subject to non-truthful payment sematics is not trivial. See Appendix A for more details.

[3] These results are more general than the results of Devanur et al. (2016), who require downward closure.

an i.i.d. position auction the agents' values are drawn from the same distribution. Equilibria in these auctions are unique and efficient (Chawla and Hartline, 2013), and are simple to compute using standard characterizations of Bayes-Nash equilibrium. One way to view our results is as a reduction from sample complexity in general single-parameter environments with non-identically distributed agents to inference and design in i.i.d. position auctions.

To implement this reduction, our mechanisms use run-time samples from their own bid distribution. Doing so effectively replaces competition between agents from distinct distributions with competition between agents with identical distributions (from the run-time samples). We show that to any agent, this is strategically equivalent to bidding in an i.i.d. position auction. Our mechanism therefore inherits the simple equilibrium structure of the latter mechanisms. We further show that mechanisms in our family exist that are approximately optimal, i.e., the representation error of our family is small, and we reduce the problem of analyzing the generalization error to the problem of estimating appropriate expected order statistics of design-time samples from bids in any i.i.d. position auction with the same agent distributions. For i.i.d. position auctions with standard payment formats like winner-pays-bid and all-pay, Chawla et al. (2017) solve this inference problem. For the truthful format, we give a straightforward solution and analysis.

Related Work Typical welfare and revenue analyses of non-truthful mechanisms take standard auctions and analyze their welfare in worst-case equilibrium. Notable examples include Syrgkanis and Tardos (2013), who prove that first-price and all-pay auctions have welfare within a constant fraction of optimal, and Hartline et al. (2014), who derive a similar result for revenue. See Roughgarden et al. (2017) for a complete survey. Unfortunately, except in settings such as i.i.d. position auctions, where equilibrium is efficient (Chawla and Hartline, 2013), asymmetries in the distributions and set of feasible allocations seem to guarantee a nontrivial fraction of welfare or revenue is lost in worst-case instances. In fact, Dütting and Kesselheim (2015) show that standard analysis techniques cannot prove near-optimal revenue results for non-truthful mechanisms in a variety of common settings. Notably, these results only apply to designs and analyses which are agnostic to agents' distributions.

We consider mechanism design from data to circumvent these negative results. Design of truthful mechanisms from data has been studied extensively. For single-parameter mechanism design, this includes the learned finite support auctions of Elkind (2007), the learned monopoly reserve of Dhangwatnotai et al. (2010), and Cole and Roughgarden (2014) as well as followups (Devanur et al., 2016; Roughgarden and Schrijvers, 2016; Morgenstern and Roughgarden, 2015; Gonczarowski and Nisan, 2017; Guo et al., 2019). A robust literature on the sample complexity of multi-parameter mechanism design also exists. As we only consider single-parameter agents, this literature is beyond the scope of our discussion. Relatively little work exists on the design of non-truthful mechanisms from data. A notable exception is Chawla et al. (2014), whose inference methodology for i.i.d. position auctions are used in our own framework. All papers mentioned above use only design-time samples.

Our design approach has close connections to the work on reducing Bayesian incentive compatible mechanism design to Bayesian algorithm design. This work relies on both design-time and run-time samples. The reductions of Hartline and Lucier (2015), Hartline et al. (2011), and Bei and Huang (2011) show that given an algorithm $\mathcal{A}$ and design-time sample access to agents' value distributions, one can construct an ϵ-Bayesian incentive compatible mechanism with at most ϵ less expected welfare than $\mathcal{A}$ in polynomial time. The constructed mechanisms can be implemented without reliance on additional samples, but fall short of exact incentive compatibility. For limited families of preferences, the former two papers show how to use run-time samples to achieve exact incentive compatibility, and Dughmi et al. (2017) recently showed how to use run-time samples to produce a fully general exactly incentive compatible reduction. Achieving exact incentive compatibility in blackbox reductions with only design-time samples is an open problem for even single-parameter agents. Indeed, eliminating the dependence on run-time samples in our setting would likely imply a breakthrough in this regard.

Organization In Section 2, we lay out notation and preliminary results. We present our parametrized family of mechanisms, which we term surrogate-ranking mechanisms, in Section 3, and analyze their equilibria. In Section 4, we show that for any set of value distributions, there exists a surrogate-ranking mechanism that uses polynomially many run-time samples and obtains nearly-optimal welfare or revenue. Finally, in Section 5, we show how to learn such a mechanism using polynomially many design-time samples.

2 Preliminaries

This work considers the *single-parameter independent private value* model of mechanism design. We describe this model in *quantile* space where the geometry of approximation mechanisms is more transparent (cf. Hartline, 2013). There are n agents drawn independently and uniformly at random from n populations. Agents are distinguished by their quantile with respect to their own population. The *quantile* q_i of agent i is the measure of population i with higher values. The *value function* v_i of population i maps agent i's quantile to her value as $v_i(q_i)$ and, with a uniformly drawn quantile, induces a value distribution. Profiles of agent values and quantiles are denoted by $\boldsymbol{v} = (v_1, \ldots, v_n)$ and $\boldsymbol{q} = (q_1, \ldots, q_n)$, respectively.

An *allocation* is $\boldsymbol{x} = (x_1, \ldots, x_n)$ where $x_i \in \{0, 1\}$ is an indicator for agent i being served. The space of feasible alloctions is given by $\mathcal{X} \subset \{0, 1\}^n$. (Notably, we do not require that $\mathcal{X}$ be downward closed.) Agent i can be assigned a non-negative payment denoted p_i and her utility is linear in allocation and payment as $v_i(q_i)\, x_i - p_i$.

A mechanism takes as input a profile of bids $\boldsymbol{b} = (b_1, \ldots, b_n)$ and outputs a feasible allocation $\boldsymbol{x} \in \mathcal{X}$ and agent payments $\boldsymbol{p}$. A mechanism consists of an *allocation algorithm* $\tilde{\boldsymbol{x}}(\boldsymbol{b})$, which maps bid profiles to a feasible allocation, and a payment rule $\tilde{\boldsymbol{p}}(\boldsymbol{b})$, which maps bid profiles to a non-negative payment for each agent. A standard allocation algorithm is *highest-bids-win* which is defined by $\tilde{\boldsymbol{x}}(\boldsymbol{b}) \in \operatorname{argmax}_{\boldsymbol{x} \in \mathcal{X}} \sum_i b_i\, x_i$. We consider payment rules defined directly from the allocation algorithm according to standard payment formats. The *winner-pays-bid* format has payment rule $\tilde{p}_i(\boldsymbol{b}) = b_i\, \tilde{x}_i(\boldsymbol{b})$, and the *all-pay* format has payment rule $\tilde{p}(\boldsymbol{b}) = b_i$. Mechanisms with these payment formats do not have truth-telling as an equilibrium. The truthful payment format is defined according to the payment identity (below, Theorem 1) and can be implemented as an integral or with any of a number of unbiased estimators with expectation equal to the integral (see, e.g., Hartline and Lucier, 2015).

We analyze non-truthful mechanisms in Bayes-Nash equilibrium (BNE): each agent's report to the mechanism is a best response to the distribution of bids induced by other agents' strategies. The strategy of agent i is denoted s_i and maps the agent's quantile to a bid and, with a uniformly drawn quantile, induces a bid distribution. The mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$, the agents' strategies $\boldsymbol{s}$, and the distribution over quantiles induce interim allocation and payment rules. Agent i's *interim allocation rule* is $x_i(q_i) = \mathbf{E}_{\boldsymbol{q}_{-i}}[\tilde{x}_i(\boldsymbol{s}(\boldsymbol{q}))]$ and *interim payment rule* $p_i(q_i) = \mathbf{E}_{\boldsymbol{q}_{-i}}[\tilde{p}_i(\boldsymbol{s}(\boldsymbol{q}))]$.

Myerson (1981) characterized the interim allocation and payment rules that arise in BNE when agents' values are independently distributed.

Theorem 1 (Myerson, 1981)**.** *For independently distributed agents, interim allocation and payment rules are induced by a Bayes-Nash equilibrium with onto strategies if and only if for each agent i,*

1. *(monotonicity) allocation rule $x_i(q)$ is monotone non-increasing in q_i, and*
2. *(payment identity) payment rule $p_i(q_i)$ satisfies $p_i(q_i) = v_i(q_i)\, x_i(q_i) + \int_{q_i}^{1} x_i(r)\, v_i'(r)\, dr + p_i(1)$.*

This paper studies the objectives of welfare and revenue. The *welfare* of a mechanism is $\mathbf{E}[\sum_i v_i(q)\, x_i(q_i)]$. The optimal mechanism for welfare allocates the value-maximizing feasible set, which is monotone and therefore implementable with payments via Theorem 1. The *revenue* of a mechanism is given by $\mathbf{E}[\sum_i p_i(q_i)]$. The revenue of a mechanism is easily analyzed in quantile space in terms of revenue curves and marginal revenue as follows.

Lemma 2 (Myerson, 1981; Bulow and Roberts, 1989)**.** *In BNE, the expected payment of an agent i satisfies*

$$\mathbf{E}_{q_i}[p_i(q_i)] = \mathbf{E}_{q_i}[-x_i'(q_i)\, R_i(q_i)] + R_i(1)\, x_i(1) = \mathbf{E}_{q_i}[R_i'(q_i)\, x_i(q_i)] + R_i(0)\, x_i(0)$$

where the revenue curve $R_i(q_i) = v_i(q_i)\, q_i$ *gives the revenue from posting price q_i and the* marginal revenue $R_i'(q_i) = v_i(q_i) + v_i'(q_i) q_i$ *is its derivative. (Note that the derivatitives of the allocation rule $x_i'(\cdot)$ and value function $v_i'(\cdot)$ are non-positive.)*

The first equality follows from revenue equivalence and noting that the allocation rule x_i is equivalent to offering a randomized posted price with price distributed according to the density function $-x_i'(\cdot)$ with a pointmass of $x_i(1)$ at price $v_i(1)$. The second equality follows from integration by parts. The optimal mechanism can be easily identified from the second equality as maximizing the surplus of marginal revenue. The value distributions are called *regular* when the revenue curves are concave, or equivalently, the marginal revenues are monotonically non-increasing.

In many environments of interest, the additive terms $R_i(0)\, x_i(0)$ and $R_i(1)\, x_i(1)$ are zero. For example, when the strongest agent $q_i = 0$ in the population has finite value $v_i(0)$, then the revenue when we post price $v_i(0)$ is $R_i(0) = 0$ as only a zero measure of the population will buy at such a price. When the weakest agent in the population $q_i = 1$ has value $v_i(1) = 0$ then $R_i(1) = 0$ as the revenue from posting price 0 is zero.

Position Auctions *I.i.d. rank-by-bid position auctions* play a fundamental role in our analysis. In i.i.d. environments the agents' value functions are identical $v_i = v_j$ for all agents i and j. An n-agent position auction is defined by n position weights $w_1 \geq \ldots \geq w_n \in [0, 1]$ and an outcome is an assignment of agents to positions. If agent i is assigned to position j her allocation is $x_i = 1$ with probability w_j and zero otherwise, i.e., $\mathbf{E}[x_i \mid \text{agent } i \text{ is assigned slot } j] = w_j$. The rank-by-bid allocation algorithm assigns agents to positions assortatively by bid. The following theorem shows that Bayes-Nash equilibria in rank-by-bid position auctions are straightforward.

Theorem 3 (Chawla and Hartline, 2013)**.** *In i.i.d. position environments, the rank-by-bid winner-pays-bid and all-pay auctions have a unique and welfare-maximizing Bayes-Nash equilibrium (in which agents are assigned to positions in order of their true values), i.e., $s_i(\cdot) = s_j(\cdot)$ for all agents i and j.*

3 Surrogate-Ranking Mechanisms

In this section, we describe the parameterized family of mechanisms for which we demonstrate polynomial sample complexity. A mechanism in this family has run-time sample access to the equilibrium bid distribution of each agent. An agent's bid can be compared to these samples to estimate the agent's strength relative to their value distribution. The mechanism then allocates solely on the basis of the agents' ranks. The choice of parameters will determine the exact mapping between ranks and allocations. This approach can be paired with any of the standard payment formats: winner-pays-bid, all-pay, or truthful.

Definition 4. *A* surrogate-ranking mechanism *(SRM) is parameterized by nT* surrogate values *$\mathbf{\Psi}$, with $\Psi_i = \{\psi_i^1 \geq \ldots \geq \psi_i^T\}$ for each agent i. The input to the mechanism is a profile of bids.*

1. *A surrogate value is calculated for each agent i as:*
   (a) *draw $T-1$ run-time samples from the agent's bid distribution,*
   (b) *calculate the rank r_i of the agent's bid relative to these samples,*
   (c) *select the agent's surrogate value $\psi_i = \psi_i^{r_i}$ according to the agent's sample rank.*
2. *For space $\mathcal{X}$ of feasible allocations, the algorithm allocates to maximize the* surrogate surplus $\operatorname{argmax}_{\boldsymbol{x}\in\mathcal{X}} \sum_i \psi_i\, x_i$.
3. *Payments are assigned according to any standard payment format, e.g., winner-pays-bid, all-pay, or truthful.*

In the paper we will focus on surrogate ranking mechanisms as defined above where the allocation is chosen to maximize the surrogate surplus, i.e., $\sum_i \psi_i\, x_i$. Our methods extend in a straightforward manner to settings where computing such an allocation is intractable. For approximation algorithms where surrogate allocations is monotone in surrogate values, all analyses in this paper hold with an additional multiplicative performance loss equal to the approximation factor of the algorithm. Non-monotone algorithms can be made monotone via the methods of Hartline and Lucier (2015) or Hartline et al. (2015).

Subsequently in Section 5, we will show how to identify good surrogate values from design-time samples. The remainder of this section is devoted to characterizing equilibria in surrogate ranking mechanisms.

3.1 Equilibria of Surrogate-Ranking Mechanisms

We now analyze the equilibrium of winner-pays-bid and all-pay surrogate-ranking mechanisms.[4] To do so, we first give a natural generalization of Bayes-Nash equilibrium to mechanisms with run-time samples from its own bid distribution.

Definition 5. *A* stationary equilibrium (with samples) *in a mechanism (with samples) is a strategy profile $\boldsymbol{s}$ where the strategy of each agent is in best response to distribution of bids induced by the strategies in the mechanism with sample access to the same bid distributions.*

[4] In Section 5, we also consider truthful SRMs. Equilibrium analysis for truthful SRMs is trivial.

We will show that stationary equilibria of surrogate ranking mechanisms are easy to characterize. Each agent plays the unique Bayes-Nash equilibrium strategy of an i.i.d. position auction for their distribution in a position environment derived from the choice of surrogate values. Specifically, rather than competing with other agents in the mechanism, an agent i's bid competes with other bids from her bid distribution which gives an outcome which is equivalent to the equilibrium of a position auction with agents with values drawn only from population i, who share a distribution.

We begin by analyzing the distribution of assigned surrogate values in a stationary equilibrium. Recall, each agent's strategy s_i, on a uniform quantile, induces a distribution over bids. Notice that, in the surrogate-ranking mechanism with sample access to this bid distribution, the surrogate value assigned to i will be uniformly distributed from the set Ψ_i of i's surrogate values. This implies:

Lemma 6. *In any stationary equilibrium of a surrogate ranking mechanism, and any agent i and surrogate value $\psi_i^j \in \Psi_i$, the ex ante probability agent i is assigned ψ_i^j is $1/T$.*

Lemma 6 implies that in a stationary equilibrium, the probability of allocation associated with a particular surrogate value is fully determined by the other agents' sets of surrogate values, and not by the form of the equilibrium bidding strategies, or even by the agents' value distributions. This characterization of outcomes can be formalized as follows.

Definition 7. *For each agent i, let $\Psi_i = \{\psi_i^1 \geq \ldots \geq \psi_i^T\}$ be agent i's set of surrogate values, and let $\hat{\boldsymbol{x}}$ denote the surrogate surplus maximizing allocation rule $\hat{\boldsymbol{x}}(\boldsymbol{\psi}) = \max_{\boldsymbol{x} \in \mathcal{X}} \sum_i \psi_i\, x_i$. The* characteristic weights *$W_i = \{w_i^1 \geq \ldots \geq w_i^T\}$ for agent i are defined by calculating the allocation probability associated with each surrogate when the surrogates of other populations are drawn uniformly at random, i.e., $w_i^j = \mathbf{E}[\hat{x}_i(\psi_i^j, \boldsymbol{\psi}_{-i})]$ for each surrogate j and uniform random $\boldsymbol{\psi}_{-i}$ from $\boldsymbol{\Psi}_{-i}$.*

We now show that from each agent's perspective, stationary equilibria in surrogate-ranking mechanisms look like a position auction among agents with the same value function. These agents compete for the characteristic weights of their population's surrogate values. Under the pay-your-bid or all-pay payment formats, they therefore inherit the equilibrium of rank-based position auctions, which is shown by Chawla and Hartline (2013) to be efficient (i.e., to rank agents by values) and unique.

Theorem 8. *For any profile of value functions $\boldsymbol{v}$, surrogate values $\boldsymbol{\Psi}$, and characteristic weights $\boldsymbol{W}$; the unique stationary equilibrium of the winner-pays-bid (resp. all-pay) SRM is given by each agent i bidding according to the unique and efficient BNE s_i of the i.i.d. winner-pays-bid (resp. all-pay) position auction with position weights W_i and value function v_i.*

Proof. Assume an arbitrary stationary equilibrium and consider an agent i. By Lemma 6, the stationary equilibrium induces a uniform distribution over each other agent's assigned surrogate values. It follows that if agent i is assigned surrogate value ψ_i^j, then they are allocated with probability w_i^j. Moreover, since the surplus-maximizing allocation algorithm is monotone, characteristic weights for agent i are monotone as well. Hence, placing the jth highest bid among the run-time samples will cause i to be assigned the jth highest characteristic weight. Thus, agent i faces the same bidding problem as if they played in the equilibrium of the i.i.d. position auction with position weights $w_i^1, \ldots, w_i^T$ and value function v_i. Thus, agent i bids according to the BNE of this i.i.d. position auction. This BNE is efficient, i.e., bids are in the same order as values, and unique.

Uniqueness of the stationary equilibrium follows by uniqueness of characteristic weights under any stationary equilibrium (Definition 7), which are determined only by the set of surrogate values $\mathbf{\Psi}$, and the uniqueness of symmetric Bayes-Nash equilibrium in i.i.d. position auctions, which follows from revenue equivalence. □

One consequence of Theorem 8 is that the bidding problem faced by agents is strategically simple. Given accurate estimates of the characteristic weights, the symmetric equilibrium bids for the corresponding position environment for each population can be computed using Theorem 1.

3.2 Equivalence of Surrogate Ranking Mechanisms

Surrogate ranking mechanisms are equivalent for revenue and welfare, irrespective of their payment format. It is helpful to relate this equivalence to the famous revenue equivalence result of Myerson (1981). In the latter, two mechanisms with the same equilibrium outcome (and with the same expected payment of the agent with the lowest value in the support of the distribution, usually zero) have the same expected revenue. For instance, with i.i.d. distributions the single-item first-price and second-price auctions have the same equilibrium outcome, i.e., the highest valued agent wins, and thus, by revenue equivalence, the same expected revenue. With non-identical distributions, these auctions do not have the same equilibrium outcome and, thus, do not generally have the same expected revenue. Our equivalence result, in contrast, holds for surrogate-ranking mechanisms in asymmetric environments (distributions and feasibility constraints).

Theorem 9. *For any fixed surrogate values and value functions, the expected welfare (resp. revenue) of the winner-pays-bid, all-pay, and truthful surrogate-ranking mechanisms in stationary equilibrium with samples are equal.*

Proof. This theorem follows because each agent is playing the symmetric BNE of an i.i.d. position auction with characteristic weights that are independent of the payment format. Such BNE are welfare-equivalent for each agent. Welfare equivalence implies, by the usual argument, revenue equivalence. □

Theorem 9 gives a revelation principle for surrogate ranking mechanisms. Bounds on the revenue and welfare of the truthful surrogate ranking mechanism implies the same bounds on that of the non-truthful ones because their equilibrium outcomes are the same in expectation.

4 Representation Error of Surrogate Ranking Mechanisms

In this section, we show that the representation error of surrogate-ranking mechanisms is small. In other words, there always exists some choice of surrogate values which induces near-optimal expected welfare or revenue. Hence, SRMs satisfy condition C3 in our statement of the problem of non-truthful sample complexity. More precisely, we prove the following guarantee:

Theorem 10. *There exists a surrogate-ranking mechanism with winner-pays-bid, all-pay, or truthful payment semantics which attains a $(1 - O(\sqrt[3]{n/T}))$-fraction of the optimal welfare in* *stationary equilibrium. With regular distributions, there exists such a mechanism which attains a $(1 - O(\sqrt[3]{n/T}))$-fraction of the optimal revenue in stationary equilibrium.*

In Section 5, we show how to estimate a SRM which is nearly welfare- or revenue-optimal among all mechanisms in the family. This mechanism consequently inherits the guarantees of Theorem 10, up to error from estimation.

Before discussing the proof of Theorem 10 we observe that a less general result follows from a theorem of Hartline et al. (2011). In this paper, the authors consider an allocation procedure which can be interpreted as a surrogate ranking mechanism. They show that for agents whose values are distributed on the interval $[0, 1]$, their mechanism has an additive welfare loss of at most $n/(4\sqrt{T})$. We improve on this welfare guarantee in three ways. First, we derive a guarantee for arbitrary distributions, even those with unbounded support. Second, our bounds will be multiplicative. Finally, our bounds apply to revenue in addition to welfare.[5]

To prove Theorem 10 we define a family of mechanisms, *surrogate binning mechanisms*, which coarsen quantile space into uniform bins, assign agents to surrogate values based on their bin, and maximize surplus with respect to the assigned surrogate values. Binning mechanisms interpolate between the structure of the optimal mechanism and SRMs; they treat agents coarsely based on their location in their distribution, whereas the optimal mechanism treats agents based on their exact location in their distribution, and SRMs treat agents coarsely based on their rank among runtime samples. Theorem 10 follows from identifying a set of surrogate values for which the surrogate binning mechanism is nearly optimal and the SRM is not much worse.

We formally define surrogate binning mechanisms in Section 4.1, and show that for any set of surrogate values where each agent's highest k and lowest k surrogate values are each identical, surrogate ranking and binning mechanisms perform comparably, with the loss depending on k. It therefore suffices to find a set of surrogate values satisfying this property that yields a near-optimal surrogate binning mechanism. We do so in Section 4.2.

4.1 Surrogate-Binning Mechanisms

We now discuss our intermediate family of mechanisms, surrogate-binning mechanisms. Recall that surrogate-ranking mechanisms assign agents surrogate values from a fixed set based on their bid's rank among run-time samples, and allocates the feasible set with the highest total surrogate value. Surrogate-binning mechanisms perform the same procedure, but assign surrogate values based on a coarsening of each distribution's quantile space, rather than ranking among samples. Formally:

Definition 11. *A* surrogate binning mechanism *(SBM) is parameterized by nT* surrogate values *$\boldsymbol{\Psi}$, with $\Psi_i = \{\psi_i^1 \geq \ldots \geq \psi_i^T\}$ for each agent i. The input to the mechanism is a profile of quantiles for each agent.*

1. *A surrogate value is calculated for each agent i as $\psi_i = \psi_i^j$, where $q_i \in ((j-1)/T, j/T]$.*

2. *For space $\mathcal{X}$ of feasible allocations, the algorithm allocates to maximize the* surrogate surplus $\operatorname{argmax}_{\boldsymbol{x} \in \mathcal{X}} \sum_i \psi_i \, x_i$.

3. *Charge truthful payments.*

[5] With the additional assumption that virtual values are bounded below by $-\underline{\phi}$, the guarantee of Hartline et al. (2011) applies to revenue as well, with an additional factor of $1 + \underline{\phi}$ applied to the loss.

We refer to the interval $((j-1)/T, j/T]$ in agent i's quantile space as the *jth bin* for that agent. SRMs and SBMs are similar in their use of surrogate values, but differ in that SRMs use ranking to discriminate between high- and low-valued agents, while SBMs compare agents' bins. In this section, we show that as long as neither mechanism discriminates at the top or bottom of the value distribution, i.e. the highest surrogate values for each agent are the same and the lowest surrogate values for each agent are the same, then binning and ranking discriminate comparably well. Formally:

Theorem 12. *For surrogate values $\psi_i^1 \geq \psi_i^2 \geq \ldots \geq \psi_i^T$ with $\psi_i^1 = \psi_i^2 = \ldots = \psi_i^{\underline{k}}$ and $\psi_i^{\overline{k}} = \psi_i^{\overline{k}+1} = \ldots = \psi_i^T$ for each population i, the SRM for $\mathbf{\Psi}$ with all-pay, winner-pays-bid, or truthful payments in stationary equilibrium attains a $(1 - O(\underline{k}^{-1/2}))$-fraction of the welfare of the SBM for $\mathbf{\Psi}$. If distributions are regular, then the SRM attains a $(1 - O(\min(\underline{k}, T - \overline{k})^{-1/2}))$-fraction of the SBM's virtual surplus as well.*

To derive Theorem 18, first prove in Section 4.1.1 that the simplest ranking mechanism, a k-unit, highest-bids-win auction, approximates the performance of posted pricing, which can be seen as the simplest binning mechanism. In Section 4.1.2, we then show that SRMs and SBMs appear as distributions over these simpler single-agent mechanisms, which implies Theorem 18.

4.1.1 Ranking Versus Pricing

In this section, we establish a lemma relating the performance of simple mechanisms which allocate agents based on their rank among samples and those that allocate based on their quantiles. This result will serve as a technical building block for comparing surrogate ranking and surrogate binning mechanisms.

Consider the following two methods for selling an unlimited supply of items to T identically distributed agents. First, the seller could impose an *ex post* supply limit of k units, and sell them to the k highest-valued agents. Second, the seller could instead relax to a k-unit supply limit *in expectation* by selling to all agents with quantile below k/T (i.e. posting a price of $v(k/T)$). Note that each approach can be implemented as a surrogate-ranking or surrogate-binning mechanism, respectively. Formally:

Definition 13. *The* k/T-price posting mechanism *allocates agents if and only if their quantile is below k/T, for some integer k. This can be achieved by posting the price with quantile k/T in an environment with unlimited supply.*

Definition 14. *The* top k-of-T mechanism *for T agents ranks agents by value and allocates the k agents with the highest values. It charges winners the $k+1$st highest value.*

As the law of large numbers might suggest, these two mechanisms perform comparably for large T when k is bounded away from the extremes (1 and $T-1$).

Lemma 15. *The top k-of-T mechanism attains at least a $(1 - O(k^{-1/2}))$-fraction of the welfare of the k/T-price posting algorithm with T i.i.d. agents. If values are regularly distributed, then it attains at least a $(1 - O(\min(k, T-k)^{-1/2}))$-fraction of the revenue of the k/T-price posting mechanism.*

A proof can be found in Appendix B. The result follows from an explicit characterization of the distributions for which the ratio in performance between the two mechanisms is largest. Hence, the bounds of the lemma are tight.

One immediate generalization of Lemma 15 is to distributions over pricing mechanisms and analogous distributions of top-k mechanisms.

Lemma 16. *Consider a distribution over k/T-price posting mechanisms for T i.i.d. agents, where the highest price is at quantile $\underline{k}/T$ and the lowest price is at quantile $\overline{k}/T$. The same distribution over corresponding top-k-of-T mechanisms attains a $(1 - O(\underline{k}^{-1/2}))$-fraction of the welfare of the distribution over price-posting mechanisms. If values are regularly distributed, then the distribution over top-k-of-T algorithms attains a $(1 - O(\min(\underline{k}, T - \overline{k})^{-1/2}))$-fraction of the price-posting revenue as well.*

Proof. Lemma 15 implies that for each price in the distribution of the price-posting mechanism, there is a top-k-of-T algorithm which approximates it and which appears with the same probability. The approximation ratio of a distribution over pairwise approximations is at least the approximation from the worst pair. □

4.1.2 Ranking Versus Binning

Lemma 16 shows that in simple settings, ranking discriminates almost as effectively as pricing. We now extend this idea to surrogate-ranking and surrogate-binning mechanisms.

Any allocation rule can be viewed by agents as a distribution over posted prices. For surrogate binning mechanisms, this distribution is especially easy to analyze. Because surrogate-binning mechanisms treat agents solely based on the bin into which their quantile falls, and because bins are distributed evenly in quantile space, it follows that for surrogate-binning mechanisms, the allocation rule is piecewise-constant, with break points occurring at multiples of $1/T$. It follows that the corresponding distribution over posted prices can actually be viewed as a distribution over k/T-quantile price posting mechanisms. Moreover, since quantiles (and agents' bins) are distributed uniformly, the probability of allocation associated with a particular surrogate value is exactly that surrogate value's characteristic weight. The probability of the price with quantile j/T being offered is therefore the marginal characteristic weight, $w_i^j - w_i^{j+1}$.

Meanwhile, a standard fact from the study of position auctions is that any rank-based position auction can be represented as a convex combination of k-unit auctions. Since the allocation rule agents face in the stationary equilibrium of a surrogate-ranking mechanism is identical to that of a rank-based position auction, a similar analysis applies. We summarize the above discussion with the following lemma:

Lemma 17. *Any surrogate-binning (resp. surrogate-ranking) mechanism with surrogate values $\{\psi_i^j\}_{i=1,\ldots,n}^{j=1,\ldots,T}$ appears to each agent i as a distribution over price-posting (resp. top-k) mechanisms. The probability of offering the price with quantile $\frac{j}{T}$ (resp. of allocating j units) is given by $w_i^j - w_i^{j+1}$, where $w_i^0 = 1$, $w_i^{T+1} = 0$, and w_i^j is the characteristic weight for ψ_i^j for $j = 1, \ldots, T$.*

Notice that if $\psi_i^1 = \ldots = \psi_i^{\underline{k}}$ for some positive integer $\underline{k}$, the binning mechanism's allocation rule on the first $\underline{k}$ intervals of distribution i's quantile space will be constant. If $\psi_i^{\overline{k}} = \ldots = \psi_i^T$ for some positive integer $\overline{k}$, then the allocation rule on the last $T - \overline{k}$ intervals will be constant. In terms of distribution i's randomization over posted pricings, the highest nontrivial price offered has quantile $\underline{k}/T$, and the lowest has quantile $\overline{k}/T$. These extremal quantiles drive the approximation guarantees relating pricing to ranking. Consequently, we obtain the following theorem:

Theorem 18. *For surrogate values $\psi_i^1 \geq \psi_i^2 \geq \ldots \geq \psi_i^T$ with $\psi_i^1 = \psi_i^2 = \ldots = \psi_i^{\underline{k}}$ and $\psi_i^{\overline{k}} = \psi_i^{\overline{k}+1} = \ldots = \psi_i^T$ for each population i, the surrogate ranking mechanism attains a $(1 - O(\underline{k}^{-1/2}))$-fraction of the welfare of the binning mechanism. If distributions are regular, then the surrogate ranking mechansim attains a $(1-O(\min(\underline{k}, T-\overline{k})^{-1/2}))$-fraction of the binning mechanism's virtual surplus.*

4.2 Approximately Optimal Binning Mechanisms

Surrogate binning mechanisms are parametrized by their surrogate values. In this section, we show how to choose a set of surrogate values that yields an SBM which (1) attains near-optimal welfare or revenue, and (2) permits approximation by a surrogate ranking mechanism via Theorem 18. More precisely:

Theorem 19. *For every $k \in \{1, \ldots, T/2\}$, there exists a choice of surrogate values which yields a surrogate binning mechanism such that:*

- *The welfare of the corresponding surrogate binning mechanism is at least a $(1-O(1/k))(1-O(nk/T))$-fraction of the optimal welfare.*
- *For each agent, the highest k surrogate values are identical and the lowest k surrogate values are identical.*

If value distributions are regular, an identical result holds for the objective of revenue.

Note that combining Theorem 18 with Theorem 19 and choosing $k = (n/T)^{2/3}$ immediately implies the representation error bound of Theorem 10. We devote this section to deriving surrogate values which prove Theorem 19. The high-level approach will be resampling. To illustrate the idea, consider mapping the jth surrogate of agent i, which corresponds to an agent with quantile in $[\frac{j-1}{T}, \frac{j}{T}]$ to a redrawn value from their value distribution conditioned to this interval. Such resampling does not change the induced allocation rule for any other agents, and replaces the allocation rule for agent i on their jth quantile interval with its average.

This naïve resampling procedure does not directly lead to an approximately optimal mechanism. For an example where this fails, note that for the top quantile interval, the optimal mechanism's allocation probability at the very top of the interval may be much higher than its average across the interval, while the highest values on the interval may be much higher than the interval's average. For example, if the value and allocation rule are both 1 for an ϵ measure and zero otherwise, then the optimal welfare is ϵ and the welfare from resampling is ϵ^2. A second issue is that we wish to be able to apply Theorem 18 to get a good approximation bound, and therefore need the k highest and k lowest surrogate values to each be the same for each agent. Resampling naïvely for every bin does not achieve this objective.

To resolve both these issues, we will artificially inflate the values of agents with low quantiles, treating them as if they had the highest value in the support of their distributions, and deflate the values of high quantiles, treating these agents as if they had the lowest value in the support of their distributions. We formalize this approach in Section 4.2.1 and Section 4.2.2. Section 4.2.1 first treats the process of exaggerating extreme quantiles in isolation, analyzing a procedure that can be applied to any mechanism. In Section 4.2.2, we combine this with the resampling procedure discussed above to produce a surrogate binning mechanism satisfying the conditions of Theorem 19.

4.2.1 Extremal Buffering

In this section, we design and analyze a procedure for privileging agents with low quantiles and penalizing agents with high quantiles. We will show that our procedure does not significantly harm welfare or revenue, and in the next section, we will use this analysis as a foundation for proving Theorem 19. Our procedure will take as input an arbitrary truthful mechanism (e.g. a revenue-optimal mechanism), and modify its allocation rule. We will compare the virtual surplus of the new rule to the original mechanism. In the next section, we will combine this with a resampling procedure to obtain a surrogate binning mechanism.

Formally, consider an arbitrary truthful mechanism with allocation rule $\boldsymbol{x} = (x_1, \ldots, x_n)$. We will define a new allocation rule by remapping agents' quantiles to exaggerate extreme quantiles before applying $\boldsymbol{x}$. In particular, we consider the following transformation of $\boldsymbol{x}$:

Definition 20. *Given a monotone allocation rule $\boldsymbol{x}$ and a quantile $q \in [0,1]$, the q-*buffering *rule for $\boldsymbol{x}$ runs $\boldsymbol{x}$ on agents with quantiles transformed for each agent as follows:*

- *For any $q_i \in [0, q]$, return 0.*
- *For any $q_i \in [q, 1-q]$, return $(q_i - q)/(1-2q)$.*
- *For any $q_i \in [1-q, 1]$, return 1.*

The q-buffering rule for $\boldsymbol{x}$ treats agents at the top of the distribution as if they had quantile 0 and agents at the bottom as if they had quantile 1. Moreover, for any agent i, conditioned on $q_i \in [q, 1-q]$, the distribution of remapped quantiles for i is uniform. We show that the q-buffering procedure approximately preserves both welfare and revenue, as well as any other virtual surplus quantity that satisfies two basic properties satisfied by values or regular virtual values. We will only consider q-buffering as an analysis tool, so it will suffice to analyze the virtual surplus generated by the resulting allocation rules independent of incentives. Note that the q-buffering procedure preserves monotonicity, so truthful payments could be found if desired. Formally:

Lemma 21. *For each agent i, let $\phi_i : [0,1] \to \mathbb{R}$ be an arbitrary nonincreasing virtual value function satisfying $\int_0^1 \phi_i(q)\, dq \geq 0$. The q-buffering rule for $\boldsymbol{x}$ attains at least a $(1 - \frac{q}{1-q})(1 - q)(1 - 2(n-1)q)$-fraction of the virtual surplus of $\boldsymbol{x}$.*

The proof of Lemma 21 can be found in Appendix C. Note that choosing $q = k/T$ for some $k \in \{1, \ldots, T/2\}$ yields an approximation factor of $(1 - O(1/k))(1 - O(nk/T))$. The proof of the lemma follows from several observations about the similarity of $\boldsymbol{x}$ and the q-buffering rule for $\boldsymbol{x}$. One particular step in this analysis may be of independent interest. Recall the characterization of expected revenue in terms of revenue curves and marginal revenue from Lemma 2:

$$\mathbf{E}_q[p(q)] = \mathbf{E}_q[-x'(q)\, R(q)] + R(1)\, x(1) = \mathbf{E}_q[R'(q)\, x(q)] + R(0)\, x(0). \tag{1}$$

The first equality enables a geometric understanding of revenue. Given an fixed allocation rule x in quantile space, for two value functions v_1 and v_2 where $R_i(q) = q\, v_i(q)$ satisfies $R_1(q) \geq R_2(q)$, then the revenue from v_1 on x is at least the revenue of v_2 on x. This follows from the first equality of equation (1), where the expressions for revenue of both value functions are weighted integrals over $q \in [0,1]$ with non-negative weights $-x_i'(q)$. Approximation bounds hold as well; specifically, if R_2 approximates R_1 at all $q \in [0,1]$ then the same approximation holds for the revenue of any fixed allocation rule x. Note however that a similar result, with a fixed distribution and similar allocation rules is not implied by the second equation, as the weights $R'(q)$ are not generally all the same sign. Instead, the following lemma shows that two allocation rules with *inverses* that are approximately close have approximately the same revenue. We state the result for general virtual value functions $\phi(\cdot)$ and cumulative virtual value curves $\Phi(q) = \int_0^q \phi(r)\,dr$. The assumption in the lemma on the cumulative virtual value curve is that lines from the origin pass from below to above the curve. This assumption, for example, is satisfied by any revenue curve, and it does not require regularity.

Lemma 22. *For virtual value function $\phi(\cdot)$ and cumulative virtual value $\Phi(q) = \int_0^q \phi(r)\,dr$ satisfying $\Phi(\alpha\, q) \geq \alpha\, \Phi(q)$ for all quantiles q and $\alpha \in [0, 1]$, and any two allocation rules x_1 and x_2 that satisfy $x_1^{-1}(z) \geq x_2^{-1}(z) \geq \frac{1}{\alpha} x_1^{-1}(z)$, the expected virtual surpluses satisfy*

$$\mathbf{E}_q[\phi(q)\, x_2(q)] + \Phi(0)\, x_2(0) \geq \tfrac{1}{\alpha}[\mathbf{E}_q[\phi(q)\, x_1(q)] + \Phi(0)\, x_1(0)].$$

Proof. The virtual surplus of any allocation rule x can be rewritten as $\int_0^1 \phi(q)\, x(q)\, dq + \Phi(0)\, x(0) = \int_0^1 \Phi(x^{-1}(z))\, dz$. This follows by the first equality of equation (1) and a change of variables to integrate the vertical axis rather than the horizontal axis as follows:

$$\begin{aligned} \int_0^1 -x'(q)\, \Phi(q)\, dq + \Phi(1)\, x(1) &= \int_{x(1)}^{x(0)} \Phi(x^{-1}(z))\, dz + \int_0^{x(1)} \Phi(1)\, dz \\ &= \int_0^{x(0)} \Phi(x^{-1}(z))\, dz. \end{aligned} \tag{2}$$

Notice that the second line follows from the first line because $x^{-1}(z) = 1$ for $z \in [0, x(1)]$.

Now consider two arbitrary quantiles q_1 and q_2 satisfying $\frac{1}{\alpha} q_1 \leq q_2 \leq q_1$. By assumption, we have $\Phi(q_2) \geq q_2 \Phi(q_1)/q_1 \geq \frac{1}{\alpha}\Phi(q_1)$. The assumption on the approximation of the two allocation rules, namely $x^{-1}(z) \geq \hat{x}^{-1}(z) \geq \frac{1}{\alpha} x^{-1}(z)$ for all $z \in [0, 1]$, and the expected virtual surplus written as rewritten in equation (2) both both x_1 and x_2, then, suffice to prove the lemma. □

4.2.2 Approximately Optimal Reasmpling

In the previous section, we analyzed the effect of exaggerating extreme quantiles. We now introduce an additional resampling transformation, and show that the performance loss is not significant when combined with the buffering procedure. Furthermore, the mechanism produced will be a surrogate binning mechanism. This will imply Theorem 19.

As we did in the previous section, we will argue for an arbitrary monotone allocation rule $\boldsymbol{x}$. Let $k \in \{1, \ldots, T/2\}$ be given, and let $\hat{\boldsymbol{x}}$ denote the k/T-buffering rule for $\boldsymbol{x}$. Our resampling procedure is defined as follows:

Definition 23. *Given an allocation rule $\boldsymbol{x}$, the* resampling rule for $\boldsymbol{x}$ *allocates according to the following randomized procedure:*

- *For each agent i:*
  - *Compute j such that i's quantile $q_i \in [(j-1)/T, j/T]$.*
  - *Resample a quantile uniformly from $[(j-1)/T, j/T]$.*
- *Run $\boldsymbol{x}$ on the resampled quantiles.*

This algorithm formalizes the procedure discussed at the beginning of Section 4.2. Because of the problems discussed earlier, the resampling procedure alone does not necessarily preserve the performance of the original allocation rule. However, composing the resampling procedure with the buffering procedure of the previous section eliminates these pathologies, yielding:

Lemma 24. *For each agent i, let $\phi_i(q_i)$ be a nonincreasing virtual value function, and let $\mathbf{x}$ be an a monotone allocation algorithm. Further, for some $k \leq T/2$, let $\hat{\mathbf{x}}$ denote the k/T-buffering algorithm for $\mathbf{x}$. Then the resampling algorithm for $\hat{\mathbf{x}}$ obtains at least a $\frac{k}{k+1}(1 - \frac{q}{1-q})(1-q)(1 - 2(n-1)q)$ the expected virtual surplus under $\mathbf{x}$, with $q = k/T$.*

A proof can be found in Appendix D. The main observations are that resampling the buffered allocation rule does not change the allocation of agents with extreme quantiles, and that the nonincreasing nature of ϕ_i implies that the loss from resampling moderate quantiles is small.

Proof of Theorem 19. We argue for the objective of revenue. The argument for welfare is essentially identical. Let $\boldsymbol{x}$ be the allocation rule of the revenue-optimal mechanism, let $\hat{\boldsymbol{x}}$ be the k/T-buffering rule based on $\boldsymbol{x}$. We first argue that the resampling rule for $\hat{\boldsymbol{x}}$ can be implemented as a surrogate binning mechanism with random surrogate values. In particular, the jth surrogate value for agent i, ψ_i^j, is produced in in the following way:

- If $j \in \{1, \ldots, k\}$, set $\psi_i^j = R_i'(0)$
- If $j \in \{T-k+1, \ldots, T\}$, set $\psi_i^j = R_i'(1)$
- Otherwise:
  - Resample q_i' uniformly from $((j-1)/(T-2k), j/(T-2k)]$.
  - Set $\psi_i^j = R_i'(q_i')$.

These surrogate values satisfy the properties required for Theorem 19, but are randomized. Note that our performance guarantees are in expectation over both the random choice of surrogate values and the quantiles of agents. This implies the existence of a deterministic choice of surrogate values for which the performance guarantees hold in expectation over the quantiles of agents. The result follows. □

5 Reduction from Sample Complexity to Rank-Based Inference

We have shown that surrogate-ranking mechanisms possess a unique stationary equilibrium (Theorem 8), and that this equilibrium may be analyzed as if it was truthful (Theorem 9). In this section, we show how to use bid data to design a surrogate-ranking mechanism with near-optimal welfare or revenue in stationary equilibrium. Specifically, we reduce this design problem to an inference problem which is better-understood: estimating expected order statistics from bid data in i.i.d. position auctions. This inference problem was solved by Chawla et al. (2017) for first-price and all-pay mechanisms and is straighforward for truthful mechanisms.

Before giving details, we describe the high-level approach. Recall from Section 4 that, for revenue with regular distributions or welfare with general distributions, polynomially many surrogate values T per agent suffice to obtain a $(1-\epsilon)$-fraction of the optimal revenue or welfare via a surrogate-ranking mechanism. Consider a surrogate-ranking mechanism with T surrogate values per agent. We first show in Section 5.1 that the revenue- or welfare-optimal choice of these nT surrogate values requires only knowledge of the expected order statistics of the value or virtual value distribution. In Section 5.2 we observe that this design approach is robust to error from inference: if one uses imperfect estimates of expected order statistics to design a SRM, then estimation error will propagate cleanly to revenue or welfare loss. Composing these three observations above yields the desired reduction.

Section 5.3 concludes by instantiating the reduction with an estimator for the requisite order statistics. Specifically, Chawla et al. (2017) show how to estimate expected order statistics using bid data from all-pay and winner-pays-bid position auctions with bounded distributions, and we show in Appendix F how to estimate expected order statistics with truthful data from unbounded regular distributions. These results imply polynomial sample complexity for winner-pays-bid, all-pay, and truthful mechanism design.

5.1 Optimal Surrogate-Ranking Mechanisms

Surrogate-ranking mechanisms are parametrized by nT surrogate values. Each choice of surrogate values induces a different allocation rule in stationary equilibrium, but by Theorem 9, this equilibrium allocation rule is the same under any of the standard payment formats. Hence the optimal surrogate values, by revenue equivalence, do not depend on the payment format and we may as well identify the optimal surrogate values for the truthful payment format. In this section, we characterize the welfare- and revenue-optimal choices of surrogate values. To choose our surrogate values optimally, we consider a relaxed problem of maximizing a generic virtual surplus quantity subject to the constraint that the allocation rule depend only on each agent's rank among $T-1$ other truthful run-time samples. The solution to this problem is straightforward given the observation that if the only information we have to make decisions on is the rank of an agent against samples from her value distribution, then decisions should be made based on expected order statistics. The right choice of surrogate values is the expected order statistics of the quantity of interest for the objective, i.e., for welfare maximization, it's order statistics for the value distribution for revenue maximization it's order statistics for the distribution of marginal revenues (via the characterization of expected revenue in Lemma 2). Formally, given $T-1$ sampled quantiles for each agent, let r_i denote the rank of the quantile q_i of agent i among these samples. We have:

Theorem 25. *The welfare-optimal surrogate-ranking mechanism uses surrogate values given by* $\psi_i^j = \mathbf{E}_{q_i}[v_i(q_i) \mid r_i = j]$. *For regular distributions, the revenue-optimal surrogate-ranking mechanism uses surrogate values* $\psi_i^j = \mathbf{E}_{q_i}[R_i'(q_i) \mid r_i = j]$ *where* $R_i'(q_i) = v_i(q_i) + q_i\, v_i'(q_i)$ *is the marginal revenue of agent* i *at quantile* q_i.

A formal proof of this theorem is in Appendix E. Because the welfare- and revenue-optimal surrogate ranking mechanisms are at least as good as any other surrogate ranking mechanism, it follows that they inherit the welfare and revenue guarantees of any other such mechanism. In particular, we obtain the following corollary to Theorem 10:

Corollary 26. *The welfare- (resp. revenue-) optimal surrogate ranking mechanism obtains a* $(1 - O(\sqrt[3]{n/T}))$*-fraction of the optimal welfare (resp. a* $(1 - O(\sqrt[3]{n/T}))$*-fraction of the optimal revenue with regular distributions) in stationary equilibrium.*

5.2 Propagation of Error

The optimal surrogate ranking mechanism for welfare (resp. revenue) uses surrogate values equal to the expected order statistics of each agent's value (resp. virtual value) distribution. We now show that a designer can in fact use noisy estimates of these quantities, and the performance will degrade smoothly with the estimation error. As before, we present our result for an arbitrary monotonic virtual value function ϕ_i for each agent.

For monotonic ϕ_i, though the expected order statistics are monotone, estimates of these expected order statistics may not be. However, if the estimates of agent i's expected order statistics are within an agent-specific error ϵ_i of being correct, then any natural method or making these estimates monotone will be similarly close, e.g., using the jth estimate of $\max_{j' \leq j} \hat{\psi}_i^j$ instead of $\hat{\psi}_i^j$. Of course, the recommended method is the standard approach of ironing, which for quantities like order statistics is formally described in Devanur et al. (2015) and Chawla et al. (2017). In fact ironing is equivalent in this setting to isotonic regression. An advantage of ironing is that it is the correct approach when the original ϕ_i is non-monotonic and, omitting the details and consequences, our results for revenue can be extended to *tail regular* distributions using this approach, cf. Devanur et al. (2015). For the remainder of the discussion, without loss of generality, we assume that the error estimates are monotonic.

The following theorem shows that errors in estimated order statistics propagate in a well-behaved fashion in surrogate-ranking mechanisms.

Theorem 27. *For all i and j, let ψ_i^j be the expected jth order statistic of agent i's virtual value distribution, and let $\hat{\psi}_i^j$ be an estimate of ψ_i^j satisfying $|\hat{\psi}_i^j - \psi_i^j| < \epsilon_i$, where ϵ_i is an agent-specific error bound. The difference between the expected virtual surplus of the surrogate ranking mechanisms with the true expected order statistics and estimated order statistics is at most $2\sum_i \epsilon_i$.*

Proof. Let $\boldsymbol{x}$ and $\hat{\boldsymbol{x}}$ denote the allocation rule of the surrogate-ranking mechanism as a function of agents' ranks $\boldsymbol{r}$ among their run-time samples with optimal surrogate values $\boldsymbol{\Psi}$ and estimated and ironed surrogate values $\hat{\boldsymbol{\Psi}}$, respectively. The theorem follows from:

$$\begin{aligned}
\mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i \mathbf{E}_{q_i}\left[\psi_i(q_i) \mid r_i\right] \hat{x}_i(\boldsymbol{r})\right] &= \mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i \psi_i^{r_i} \hat{x}_i(\boldsymbol{r})\right] \\
&\geq \mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i (\hat{\psi}_i^{r_i} - \epsilon_i) \hat{x}_i(\boldsymbol{r})\right] \\
&\geq \mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i \hat{\psi}_i^{r_i}\, \hat{x}_i(\boldsymbol{r})\right] - \sum\nolimits_i \epsilon_i \\
&\geq \mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i \hat{\psi}_i^{r_i}\, x_i(\boldsymbol{r})\right] - \sum\nolimits_i \epsilon_i \\
&\geq \mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i (\psi_i^{r_i} - \epsilon_i)\, x_i(\boldsymbol{r})\right] - \sum\nolimits_i \epsilon_i \\
&\geq \mathbf{E}_{\boldsymbol{r}}\left[\sum\nolimits_i \psi_i^{r_i}\, x_i(\boldsymbol{r})\right] - 2\sum\nolimits_i \epsilon_i.
\end{aligned}$$

The second and fifth lines follow from the assumption that $|\hat{\psi}_i^j - \psi_i^j| < \epsilon_i$, and the fourth line follows from the fact that $\hat{\boldsymbol{x}}$ is the allocation rule that maximizes $\mathbf{E}_{\boldsymbol{r}}\left[\sum_i \hat{\psi}_i^{r_i}\, \hat{x}_i(\boldsymbol{r})\right]$. The last line is the expected virtual surplus of the optimal surrogate-ranking mechanism $\boldsymbol{x}$, which implies the result. □

5.3 Sample Complexity of Non-Truthful Mechanisms

We can now formalize the reduction from non-truthful sample complexity to inference in rank-based position auctions. In Section 3, we observed that the data generated by an agent in a SRM is distributed according to the unique BNE of an i.i.d. position auction. In Section 4, we demonstrated the existence of SRMs with near-optimal welfare and revenue, and in Section 5.1 and Section 5.2, we showed that it was possible to construct such a mechanism from noisy estimates of expected order statistics. We have therefore reduced the problem designing a near-optimal non-truthful mechanism from data to the problem of estimating expected order statistics in an i.i.d. position auction. We now instantiate the reduction in two settings: bounded distributions with the winner-pays-bid or all-pay format, where we seek additive error bounds, and unbounded distributions with truthful payments, where we seek multiplicative error bounds.

We first consider bounded distributions with the winner-pays-bid or all-pay format. Existing literature shows how to one use bid data from SRMs to infer these parameters efficiently. Chawla et al. (2017) study the problem of inferring order statistics from bid data in all-pay and winner-pays-bid i.i.d. position auctions. They show that for any non-trivial position weights, it is possible to efficiently infer order statistics for both the values and marginal revenues. We summarize their results below:

Theorem 28 (Chawla et al., 2017)**.** *Consider a T-agent all-pay or winner-pays-bid i.i.d. position auction with arbitrary position weights and values in $[0,1]$. There exists an estimator $\hat{V}_k$ for the expected kth order statistic of the value distribution V_k such that with $N \geq O(T^4(\log^2(1/\delta+T)+\log^2(1/\epsilon+T))\epsilon^{-2}\delta^{-2})$ sampled bids from the unique BNE, $|\hat{V}_k - V_k| \leq \epsilon$ with probability at least $1-\delta$.*

Theorem 29 (Chawla et al., 2017)**.** *Consider a T-agent all-pay or winner-pays-bid i.i.d. position auction with arbitrary position weights and values in $[0,1]$. There exists an estimator $\hat{\Psi}_k$ for the expected kth order statistic of the virtual value distribution Ψ_k such that with $N \geq O(T^4(\log^2(1/\delta)+\log^2(1/\epsilon))\epsilon^{-2}\delta^{-2})$ sampled bids from the unique BNE, $|\hat{\Psi}_k - \Psi_k| \leq \epsilon$ with probability at least $1-\delta$.*

We note that the above results combine with the incentive analysis of Section 3, approximation analysis of Section 4, and the analysis of error propagation in Section 5.1 and Section 5.2 to imply a solution to the non-truthful sample complexity problem for agents with values in $[0,1]$ and additive loss. We summarize below:

Theorem 30. *For agents with bounded values in $[0,1]$, there are families of winner-pays-bid and all-pay mechanisms that satisfy C1, C2, and C3 with $p_{run}(n,\epsilon^{-1}) = O(n^4\epsilon^{-3})$ and $p_{design}(n,\epsilon^{-1},\delta^{-1}) = \tilde{O}(n^{28}\epsilon^{20}\delta^2)$ for additive loss and the welfare objective. If the agents' distributions are regular and bounded, then the same result also holds for the revenue objective.*

Proof. Because values are bounded, a multiplicative error of ϵ corresponds to an additive error of at most $n\epsilon$. Hence, choosing $T = n^4\epsilon^{-3}$ guarantees additive representation error ϵ, by Theorem 10. Next, note that to obtain an additive estimation error of ϵ, Theorem 27 implies that estimating each surrogate value to additive error at most ϵ/n suffices. Finally, to obtain these guarantees with probability at least $1-\delta$, the union bound implies that a failure probability of $n^{-1}T^{-1}\delta = n^{-5}\epsilon^3\delta$ per surrogate value suffices. Theorems 28 and 29 therefore imply that $\tilde{O}(n^4\epsilon^{-3})^4(\epsilon/n)^{-2}(n^{-5}\epsilon^3\delta)^{-2}) = \tilde{O}(n^{28}\epsilon^{20}\delta^2)$ design-time samples suffice. □

We conclude by discussing our results' implications for the truthful sample complexity literature. Our results yield polynomial sample complexity for revenue maximization with unbounded regular distributions and general feasibility settings. This extends the result of Devanur et al. (2016) by dropping the downward-closure requirement on the feasibility constraint. Details can be found in Appendix F.

Theorem 31. *For agents with regularly distributed (but potentially unbounded) values, there is a family of truthful mechanisms that satisfies C1, C2, and C3 with* $p_{run}(n, \epsilon^{-1}) = O(n\epsilon^{-3})$ *and* $p_{design}(n, \epsilon^{-1}, \delta^{-1}) = \tilde{O}(n^{10}\epsilon^{-22})$ *for multiplicative approximation and the revenue objective.*

We conclude by noting that the polynomials in the sample complexity guarantees of Theorem 30 and Theorem 31 are clearly impractical. We leave optimizing the polynomials and deriving tight tradeoffs between run-time and design-time samples to future work.

References

Ausubel, L. M. and Milgrom, P. (2006). The lovely but lonely Vickrey auction. *Combinatorial auctions*, pages 17–40.

Bei, X. and Huang, Z. (2011). Bayesian incentive compatibility via fractional assignments. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 720–733.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–1090.

Chawla, S., Hartline, J., and Nekipelov, D. (2014). Mechanism design for data science. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 711–712. ACM.

Chawla, S. and Hartline, J. D. (2013). Auctions with unique equilibria. In *ACM Conference on Electronic Commerce*, pages 181–196.

Chawla, S., Hartline, J. D., and Nekipelov, D. (2017). Mechanism redesign. *arXiv preprint arXiv:1708.04699*.

Chawla, S., Hartline, J. D., and Sivan, B. (2015). Optimal crowdsourcing contests. *Games and Economic Behavior*.

Cole, R. and Roughgarden, T. (2014). The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 243–252. ACM.

Dasgupta, A. and Ghosh, A. (2013). Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330. ACM.

Devanur, N. R., Hartline, J. D., and Yan, Q. (2015). Envy freedom and prior-free mechanism design. *Journal of Economic Theory*, 156:103–143.

Devanur, N. R., Huang, Z., and Psomas, C. (2016). The sample complexity of auctions with side information. In *STOC 2016*, pages 426–439.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.

Dughmi, S., Hartline, J. D., Kleinberg, R., and Niazadeh, R. (2017). Bernoulli factories and black-box reductions in mechanism design. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 158–169. ACM.

Dütting, P. and Kesselheim, T. (2015). Algorithms against anarchy: Understanding non-truthful mechanisms. In *16th ACM Conference on Economics and Computation*.

Elkind, E. (2007). Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745. Society for Industrial and Applied Mathematics.

Feng, Y. and Hartline, J. D. (2018). An end-to-end argument in mechanism design (prior-independent auctions for budgeted agents). In *IEEE 59th Annual Symposium on Foundations of Computer Science*.

Gonczarowski, Y. A. and Nisan, N. (2017). 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. ACM.

Guo, C., Huang, Z., and Zhang, X. (2019). Settling the sample complexity of single-parameter revenue maximization. *arXiv preprint arXiv:1904.04962*.

Hartline, J., Hoy, D., and Taggart, S. (2014). Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 693–710. ACM.

Hartline, J. D. (2013). Bayesian mechanism design. *Foundations and Trends® in Theoretical Computer Science*, 8(3):143–263.

Hartline, J. D., Kleinberg, R., and Malekian, A. (2011). Bayesian incentive compatibility via matchings. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 734–747.

Hartline, J. D., Kleinberg, R., and Malekian, A. (2015). Bayesian incentive compatibility via matchings. *Games and Economic Behavior*, 92:401–429.

Hartline, J. D. and Lucier, B. (2015). Non-optimal mechanism design. *American Economic Review*, 105(10):3102–24.

Jansen, B. J. and Mullen, T. (2008). Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, 6(2):114–131.

Morgenstern, J. and Roughgarden, T. (2015). On the pseudo-dimension of nearly optimal auctions. In *Annual Conference on Neural Information Processing Systems 2015*, pages 136–144.

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

Osband, K. (1989). Optimal forecasting incentives. *Journal of Political Economy*, 97(5):1091–1112.

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

Roughgarden, T., Syrgkanis, V., and Tardos, E. (2017). The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101.

Syrgkanis, V. and Tardos, E. (2013). Composable and efficient mechanisms. In *ACM Symposium on Theory of Computing*, pages 211–220.

Yan, Q. (2011). Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719.

A Undoing the Revelation Principle

Good first-price and all-pay mechanisms for a given environment can be found by undoing the revelation principle (ignoring computational complexity). This construction applies to any revelation mechanism $\mathcal{M}$. For concreteness, imagine a applying this approach to a single-minded combinatorial auction problem where $\mathcal{M}$ is the Vickrey-Clarke-Groves (VCG) mechanism. We give the all-pay version of the construction which is slightly simpler, but exhibits the same issues.

Definition 32. The all-pay unrevelation mechanism *for a revelation mechanism $\mathcal{M}$ is:*

1. *For each agent i and value v_i, calculate $s_i(v_i)$ as the expected payment in $\mathcal{M}$ when the agent's value is v_i and other agents' values are drawn from the distribution.*

2. *For each agent i, given bid b_i in the un-revelation mechanism, calculate the agent's value as $v_i = s_i^{-1}(b_i)$.*

3. *Serve the agents who are served by $\mathcal{M}$ on values $\mathbf{v} = (v_1, \ldots, v_n)$; all agents pay their bids.*

The characterization of Bayes-Nash equilibrium (Theorem 1) implies that s_i is the strategy that agents will employ in equilibrium of the constructed all-pay mechanism. Thus, the all-pay mechanism has the same equilibrium outcome.

From this definition we can see why symmetric and ordinal environments (i.e., IID position environments) are special. For these environments all agents will have the same strategy function, this strategy function will order higher valued bidders higher (by monotonicity), and the ordinal environment then implies that all that is needed to select an outcome is the order of values not their cardinal values. Thus, the mechanism simplifies to simply ordering the bids and the strategy function does not need to be calculated.

Even absent computational issues in estimating the strategy functions so as to implement this mechanism, it is clear that very detailed distributional information is needed to run the unrevelation mechanism. Moreover, the resulting outcomes may be very sensitive to small errors with the inversion of the strategy function. This unrevelation mechanism is not to be considered practical.

B Proof of Lemma 15

Lemma 15. *The top k-of-T mechanism attains at least a $(1 - O(k^{-1/2}))$-fraction of the welfare of the k/T-price posting algorithm with T i.i.d. agents. If values are regularly distributed, then it attains at least a $(1 - O(\min(k, T - k)^{-1/2}))$-fraction of the revenue of the k/T-price posting mechanism.*

Proof. We argue separately for the objectives of welfare and virtual surplus, but in both cases, the proof strategy will be the same. We will first explicitly characterize the worst-case distribution for each objective, and then we will analyze the performance ratio of the top-k-of-T mechanism and k/T-price posting mechanism using the correlation gap approach of Yan (2011) or a similar analysis. In what follows, we will suppress subscripts denoting a particular agent when the agent's identity is irrelevant.

Key to the analysis will be two formulae for the expected surplus of an mechanism, in terms of its interim allocation rule $x(\cdot)$ and the distribution's value function $v(\cdot)$. We have that a mechanism's surplus is:

$$\mathbb{E}_{q \sim U[0,1]}[x(q)v(q)] = \mathbb{E}_{q \sim U[0,1]}[-x'(q)V(q)], \tag{3}$$

where $V(q) = \int_0^q v(z)\,dz$, and the equality follows from integration by parts. This is the welfare analog of Lemma 2. We will make use of the latter result for revenue analysis. The only real difference between the two objectives is the fact that values are always positive, whereas virtual values may be negative. This will change the nature of the approximation, as allocating the wrong agent becomes actively harmful to the performance of the algorithm in the case of revenue.

Welfare. We begin by normalizing the per-agent surplus of the price-posting mechanism to 1. Note that for the k/T-price posting algorithm, the allocation rule is 1 until quantile k/T, and then drops to 0. It follows from equation (3) that our normalization is equivalent to the assumption that $V(k/T) = 1$. Next, we note that because $v(\cdot)$ is positive and decreasing, $V(\cdot)$ is increasing and concave, with $V(0) = 0$. Let $x(\cdot)$ be the allocation rule of the top-k-of-T algorithm. Given our normalization, the problem of finding the worst-case distribution then becomes:

$$\begin{aligned} \min_{V(\cdot)} \quad & \mathbb{E}_{q \sim U[0,1]}[-x'(q)V(q)] \\ \text{subject to} \quad & V(0) = 0 \\ & V(k/T) = 1 \\ & V(\cdot) \text{ concave} \\ & V(\cdot) \text{ increasing} \end{aligned}$$

This program can be solved by inspection by noticing that there is pointwise minimal function satisfying the constraints of the program: namely, the optimal $V(q)$ is linear with slope $v(q) = T/k$ for $q \leq k/T$, and constant at 1 for $q \geq k/T$. This corresponds to the distribution with k/T mass on the value T/k, and the rest on 0. The result then immediately follows from the correlation gap for k-uniform matroids (Yan, 2011).

Virtual Surplus. We now adapt the above proof to virtual surplus. We will again characterize the worst-case distribution, but this time, we cannot simply apply the correlation gap as we did for welfare. As before, we normalize the per-agent virtual surplus from price-posting to 1. This corresponds with setting $R(q) = 1$. Subject to normalization, we use properties of revenue curves to derive the worst-case distribution for virtual surplus. We assume values are regularly distributed, which implies that $R(q)$ is concave. Moreover, since $R(q) = v(q)q$, we have that $R(0) = R(1) = 0$. These properties yield the following program for the worst-case distribution:

$$\begin{aligned} \min_{R(\cdot)} \quad & \mathbb{E}_{q\sim U[0,1]}[-x'(q)R(q)] \\ \text{subject to} \quad & R(0) = R(1) = 0 \\ & R(k/T) = 1 \\ & R(\cdot) \text{ concave} \end{aligned}$$

Again, this may be solved by inspection. The worst-case $R(\cdot)$ is triangular, with its apex at $(k/n, 1)$. That is, on $[0, k/T]$, $R(q)$ has slope T/k, and on $[k/T, 1]$, it has slope $-T/(T-k)$. In other words, the worst-case distribution for virtual surplus has virtual value T/k with probability k/T, and virtual value $-T/(T-k)$ otherwise.

Let X^+ (resp. X^-) denote the number of agents with positive (resp. negative) virtual value. Further let Y^+ (resp. Y^-) denote the number of agents with positive (resp. negative) virtual value allocated by the top-k-of-T mechanism, and $Z^+ = X^+ - Y^+$ (resp. $Z^- = X^- - Y^-$) denote the number of agents with positive (resp. negative) virtual value who go unallocated. Since the expected virtual value of any individual agent is 0, we may write:

$$\mathbb{E}\left[\tfrac{T}{k}X^+ - \tfrac{T}{T-k}X^-\right] = 0.$$

This yields two equivalent expressions for the virtual surplus of the top-k-of-T mechanism:

$$\begin{aligned} \mathbb{E}\left[\text{Rev}(k,T)\right] &= \tfrac{T}{k}\mathbb{E}\left[Y^+\right] - \tfrac{T}{T-k}\mathbb{E}\left[Y^-\right] && (4) \\ &= \tfrac{T}{T-k}\mathbb{E}\left[Z^-\right] - \tfrac{T}{k}\mathbb{E}\left[Z^+\right]. && (5) \end{aligned}$$

We can break our analysis into two cases, based on whether $k \leq T-k$ or $k > T-k$. In the former case, we analyze eq. (4), and in the latter, we analyze eq. (5). In either case, we can lower bound the first term using the correlation gap as we did for welfare. For the second term, we bound the loss by a new analysis. We argue the $k \leq T-k$ case below. The other case follows by a symmetric argument.

Assume $k \leq T-k$. We must first lower bound $\frac{T}{k}\mathbb{E}\left[Y^+\right]$. Noting that the setup is identical to the welfare analysis implies that we may again apply the correlation gap result of Yan (2011) to get a lower bound of $(1 - O(k^{-1/2}))T$. We next upper bound $\frac{T}{T-k}\mathbb{E}\left[Y^-\right]$. The expected value of Y^- can be written as:

$$\frac{T}{T-k}\sum_{j=0}^{k-1}\binom{T}{j}\left(\frac{k}{T}\right)^j\left(\frac{T-k}{T}\right)^{T-j}(k-j). \qquad (6)$$

For any fixed k and j, the summand in eq. (6) is an increasing function of T. We may therefore upper bound it by its limit as $T \to \infty$ to obtain:

$$\begin{aligned}
\frac{T}{T-k}\mathbb{E}\left[Y^-\right] &\le \frac{T}{T-k}\sum_{j=0}^{k-1}\frac{k^j(k-j)}{e^k j!} \\
&= \frac{T}{T-k}\left[\sum_{j=0}^{k-1}\frac{k^{j+1}}{e^k j!} - \sum_{j=1}^{k-1}\frac{k^j}{e^k(j-1)!}\right] \\
&= \frac{T}{T-k}\cdot\frac{k^k}{e^k(k-1)!} \\
&\le T\Big(\frac{k^k}{e^k k!}\Big) = T\cdot O(k^{-1/2})
\end{aligned}$$

The last inequality follows from the assumption that $k \le T-k$, and the last equality from Stirling's approximation. □

C Proof of Lemma 21

Lemma 21. *For each agent i, let $\phi_i : [0,1] \to \mathbb{R}$ be an arbitrary nonincreasing virtual value function satisfying $\int_0^1 \phi_i(q)\,dq \ge 0$. The q-buffering rule for $\boldsymbol{x}$ attains at least a $(1-\frac{q}{1-q})(1-q)(1-2(n-1)q)$-fraction of the virtual surplus of $\boldsymbol{x}$.*

The proof of Lemma 21 will proceed in two main steps. First, we will show that applying the quantile remapping procedure in Definition 20 to a single agent i (leaving other agents' quantiles untouched) cannot reduce the expected virtual surplus from that agent by too much. This will follow from Lemma 22, which relates the virtual surpluses of allocation rules with inverses that are multiplicatively close. Second, we will show that subsequently applying the quantile remapping procedure to the populations other than i also does not significantly reduce the expected virtual surplus from i. This will follow from the fact that the distribution of quantiles input to the base allocation algorithm is identical, conditioned on no agents having extreme quantiles.

We begin with the single-agent analysis. Note that for a single agent, the q-buffering procedure can be thought of as two composed steps. First is a *top promotion* procedure, which remaps sufficiently low quantiles to 0 while remapping the remaining quantiles to induce a uniform distribution over $[0,1]$. Top promotion is then composed with *bottom demotion*, which performs analogous transformation, mapping high quantiles to 1 and mapping the rest of the interval to $[0,1]$. We formalize this as follows:

Definition 33. *Given a monotone single-agent allocation rule x and quantile $\underline{q}$, the* top promotion algorithm *for x and $\underline{q}$ runs x on the agent with quantiles transformed as follows:*

- *For any $q \in [0,\underline{q}]$, return 0.*
- *For any $q \in [\underline{q},1]$, return $(q-\underline{q})/(1-\underline{q})$.*

Definition 34. *Given a monotone single-agent allocation rule x and quantile $\overline{q}$, the* bottom demotion algorithm *for x and $\overline{q}$ runs x on the agent with quantiles transformed as follows:*

- *For any $q \in [0,\overline{q}]$, return $q/\overline{q}$.*

- *For any $q \in [\overline{q}, 1]$, return 1.*

The interim allocation rule faced by an agent i after applying the extremal buffering procedure to just i is the composition of the bottom demotion algorithm for quantile $1-q$ composed with the top promotion algorithm for original allocation rule x_i and quantile $q/1-q$. Consequently, we may analyze the loss from applying these two transformations separately and multiply the losses.

We first analyze bottom demotion. While bottom demotion does not produce an allocation rule which is multiplicatively close to the original rule, it does produce one which is close in the sense that its inverse is close to the inverse of the original rule. We may therefore apply Lemma 22 from Section 4, which we restate here for convenience. Recall, a virtual value function is $\phi(\cdot)$ and has cumulative virtual curve $\Phi(q) = \int_0^q \phi(r)\,dr$.

Lemma 22. *For virtual value function $\phi(\cdot)$ and cumulative virtual value $\Phi(q) = \int_0^q \phi(r)\,dr$ satisfying $\Phi(\alpha\, q) \geq \alpha\, \Phi(q)$ for all quantiles q and $\alpha \in [0, 1]$, and any two allocation rules x_1 and x_2 that satisfy $x_1^{-1}(z) \geq x_2^{-1}(z) \geq \frac{1}{\alpha} x_1^{-1}(z)$, the virtual surpluses satisfy*

$$\mathbf{E}_q[\phi(q)\, x_2(q)] + \Phi(0)\, x_2(0) \geq \tfrac{1}{\alpha}[\mathbf{E}_q[\phi(q)\, x_1(q)] + \Phi(0)\, x_1(0).$$

Lemma 35. *Let $\phi : [0, 1] \to \mathbb{R}$ be an arbitrary nonincreasing virtual value function. Given a monotone single-agent allocation rule x and quantile $\overline{q}$, the bottom demotion algorithm for x and $\overline{q}$ obtains at least a $\overline{q}$-fraction of the expected virtual surplus of x.*

Proof. The lemma will follow from a straightforward application of Lemma 22. For a quantile q receiving allocation $x(q)$ from the base algorithm, the quantile receiving this probability of allocation under the bottom demotion algorithm will be $\overline{q}q$. Hence, $x^{-1}(z) \geq \hat{x}^{-1}(z) = \overline{q}x^{-1}(z)$. Since ϕ is nonincreasing in q, we have that $R(q) = \int_0^q \phi(r)\,dr$ satisfies $R(\alpha q) \geq \alpha R(q)$ for all $\alpha \in [0, 1]$. Lemma 22 therefore implies the desired result. □

We have shown that bottom demotion results in an allocation rule which has an inverse close to that of the original rule on which it is based. To derive an approximation result for the top promotion procedure requires a more nuanced version of the same approach, based on two observations. First, the "unallocation rules", i.e., $y(q) = 1 - x(1-q)$ for allocation rule $x(q)$, satisfy the inverse-approximation condition of the lemma. Second, the virtual surplus of the unallocation rule is given by the expected virtual value plus the negative virtual surplus of the unallocation rule. Specifically $\mathbb{E}_q[\phi(q)\, x(q)] = \mathbb{E}_q[\phi(q)] + \mathbb{E}_q[(-\phi(1-q))\, y(q)]$. While virtual values for revenue always satisfy the property that rays from the origin cross the cumulative virtual value curve from below, this property does not generally hold for the negative virtual values $-\phi(1-q)$. Regularity, i.e., monotonicity of the original virtual value function, however, implies the property for negative virtual values. These observations are formally summarized in the subsequent lemma:

Lemma 36. *Let $\phi : [0, 1] \to \mathbb{R}$ be a nonincreasing virtual value function with $\int_0^1 \phi(q)\,dq \geq 0$. Given a monotone single-agent allocation rule x and quantile $\underline{q}$, the top promotion algorithm for x and $\underline{q}$ obtains at least a $(1-\underline{q})$-fraction of the expected virtual surplus of x.*

Proof. Note that the expected virtual surplus from any allocation rule x is can be written as $\int_0^1 \phi(q)x(q)\,dq = \int_0^1 \phi(q)\,dq - \int_0^1 \phi(q)(1-x(q))\,dq$. Specifically, let $\hat{x}$ be the allocation rule of the top promotion algorithm, and x the allocation rule of the original algorithm. Moreover, define

$\hat{y}(q) = 1 - \hat{x}(1-q)$ and $y(q) = 1 - x(1-q)$ to be the corresponding "unallocation rules." We will show that

$$\int_0^1 -\phi(1-q)\hat{y}(q)\,dq \geq (1-\underline{q})\int_0^1 -\phi(1-q)y(q)\,dq. \tag{7}$$

Since, $\int_0^1 \phi(q)\,dq \geq 0$, this will prove that $\int_0^1 \phi(q)x(q)\,dq \geq (1-\underline{q})\int_0^1 \phi(q)\hat{x}(q)\,dq$.

To prove (7), note that the definition of the top promotion algorithm can be manipulated to obtain $\hat{x}^{-1}(z) = x^{-1}(z)(1-\underline{q}) + \underline{q}$. Moreover, by the definition of y and $\hat{y}$, we have $y^{-1} = 1 - x^{-1}(1-z)$ and $\hat{y}^{-1} = 1 - \hat{x}^{-1}(1-z)$. Combining these three equations yields that $y^{-1}(z) \geq \hat{y}^{-1}(z) = (1-\underline{q})y^{-1}(z)$ for all $z \in [0,1]$. Moreover, note that $-\phi(1-q)$ is decreasing in q. This implies that $R(q)/q \geq -\phi(1-q)$, where $R(q) = \int_0^q -\phi(1-q)\,dq$. We may therefore apply Lemma 22, which yields (7). □

Combining Lemmas 35 and 36 yields the following:

Lemma 37. *Let $\phi : [0,1] \to \mathbb{R}$ be an arbitrary nonincreasing virtual value function satisfying $\int_0^1 \phi(q)\,dq \geq 0$, and consider an arbitrary agent i. The q-buffering algorithm for $\mathbf{x}$ and $\hat{Q}$, when applied only to agent i, attains at least a $(1 - \frac{q}{1-q})(1-q)$-fraction of the expected virtual surplus for i.*

Having derived a single-agent guarantee, we now show that applying the q-buffering procedure to all agents at once, rather than just to one agent, yields only a small additional loss. Intuitively, for each agent, the mechanism only appears different when another agent has an extreme quantile which is promoted or demoted by the buffering procedure. The probability of such an event can be controlled using the union bound. Formally, we have:

Proof of Lemma 21. Lemma 37 states that the virtual surplus lost from applying the q-buffering procedure to a single agent is small. We now argue that applying the algorithm to all agents at once does not incur much additional loss. We argue from the perspective of some agent i.

The key observation in our analysis is that the distribution of the quantiles of other agents input to $\boldsymbol{x}$ is nearly unchanged by the extremal buffering procedure. In particular, note that the probability that one or more agents other than i with quantiles set to 0 or 1 by the extremal buffering procedure is at most $(n-1)(1-2q)$, by the union bound. Conditioned on there being no such agents, the distribution of quantiles input to $\boldsymbol{x}$ remains uniform. It follows that the virtual surplus from agent i conditioned on this event is identical to the virtual surplus from the extremal buffering procedure applied only to i.

In the event that there are one or more agents from populations other than i who have top quantiles (which are promoted) or bottom quantiles (which are demoted), we note that the conditional virtual surplus from population i is nonnegative. To see this, let $\tilde{x}_i$ be the q-buffered interim allocation rule for agent i, conditioned on the event $\mathcal{E}$ that at least one agent j other than i has a quantile in $[0,q] \cup [1-q,1]$. Since $\boldsymbol{x}$ is a monotone function of its inputs, it must be that $\tilde{x}_i$ is nondecreasing. The expected virtual surplus from agent i conditioned on $\mathcal{E}$ is $\int_0^1 \phi_i(q)\tilde{x}_i(q)\,dq$. By assumption, $\int_0^1 \phi_i(q)\,dq \geq 0$, so it must also be the case that $\int_0^1 \phi_i(q)\tilde{x}_i(q)\,dq \geq 0$.

To conclude the proof, let $\hat{x}_i$ denote the interim allocation rule after extremal buffering, conditioned on the event $\overline{\mathcal{E}}$. The total virtual surplus from agent i is:

$$\Pr(\mathcal{E})\int_0^1 \phi_i(q)\tilde{x}_i(q)\,dq + \Pr(\overline{\mathcal{E}})\int_0^1 \phi_i(q)\hat{x}_i(q)\,dq$$

By the union bound, $\Pr(\overline{\mathcal{E}}) = 1 - \Pr(\mathcal{E}) \geq 1 - (n-1)(1-2q)$. By Lemma 37 and the fact that, conditioned on $\overline{\mathcal{E}}$, the distribution of quantiles i perceives from other agents is uniform implies that

$$\int_0^1 \phi_i(q)\hat{x}_i(q)\,dq \quad \geq (1 - \tfrac{q}{1-q})(1-q)\textstyle\int_0^1 \phi_i(q)x_i(q)\,dq$$
$$\geq (1 - \tfrac{q}{1-q})(1-q)\textstyle\int_0^1 \phi_i(q)x_i(q)\,dq.$$

Since we have shown that $\int_0^1 \phi_i(q)\hat{x}_i(q)\,dq \geq 0$, we can combine the above to obtain:

$$\Pr(\mathcal{E})\int_0^1 \phi_i(q)\tilde{x}_i(q)\,dq + \Pr(\overline{\mathcal{E}})\int_0^1 \phi_i(q)\hat{x}_i(q)\,dq$$
$$\geq \big(1 - \frac{q}{1-q}\big)(1-q)(1-2(n-1)q)\int_0^1 \phi_i(q)x_i(q)\,dq.$$

Summing over agents proves the lemma. □

Since we have reasoned about abstract virtual surplus, which could be value or Myerson virtual value, we obtain revenue and welfare approximation results for extremal buffering.

D Proof of Lemma 24

Lemma 24. *For each agent i, let $\phi_i(q_i)$ be a nonincreasing virtual value function, and let $\mathbf{x}$ be an a monotone allocation algorithm. Further, for some $k \leq T/2$, let $\hat{\mathbf{x}}$ denote the k/T-buffering algorithm for $\mathbf{x}$. Then the resampling algorithm for $\hat{\mathbf{x}}$ obtains at least a $\frac{k}{k+1}(1 - \frac{q}{1-q})(1-q)(1-2(n-1)q)$ the expected virtual surplus under $\mathbf{x}$, with $q = k/T$.*

To prove Lemma 24, we first rephrase a useful lemma from Roughgarden and Schrijvers (2016), which characterizes the relationship between the virtual surplus of resampling rules and their base allocation rules. Intuitively, it states that resampling can be treated as a distributional transformation - rather than replacing the allocation on each interval with its average, you could think about replacing the virtual values on each interval with their average. Formally:

Lemma 38 (Roughgarden and Schrijvers, 2016)**.** *For every agent i, let ϕ_i be a virtual value function for each agent, with cumulative virtual surplus $R_i(q) = \int_0^q \phi_i(r)\,dr$. For any allocation algorithm $\boldsymbol{x}$, let $\overline{\boldsymbol{x}}$ denote the resampling algorithm based on $\boldsymbol{x}$. Then for each agent i we have:*

$$\mathbb{E}[-\overline{x}_i'(q_i)R_i(q_i)] = \mathbb{E}[-\overline{x}_i'(q_i)\overline{R}_i(q_i)] = \mathbb{E}[-x_i'(q_i)\overline{R}_i(q_i)],$$

where $\overline{R}_i(q_i)$ is a piecewise linear approximation to R_i given by:

$$\overline{R}_i(q) = T(q - \tfrac{j-1}{T})R_i(\tfrac{j}{T}) + T(\tfrac{j}{T} - q)R_i(\tfrac{j-1}{T}) \text{ for } q \in [\tfrac{j-1}{T}, \tfrac{j}{T}].$$

Proof of Lemma 24. Let $\hat{\boldsymbol{x}}$ denote the k/T-buffering rule, and $\overline{\boldsymbol{x}}$ denote the resampling rule based on $\hat{\boldsymbol{x}}$. Define

$$\hat{R}_i(q_i) = \begin{cases} R_i(q_i) & \text{for } q_i \in [\frac{k}{T}, \frac{T-k}{T}] \\ \frac{Tq_i}{k}R_i(\frac{k}{T}) & \text{for } q_i \in [0, \frac{k}{T}] \\ R_i(\frac{T-k}{T}) + \frac{T}{k}(q - \frac{T-k}{T})(R_i(1) - R_i(\frac{T-k}{T})) & \text{for } q_i \in [\frac{T-k}{T}, 1]. \end{cases}$$

That is, $\hat{R}_i$ is equal to R_i except on $[0, \frac{k}{T}]$ and $[\frac{T-k}{T}, 1]$, where it is a linear interpolation between the values of the R_i at the endpoints of those intervals.

We will argue the following sequence of inequalities for each agent:

$$\begin{aligned}
\mathbb{E}[-\overline{x}_i'(q_i)R_i(q_i)] &= \mathbb{E}[-\hat{x}_i'(q_i)\overline{R}_i(q_i)] \\
&\geq \tfrac{k}{k+1}\mathbb{E}[-\hat{x}_i'(q_i)\hat{R}_i(q_i)] \\
&= \tfrac{k}{k+1}\mathbb{E}[-\hat{x}_i'(q_i)R_i(q_i)].
\end{aligned}$$

The first equality follows from Lemma 38. We will argue the second and third lines shortly. Since the first and last expressions in the above chain are the respective virtual surpluses of the resampling and buffering rules, respectively, the result will follow.

To see that $\overline{R}_i(q_i) \geq \frac{k}{k+1}\hat{R}_i(q_i)$, note that $\overline{R}_i(q_i) = \hat{R}_i(q_i)$ for all quantiles in $[0, \hat{q}_i^k] \cup [\hat{q}_i^{T-k}, 1]$. Otherwise, consider $q_i \in [\frac{j}{T}, \frac{j+1}{T}]$ for $j \in \{k, \ldots, T-k-1\}$. Assume without loss of generality that $R_i(\frac{j}{T}) \leq R_i(\frac{j+1}{T})$; a symmetric argument will apply to the case where $R_i(\frac{j}{T}) \geq R_i(\frac{j+1}{T})$. The concavity of R_i implies that for all $q \in [\frac{j}{T}, \frac{j+1}{T}]$, $R_i(q) \leq \frac{Tq}{j}R_i(\frac{j}{T})$. Moreover, note that $\overline{R}_i(q) \geq R_i(\frac{j}{T})$ for all $q \in [\frac{j}{T}, \frac{j+1}{T}]$. Since $\frac{Tq}{j} \leq \frac{k+1}{k}$, it follows that $\overline{R}_i(q_i) \geq \frac{k}{k+1}R_i(q_i) = \frac{k}{k+1}\hat{R}_i(q_i)$ for all $q_i \in [\frac{k}{T}, \frac{T-k}{T}]$, and $\overline{R}_i(q_i) = \hat{R}_i(q_i)$ elsewhere.

Finally, we argue that $\mathbb{E}[-\hat{x}_i'(q_i)\hat{R}_i(q_i)] = E[-\hat{x}_i'(q_i)R_i(q_i)]$. Note that since $\hat{R}_i(q_i) = R_i(q_i)$ for $q_i \in [\frac{k}{T}, \frac{T-k}{T}]$, it follows that

$$\int_{\frac{k}{T}}^{\frac{T-k}{T}} -\hat{x}_i(q_i)R_i(q_i)\, dq_i = \int_{\frac{k}{T}}^{\frac{T-k}{T}} -\hat{x}_i(q_i)\hat{R}_i(q_i)\, dq_i.$$

To prove the claim, notice that $\hat{x}_i'(q_i) = 0$ on $[0, \frac{k}{T}]$ and $[\frac{T-k}{T}, 1]$. Hence,

$$\int_0^{\frac{k}{T}} -\hat{x}_i'(q_i)R_i(q_i)\, dq_i = \int_0^{\frac{k}{T}} -\hat{x}_i'(q_i)\hat{R}_i(q_i)\, dq_i.$$

and

$$\int_{\frac{T-k}{T}}^1 -\hat{x}_i'(q_i)R_i(q_i)\, dq_i = \int_{\frac{T-k}{T}}^1 -\hat{x}_i'(q_i)\hat{R}_i(q_i)\, dq_i.$$

This proves the lemma. □

E Proof of Theorem 25

Theorem 25. *The welfare-optimal surrogate-ranking mechanism uses surrogate values given by $\psi_i^j = \mathbf{E}_{q_i}[v_i(q_i) \mid r_i = j]$. For regular distributions, the revenue-optimal surrogate-ranking mechanism uses surrogate values $\psi_i^j = \mathbf{E}_{q_i}[R_i'(q_i) \mid r_i = j]$ where $R_i'(q_i) = v_i(q_i) + q_i\, v_i'(q_i)$ is the marginal revenue of agent i at quantile q_i.*

Proof. We define the *rank-based allocation problem* as follows: the designer must choose an allocation rule $\bar{x}$ which takes as input the rank $\boldsymbol{r} = (r_1, \ldots, r_n)$ of each agent among $T-1$ runtime samples for their distributions and outputs a (possibly randomized) feasible allocation $\bar{\boldsymbol{x}}(\boldsymbol{r})$. As a constraint, $\bar{\boldsymbol{x}}$ must be monotone in the ranks of each agent. The objective is to maximize $\mathbf{E}[\sum_i \phi_i(q_i)\bar{x}_i(\boldsymbol{r})]$ for some given virtual value function $\phi_i(\cdot)$, where the expectation is over agents' quantiles being uniformly distributed and over the runtime samples used to compute $\boldsymbol{r}$. For example, $\phi_i(q_i) = v_i(q_i)$ corresponds to welfare maximization and $\phi_i(q_i) = R_i'(q_i)$ corresponds to revenue maximization.

The rank-based allocation problem can be solved by inspection. Fixing an allocation rule, the objective can be rewritten as $\sum_i \mathbf{E}[\phi_i(q_i) \mid r_i]\ \bar{x}_i(\boldsymbol{r})$ by linearity of expectation. From this expression, it becomes clear that the optimal solution chooses the allocation which maximizes the quantity $\sum_i \mathbf{E}[\phi_i(q_i) \mid r_i]\ \bar{x}_i(\boldsymbol{r})$. Note that if $\phi_i(\cdot)$ is monotone, then this allocation rule will be monotone as well, and therefore feasible.[6] Setting these expected order statistics as surrogate values, the surrogate-ranking mechanism (Definition 4) optimizes this quantity. □

F Sample Complexity in General Feasibility Environments

In this appendix, we show how to use polynomially many truthfully sampled values to estimate the revenue of the k-unit, T-buyer auction for all k from 1 to $T-1$ simultaneously. We assume the value distribution is regular, but may have unbounded support.

Theorem 39. *Given a regular value distribution, let $R^* = \max_q R(q)$ be the* monopoly revenue. *For any $\epsilon, \delta \in (0,1)$, $O(T^6\epsilon^{-4}\log(T/\delta))$ samples suffice to estimate the expected revenue of a k-unit, T-bidder auction up to additive error ϵR^* for all $k \in \{1, \dots, T\}$ simultaneously with probability at least $1-\delta$.*

We may combine Theorem 39 our bounds on the propagation of error and the representation error of surrogate ranking mechanisms (Theorem 27 and Theorem 10, respectively) to obtain the desired sample complexity result for truthful mechanisms, restated below.

Theorem 31. *For agents with regularly distributed (but potentially unbounded) values, there is a family of truthful mechanisms that satisfies C1, C2, and C3 with $p_{run}(n, \epsilon^{-1}, \delta^{-1}) = \tilde{O}(n\epsilon^{-3})$ and $p_{design}(n, \epsilon^{-1}, \delta^{-1}) = \tilde{O}(n^{10}\epsilon^{-22})$ for multiplicative approximation and the revenue objective.*

Proof. First, note that $T = n\epsilon^{-3}$ surrogate values per agent suffice to obtain multiplicative representation error at most ϵ, by Theorem 10. Now let R_i^* denote the monopoly revenue for agent i's distribution, and let OPT denote the optimal revenue. Estimating each surrogate value to an additive $\epsilon n^{-1} R_i$ suffices to obtain multiplicative estimation error of $O(\epsilon)$, as $\sum_i R_i^* \leq n\mathrm{OPT}$. Finally, failure probability of δ/n suffices for each agent to obtain these surrogate value estimates with probability at least $1-\delta$, by the union bound. Hence, $\tilde{O}((n\epsilon^{-3})^6(\epsilon/n)^{-4}) = \tilde{O}(n^{10}\epsilon^{-22})$ samples suffice to obtain multiplicative loss of ϵ with probability at least $1-\delta$. □

We now outline the high-level strategy for proving Theorem 39. First, for any $j \in \{0, \dots, T\}$, let P_j denote the expected revenue of a j-unit auction with T agents, and let ψ^k denote the expected kth order statistic of the virtual value distribution. Then we may write: $\psi^k = P_k - P_{k-1}$. It follows that to estimate ψ^k with additive error ϵR^*, it suffices to estimate P_k and P_{k-1} with error $\epsilon R^*/2$.

To estimate the k-unit revenue P_k, we will estimate the revenue contribution from a single agent, $\mathbb{E}_q[p(q)]$. Note that from an agent's perspective, playing in a k-unit auction is equivalent to facing a posted price distributed according to $v(q_{k:T-1})$, where $q_{k:T-1}$ denotes the kth lowest order statistic of $T-1$ $U[0,1]$ random variables. A basic property of the order statistics of

[6] If it is not monotone, then then the resulting surrogate values may not be monotone, if the surrogate values by this approach are not monotone, they can be ironed using the standard procedure.

uniform variables is that $q_{k:T-1}$ is distributed according to Beta$(k, T-k)$. Let $f_{k:T-1}$ denote the density of $q_{k:T-1}$, and let $R = v(q)q$ denote the price-posting revenue curve. We have:

Lemma 40. *For any $k \in \{1, \dots, T-1\}$, $P_k = T \int_0^1 f_{k:T-1}(q) R(q)\, dq$.*

In what follows, we will show how to estimate $R(q)$ for all $q \in [1/T^2, 1 - 1/T^2]$. We will further show that the loss from misestimating $R(q)$ on $[0, 1/T^2] \cup [1 - 1/T^2, 1]$ is minimal. This will immediately imply Theorem 39.

F.1 Estimation on a Grid

To create a skeleton for our estimated revenue curve, we first estimate $R(q)$ for $q \in \{1/K, \dots, 1 - 1/K\}$, for some large K to be determined later. The concavity of R will imply that the rest of the revenue can be estimated with low error via interpolation.

Lemma 41. *Let $R^* = \max_q R(q)$. For any $1 > \epsilon > 0$ and δ, $O(K^2 \epsilon^{-2} \log(K/\delta))$ samples suffice to guarantee that $K\hat{v}_j/j \in [R(j/K) - \epsilon R^*, R(j/K) + \epsilon R^*]$ for all $j \in \{1, \dots, K-1\}$ simultaneously with probability at least $1 - \delta$.*

Consider drawing $N = Km - 1$ samples, for some positive integer m. Note that the jmth smallest sample has mean j/K. Let $\hat{v}_j$ denote the value of this sample. We will use $\hat{v}_j$ as an estimator for $v(j/K)$, and $K\hat{v}_j/j$ as an estimator for $R(j/K) = Tv(j/K)/j$. The proof will proceed in two steps. First, we will use a Chernoff bound to show that with high probability, $q(\hat{v}_j)$ is close to j/K. We will then use the concavity of R to show that $\hat{v}_j$ is close to $v(j/K)$.

Lemma 42. *For any $1 > \epsilon > 0$ and δ, $Km = O(K\epsilon^{-2} \log(1/\delta))$ samples suffice to guarantee that $q_j \in [(1-\epsilon)j/K, (1+\epsilon)j/K]$ with probability at least $1 - \delta$.*

Proof. We now bound the probability of significantly misestimating $q(\hat{v}_j)$. Let $q_j = q(\hat{v}_j)$. Note that for any $\epsilon \in (0, 1)$, the number of samples with quantile at most $(1+\epsilon)j/K$ is the sum of N iid Bernoulli random variables with mean $(1+\epsilon)j/K$. Moreover, note that $q_j > (1+\epsilon)j/K$ only if at most $jm - 1$ samples overall have quantile at most $(1+\epsilon)j/K$. Chernoff then gives us that

$$\Pr[q_j \geq (1+\epsilon)j/K] \leq e^{-\left(1 - \frac{(jm-1)K}{(1+\epsilon)Nj}\right)^2 \frac{(1+\epsilon)Nj}{2K}} = e^{-\Omega(\epsilon^2 mj)}$$

Similarly, the the number of samples with quantile at most $(1-\epsilon)j/K$ is the sum of N iid Bernoullis with mean $(1-\epsilon)j/K$. We have $q_j < (1-\epsilon)j/K$ only if at least jm samples have quantile at most $(1-\epsilon)j/K$. Chernoff then gives us:

$$\Pr[q_j \leq (1-\epsilon)j/K] \leq e^{-\Omega(\epsilon^2 mj)}$$

It follows that as $m = O(\epsilon^{-2} j^{-1} \log(1/\delta))$ suffices for suffices for $q_j \in [(1-\epsilon)j/K, (1+\epsilon)j/K]$ with probability at least $1-\delta$. This bound is worst when $j = 1$. Hence $mK = O(K\epsilon^{-2} \log(1/\delta))$ samples suffice overall. □

We next show that if q_j is close to j/K, then $K\hat{v}_j/j$ will be a close estimate of $R(j/K)$.

Lemma 43. *Assume $q_j \in [(1-\epsilon)j/K, (1+\epsilon)j/K]$. Then $K\hat{v}_j/j \in [(1-\epsilon K)R(j/K), (1 + \epsilon K)R(j/K)]$.*

Proof. We will show that $\hat{v}_j \in [(1-\epsilon K)v(j/K), (1+\epsilon K)v(j/K)]$. The result follows from multiplying by K/j. We first argue the case where $q_j < j/K$. That is, $\hat{v}_j \geq v(j/K)$. Concavity of the revenue curve implies that $R(q_j) \leq \frac{1-q_j}{1-j/K}R(j/K)$. In the event that $q_j \geq (1-\epsilon)j/K$, we have:

$$\hat{v}_j q_j \leq \frac{1-q_j}{1-j/K}\frac{v(j/K)j}{K} \leq (1+\epsilon K)v(j/K).$$

Dividing the inequality by q_j and using the fact that $q_j \geq (1-\epsilon)j/K$ yields:

$$\hat{v}_j \leq \frac{(1-(1-\epsilon)j/K)}{(1-j/K)(1-\epsilon)}v(j/K).$$

A symmetric argument applies when $q_j > j/K$. By concavity, we have $R(q_j) \geq \frac{1-q_j}{1-j/K}R(j/K)$. In the event that $q_j \leq (1+\epsilon)j/K$, we have

$$v_j \geq \frac{(1-(1+\epsilon)j/K)}{(1-j/K)(1+\epsilon)}v(j/K) \geq (1-\epsilon K)v(j/K).$$

□

Lemma 41 follows from the above results by applying the union bound.

Proof of Lemma 41. Lemmas 42 and 43 imply that $O(K^2\epsilon^{-2}\log(K/\epsilon))$ samples suffice to guarantee that $K\hat{v}_j/j \in [(1-\epsilon)R(j/K), (1+\epsilon)R(j/K)]$ with probability at least $1-\delta/K$. Applying the union bound, it follows that $O(K^2\epsilon^{-2}\log(K/\epsilon))$ samples suffice to guarantee that this guarantee holds for all $j \in \{1, \dots, K-1\}$ simultaneously with probability at least $1-\delta$. Noting that $R(j/K) \leq R^*$ implies the lemma. □

F.2 Estimating The Interior Revenue Curve

In the previous section, we showed how to estimate $R(j/K)$ up to an additive ϵR^*. Pick some $q \in [(j-1)/K, j/K]$ with $j \in \{2, \dots, K-1\}$. We may linearly interpolate between our estimates of $R((j-1)/K)$ and $R(j/K)$ to estimate $R(q)$. More formally, let $\hat{R}_j$ denote the estimator for $R(j/K)$ analyzed in the previous section. We will estimate $R(q)$ as

$$\hat{R}(q) = (q - \tfrac{j-1}{K})K\hat{R}_j + (\tfrac{j}{K} - q)K\hat{R}_{j-1}$$

We will use concavity of R to bound the error from this estimate.

Lemma 44. *Assume that for all $j \in \{1, \dots, K-1\}$, $K\hat{v}_j/j \in [R(j/K) - \epsilon R^*, R(j/K) + \epsilon R^*]$ for all $j \in \{1, \dots, K-1\}$. Then for all $q \in [1/K, 1-1/K]$:*

$$\hat{R}(q) \in [R(q) - (\epsilon + K^{-1})R^*, R(q) + \epsilon R^*]$$

Proof. We bound the overestimation and underestimation error in turn. First, not that since R is concave, it must be that $R(q)$ lies above the line between $((j-1)/K, \hat{R}_{j-1} - \epsilon R^*)$ and $(j/K, \hat{R}_j - \epsilon R^*)$. It follows that $\hat{R}(q)$ can only overestimate by at most ϵR^*. Next, note that a lower bound on $\hat{R}(q)$ is

$$\hat{R}(q) \geq (q - \tfrac{j-1}{K})KR(\tfrac{j-1}{K}) + (\tfrac{j}{K} - q)KR(\tfrac{j}{K}) - \epsilon R^*$$

In other words, the worst-case underestimation is ϵR^*, plus the worst-case underestimation of the piecewise linear curve through the points $(j/K, R(j/K))$ for all $j \in \{0, K\}$. Using the facts that R is concave, $R(0) = 0$, and $R(1) = 0$, this underestimation can be shown to be $O(R^*/K)$. The lemma follows.

□

F.3 Proof of Theorem 39

We have shown how to estimate the revenue curve to low additive error for quantiles bounded away from 0 and 1. We now use this estimator to prove the main sample complexity result of this appendix: that $O(T^6\epsilon^{-4}\log(T/\epsilon))$ samples suffice to estimate estimate ψ^k for all $k \in \{1, \ldots, T-1\}$ to within additive error ϵR^* with probability at least $1-\delta$. To do so, we will consider estimating P_k for $k \in \{1, \ldots, T-1\}$ as $T\int_{\epsilon/T^2}^{1-\epsilon/T^2} f_{k:T-1}(q)\hat{R}(q)\,dq$.

First we show that ignoring the revenue contribution from the intervals $[0, \epsilon/T^2]$ and $[1-\epsilon/T^2, 1]$ cannot hurt our estimate by much. Let $F_{k:T-1}$ denote the CDF of the kth lowest order statistic of $T-1$ $U[0,1]$ random variables. Then using properties of the Beta distribution, we have that $F_{k:T-1}(\epsilon/T^2) \leq \epsilon/T$ and $1 - F_{k:T-1}(1-\epsilon/T^2) \leq \epsilon/T$ for all $k \in \{1, \ldots, T-1\}$. Since for $q \in [0, \epsilon/T^2] \cup [1-\epsilon/T^2, 1]$, $R(q) \leq R^*$, it follows from Lemma 40 that

$$P_k - T\int_{\epsilon/T^2}^{1-\epsilon/T^2} f_{k:T-1}(q)R(q)\,dq \leq \epsilon R^*$$

Now assume that for all $q \in [\epsilon/T^2, 1-\epsilon/T^2]$, $|\hat{R}(q) - R(q)| \leq \epsilon T^{-1}R^*$. If this is the case, then we have

$$\left|T\int_{\epsilon/T^2}^{1-\epsilon/T^2} f_{k:T-1}(q)R(q)\,d - T\int_{\epsilon/T^2}^{1-\epsilon/T^2} f_{k:T-1}(q)\hat{R}(q)\,dq\right| \leq \epsilon R^*.$$

Choosing $K = T^2/\epsilon$ in Lemma 44 therefore yields the result.