Sample Complexity for Non-Truthful Mechanisms*

JASON HARTLINE, Northwestern University

SAMUEL TAGGART, Oberlin College

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.

CCS Concepts: • **Applied computing** → **Economics**; • **Theory of computation** → **Algorithmic game theory**; **Computational pricing and auctions**; *Sample complexity and generalization bounds*.

Additional Key Words and Phrases: mechanism design; sample complexity; non-truthful mechanisms; first-price auctions; all-pay auctions; position auctions

ACM Reference Format:

Jason Hartline and Samuel Taggart. 2019. Sample Complexity for Non-Truthful Mechanisms. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19), June 24–28, 2019, Phoenix, AZ, USA*. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3328526.3329632>

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.

*This work was supported in part by NSF CCF-1101717 and CCF-1618502. Full paper: <https://arxiv.org/abs/1608.01875>

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

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329632>

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] and Gonczarowski and Nisan [2017].

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 $m_{\text{design}} = p_{\text{design}}(n, 1/\epsilon)$ 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 $m_{\text{run}} = p_{\text{run}}(n, 1/\epsilon)$ 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 $1 - \epsilon$, 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 and $1/\epsilon$ design-time samples are sufficient to identify a mechanism that, with polynomially many run-time samples, approximates the performance of the optimal mechanism to within precision ϵ in the following environments:

¹Multiplicative versions of C3 are also interesting. In the multiplicative version a $(1 + \epsilon)$ approximation is required.

- (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 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 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; Gonczarowski and Nisan 2017; Guo et al. 2019; Morgenstern and Roughgarden 2015; Roughgarden and Schrijvers 2016]. 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

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

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

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 $\mathbf{v} = (v_1, \dots, v_n)$ and $\mathbf{q} = (q_1, \dots, q_n)$, respectively.

An *allocation* is $\mathbf{x} = (x_1, \dots, x_n)$ where $x_i \in \{0, 1\}$ is an indicator for agent i being served. The space of feasible allocations 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 $\mathbf{b} = (b_1, \dots, b_n)$ and outputs a feasible allocation $\mathbf{x} \in \mathcal{X}$ and agent payments $\mathbf{p}$. A mechanism consists of an *allocation algorithm* $\tilde{\mathbf{x}}(\mathbf{b})$, which maps bid profiles to a feasible allocation, and a *payment rule* $\tilde{\mathbf{p}}(\mathbf{b})$, which maps bid profiles to a non-negative payment for each agent. A standard allocation algorithm is *highest-bids-win* which is defined by $\tilde{\mathbf{x}}(\mathbf{b}) \in \arg\max_{\mathbf{x} \in \mathcal{X}} \sum_i b_i x_i$. We consider payment rules defined directly from the allocation algorithm according to standard payment formats. The *winner-pays-bid* format has payment rule $\tilde{p}_i(\mathbf{b}) = b_i \tilde{x}_i(\mathbf{b})$, and the *all-pay* format has payment rule $\tilde{p}_i(\mathbf{b}) = b_i$. 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 $(\tilde{x}, \tilde{p})$, the agents' strategies s , and the distribution over quantiles induce interim allocation and payment rules. Agent i 's *interim allocation rule* is $x_i(q_i) = \mathbf{E}_{q_{-i}}[\tilde{x}_i(s(q))]$ and *interim payment rule* $p_i(q_i) = \mathbf{E}_{q_{-i}}[\tilde{p}_i(s(q))]$. 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 $x_i(q_i)$ 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 derivatives 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 $w_1 \geq \dots \geq w_n \in [0, 1]$ 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 Ψ , with $\Psi_i = \{\psi_i^1 \geq \dots \geq \psi_i^T\}$ 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 X of feasible allocations, the algorithm allocates to maximize the surrogate surplus $\arg\max_{x \in X} \sum_i \psi_i x_i$.
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 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.

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,

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

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 $\Psi_i = \{\psi_i^1 \geq \dots \geq \psi_i^T\}$ be agent i 's set of surrogate values, and let $\hat{x}$ denote the surrogate surplus maximizing allocation rule $\hat{x}(\psi) = \max_{x \in X} \sum_i \psi_i x_i$. The characteristic weights $W_i = \{w_i^1 \geq \dots \geq w_i^T\}$ 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., $w_i^j = E[\hat{x}_i(\psi_i^j, \psi_{-i})]$ for each surrogate j and uniform random ψ_{-i} from Ψ_{-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 v , surrogate values Ψ , and characteristic weights 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 j th highest bid among the run-time samples will cause i to be assigned the j th 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 $w_i^1, \dots, w_i^T$ 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 Ψ , and the uniqueness of symmetric Bayes-Nash equilibrium in i.i.d. position auctions, which follows from revenue equivalence. $\square$

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. $\square$

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 APPROXIMATION ANALYSIS

In this section, we state and briefly discuss our approximation analysis for surrogate ranking mechanisms. Condition C3 requires that there exist some choice of surrogate values which induces near-optimal expected performance. In the full version of the paper, we show how to construct a choice of surrogate values which proves 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.*

Theorem 10 implies that the representation error of SRMs is low. 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.

The proof of Theorem 10 is technically involved, and we defer details to the full version of the paper. The high level approach involves two main steps. We first argue that there exists an approximately-optimal mechanism which coarsens quantile space into bins, and treats agents based solely on their bin. We then show that this bin-based mechanism can be transformed into a surrogate-ranking mechanism with only polynomial loss in performance.

A main challenge in proving approximation bounds in sample complexity is arguing that agents at the extremes, i.e., quantiles close to 0 and 1, are treated appropriately. For revenue, allocation to quantiles close to 0 could have high positive contribution to the revenue while allocation to quantiles close to 1 could have high negative contribution to the revenue. Our approach is to treat all agents that have quantile close to 0 as if they have quantile 0 and all who have quantile close to 1 as if they have quantile 1. The key step in the analysis is then to show that the loss from this

change, which disrupts the allocation to other agents, is bounded. The subsequent analysis is then relatively straightforward.

A key lemma in the analysis bounds virtual surplus from allocation rules that are approximately similar and may be of independent interest. Recall the characterization of expected revenue in terms of revenue curves and marginal revenue from Lemma 2:

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

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 curves $\Phi(q) = \int_0^q \phi(r) dr$. The assumption in the lemma on the cumulative virtual curve is that lines from the origin pass from below to above the curve. This assumption, for example, is satisfied by any revenue curve, specifically, it does not require regularity.

Lemma 11. *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*

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

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:

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

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. $\square$

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 straightforward 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 C 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 12. *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 B. 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 13. *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 of making these estimates monotone will be similarly close, e.g., using the j th 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 14. *For all i and j , let ψ_i^j be the expected j th 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 $\mathbf{x}$ and $\hat{\mathbf{x}}$ denote the allocation rule of the surrogate-ranking mechanism as a function of agents' ranks $\mathbf{r}$ among their run-time samples with optimal surrogate values Ψ and estimated and ironed surrogate values $\hat{\Psi}$, respectively. The theorem follows from:

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

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 $\hat{\mathbf{x}}$ is the allocation rule that maximizes $\mathbf{E}_{\mathbf{r}} [\sum_i \hat{\psi}_i^{r_i} \hat{\mathbf{x}}_i(\mathbf{r})]$. The last line is the expected virtual surplus of the optimal surrogate-ranking mechanism $\mathbf{x}$, which implies the result. $\square$

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 by noting that for bounded distributions, existing literature shows how to 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 15 (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 k th 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 16 (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 k th 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 17. *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 $p_{\text{run}}(n, 1/\epsilon) = \tilde{O}(n\epsilon^{-3})$ and $p_{\text{design}}(n, 1/\epsilon) = \tilde{O}(n^8\epsilon^{-20})$ 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.*

Our results have new implications for the truthful sample complexity literature as well. Specifically, we give 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 of this instantiation of the reduction are in Appendix C.

Theorem 18. *For agents with regularly distributed values (but potentially unbounded), there is a family of truthful mechanisms that satisfies C1, C2, and C3 with $p_{\text{run}}(n, 1/\epsilon) = \tilde{O}(n\epsilon^{-3})$ and $p_{\text{design}}(n, 1/\epsilon) = \tilde{O}(n^{10}\epsilon^{-22})$ for multiplicative approximation and the revenue objective.*

We conclude by noting that the polynomials in the sample complexity guarantees of Theorem 17 and Theorem 18 are quite unwieldy. We leave the derivation of tight tradeoffs between run-time and design-time samples to future work.

REFERENCES

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

- Xiaohui Bei and Zhiyi Huang. 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*. 720–733. <https://doi.org/10.1137/1.9781611973082.57>
- Jeremy Bulow and John Roberts. 1989. The Simple Economics of Optimal Auctions. *Journal of Political Economy* 97, 5 (October 1989), 1060–1090.
- Shuchi Chawla, Jason Hartline, and Denis Nekipelov. 2014. Mechanism design for data science. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 711–712.
- Shuchi Chawla and Jason D Hartline. 2013. Auctions with unique equilibria. In *ACM Conference on Electronic Commerce*. 181–196.
- Shuchi Chawla, Jason D Hartline, and Denis Nekipelov. 2017. Mechanism Redesign. *arXiv preprint arXiv:1708.04699* (2017).
- Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. 2015. Optimal crowdsourcing contests. *Games and Economic Behavior* (2015).
- Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM, 243–252.
- Anirban Dasgupta and Arpita Ghosh. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 319–330.
- Nikhil R Devanur, Jason D Hartline, and Qiqi Yan. 2015. Envy freedom and prior-free mechanism design. *Journal of Economic Theory* 156 (2015), 103–143.
- Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. 2016. The sample complexity of auctions with side information. In *STOC 2016*. 426–439. <https://doi.org/10.1145/2897518.2897553>
- Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. 2010. Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*. 129–138.
- Shaddin Dughmi, Jason D Hartline, Robert Kleinberg, and Rad Niazadeh. 2017. Bernoulli factories and black-box reductions in mechanism design. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 158–169.
- Paul Dütting and Thomas Kesselheim. 2015. Algorithms against Anarchy: Understanding Non-Truthful Mechanisms. In *16th ACM Conference on Economics and Computation*.
- Edith Elkind. 2007. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 736–745.
- Yiding Feng and Jason D Hartline. 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*.
- Yannai A Gonczarowski and Noam Nisan. 2017. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 856–868.
- Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. 2019. Settling the Sample Complexity of Single-parameter Revenue Maximization. *arXiv preprint arXiv:1904.04962* (2019).
- Jason Hartline, Darrell Hoy, and Sam Taggart. 2014. Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 693–710.
- Jason D Hartline. 2013. Bayesian mechanism design. *Foundations and Trends® in Theoretical Computer Science* 8, 3 (2013), 143–263.
- Jason D Hartline, Robert Kleinberg, and Azarakhsh Malekian. 2011. Bayesian incentive compatibility via matchings. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. 734–747.
- Jason D Hartline, Robert Kleinberg, and Azarakhsh Malekian. 2015. Bayesian incentive compatibility via matchings. *Games and Economic Behavior* 92 (2015), 401–429.
- Jason D Hartline and Brendan Lucier. 2015. Non-optimal mechanism design. *American Economic Review* 105, 10 (2015), 3102–24.
- Bernard J Jansen and Tracy Mullen. 2008. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business* 6, 2 (2008), 114–131.
- Jamie Morgenstern and Tim Roughgarden. 2015. On the Pseudo-Dimension of Nearly Optimal Auctions. In *Annual Conference on Neural Information Processing Systems 2015*. 136–144.
- Roger Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- Kent Osband. 1989. Optimal forecasting incentives. *Journal of Political Economy* 97, 5 (1989), 1091–1112.
- Tim Roughgarden and Okke Schrijvers. 2016. Ironing in the Dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. 1–18.
- Tim Roughgarden, Vasilis Syrgkanis, and Eva Tardos. 2017. The price of anarchy in auctions. *Journal of Artificial Intelligence Research* 59 (2017), 59–101.
- Vasilis Syrgkanis and Eva Tardos. 2013. Composable and Efficient Mechanisms. In *ACM Symposium on Theory of Computing*. 211–220. <https://doi.org/10.1145/2488608.2488635>

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 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 19. 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 $\mathbf{v} = (v_1, \dots, v_n)$; 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 Theorem 12

Theorem 12. 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 $\mathbf{r} = (r_1, \dots, r_n)$ of each agent among $T - 1$ runtime samples for their distributions and outputs a (possibly randomized) feasible allocation $\bar{x}(\mathbf{r})$. As a constraint, $\bar{x}$ must be monotone in the ranks of each agent. The objective is to maximize $\mathbf{E}[\sum_i \phi_i(q_i) \bar{x}_i(\mathbf{r})]$ 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 $\mathbf{r}$. 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 $\sum_i \mathbf{E}[\phi_i(q_i) \mid r_i] \bar{x}_i(\mathbf{r})$ by linearity of expectation. From this expression, it becomes clear that the optimal solution chooses the allocation which maximizes the quantity $\sum_i \mathbf{E}[\phi_i(q_i) \mid r_i] \bar{x}_i(\mathbf{r})$. Note that if $\phi_i(\cdot)$ is monotone, then this allocation rule will be monotone

as well, and therefore feasible.⁵ Setting these expected order statistics as surrogate values, the surrogate-ranking mechanism (Definition 4) optimizes this quantity. $\square$

C 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. Formally, we prove the following guarantee:

Theorem 20. *Let $R^* = \max_q R(q)$ be the monopoly revenue. For any $\delta, \gamma \in (0, 1)$, $O(T^6 \delta^{-4} \log(T/\gamma))$ samples suffice to estimate the expected revenue of a k -unit, T -bidder auction up to additive error δR^* for all $k \in \{1, \dots, T - 1\}$ simultaneously with probability at least $1 - \gamma$.*

Theorem 20 gives a sample complexity result as follows. Since $\sum_i R_i^*$ is at most n times the revenue of the optimal mechanism, the additive propagation of error guarantee of Theorem 14 also implies a multiplicative $(1 - n\epsilon)$ guarantee, which can be composed with the multiplicative guarantee of Theorem 10. The result, restated from Section 5.3, is:

Theorem 18. *Consider an arbitrary feasibility environment. Assume agents have regular (possibly unbounded value distributions). Then polynomially many sampled value profiles and polynomial run-time samples suffice to estimate a mechanism which obtains at least a $(1 - \epsilon)$ fraction of the revenue of the optimal mechanism with probability at least $1 - \epsilon$.*

We first outline the high-level strategy for proving Theorem 20. 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 k th 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 δR^* , it suffices to estimate P_k and P_{k-1} with error $2\delta R^*$.

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 $q_{k:T-1}$, where $q_{k:T-1}$ denotes the k th lowest order statistic of $T - 1$ $U[0, 1]$ random variables. Note that $q_{k:T-1}$ is distributed according to $\text{Beta}(k, T - k)$. 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 21. *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 20.

C.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 22. *Let $R^* = \max_q R(q)$. For any $1 > \delta > 0$ and γ , $O(K^2 \delta^{-2} \log(K/\gamma))$ samples suffice to guarantee that $K\hat{v}_j/j \in [R(j/K) - \delta R^*, R(j/K) + \delta R^*]$ for all $j \in \{1, \dots, K - 1\}$ simultaneously with probability at least $1 - \gamma$.*

⁵If 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.

Consider drawing $N = Km - 1$ samples, for some positive integer m . Note that the j th 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 23. *For any $1 > \delta > 0$ and γ , $Km = O(K\delta^{-2} \log(1/\gamma))$ samples suffice to guarantee that $q_j \in [(1 - \delta)j/K, (1 + \delta)j/K]$ with probability at least $1 - \gamma$.*

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

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

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

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

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

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

PROOF. We will show that $\hat{v}_j \in [(1 - \delta K)v(j/K), (1 + \delta K)v(j/K)]$. 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 $q_j \geq (1 - \delta)j/K$, we have:

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

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

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

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 $q_j \leq (1 + \delta)j/K$, we have

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

$\square$

Lemma 22 follows from the above results by applying the union bound.

PROOF OF LEMMA 22. Lemmas 23 and 24 imply that $O(K^2 \delta^{-2} \log(K/\delta))$ samples suffice to guarantee that $K\hat{v}_j/j \in [(1 - \delta)R(j/K), (1 + \delta)R(j/K)]$ with probability at least $1 - \gamma/K$. Applying the union bound, it follows that $O(K^2 \delta^{-2} \log(K/\delta))$ samples suffice to guarantee that this guarantee holds for all $j \in \{1, \dots, K - 1\}$ simultaneously with probability at least $1 - \gamma$. Noting that $R(j/K) \leq R^*$ implies the lemma. $\square$

C.2 Estimating The Interior Revenue Curve

In the previous section, we showed how to estimate $R(j/K)$ up to an additive δR^* . 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

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

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

Lemma 25. *Assume that for all $j \in \{1, \dots, K-1\}$, $K\hat{v}_j/j \in [R(j/K) - \delta R^*, R(j/K) + \delta R^*]$ for all $j \in \{1, \dots, K-1\}$. Then for all $q \in [1/K, 1 - 1/K]$:*

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

PROOF. We bound the overestimation and underestimation error in turn. First, note that since R is concave, it must be that $R(q)$ lies above the line between $((j-1)/K, \hat{R}_{j-1} - \delta R^*)$ and $(j/K, \hat{R}_j - \delta R^*)$. It follows that $\hat{R}(q)$ can only overestimate by at most δR^* . Next, note that a lower bound on $\hat{R}(q)$ is

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

In other words, the worst-case underestimation is δR^* , 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. $\square$

C.3 Proof of Theorem 20

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 $O(T^6 \delta^{-4} \log(T/\delta))$ samples suffice to estimate estimate ψ^k for all $k \in \{1, \dots, T-1\}$ to within additive error δR^* with probability at least $1 - \gamma$. To do so, we will consider estimating P_k for $k \in \{1, \dots, T-1\}$ as $T \int_{\delta/T^2}^{1-\delta/T^2} f_{k:T-1}(q) \hat{R}(q) dq$.

First we show that ignoring the revenue contribution from the intervals $[0, \delta/T^2]$ and $[1 - \delta/T^2, 1]$ cannot hurt our estimate by much. Let $F_{k:T-1}$ denote the CDF of the k th lowest order statistic of $T-1$ $U[0, 1]$ random variables. Then using properties of the Beta distribution, we have that $F_{k:T-1}(\delta/T^2) \leq \delta/T$ and $1 - F_{k:T-1}(1 - \delta/T^2) \leq \delta/T$ for all $k \in \{1, \dots, T-1\}$. Since for $q \in [0, \delta/T^2] \cup [1 - \delta/T^2, 1]$, $R(q) \leq R^*$, it follows from Lemma 21 that

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

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

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

Choosing $K = T^2/\delta$ in Lemma 25 yields the result.