
Refined Mechanism Design for Approximately Structured Priors via Active Regression

Christos Boutsikas
Purdue University
cboutsik@purdue.edu

Petros Drineas
Purdue University
pdrineas@purdue.edu

Marios Mertzaniadis
Purdue University
mmertzan@purdue.edu

Alexandros Psomas
Purdue University
apsomas@cs.purdue.edu

Paritosh Verma
Purdue University
verma136@purdue.edu

Abstract

We consider the problem of a revenue-maximizing seller with a large number of items m for sale to n strategic bidders, whose valuations are drawn independently from high-dimensional, unknown prior distributions. It is well-known that optimal and even approximately-optimal mechanisms for this setting are notoriously difficult to characterize or compute, and, even when they can be found, are often rife with various counter-intuitive properties. In this paper, following a model introduced recently by Cai and Daskalakis [CD22], we consider the case that bidders' prior distributions can be well-approximated by a topic model. We design an active learning component, responsible for interacting with the bidders and outputting low-dimensional approximations of their types, and a mechanism design component, responsible for robustifying mechanisms for the low-dimensional model to work for the approximate types of the former component. On the active learning front, we cast our problem in the framework of Randomized Linear Algebra (RLA) for regression problems, allowing us to import several breakthrough results from that line of research, and adapt them to our setting. On the mechanism design front, we remove many restrictive assumptions of prior work on the type of access needed to the underlying distributions and the associated mechanisms. To the best of our knowledge, our work is the first to formulate connections between mechanism design, and RLA for active learning of regression problems, opening the door for further applications of randomized linear algebra primitives to mechanism design.

1 Introduction

The design of revenue-optimal auctions is a central problem in Economics and Computer Science. In this problem, a revenue-maximizing seller has m heterogeneous items for sale to n strategic bidders. Each bidder i has a type $\mathbf{t}_i \in \mathbb{R}^d$ which contains enough information to encode the bidder's willingness to pay for every subset of items. Bidders' types are private information, and thus, in order to provide meaningful guarantees on the seller's revenue, the standard approach in Economics is to make a Bayesian assumption: types are drawn from a joint distribution \mathcal{D} .

Assuming a single item for sale and bidders' types that are drawn independently from known distributions, Myerson's seminal work [Mye81] provides a closed-form solution for the revenue-optimal mechanism. Beyond this single-item case, however, multi-item mechanism design remains an active research agenda, even forty years later. Optimal mechanisms are no longer tractable, in any sense of the word, as well as exhibit various counter-intuitive properties [MV07, DDT13, DDT15, BCKW15, HR15, HN19, PSCW22]; see [Das15] for a survey. On the other

hand, significant research effort has culminated in numerous compelling positive results, such as simple and approximately optimal auctions [CHK07, CHMS10, CMS15, Yao15, RW15, CM16, CDW16, CZ17, KW19, BILW20], as well as efficient algorithms for computing near-optimal auctions [AFH⁺12, CDW12a, CDW12b, CDW13a, CDW13b], even with just sampling access to the type distribution [CR14, HMR15, DHP16, MR16, CD17, GHZ19, BCD20, GW21].

Despite all this progress, however, there are key challenges in applying these results in practice. First, the computational complexity, sample complexity, approximation guarantees, and communication complexity (i.e., the amount of information the bidder should communicate to the seller) of these results often depend on the number of items m , which could be prohibitively large (e.g., think of m as the number of items on Amazon.com). Second, bidders' type distributions are typically high-dimensional, or otherwise complex objects, that are not known nor can they be sampled. Instead, the designer might have an estimated distribution $\hat{\mathcal{D}}$, e.g., through market research, that is close to the real distribution \mathcal{D} . Motivated by these issues, Cai and Daskalakis [CD22] introduce a model where the true type distribution \mathcal{D}_i of bidder i is close to a structured distribution $\hat{\mathcal{D}}_i$. Specifically, they assume that there is an underlying design matrix $\mathbf{A} \in \mathbb{R}^{m \times k}$ of k "archetypes," with $k \ll m$. Intuitively, bidder i can be approximated by a linear combination of k archetypal bidders. This same assumption has been central in the study of recommender systems. In this model, [CD22] give a framework for transforming a mechanism $\hat{\mathcal{M}}$ for the low-dimensional distribution $\hat{\mathcal{D}}_z$ into a mechanism for the true type distribution with good revenue guarantees, and whose computational and communication complexity does not depend on the number of items m .

The impact of their work notwithstanding, the above results require strong structural assumptions on: the design matrix \mathbf{A} ; the bidders' valuation functions; and very specific access to (or exact knowledge of) the structured distribution $\hat{\mathcal{D}}_z$ and the mechanism $\hat{\mathcal{M}}$. *Our work connects the recommender system approaches for mechanism design with recent progress on Randomized Linear Algebra for active learning for regression problems. We relax, and even remove these restrictive assumptions, and open the door to future exploration of more elaborate recommender system models in the context of mechanism design, using randomized linear algebra primitives.*

The framework and results [CD22]. To place our results in context, we start with a brief overview of [CD22], which considers a setting where the type distribution \mathcal{D}_i of bidder i is close to a distribution $\hat{\mathcal{D}}_i$ in the Prokhorov distance, under the ℓ_∞ norm (Definition 1). Here, $\hat{\mathcal{D}}_i$ first samples a vector $\mathbf{z} \in [0, 1]^k$ from a low-dimensional distribution $\hat{\mathcal{D}}_{z,i}$, and then outputs $\mathbf{A}\mathbf{z}$, where $\mathbf{A} \in \mathbb{R}^{m \times k}$ is a known matrix. The proposed framework has a learning and a mechanism design component.

Their learning component consists of a communication-efficient query protocol \mathcal{Q} for interacting with each bidder i such that, if the type \mathbf{t}_i of bidder i satisfies $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}\|_\infty \leq \varepsilon$, the query protocol outputs a vector $\mathcal{Q}(\mathbf{t}_i)$ such that $\|\mathcal{Q}(\mathbf{t}_i) - \mathbf{z}\|_\infty \leq \zeta_\infty$.¹ [CD22] give such query protocols under strong conditions on \mathbf{A} , and specifically when \mathbf{A} : (i) satisfies an assumption similar to the separability condition of [DS03]; (ii) is generated from a distribution where each archetype is an independent copy of an m -dimensional Gaussian; and (iii) is generated from a distribution where each archetype is an independent copy of an m -dimensional, bounded distribution with weak dependence. We discuss these restrictions in more detail in Appendix D.4.1. The query complexity, as well as the error ζ_∞ , depend on them, but, importantly, are independent of the number of items m .

Their mechanism design component is a refinement of a robustification result of [BCD20], and transforms a mechanism $\hat{\mathcal{M}}$ for the low-dimensional distribution into a mechanism for \mathcal{D} , while being agnostic about \mathcal{D} , and in a way that the overall error does not depend on m . For this transformation to work, one needs to interact with $\hat{\mathcal{M}}$ and $\hat{\mathcal{D}}_z$ using highly non-trivial operations, and specifically, the construction requires an oracle for computing $\arg\max_{\mathbf{t}'_i \in \text{supp}(\tilde{\mathcal{D}}_{z,i})} \mathbb{E}_{\mathbf{b}_{-i} \sim \tilde{\mathcal{D}}_{z,-i}} [u_i(v_i, \hat{\mathcal{M}}(\mathbf{t}'_i, \mathbf{b}_{-i}))]$, where $\tilde{\mathcal{D}}_{z,i}$ is a rounded down version of $\hat{\mathcal{D}}_{z,i}$. This oracle is therefore solving a computationally demanding problem,² as well as requires exact knowledge of bidders' valuation functions.

¹Recall that for any (integer) $1 \leq p < \infty$ and a vector $\mathbf{x} \in \mathbb{R}^d$, $\|\mathbf{x}\|_p^p = \sum_{i=1}^d |\mathbf{x}_i|^p$; for $p = \infty$, $\|\mathbf{x}\|_\infty = \max_{i=1 \dots d} |\mathbf{x}_i|$. See [CD22] for an exact expression for ζ_∞ .

²For an arbitrary mechanism $\hat{\mathcal{M}}$ and discrete $\tilde{\mathcal{D}}_{z,i}$, this task requires $\Pi_{i \in [n]} |\text{supp}(\tilde{\mathcal{D}}_{z,i})|$ operations.

Combining the two components, [CD22] obtain mechanisms for \mathcal{D} , for constrained-additive bidders.³ In these mechanisms, each bidder is required to answer a small number of simple Yes/No queries of the form “are you willing to pay p for item j ?”, such that the loss in revenue and the violation in incentive compatibility do not depend on the number of items m .

1.1 Our contributions

Randomized Linear Algebra (RLA) for active learning. RLA for active learning has focused on solving regression problems of the form

$$\mathbf{z} = \arg \min_{\mathbf{x}} \|\mathbf{t} - \mathbf{A}\mathbf{x}\|_p,$$

for ℓ_p norms with $1 \leq p < \infty$, by querying *only* a subset of elements of the vector \mathbf{t} . In prior work on RLA for active learning, the focus has been on recovering an approximate solution that achieves a relative error or constant factor approximation to the optimum loss. We adapt these bounds to our setting, which could include noisy instead of exact queries, and prove bounds for the ℓ_p norm error of the exact minus the approximate solution. Specifically, we bound $\|\mathcal{Q}(\mathbf{t}_i) - \mathbf{z}\|_p \leq \zeta_p$. We provide bounds on ζ_p for all integers $1 \leq p < \infty$. Our bounds depend on the modelling error $\|\mathbf{t} - \mathbf{A}\mathbf{x}\|_p$ and some measure of the query noise (see Definitions 4 and 5); both dependencies are expected. Importantly, our bounds hold for *arbitrary archetype matrices*, very much unlike the work of [CD22], which focused on very restricted classes of matrices. A single property of the archetype matrix, the smallest singular value with respect to the ℓ_p induced matrix norm, $\sigma_{\min,p}(\mathbf{A})$, can characterize the quality of the error. As $\sigma_{\min,p}(\mathbf{A})$ decreases, the error ζ_p grows. Intuitively, $\sigma_{\min,p}(\mathbf{A})$ is a measure of independence of the archetypes, with small values corresponding to linearly dependent archetypes. The query complexity needed to achieve error ζ_p is almost linear (up to logarithmic factors) on k for $p = 1$ and $p = 2$, and grows with $k^{p/2}$ for $p \geq 3$. Our query complexity bounds have no dependency on m (number of items) or d (dimensionality of a type) enabling us to produce results way beyond constrained-additive valuations, as $d = 2^m$ dimensions suffice to encode *arbitrary* valuation functions.

It is critical to highlight that our ability to provide bounds on the approximation error for arbitrary archetype matrices is, at least partly, due to leveraging information from the archetype matrix \mathbf{A} . Specifically, we use this information to select which bidder types to query, instead of just querying types uniformly at random. This information involves the computation or approximation of the well-studied leverage scores of \mathbf{A} for $p = 2$ and of the so-called Lewis weights for all other values of p (see Section 3.1). We do note that in our framework, the errors due to the modeling of the bidder type \mathbf{t} as the product $\mathbf{A}\mathbf{z}$ and the query noise are always bounded by the respective ℓ_p norm. Thus, our models are more restrictive compared to the ℓ_∞ norm models of [CD22]. However, to the best of our knowledge, even assuming such restrictive models, the results and tools of prior work do not extend to arbitrary archetype matrices. This is precisely the gap that is bridged by our work, using RLA for active learning of ℓ_p norm regression for $1 \leq p < \infty$.

Mechanism Design. On the mechanism design front, our main contribution is relaxing the assumptions of [CD22, BCD20] on the type of access needed to the low-dimensional distribution and the mechanism for it. Specifically, we further refine the robustification result of Brustle et al. [BCD20] and remove the need for the aforementioned strong oracle.

The main difficulty of transforming mechanisms for one distribution into mechanisms for another distribution is that the two distributions might not share the same support. The crux of the issue is that the incentive constraints are very delicate; a small change in the underlying distribution may drastically change the agents’ *valuation distribution* over the mechanisms’ outcomes. One way to tame the distribution of outcomes is to map bids that are not in the support of the initial distribution, to bids that are. Brustle et al. [BCD20] do this by “optimally misreporting” on behalf of the bidder, by calculating $\arg \max_{\mathbf{t}'_i \in \text{supp}(\tilde{\mathcal{D}}_{z,i})} \mathbb{E}_{\mathbf{b}_{-i} \sim \tilde{\mathcal{D}}_{z,-i}} [u_i(v_i, \hat{\mathcal{M}}(\mathbf{t}'_i, \mathbf{b}_{-i}))]$, where $\tilde{\mathcal{D}}_{z,i}$ is a rounded-down version of $\hat{\mathcal{D}}_{z,i}$. As we’ve discussed, for this operation to be viable, many things need to be assumed about what the designer knows and can compute about $\hat{\mathcal{M}}$, $\hat{\mathcal{D}}_{z,i}$, and the bidder’s valuation function. Our approach, instead, leverages the fact that when two distributions are close in Prokhorov distance, under any ℓ_p norm, *any* point on the support of one distribution is close to a point on the support

³A function is constrained-additive if $v(t, S) = \max_{T \in \mathcal{I} \cap 2^S} \sum_{j \in T} t_j$, where \mathcal{I} is a downward-closed set system.

of the other, with high probability. Our construction simply maps a report w_i to the “valid” report (approximately) closest to w_i in ℓ_p distance. This operation is linear on the support size. Furthermore, our overall robustification result holds for all norms, not just ℓ_∞ , and our construction is completely agnostic to bidders’ valuation functions.

Combining the components. Combining the two components we can, without any assumptions on \mathbf{A} , given a mechanism for the low-dimensional prior, design mechanisms with comparable revenue guarantees, where each bidder is required to answer a small number of queries. Our queries ask a bidder her value for a subset of items, and our mechanism can accommodate *any* valuation function, significantly extending the scope of our results.

Related Work. Aside from the work of [CD22], on the mechanism design front, the most relevant works are [CD17], that consider learning multi-item auctions given “approximate” distributions, and [BCD20], that consider learning multi-item auctions when type distributions are correlated, yet admit special structure. On the RLA front, we leverage and adapt multiple recent results on approximating ℓ_p regression problems in an active learning setting. We discuss prior work on RLA for active learning and its connections to our setting in Section 3.1 and Appendix D.

2 Preliminaries

Let $[n] := \{1, 2, \dots, n\}$ denote the first n natural numbers. A revenue-maximizing seller has a set $[m]$ of m heterogeneous items for sale to n strategic bidders indexed by $[n]$. Bidder i has a private type vector $\mathbf{t}_i \in \mathbb{R}^d$, and the types of all bidders are represented by a *type profile* $\mathbf{t} = (\mathbf{t}_1, \dots, \mathbf{t}_n)$. Bidder i has a valuation function $v_i : \mathbb{R}^d \times [0, 1]^m \rightarrow \mathbb{R}_+$, that takes as input the bidder’s type and a (possibly randomized) set of items $S \subseteq [m]$ and outputs the bidder’s value for S . Note that $d \leq 2^m$, since expressing a valuation function requires at most one real number per subset of items. Types are drawn independently. Let $\mathcal{D} = \times_{i \in [n]} \mathcal{D}_i$ be the distribution over bidders’ types, $\mathcal{D}_{-i} = \times_{j \in [n]/\{i\}} \mathcal{D}_j$ be the distribution of all bidders excluding i , and $\text{supp}(\mathcal{D})$ be the support of a distribution \mathcal{D} .

Mechanisms. A mechanism $\mathcal{M} = (x, p)$ is a tuple where $x : \mathbb{R}_+^d \rightarrow [0, 1]^{nm}$ is the allocation rule, and $p : \mathbb{R}_+^d \rightarrow \mathbb{R}_+^n$ is the payment rule, which map *reported* types to allocations of the items and payments, respectively. Specifically, $x_{i,j}(\mathbf{b}) := (x(\mathbf{b}))_{i,j}$ denotes the probability that bidder i receives item j for input valuation profile \mathbf{b} , and $p_i(\mathbf{b}) := (p(\mathbf{b}))_i$ denotes the amount bidder i has to pay. Let $u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{b}))$ be the utility of bidder i with type \mathbf{t}_i for participating in mechanism \mathcal{M} , under reports \mathbf{b} . Bidders are risk-free and quasi-linear i.e., $u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{b})) = \mathbb{E}[v_i(\mathbf{t}_i, x(\mathbf{b})) - p_i(\mathbf{b})]$, where the expectation is taken over the randomness of the allocation rule. Since we only consider truthful mechanisms, unless stated otherwise, reported types will be the same as the true types.

A bidder’s objective is to maximize her utility. The seller strives to design mechanisms that incentivize bidders to report truthfully. We use the following notions of truthfulness. A mechanism \mathcal{M} is ε -*Bayesian Incentive Compatible* (ε -*BIC*), if for each bidder $i \in [n]$, any type \mathbf{t}_i and misreport \mathbf{t}'_i we have that: $\mathbb{E}_{\mathbf{t}_{-i} \sim \mathcal{D}_{-i}}[u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{t}_i, \mathbf{t}_{-i}))] \geq \mathbb{E}_{\mathbf{t}_{-i} \sim \mathcal{D}_{-i}}[u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{t}'_i, \mathbf{t}_{-i}))] - \varepsilon$. A mechanism \mathcal{M} is (ε, δ) -*BIC* if for each bidder $i \in [n]$, and any misreport \mathbf{t}'_i we have that:

$$\mathbb{P}_{\mathbf{t}_i \sim \mathcal{D}_i} \left[\mathbb{E}_{\mathbf{t}_{-i} \sim \mathcal{D}_{-i}} [u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{t}_i, \mathbf{t}_{-i}))] \geq \mathbb{E}_{\mathbf{t}_{-i} \sim \mathcal{D}_{-i}} [u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{t}'_i, \mathbf{t}_{-i}))] - \varepsilon \right] \geq 1 - \delta.$$

A $(\varepsilon, 0)$ -BIC mechanism is a ε -BIC mechanism; a 0-BIC mechanism is simply BIC. Finally, a mechanism \mathcal{M} is *ex-post Individually Rational* (*IR*) if for all type profiles \mathbf{t} and all bidders $i \in [n]$, $u_i(\mathbf{t}_i, \mathcal{M}(\mathbf{t}_i, \mathbf{t}_{-i})) \geq 0$. The seller’s objective is to maximize her expected *revenue*. For a mechanism \mathcal{M} and distribution \mathcal{D} we denote the expected revenue as $\text{Rev}(\mathcal{M}, \mathcal{D}) = \mathbb{E}_{\mathbf{t} \sim \mathcal{D}} \left[\sum_{i \in [n]} p_i(\mathbf{t}) \right]$. Note that, we are calculating revenue assuming truthful reports, even for, e.g., (ε, δ) -BIC mechanisms.

Statistical Distance. In this work we design mechanisms that work well, as long as they are evaluated on distributions that are “close” to \mathcal{D} . Here, we define the notion of distance between two probability measures that we use throughout the paper.

Definition 1 (Prokhorov Distance). *Let (\mathcal{U}, d) be a metric space and \mathcal{B} be a σ -algebra on \mathcal{U} . For $A \in \mathcal{B}$, let $A^\varepsilon = \{x : \exists y \in A \text{ s.t. } \pi(x, y) < \varepsilon\}$ where π is some distance metric. Two probability measures P, Q on \mathcal{B} have Prokhorov distance: $\inf\{\varepsilon > 0 : P(A) \leq Q(A^\varepsilon) + \varepsilon \text{ and } Q(A) \leq$*

$P(A^\varepsilon) + \varepsilon, \forall A \in \mathcal{B}$. We choose π to be the ℓ_p -distance, and we denote the Prokhorov distance between measures P, Q as $\pi_p(P, Q)$.

The following is an equivalent characterization of Prokhorov distance due to Strassen [Str65].

Lemma 1 ([Str65]). *Let \mathcal{D} and \mathcal{D}' be two distributions supported on \mathbb{R}^n . $\pi_p(\mathcal{D}, \mathcal{D}') \leq \varepsilon$ iff there exists coupling γ of \mathcal{D} and \mathcal{D}' , such that $\mathbb{P}_{(\mathbf{x}, \mathbf{y}) \sim \gamma}[\|\mathbf{x} - \mathbf{y}\|_p > \varepsilon] \leq \varepsilon$.*

Recommendation system-inspired model. We assume that, for each bidder, there exists a known design matrix $\mathbf{A} \in \mathbb{R}^{d \times k}$, where the columns of \mathbf{A} represent k ‘‘archetypes’’, for a constant k . Our results hold if these matrices are different for each bidder, however, for ease of notation we will assume all bidders have the same design matrix. For each bidder i there exists a distribution $\hat{\mathcal{D}}_{z,i}$ supported on the latent space $[0, 1]^k$. Let $\hat{\mathcal{D}}_i = \mathbf{A} \circ \hat{\mathcal{D}}_{z,i}$ be the distribution induced by multiplying a sample from $\hat{\mathcal{D}}_{z,i}$ with the design matrix \mathbf{A} , i.e., $\hat{\mathcal{D}}_i$ is the distribution of $\mathbf{A}\mathbf{y}$ where $\mathbf{y} \sim \hat{\mathcal{D}}_{z,i}$.

The valuation function over the latent types is defined as $v_i^{\mathbf{A}}(\mathbf{z}_i, S) := v_i(\mathbf{A}\mathbf{z}_i, S)$ for any bundle $S \subseteq [m]$. We will use the following notion of Lipschitz continuity for valuation functions.

Definition 2 (Lipschitz Valuation). *A valuation function $v(\cdot, \cdot) : \mathbb{R}^d \times [0, 1]^m \rightarrow \mathbb{R}_+$ is \mathcal{L} -Lipschitz, if for any two types $\mathbf{t}, \mathbf{t}' \in \mathbb{R}^d$ and any bundle $S \subseteq [m]$, $|v(\mathbf{t}, S) - v(\mathbf{t}', S)| \leq \mathcal{L}\|\mathbf{t} - \mathbf{t}'\|_\infty$.*

We include a table, Table 1, with all the notation used throughout the paper in the appendix.

3 Active Learning for Regression and Mechanism Design

In this section, we state our main results, deferring all technical proofs to the appendix. We present mechanisms that are completely agnostic with respect to \mathcal{D} , the distribution from which bidders’ types are drawn from. However, we have limited access (described later in this section) to (i) a design matrix \mathbf{A} ; (ii) distributions $\hat{\mathcal{D}}_{z,i}$ over \mathbb{R}^k , where for all $i \in [n]$, $\pi_p(\mathcal{D}_i, \mathbf{A} \circ \hat{\mathcal{D}}_{z,i}) \leq \varepsilon_{\text{mdl},p}$ for some $\varepsilon_{\text{mdl},p} > 0$; and (iii) a mechanism \mathcal{M} for $\hat{\mathcal{D}}_z = \times_{i \in [n]} \hat{\mathcal{D}}_{z,i}$. Our approach is modular and starts by building an active learning component (Section 3.1) followed by the mechanism design component (Section 3.2). We combine the two components to get an overall mechanism for \mathcal{D} in Section 3.3.

3.1 Active learning for regression via Randomized Linear Algebra

Our objective is to design a communication-efficient, active learning query protocol for the seller that interacts with each bidder i , and infers their type $\mathbf{t}_i \in \mathbb{R}^d$ by accessing only a small subset of elements of the type vector (as d is very large). We use \mathcal{Q} to denote the query protocol, whose output is a vector in the low-dimensional latent space \mathbb{R}^k . A bidder interacts with the query protocol *truthfully* if it is in her best interest to evaluate functions requested by the protocol on her true private type \mathbf{t}_i . We use $\mathcal{Q}(\mathbf{t}_i) \in \mathbb{R}^k$ to denote the output of \mathcal{Q} when interacting with a truthful bidder with type \mathbf{t}_i . We now define the notion of an $(\varepsilon_{\text{mdl},p}, \zeta_p, p)$ -query protocol and the notion of *query noise*.

Definition 3 ($(\varepsilon_{\text{mdl},p}, \zeta_p, p)$ -query protocol). *\mathcal{Q} is called an $(\varepsilon_{\text{mdl},p}, \zeta_p, p)$ -query protocol, if, for all $\mathbf{t} \in \mathbb{R}^d$ and $\mathbf{z} \in \mathbb{R}^k$ satisfying $\|\mathbf{t} - \mathbf{A}\mathbf{z}\|_p \leq \varepsilon_{\text{mdl},p}$, we have $\|\mathbf{z} - \mathcal{Q}(\mathbf{t})\|_p \leq \zeta_p$.*

Definition 4 (Query noise). *Let \mathbf{t}_i be the true type of a bidder. Our query protocol can access elements of $\mathbf{t}_i + \epsilon_{\text{ng},p}$, where $\epsilon_{\text{ng},p}$ is an (unknown) vector. The query noise $\epsilon_{\text{ng},p}$ satisfies $\|\epsilon_{\text{ng},p}\|_p \leq \epsilon_{\text{ng},p}$.*

The query noise depends on the specifics of the interactions between the seller and the bidder. For example, if the seller is only allowed to ask queries of the form ‘‘what is your value for the subset S ?’’, the query noise $\epsilon_{\text{ng},p}$ is equal to zero. Our bounds will also depend on the *model error*.

Definition 5 (Model error). *Given a type profile $\mathbf{t} \in \mathbb{R}^{nd}$, the model error is $\varepsilon_{\text{mdl},p}$ if, for all $i \in [n]$, there exists a $\mathbf{z}_i \in \mathbb{R}^k$ such that $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_p \leq \varepsilon_{\text{mdl},p}$.*

Note that, we don’t have bounds of the form ‘‘ $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\| \leq \varepsilon$ ’’ for individual types, but for the distributions \mathcal{D}_i and $\mathbf{A} \circ \hat{\mathcal{D}}_{z,i}$. The characterization of Prokhorov distance (Lemma 1) allows us to relate the two quantities in the proofs that follow.

We now rephrase the above discussion in order to cast it in the framework of Randomized Linear Algebra (RLA) and active learning. Dropping the index i for notational simplicity, we assume that

$\mathbf{Az} \approx \mathbf{t}$ and we seek to recover an approximate solution vector $\mathcal{Q}(\mathbf{t})$ such that the ℓ_p norm error between the approximate and the optimal solution is bounded. *Importantly*, the query protocol \mathcal{Q} is *not* allowed full access to the vector \mathbf{t} in order to construct the approximate solution vector. This is a well-studied problem in the context of RLA: the learner is given a large collection of k -dimensional data points (the $d \gg k$ rows of the design matrix $\mathbf{A} \in \mathbb{R}^{d \times k}$), but can only query a small subset of the real-valued labels associated with each data point (elements of the vector $\mathbf{t} \in \mathbb{R}^d$). Prior work in RLA and active learning has studied this problem in order to identify the optimal number of queries that allow efficient, typically relative error, approximations of the loss function. In our parlance, prior work has explored the existence of query protocols that construct a vector $\mathcal{Q}(\mathbf{t})$ such that

$$\|\mathbf{t} - \mathbf{A}\mathcal{Q}(\mathbf{t})\|_p \leq \gamma_p \|\mathbf{t} - \mathbf{Az}\|_p, \quad (1)$$

where $\gamma_p > 1$ is an error parameter that controls the approximation accuracy. Of particular interest in the RLA literature are *relative error* approximations, with $\gamma_p = 1 + \epsilon$, for some small $\epsilon > 0$; see [MMWY21, MMWY22] for a detailed discussion. However, relative error approximations are less important in our setting, since our protocols in Section 3.2 necessitate $\zeta_p \geq \epsilon_{\text{mdl},p}$. For $p = 2$, the underlying problem is active learning for least-squares regression: [DMMS11] analyzed its complexity (namely, the number of queries) of query protocols in this setting, eventually providing matching upper and lower bounds. Similarly, for $p = 1$, the underlying problem is active learning for least absolute deviation regression, a robust version of least-squares regression: [CD21] analyzed the complexity of query protocols in this setting. The query protocols of [DMMS11, CD21] are straightforward: they sample a small set of labels (i.e., bidder types) and elicit the bidder’s preferences for this set. Then, the respective ℓ_p norm regression problem is solved on the smaller set and the resulting solution is returned as an approximation to the original solution.⁴ The types to be sampled (see Appendix D.1 for details) are selected using distributions that can be constructed by accessing *only* the design matrix \mathbf{A} . Specifically, for the $p = 2$ case, one needs to compute or approximate the *leverage scores* of the rows of the matrix \mathbf{A} . For the $p = 1$ case, one needs to compute or approximate the *Lewis weights* of the design matrix \mathbf{A} . (The Lewis weights are an extension of the leverage scores to ℓ_p norms for $p \neq 2$.) The work of [DW17, DWH18, DM21, DW18, CP19] for the $p = 2$ case involves more elaborate query protocols, using primitives such as volume sampling and the Batson-Spielman-Srivastava sparsifier to improve the query complexity. Finally, the $p > 2$ case for active learning for regression problems was recently resolved in [MMWY22, MMWY21]; we discuss their approach in our context in Appendix D.2.

To the best of our knowledge, our work is the first one to formulate connections between mechanism design and Randomized Linear Algebra for active learning. Two technical points of departure that are needed in order to adapt the RLA work for active learning to the mechanism design framework are: (i) we need to derive bounds of the form of eqn. (1) for the ℓ_p norm distance between the exact and approximate solutions, whereas prior work typically bounds the error of the *loss* function when an approximate solution is used; and (ii) the entries of the bidder’s type vector \mathbf{t} might not be known exactly, but only up to a small error. The latter assumption corresponds to the use of *noisy queries* in the model of [CD22] and is known to be equivalent, up to logarithmic factors, to *threshold queries* via binary search. Our work addresses both technicalities and seamlessly combines the RLA work for active learning with mechanism design.

Prior to stating our main result, we need to define a fundamental property of the design matrix $\mathbf{A} \in \mathbb{R}^{d \times k}$ that will affect the approximation error. Let

$$\sigma_{\min,p}(\mathbf{A}) = \min_{\mathbf{x} \in \mathbb{R}^k, \|\mathbf{x}\|_p=1} \|\mathbf{Ax}\|_p. \quad (2)$$

For $p = 2$, this is simply the smallest singular value of the matrix \mathbf{A} . For other values of p , the above definition is the standard generalization of the smallest singular value of \mathbf{A} for the induced matrix ℓ_p norm. Notice that $\sigma_{\min,p}(\mathbf{A})$ is a property of the matrix \mathbf{A} and can be computed *a priori* via, say, the QR factorization or the Singular Value Decomposition (SVD) for $p = 2$ and via linear programming for $p = 1$. As we will see in Theorem 1 below, smaller values of $\sigma_{\min,p}(\mathbf{A})$ result in increased sample complexity for our query protocols.

Theorem 1. *Let $\mathbf{A} \in \mathbb{R}^{d \times k}$ be the design matrix, and recall the definitions of the model error $\epsilon_{\text{mdl},p}$ (Definition 5) and the query noise (Definition 4). For all integers $1 \leq p < \infty$, there exist query*

⁴To be precise, multiple smaller problems have to be solved and a “good enough” solution has to be chosen in order to boost the success probability. See Appendix D for details.

protocols \mathcal{Q} using s_p queries for each bidder $i \in [n]$, such that, with probability at least $1 - \delta$,

$$\|\mathbf{z}_i - \mathcal{Q}(\mathbf{t}_i)\|_p \leq \frac{c_p(\varepsilon_{\text{mdl},p} + \varepsilon_{\text{nq},p})}{\sigma_{\min,p}(\mathbf{A})} = \zeta_p$$

holds for all n bidders $i \in [n]$. Here c_p is a small constant that depends on p .⁵ The respective query complexities for $p = 1$ and $p = 2$ are (asymptotically) identical:

$$s_1 = s_2 = O(k \cdot \ln k \cdot \ln n / \delta). \quad (3)$$

For $p \geq 3$, the query complexity is

$$s_p = O\left(k^{p/2} \cdot \ln^3 k \cdot \ln n / \delta\right). \quad (4)$$

Several comments are in order. (i) The error ζ_p is a small constant times the modelling error plus the error due to noisy queries. In the limit case where the modelling error is equal to zero and the queries are noiseless, the bidders' types can be recovered exactly in our framework. However, as the modelling error and the query noise increase, approximating user types becomes harder and less accurate. (ii) Importantly, the approximation accuracy of Theorem 1 grows linearly with the inverse of the smallest ℓ_p norm singular value of the design matrix \mathbf{A} . Our results indicate that the approximation accuracy of the query model \mathcal{Q} depends on this simple property of the design matrix \mathbf{A} . For example, focusing on the $p = 2$ case, our theorem shows that as the archetypes (columns of the matrix \mathbf{A}) become more linearly dependent and the smallest singular value approaches zero, the error of our approximation worsens. This is quite reasonable: if archetypes are linearly dependent, then it is increasingly difficult to approximate the respective entries of the vector \mathbf{z} . (iii) The query complexities s_1 and s_2 are asymptotically identical, growing linearly with $k \ln k$, where k is the number of archetypes. They both depend on the log of the number of bidders (due to a union bound) and on the log of $1/\delta$, where δ is the failure probability. The query complexity for $p \geq 3$ is larger and is dominated by the $k^{p/2}$ term. Importantly, the query complexity remains independent of d , the number of bidder types, which, in worst case, could be exponential to the number of underlying items. (iv) Improving the sampling complexities s_1 and s_2 has been a topic of intense interest in the RLA community and we defer the reader to [MMWY21, MMWY22], which has essentially provided matching upper and lower bounds for various values of p . We just note that for the well-studied $p = 2$ case, volume sampling approaches [DW17, DWH18, DM21, DW18] achieve essentially matching bounds, while the work of [CP19] removes (at least in expectation) the $\ln k$ factor from s_2 , at the expense of significant additional protocol complexity. From a technical perspective, we note that $\zeta_p \geq \varepsilon_{\text{mdl},p}$, as necessitated in Theorem 2 and that our query protocols are all one-round protocols.

Finally, notice that our theorem works for all $p \geq 1$, but not for $p = \infty$, which was the setting of [CD22]. In Appendix D.4, we present a (modest) improvement of the result of [CD22] and explain why it seems improbable that the $p = \infty$ case can be generalized to a much broader class of design matrices. This is a strong motivating factor to explore properties of mechanism design for the recommender system setting for other ℓ_p norms, as we do in this work.

3.2 The Mechanism Design component

The goal of the mechanism design component is to transform a mechanism $\hat{\mathcal{M}}$ for $\hat{\mathcal{D}}$ into a mechanism \mathcal{M} for \mathcal{D}_z . We first define exactly the type of access to $\hat{\mathcal{D}}_z$ and $\hat{\mathcal{M}}$ our construction requires.

Definition 6 (Access to $\hat{\mathcal{M}}$). By “query access to $\hat{\mathcal{M}}$ ” we mean access to an oracle which, given a valuation profile \mathbf{t} , outputs the allocation and payments of $\hat{\mathcal{M}}$ on input \mathbf{t} .

Definition 7 (Access to $\hat{\mathcal{D}}_z$). By “oracle access to $\hat{\mathcal{D}}_z$ ” we mean access to (1) a sampling algorithm \mathcal{S}_i for each $i \in [n]$, where $\mathcal{S}_i(\mathbf{x}, \delta)$ draws a sample from the conditional distribution of $\hat{\mathcal{D}}_{z,i}$ on the k -dimensional cube $\times_{j \in [k]} [x_j, x_j + \delta_j]$, and (2) an oracle which, given as input a type \mathbf{t}_i for bidder i , outputs the type in the support of $\hat{\mathcal{D}}_{z,i}$ that is closest to \mathbf{t}_i in ℓ_p distance, i.e., outputs $\text{argmin}_{\mathbf{t}'_i \in \text{supp}(\hat{\mathcal{D}}_{z,i})} \|\mathbf{t}_i - \mathbf{t}'_i\|_p$.

⁵We make no attempt to optimize constants and focus on simplicity of presentation. In our proofs, $c_1 = 2.5$; $c_2 = 7.5$; and for $p \geq 3$, $c_p = 18 \cdot (200)^{1/p} + 3$. Notice that the last constant converges to 21 as p increases.

If the allocation is randomized, our approach works even if the query to the oracle returns a (correct) deterministic instantiation of the randomized allocation.⁶

In Definition 7, the first part of our oracle access (sampling from the conditional) is also necessary in [CD22]. The second part is new to our work, and replaces a strong requirement in [CD22]. In more detail, given a type \mathbf{t}_i , Cai and Daskalakis [CD22] (as well as Brustle et al. [BCD20]) need to know if $\mathbf{t}_i \in \text{supp}(\tilde{\mathcal{D}}_{z,i})$, and, if not, need access to $\text{argmax}_{\mathbf{t}'_i \in \text{supp}(\tilde{\mathcal{D}}_{z,i})} \mathbb{E}_{\mathbf{b}_{-i} \sim \tilde{\mathcal{D}}_{z,-i}} [u_i(v_i, \hat{\mathcal{M}}(\mathbf{t}'_i, \mathbf{b}_{-i}))]$ where $\tilde{\mathcal{D}}_{z,i}$ is a rounded-down version of $\hat{\mathcal{D}}_{z,i}$. However, for arbitrary distributions and mechanisms, this task might be computationally inefficient, or simply infeasible. In our work, we need access to $\text{argmin}_{\mathbf{t}'_i \in \text{supp}(\hat{\mathcal{D}}_{z,i})} \|\mathbf{t}_i - \mathbf{t}'_i\|_p$ in the “no” case.

Given these definitions, our main theorem for this component is stated as follows.

Theorem 2. *Let $\mathcal{D} = \times_{i=1}^n \mathcal{D}_i$ be the bidders’ type distribution and $v_i : \mathbb{R}^d \times 2^{[m]} \rightarrow \mathbb{R}_+$ be a \mathcal{L} -Lipschitz valuation function for each bidder $i \in [n]$. Also, let $\mathbf{A} \in \mathbb{R}^{d \times k}$ be a design matrix and $\hat{\mathcal{D}}_z = \times_{i=1}^n \hat{\mathcal{D}}_{z,i}$, where $\hat{\mathcal{D}}_{z,i}$ is a distribution over \mathbb{R}^k for each $i \in [n]$.*

Suppose that we are given (1) query access to a mechanism $\hat{\mathcal{M}}$ that is IR and BIC w.r.t. $\hat{\mathcal{D}}_z$ and valuations $\{v_i^{\mathbf{A}}\}_{i \in [n]}$, (2) oracle access to $\hat{\mathcal{D}}_z$, and (3) any $(\varepsilon_{\text{mdl},p}, \zeta_p, p)$ -query protocol \mathcal{Q} with $\zeta_p \geq \varepsilon_{\text{mdl},p}$. Then, we can construct a mechanism \mathcal{M} that is oblivious to \mathcal{D} and $v(\cdot, \cdot)$, such that for all \mathcal{D} that satisfy $\pi_p(\mathcal{D}_i, \mathbf{A} \circ \hat{\mathcal{D}}_{z,i}) \leq \varepsilon_{\text{mdl},p}$ for all $i \in [n]$, the following hold: (1) \mathcal{M} only interacts with every bidder using \mathcal{Q} once, (2) \mathcal{M} is IR and (η, μ) -BIC w.r.t. \mathcal{D} , where $\mu = O(\sqrt{\zeta_p})$ and $\eta = O(n\|\mathbf{A}\|_\infty \mathcal{L} \sqrt{\zeta_p})$, and (3) the expected revenue of \mathcal{M} is at least $\text{Rev}(\hat{\mathcal{M}}, \hat{\mathcal{D}}_z) - O(n\eta)$.

Note that \mathcal{M} is an indirect mechanism, so it is slightly imprecise to call it (η, μ) -BIC. Formally, interacting with \mathcal{Q} truthfully is an approximate Bayes-Nash equilibrium.

In order to prove Theorem 2, we use a key lemma, Lemma 2, which establishes the robustness guarantees of Theorem 2, but in the space of latent types. Intuitively, let \mathbf{t}_i be the type of bidder i , and \mathbf{z}_i be a random variable distributed according to $\hat{\mathcal{D}}_{z,i}$. We know that $\pi_p(\mathcal{D}_i, \mathbf{A} \circ \hat{\mathcal{D}}_i) \leq \varepsilon_{\text{mdl},p} \leq \zeta_p$. Due to Lemma 1, there exists a coupling such that with probability greater than $1 - \zeta_p$, $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_p \leq \zeta_p$. Since the seller uses a $(\varepsilon_{\text{mdl},p}, \zeta_p, p)$ -query protocol, with probability at least $1 - \zeta_p$, $\|\mathcal{Q}(\mathbf{t}_i) - \mathbf{z}_i\|_p \leq \zeta_p$. Note that, this implies that \mathbf{z}_i and $\mathcal{Q}(\mathbf{t}_i)$ are distributed such that their Prokhorov distance is at most ζ_p . At this step, Lemma 2 provides us a mechanism $\tilde{\mathcal{M}}$, constructed from $\hat{\mathcal{M}}$, that we can execute on types $\mathcal{Q}(\mathbf{t}_1), \dots, \mathcal{Q}(\mathbf{t}_n)$, obtained by interacting with the bidders via the query protocol. With probability at least $1 - \zeta_p$, we have that $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_\infty \leq \|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_p \leq \varepsilon_{\text{mdl},p}$ and thus, using the fact that the query protocol ensures $\|\mathcal{Q}(\mathbf{t}_i) - \mathbf{z}_i\|_p \leq \zeta_p$ as well, we have $\|\mathbf{t}_i - \mathbf{A}\mathcal{Q}(\mathbf{t}_i)\|_\infty \leq \varepsilon_{\text{mdl},p} + k\|\mathbf{A}\|_\infty \zeta_p$. The guarantees of $\tilde{\mathcal{M}}$ for the distribution over $\mathcal{Q}(\mathbf{t}_i)$ s are therefore translated into guarantees (with a small error) of the overall mechanism for the \mathcal{D} .

The proof of Lemma 2 is quite involved and is the main focus of our analysis. Here, we sketch the key ideas behind the proof, and defer all formal arguments to Appendix C. The proof uses the following notion of a rounded distribution.

Definition 8 (Rounded Distribution). *Let \mathcal{F} be a distribution supported on $\mathbb{R}_{\geq 0}^k$. For any $\delta > 0$ and $\ell \in [0, \delta]^k$, we define function $r^{(\ell, \delta)} : \mathbb{R}_{\geq 0}^k \mapsto \mathbb{R}_{\geq 0}^k$ such that $r_i^{(\ell, \delta)}(\mathbf{x}) = \max \left\{ \lfloor \frac{x_i - \ell_i}{\delta} \rfloor \cdot \delta + \ell_i, 0 \right\}$ for all $i \in [k]$. Let \mathbf{x} be a random vector sampled from \mathcal{F} . We define $\lfloor \mathcal{F} \rfloor_{\ell, \delta}$ as the distribution for the random variable $r^{(\ell, \delta)}(\mathbf{x})$, and we call $\lfloor \mathcal{F} \rfloor_{\ell, \delta}$ the rounded distribution of \mathcal{F} .*

We follow an approach similar to Brustle et al. [BCD20]. The main idea is that arguing directly about mechanisms for distributions that are close in Prokhorov distance is difficult. On the flip side, arguing about mechanisms for distributions that are close in total variation distance is much easier, since the total variation is a more stringent (and hence more tamable) notion of distance. The key observation is that, if two distributions are close in Prokhorov distance then, in expectation over the random parameter ℓ , their rounded-down versions are close in total variation distance.

⁶That is, if, for example, $\hat{\mathcal{M}}$ allocates item j to bidder i with probability 1/2 and with the remaining probability item j is not allocated, our construction does not need to know this distribution/fractional allocation and works even if nature samples and returns an integral allocation for item j .

Our overall construction is via three reductions. First, in Lemma 3, given a mechanism for $\hat{\mathcal{F}}_z$ we design a mechanism for the rounded-down version. Second, in Lemma 4, given a mechanism for the rounded-down $\hat{\mathcal{F}}_z$ we design a mechanism for $\lfloor \mathcal{F}_z \rfloor_{\ell, \delta}$, which maintains its guarantees if $\pi_p(\mathcal{F}_z, \hat{\mathcal{F}}_z)$ is small. Third, in Lemma 5, given a mechanism for $\lfloor \mathcal{F}_z \rfloor_{\ell, \delta}$ we design a mechanism for \mathcal{F}_z .

Our proofs for Lemmas 3 and 5 are adaptations of the corresponding lemmas of [BCD20], where our main task is to flesh out the exact dependence on the dimensionality of the latent space and the ℓ_p -norm (versus the ℓ_1 -norm in [BCD20]). The novelty of our approach comes in the construction and analysis of Lemma 4. The difficulty of transforming mechanisms for $\lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}$ into mechanisms for $\lfloor \mathcal{F}_z \rfloor_{\ell, \delta}$ is that the two distributions might not share the same support. Thus, we need a way to map bids that are not in the support of $\lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}$ to bids that are. Brustle et al. [BCD20] do this by “optimally misreporting” on behalf of the bidder, by calculating $\arg\max_{\mathbf{z} \in \text{supp}(\lfloor \hat{\mathcal{F}}_{z,i} \rfloor_{\ell, \delta})} \mathbb{E}_{b_{-i} \sim \lfloor (\hat{\mathcal{F}}_{z,i})_{-i} \rfloor_{\ell, \delta}} [u_i(v_i, \hat{\mathcal{M}}(\mathbf{z}, \mathbf{b}_{-i}))]$, and then picking matching payments that make the overall mechanism IR. Our approach leverages the fact that $\hat{\mathcal{F}}_{z,i}$ and $\mathcal{F}_{z,i}$ are close in Prokhorov distance, and thus any point on the support of one distribution is close to a point on the support of the other, with high probability. An ideal construction would map a report \mathbf{w}_i to the “valid” report (i.e., a report in the support of $\lfloor \hat{\mathcal{F}}_{z,i} \rfloor_{\ell, \delta}$) that minimizes the ℓ_p distance to \mathbf{w}_i . This operation is linear on the support of $\lfloor \hat{\mathcal{F}}_{z,i} \rfloor_{\ell, \delta}$, and does not need any information on the valuation functions, nor on the actual probabilities that elements of the distribution are sampled with. Unfortunately, our assumption on what “oracle access” means does not allow us to do this operation (finding the closest point w.r.t. ℓ_p) on $\lfloor \hat{\mathcal{F}}_{z,i} \rfloor_{\ell, \delta}$, but only on $\hat{\mathcal{F}}_{z,i}$; we prove that, by incurring a small loss, our assumption suffices.

3.3 Putting everything together

Combining Theorems 1 and 2, we can give mechanisms for concrete settings. Formally, we have the following theorem.

Theorem 3. *Under the same setting as in Theorem 2, for bidders with arbitrary valuation functions, we can construct mechanism \mathcal{M} using only query access to the mechanism $\hat{\mathcal{M}}$ (Definition 6) and oracle access to distribution $\hat{\mathcal{D}}$ (Definition 7), and oblivious to the true type distribution \mathcal{D} . We consider queries (to each bidder i) of the form “What is your value for the subset of items S ?”*

Mechanism \mathcal{M} is IR and (η, μ) -BIC w.r.t. \mathcal{D} , where $\mu = O(\sqrt{\zeta_p})$ and $\eta = O(n \|\mathbf{A}\|_\infty \mathcal{L} \sqrt{\zeta_p})$, and the expected revenue of \mathcal{M} is at least $\text{Rev}(\hat{\mathcal{M}}, \hat{\mathcal{D}}_z) - O(n\eta)$. Additionally, with probability at least $1 - \delta$,

$$\zeta_p = c_p \cdot (\sigma_{\min, p}(\mathbf{A}))^{-1} \cdot \varepsilon_{\text{mdl}, p}$$

for a small constant c_p that depends on the parameter p (see footnote 5). The number of queries is $O(k \cdot \ln k \cdot \ln(n/\delta))$ and $O(k^{p/2} \cdot \ln^3 k \cdot \ln(n/\delta))$ for $p = 1, 2$ and $p \geq 3$, respectively.

The proof of Theorem 3 follows from Theorem 2 and Theorem 1, and is deferred to Appendix B.

As we’ve already discussed, the main mechanism of Cai and Daskalakis [CD22] requires bidders to have constrained-additive valuations⁷, as well \mathbf{A} to satisfy a number of restrictions. Here, we completely remove both conditions. On the flip-side, [CD22] ask bidders weaker queries, of the form “are you willing to pay price τ for item j ?” Using such queries, one can binary search over τ , and drive down the query noise (see Definition 4). For ℓ_∞ , the extra cost of such an operation would be $\ln(\|\mathbf{A}\|_\infty/\varepsilon)$, where ε is the desired accuracy. However, for other p -norms, for the same target accuracy, this operation requires an extra factor of $\Theta(\log(d^{1/p}))$ queries, giving a dependence on d .

4 Conclusions and Future Work

In this paper, we study mechanism design for prior distributions close to a topic model, inspired from the recommender systems literature. We formulate connections between mechanism design and Randomized Linear Algebra for active learning for regression problems, importing state-of-the-art

⁷A valuation function is constrained-additive if $v(t, S) = \max_{T \in \mathcal{I} \cap 2^S} \sum_{j \in T} t_j$, where \mathcal{I} is a downward-closed set system.

results from Randomized Linear Algebra to mechanism design, and alleviate or relax restrictive assumptions of prior work. Developing a deeper understanding of such connections is an important direction for future research. For example, one could study this and other topic models in the context of mechanism design for correlated bidders, two-sided markets, information structure design, etc. Additionally, another interesting open problem would be to develop a framework for proving lower bounds for mechanism design (e.g., lower bounds on the query complexity for single-round or multi-round protocols used to communicate with the bidders) using known limitations of algorithms in active learning, and vice-versa.

References

- [AFH⁺12] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi-to single-agent reduction. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 17–17, 2012.
- [BCD20] Johannes Brustle, Yang Cai, and Constantinos Daskalakis. Multi-item mechanisms without item-independence: Learnability via robustness. In *Proceedings of the 21st ACM Conference on Economics and Computation, EC '20*, page 715–761, New York, NY, USA, 2020. Association for Computing Machinery.
- [BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015.
- [BILW20] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)*, 67(4):1–40, 2020.
- [CD17] Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 516–527. IEEE, 2017.
- [CD21] Xue Chen and Michal Dereziński. Query complexity of least absolute deviation regression via robust uniform convergence. In *Conference on Learning Theory*, pages 1144–1179. PMLR, 2021.
- [CD22] Yang Cai and Constantinos Daskalakis. Recommender systems meet mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation, EC '22*, page 897–914, New York, NY, USA, 2022. Association for Computing Machinery.
- [CDW12a] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 459–478, 2012.
- [CDW12b] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 130–139. IEEE, 2012.
- [CDW13a] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 578–595. SIAM, 2013.
- [CDW13b] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Understanding incentives: Mechanism design becomes algorithm design. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 618–627. IEEE, 2013.
- [CDW16] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 926–939, 2016.
- [CHK07] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07*, pages 243–251, New York, NY, USA, 2007. ACM.
- [CHMS10] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.
- [CM16] Shuchi Chawla and J Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 579–596. ACM, 2016.

- [CMS15] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- [CP15] Michael B Cohen and Richard Peng. ℓ_p Row Sampling by Lewis Weights. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 183–192, 2015.
- [CP19] Xue Chen and Eric Price. Active regression via linear-sample sparsification. In *Conference on Learning Theory*, pages 663–695. PMLR, 2019.
- [CR14] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 243–252, 2014.
- [CZ17] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 170–183. ACM, 2017.
- [Das15] Constantinos Daskalakis. Multi-item auctions defying intuition? *ACM SIGecom Exchanges*, 14(1):41–75, 2015.
- [DDT13] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 269–286, 2013.
- [DDT15] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 449–450, 2015.
- [DHP16] Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 426–439, 2016.
- [DM16] Petros Drineas and Michael W Mahoney. RandNLA: Randomized Numerical Linear Algebra. *Communications of the ACM*, 59(6):80–90, 2016.
- [DM21] Michał Dereziński and Michael W Mahoney. Determinantal point processes in randomized numerical linear algebra. *Notices of the American Mathematical Society*, 68(1):34–45, 2021.
- [DMMS11] Petros Drineas, Michael W Mahoney, Shan Muthukrishnan, and Tamás Sarlós. Faster least squares approximation. *Numerische Mathematik*, 117(2):219–249, 2011.
- [DS03] David Donoho and Victoria Stodden. When does non-negative matrix factorization give a correct decomposition into parts? *Advances in neural information processing systems*, 16, 2003.
- [DW17] Michał Dereziński and Manfred KK Warmuth. Unbiased estimates for linear regression via volume sampling. *Advances in Neural Information Processing Systems*, 30, 2017.
- [DW18] Michał Dereziński and Manfred K. Warmuth. Reverse iterative volume sampling for linear regression. *Journal of Machine Learning Research*, 19(23):1–39, 2018.
- [DWH18] Michał Dereziński, Manfred K Warmuth, and Daniel Hsu. Tail bounds for volume sampled linear regression. *arXiv preprint arXiv:1802.06749*, 2018.
- [GHZ19] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673, 2019.
- [GW21] Yannai A Gonczarowski and S Matthew Weinberg. The sample complexity of up-to- ε multi-dimensional revenue maximization. *Journal of the ACM (JACM)*, 68(3):1–28, 2021.

- [HMR15] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 45–60, 2015.
- [HN19] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019.
- [HR15] Sergiu Hart and Philip J Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.
- [KW19] Robert Kleinberg and S Matthew Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113:97–115, 2019.
- [LPW09] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI, USA, first edition, 2009.
- [M. 11] M. W. Mahoney. *Randomized algorithms for matrices and data*. Foundations and Trends in Machine Learning. NOW Publishers, Boston, 2011.
- [MMWY21] Cameron Musco, Christopher Musco, David P Woodruff, and Taisuke Yasuda. Active linear regression for ℓ_p norms and beyond. *arXiv preprint arXiv:2111.04888*, 2021.
- [MMWY22] Cameron Musco, Christopher Musco, David P Woodruff, and Taisuke Yasuda. Active linear regression for ℓ_p norms and beyond. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 744–753. IEEE, 2022.
- [MR16] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Conference on Learning Theory*, pages 1298–1318. PMLR, 2016.
- [MV07] Alejandro M Manelli and Daniel R Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic theory*, 137(1):153–185, 2007.
- [Mye81] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [PSCW22] Alexandros Psomas, Ariel Schwartzman Cohenca, and S Weinberg. On infinite separations between simple and optimal mechanisms. *Advances in Neural Information Processing Systems*, 35:4818–4829, 2022.
- [RW15] Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394. ACM, 2015.
- [Str65] V. Strassen. The Existence of Probability Measures with Given Marginals. *The Annals of Mathematical Statistics*, 36(2):423 – 439, 1965.
- [Woo14] David P. Woodruff. Sketching as a Tool for Numerical Linear Algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2):1–157, 2014.
- [Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109. Society for Industrial and Applied Mathematics, 2015.

A Notation

Table 1: Notations

Symbol	Explanation
n	Number of bidders
m	Number of items
k	Dimensionality of latent space
d	Number of types
\mathcal{D}_i	Distribution of bidders i type vector
$\hat{\mathcal{D}}_i$	Induced noisy representation of \mathcal{D}_i
$\hat{\mathcal{D}}_{z,i}$	Distribution of latent vector type of bidder i
$\mathbf{t}_i \in \mathbb{R}^d$	Type vector of bidder i
$\mathbf{z}_i \in \mathbb{R}^k$	Sample from $\hat{\mathcal{D}}_{z,i}$ of bidder i
$\mathbf{A} \in \mathbb{R}^{d \times k}$	Design matrix of archetypes
$\mathcal{Q}(\mathbf{t}_i) = \tilde{\mathbf{z}}_i \in \mathbb{R}^k$	Output of query protocol under true type \mathbf{t}_i
$\epsilon_{\text{mdl},p} \in \mathbb{R}^d$	Modeling vector noise
$\epsilon_{\text{mdl},p} > 0$	Upper bound on $\ \epsilon_{\text{mdl},p}\ _p$
$\epsilon_{\text{nq},p} \in \mathbb{R}^d$	Query vector noise
$\epsilon_{\text{nq},p} \geq 0$	Upper bound on $\ \epsilon_{\text{nq},p}\ _p$
$v_i(\cdot, \cdot)$	Valuation function
\mathcal{M}	Mechanism
$x(\cdot)$	Allocation rule
$p(\cdot)$	Payment rule
$Rev(\mathcal{M}, \mathcal{D})$	Expected revenue of mechanism \mathcal{M} under type distributions \mathcal{D}
$\pi_P(\cdot, \cdot)$	Prokhorov distance of two probability measures
γ	A coupling between two probability measures
$v_i^{\mathbf{A}}(\mathbf{z}_i, S)$	$v_i(\mathbf{A}\mathbf{z}_i, S)$
\mathcal{L}	Lipschitz constant of valuation function
$\lfloor F \rfloor_{\lambda, \delta}$	Rounded down distribution
$r^{(\lambda, \delta)}(x)$	Rounding down mapping function where δ is a parameter chosen by the designer
$\text{supp}(F)$	Support of probability distribution F
ρ^ℓ	Total TV distance of rounded down distributions
$u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{t}_{-i})$	Change in utility of bidder i when bidder reports \mathbf{b}_i instead of \mathbf{t}_i and the remaining bidders bid \mathbf{t}_{-i}

B Proof of Theorem 3

Proof of Theorem 3. For arbitrary valuation functions, $d = 2^m$, and each type \mathbf{t}_i represents the valuation of bidder i for all possible bundles of items. Since we can ask queries of the form ‘‘What is your value for subset S ?’’, in Theorem 1, we have that $\epsilon_{\text{nq},p} = 0$. Thus from Theorem 1 with probability at least $1 - \delta$ we have a $\left(\epsilon_{\text{mdl},p}, \frac{c_p \epsilon_{\text{mdl},p}}{\sigma_{\min,p}(\mathbf{A})}, p\right)$ -query protocol that communicates once with the bidders and asks a total of s_p queries. Where $s_1 = s_2 = O(k \cdot \ln k \cdot \ln n/\delta)$ and $s_p = O(k^{p/2} \cdot \ln^3 k \cdot \ln n/\delta)$. Combining this observation with Theorem 2 proves the theorem. \square

C The Mechanism Design Component

The main goal of this section is to prove Theorem 2. In order to establish the theorem, we use the following key lemma, which essentially establishes the robustness guarantees described in Theorem 2, but in the space of latent types.

Lemma 2. Let $\mathbf{A} \in \mathbb{R}^{d \times k}$ be a design matrix. Suppose there exists a collection of distributions over latent types $\{\hat{\mathcal{F}}_{z,i}\}_{i \in [n]}$, where the support of each $\hat{\mathcal{F}}_{z,i}$ lies in $[0, 1]^k$, and a mechanism $\hat{\mathcal{M}}$ that is IR and BIC w.r.t. $\hat{\mathcal{F}}_z = \times_{i=1}^n \hat{\mathcal{F}}_{z,i}$ and valuations $\{v_i^{\mathbf{A}}\}_{i \in [n]}$, where each v_i is an \mathcal{L} -Lipschitz valuation. Let $\mathcal{F}_z = \times_{i=1}^n \mathcal{F}_{z,i}$ be any distribution such that $\pi_p(\mathcal{F}_{z,i}, \hat{\mathcal{F}}_{z,i}) \leq \zeta_p$ for all $i \in [n]$. Given query access to $\hat{\mathcal{M}}$ (as defined in Definition 6), and sampling access to $\hat{\mathcal{F}}_z$ (as defined in Definition 7), we can construct a mechanism $\tilde{\mathcal{M}}$ such that the following hold: (i) $\tilde{\mathcal{M}}$ is IR and (η, μ) -BIC w.r.t. \mathcal{F}_z , and (ii) the expected revenue of $\tilde{\mathcal{M}}$ is $Rev(\tilde{\mathcal{M}}, \mathcal{F}_z) \geq Rev(\hat{\mathcal{M}}, \hat{\mathcal{F}}_z) - O(n\eta)$, where $\mu = O(\zeta_p + \delta k^{1/p})$ and $\eta = O\left(k\mathcal{L}\|\mathbf{A}\|_\infty \left(\delta + n \left(1 + \frac{k^{1-1/p}}{\delta}\right) \zeta_p\right)\right)$, for all $\delta > 0$.

The proof of Lemma 2 is quite involved and will be the main focus of the subsequent analysis. Before proving Lemma 2, we will show that it is sufficient to prove our main result, Theorem 2.

Proof of Theorem 2. Let \mathbf{t}_i be the type of bidder i , and \mathbf{z}_i be a random variable distributed according to $\hat{\mathcal{D}}_{z,i}$. We know that $\pi_p(\mathcal{D}_i, \mathbf{A} \circ \hat{\mathcal{D}}_i) \leq \varepsilon_{\text{md1},p} \leq \zeta_p$. Due to Lemma 1, there exists a coupling such that with probability greater than $1 - \zeta_p$, $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_p \leq \zeta_p$. Since the seller uses a $(\varepsilon_{\text{md1},p}, \zeta_p, p)$ -query protocol, with probability at least $1 - \zeta_p$, $\|\mathcal{Q}(\mathbf{t}_i) - \mathbf{z}_i\|_p \leq \zeta_p$. Note that, this implies that \mathbf{z}_i and $\mathcal{Q}(\mathbf{t}_i)$ are distributed such that their Prokhorov distance is at most ζ_p .

We now invoke Lemma 2 by setting $\mathcal{F}_{z,i}$ to be distributed as $\mathcal{Q}(\mathbf{t}_i)$, $\hat{\mathcal{F}}_{z,i}$ to be distributed as $\hat{\mathcal{D}}_{z,i}$, and $\delta = \sqrt{\zeta_p}$. We run the resultant mechanism, $\tilde{\mathcal{M}}$, on types $\mathcal{Q}(\mathbf{t}_1), \dots, \mathcal{Q}(\mathbf{t}_n)$, obtained by interacting with the bidders via the query protocol. With probability at least $1 - \zeta_p$, we have that $\|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_\infty \leq \|\mathbf{t}_i - \mathbf{A}\mathbf{z}_i\|_p \leq \varepsilon_{\text{md1},p}$ and thus, using the fact that the query protocol ensures $\|\mathcal{Q}(\mathbf{t}_i) - \mathbf{z}_i\|_p \leq \zeta_p$ as well, we have $\|\mathbf{t}_i - \mathbf{A}\mathcal{Q}(\mathbf{t}_i)\|_\infty \leq \varepsilon_{\text{md1},p} + k\|\mathbf{A}\|_\infty \zeta_p$. Note that $\tilde{\mathcal{M}}$ is a (η, μ) -BIC mechanism w.r.t. $\times_{i \in [n]} \mathcal{F}_{z,i}$, with $\eta = O(n\|\mathbf{A}\|_\infty \mathcal{L} \sqrt{\zeta_p})$ and $\mu = \sqrt{\zeta_p}$. Therefore, conditioned on the aforementioned scenario having probability at least $1 - \zeta_p$, with probability at least $1 - \mu$, deviating from interacting with \mathcal{Q} truthfully can increase the expected utility by at most $O(\mathcal{L}(\varepsilon_{\text{md1},p} + k\|\mathbf{A}\|_\infty \zeta_p) + \eta) = O(\eta)$, and bidders can improve their utility by deviating with probability less than $\varepsilon_{\text{md1},p} \leq \zeta_p$, i.e., with probability at most $1 - (1 - \zeta_p)(1 - \mu) = O(\zeta_p + \mu) = O(\mu)$. Hence, the overall mechanism is (η, μ) -BIC w.r.t. \mathcal{D} , as well as IR (since $\tilde{\mathcal{M}}$ is IR). The revenue guarantee is immediate. This concludes the proof. \square

C.1 Proof of Lemma 2

For the remainder of this section, we use the following notation, where all distributions of an agent i are supported on $[0, 1]^k$: $\hat{\mathcal{F}}_z = \times_{i \in [n]} \hat{\mathcal{F}}_{z,i}$, $\left[\hat{\mathcal{F}}_z\right]_{\ell, \delta} = \times_{i \in [n]} \left[\hat{\mathcal{F}}_{z,i}\right]_{\ell, \delta}$, $\lfloor \mathcal{F}_z \rfloor_{\ell, \delta} = \times_{i \in [n]} \lfloor \mathcal{F}_{z,i} \rfloor_{\ell, \delta}$, and $\mathcal{F}_z = \times_{i \in [n]} \mathcal{F}_{z,i}$. We follow an approach similar to Brustle et al. [BCD20]. The main idea is that arguing directly about mechanisms for distributions that are close in Prokhorov distance is difficult. On the flip side, arguing about mechanisms for distributions that are close in total variation distance is much easier, since the total variation is a more stringent — and hence more well-behaved — notion of distance. The key observation here is that if two distributions are close in Prokhorov distance then, in expectation, their rounded-down versions will also be close in total variation distance, where the expectation is taken over the randomness of parameter ℓ .

Our overall construction of the mechanism $\tilde{\mathcal{M}}$ (of Lemma 2) is via three reductions. For ease of presentation, we defer some proofs (and specifically, the proofs for Lemmas 3, 4 and 5, to Appendix C.2. First, in Lemma 3, given a mechanism for $\hat{\mathcal{F}}_z$ we design a mechanism for the rounded-down version.

Lemma 3. Given query access to a mechanism \mathcal{M} that is IR and BIC w.r.t. $\hat{\mathcal{F}}_z$, and a sampling algorithm \mathcal{S}_i for each $i \in [n]$, where $\mathcal{S}_i(\mathbf{x}, \delta)$ draws a sample from the conditional distribution of $\hat{\mathcal{F}}_{z,i}$ on the k -dimensional cube $\times_{j \in [k]} [x_j, x_j + \delta_j)$, we can construct a mechanism \mathcal{M}_1^ℓ , oblivious to $\hat{\mathcal{F}}$, that is IR and $O(k\|\mathbf{A}\|_\infty \mathcal{L} \delta)$ -BIC w.r.t. $\left[\hat{\mathcal{F}}_z\right]_{\ell, \delta}$. Furthermore,

$$Rev\left(\mathcal{M}_1^\ell, \left[\hat{\mathcal{F}}_z\right]_{\ell, \delta}\right) \geq Rev(\mathcal{M}, \hat{\mathcal{F}}_z) - O(kn\|\mathbf{A}\|_\infty \mathcal{L} \delta).$$

Second, in Lemma 4, given a mechanism for $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$ we design a mechanism for $\lfloor\mathcal{F}_z\rfloor_{\ell,\delta}$, which maintains its guarantees if $\pi_p(\mathcal{F}_z, \hat{\mathcal{F}}_z)$ is small.

Lemma 4. Let $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$ and $\lfloor\mathcal{F}_z\rfloor_{\ell,\delta}$ be distributions such that, for all $i \in [n]$, it holds that (1) $\left\|\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta} - \lfloor\mathcal{F}_{z,i}\rfloor_{\ell,\delta}\right\|_{TV} \leq \varepsilon_i^\ell$, and (2) $\pi_p\left(\hat{\mathcal{F}}_{z,i}, \mathcal{F}_{z,i}\right) \leq \zeta_p$. Then, letting $\rho^\ell := \sum_i \varepsilon_i^\ell$, given a mechanism \mathcal{M}_1^ℓ that is IR and η -BIC w.r.t $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$ we can construct a mechanism \mathcal{M}_2^ℓ that is IR, $(O(k\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell + k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L} + \eta), \zeta_p + \delta k^{1/p})$ -BIC w.r.t. $\lfloor\mathcal{F}_z\rfloor_{\ell,\delta}$. Furthermore:

$$Rev\left(\mathcal{M}_2^\ell, \lfloor\mathcal{F}_z\rfloor_{\ell,\delta}\right) \geq Rev\left(\mathcal{M}_1^\ell, \left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}\right) - nk\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell - nk(\zeta_p + \delta k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L}.$$

Third, in Lemma 5, given a mechanism for $\lfloor\mathcal{F}_z\rfloor_{\ell,\delta}$ we design a mechanism for \mathcal{F}_z .

Lemma 5. Given a mechanism \mathcal{M}_2^ℓ that is IR and (η, μ) -BIC w.r.t. $\lfloor\mathcal{F}_z\rfloor_{\ell,\delta}$, we can construct a mechanism \mathcal{M}^ℓ that is IR and $(3k\mathcal{L}\|\mathbf{A}\|_\infty\delta + \eta, \mu)$ -BIC w.r.t. \mathcal{F}_z . Moreover, $Rev\left(\mathcal{M}^\ell, \mathcal{F}_z\right) \geq Rev\left(\mathcal{M}_2^\ell, \lfloor\mathcal{F}_z\rfloor_{\ell,\delta}\right) - nk\mathcal{L}\|\mathbf{A}\|_\infty\delta$.

With all the prerequisite technical lemmas at hand, we finally prove Lemma 2.

Proof of Lemma 2. The mechanism $\widetilde{\mathcal{M}}$ simply operates as follows: (i) ℓ is sampled uniformly from the interval $[0, \delta]$ (i.e., $\ell \sim \mathcal{U}[0, \delta]$) and (ii) the mechanism \mathcal{M}^ℓ of Lemma 5 is run on the input bids. Specifically, $\widetilde{\mathcal{M}}$ calls mechanism \mathcal{M}^ℓ of Lemma 5, which calls mechanism \mathcal{M}_2^ℓ of Lemma 4, which calls mechanism \mathcal{M}_1^ℓ of Lemma 3, which calls $\hat{\mathcal{M}}$ which is IR and BIC w.r.t. $\hat{\mathcal{F}}$. In order to invoke Lemma 4 we need a bound on the TV distance between $\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}$ and $\lfloor\mathcal{F}_{z,i}\rfloor_{\ell,\delta}$. We use the following lemma (whose proof can be found in Appendix C.2).

Lemma 6. Let \mathcal{F}_z and $\hat{\mathcal{F}}_z$ be two distributions supported on \mathbb{R}^k such that $\pi_p(\mathcal{F}_z, \hat{\mathcal{F}}_z) \leq \varepsilon$. For any $\delta > 0$, $\mathbb{E}_{\ell \sim \mathcal{U}[0, \delta]^k} \left[\left\| \lfloor\mathcal{F}_z\rfloor_{\ell,\delta} - \left[\hat{\mathcal{F}}_z\right]_{\ell,\delta} \right\|_{TV} \right] \leq \left(1 + \frac{k^{1-1/p}}{\delta}\right) \varepsilon$.

Letting $\rho^\ell = \sum_i \varepsilon_i^\ell$, where $\left\|\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta} - \lfloor\mathcal{F}_{z,i}\rfloor_{\ell,\delta}\right\|_{TV} \leq \varepsilon_i^\ell$, Lemma 6 implies that $\mathbb{E}_{\ell \sim \mathcal{U}[0, \delta]} [\rho^\ell] \leq n \left(1 + \frac{k^{1-1/p}}{\delta}\right) \zeta_p$, allowing us to invoke Lemma 4. The IR guarantee for $\widetilde{\mathcal{M}}$ is immediate. For the BIC guarantee, there is a trade-off in choosing δ : $\mathbb{E}_{\ell \sim \mathcal{U}[0, \delta]} [\rho^\ell]$ has a term inversely proportional to δ , while the BIC guarantees of Lemma 4 and Lemma 3 have a term proportional to δ . We set $\delta = \sqrt{\zeta_p}$ to strike a balance between these terms. Combining the BIC and revenue guarantees of Lemma 3, Lemma 4, and Lemma 5, and using the fact that k is a constant, we get that $\widetilde{\mathcal{M}}$ is a (η, μ) -BIC and IR w.r.t. \mathcal{F} where $\mu = O(\zeta_p + \delta \cdot k^{1/p})$ and $\eta = O\left(k\mathcal{L}\|\mathbf{A}\|_\infty\left(\delta + n \cdot \left(1 + \frac{k^{1-1/p}}{\delta}\right) \cdot \zeta_p\right)\right)$ and $Rev(\widetilde{\mathcal{M}}, \mathcal{F}) \geq Rev(\hat{\mathcal{M}}, \hat{\mathcal{F}}) - O(n\eta)$. \square

C.2 Proofs of Lemma 3, 4, 5, and 6

The following proposition will be useful in multiple proofs throughout this section.

Proposition 1. If v_i is a \mathcal{L} -Lipschitz valuation function, then $v_i^{\mathbf{A}}$ is a $k\|\mathbf{A}\|_\infty\mathcal{L}$ -Lipschitz valuation function.

Proof. The statement follows from the fact that $v_i^{\mathbf{A}}(\mathbf{z}) = v_i(\mathbf{A}\mathbf{z})$, where $\mathbf{z} \in [0, 1]^k$ and v_i is \mathcal{L} -Lipschitz. \square

Lemma 3. Given query access to a mechanism \mathcal{M} that is IR and BIC w.r.t. $\hat{\mathcal{F}}_z$, and a sampling algorithm \mathcal{S}_i for each $i \in [n]$, where $\mathcal{S}_i(\mathbf{x}, \delta)$ draws a sample from the conditional distribution of

$\hat{\mathcal{F}}_{z,i}$ on the k -dimensional cube $\times_{j \in [k]} [x_j, x_j + \delta_j)$, we can construct a mechanism \mathcal{M}_1^ℓ , oblivious to $\hat{\mathcal{F}}$, that is IR and $O(k\|\mathbf{A}\|_\infty \mathcal{L}\delta)$ -BIC w.r.t. $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$. Furthermore,

$$\text{Rev}\left(\mathcal{M}_1^\ell, \left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}\right) \geq \text{Rev}(\mathcal{M}, \hat{\mathcal{F}}_z) - O(kn\|\mathbf{A}\|_\infty \mathcal{L}\delta).$$

Proof. Let x be the allocation rule and p be the payment rule of \mathcal{M} . We construct \mathcal{M}_1^ℓ as follows. Upon receiving bids $\{\mathbf{w}_i\}_{i \in [n]}$ we first query \mathcal{S}_i for each bidder i to get $\mathbf{w}'_i = \mathcal{S}_i(\mathbf{w}_i, \beta(\mathbf{w}_i))$ where $\beta_j(\mathbf{w}_i) = \delta$ if $w_{i,j} \neq 0$ and $\beta_j(\mathbf{w}_i) = \ell_j$ otherwise. Then, each bidder i gets allocated items $x_i(\mathbf{w}')$ and pays $\max\{0, p_i(\mathbf{w}') - k\|\mathbf{A}\|_\infty \mathcal{L}\delta\}$.

We first prove that this mechanism is IR. Note that, by definition of \mathbf{w}_i , we have $\|\mathbf{w}_i - \mathbf{w}'_i\|_\infty \leq \delta$ for each $i \in [n]$. Since \mathcal{M} is IR then we know that for all bids \mathbf{t}_{-i} of other bidders, $v_i^{\mathbf{A}}(\mathbf{w}'_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{t}_{-i})) - p_i(\mathbf{w}'_i, \mathbf{t}_{-i}) = u_i(\mathbf{w}'_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{t}_{-i})) \geq 0$. This gives us the required inequality since,

$$\begin{aligned} 0 &\leq v_i^{\mathbf{A}}(\mathbf{w}'_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{t}_{-i})) - p_i(\mathbf{w}'_i, \mathbf{t}_{-i}) \\ &= v_i^{\mathbf{A}}(\mathbf{w}'_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{t}_{-i})) - k\|\mathbf{A}\|_\infty \mathcal{L}\delta - (p_i(\mathbf{w}'_i, \mathbf{t}_{-i}) - k\|\mathbf{A}\|_\infty \mathcal{L}\delta) \\ &\leq v_i^{\mathbf{A}}(\mathbf{w}_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{t}_{-i})). \end{aligned} \quad (\text{via Proposition 1})$$

We now prove that the mechanism is $O(k\|\mathbf{A}\|_\infty \mathcal{L}\delta)$ -BIC. From the point of view of bidder i the types of the other bidders are drawn from $(\hat{\mathcal{F}}_z)_{-i}$ (i.e. $\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}$). From the fact that \mathcal{M} is BIC w.r.t. $\hat{\mathcal{F}}_z$ we have the following:

$$\begin{aligned} \mathbb{E}_{\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{w}'_{-i}))] &\geq \mathbb{E}_{\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}} [u_i(\mathbf{w}'_i, \mathcal{M}(\mathbf{w}'_i, \mathbf{w}'_{-i}))] - k\|\mathbf{A}\|_\infty \mathcal{L}\delta \\ &\geq \max_{\mathbf{x} \in \text{supp}(\hat{\mathcal{F}}_{z,i})} \mathbb{E}_{\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}} [u_i(\mathbf{w}'_i, \mathcal{M}(\mathbf{x}, \mathbf{w}'_{-i}))] - k\|\mathbf{A}\|_\infty \mathcal{L}\delta \\ &\geq \max_{\mathbf{x} \in \text{supp}(\hat{\mathcal{F}}_{z,i})} \mathbb{E}_{\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}(\mathbf{x}, \mathbf{w}'_{-i}))] - 2k\|\mathbf{A}\|_\infty \mathcal{L}\delta. \end{aligned} \quad (5)$$

The first and the third inequalities above follow from Proposition 1. Thus if i could pick the type exactly \mathbf{w}'_i as she pleased, she could not possibly make more than $2k\|\mathbf{A}\|_\infty \mathcal{L}\delta$. However she must pick a $\mathbf{b}_i \in \text{supp}\left(\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}\right)$, which gets rounded to $\mathbf{b}'_i = \mathcal{S}_i(\mathbf{b}_i, \beta(\mathbf{b}_i))$. Specifically, $\mathbf{b}'_i \sim \hat{\mathcal{F}}_{z,i} \times_{j \in [k]} [b_{i,j}, b_{i,j} + \beta_j(\mathbf{b}_i)]$; for notational simplicity, we will simply denote this as $\mathbf{b}'_i \sim \mathcal{S}_i(\mathbf{b}_i, \beta(\mathbf{b}_i))$. Bidder i 's utility when she reports bid $\mathbf{b}_i \in \text{supp}\left(\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}\right)$ can be bounded using the following observation.

$$\begin{aligned} \mathbb{E}_{\mathbf{b}'_i \sim \mathcal{S}_i(\mathbf{b}_i, \beta(\mathbf{b}_i))} [u_i(\mathbf{w}_i, \mathcal{M}(\mathbf{b}'_i, \mathbf{w}'_{-i}))] &= \mathbb{E}_{\mathbf{b}'_i \sim \mathcal{S}_i(\mathbf{b}_i, \beta(\mathbf{b}_i))} \left[\mathbb{E}_{\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}(\mathbf{b}'_i, \mathbf{w}'_{-i}))] \right] \\ &\leq \max_{\mathbf{x} \in \text{supp}(\hat{\mathcal{F}}_{z,i})} \mathbb{E}_{\mathbf{w}'_{-i} \sim (\hat{\mathcal{F}}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}(\mathbf{x}, \mathbf{w}'_{-i}))]. \end{aligned} \quad (6)$$

On combining inequalities (5) and (6), we get that our new mechanism is $2k\|\mathbf{A}\|_\infty \mathcal{L}\delta$ -BIC w.r.t. $\times_{i \in [m]} \left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}$. Finally, note that under truthful bidding this mechanism extracts revenue at least $\text{Rev}(\mathcal{M}, \hat{\mathcal{F}}_z) - O(kn\|\mathbf{A}\|_\infty \mathcal{L}\delta)$ since \mathbf{w}' is essentially drawn from $\hat{\mathcal{F}}_z$ and each bidder gets a discount of $k\|\mathbf{A}\|_\infty \mathcal{L}\delta$. \square

Lemma 4. Let $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$ and $\left[\mathcal{F}_z\right]_{\ell,\delta}$ be distributions such that, for all $i \in [n]$, it holds that (1) $\left\|\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta} - \left[\mathcal{F}_{z,i}\right]_{\ell,\delta}\right\|_{TV} \leq \varepsilon_i^\ell$, and (2) $\pi_p\left(\hat{\mathcal{F}}_{z,i}, \mathcal{F}_{z,i}\right) \leq \zeta_p$. Then, letting $\rho^\ell := \sum_i \varepsilon_i^\ell$, given a mechanism \mathcal{M}_1^ℓ that is IR and η -BIC w.r.t $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$ we can construct a mechanism \mathcal{M}_2^ℓ that is IR, $(O(k\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell + k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L} + \eta), \zeta_p + \delta k^{1/p})$ -BIC w.r.t. $\left[\mathcal{F}_z\right]_{\ell,\delta}$. Furthermore:

$$\text{Rev}\left(\mathcal{M}_2^\ell, \left[\mathcal{F}_z\right]_{\ell,\delta}\right) \geq \text{Rev}\left(\mathcal{M}_1^\ell, \left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}\right) - nk\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell - nk(\zeta_p + \delta k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L}.$$

Proof. We will construct \mathcal{M}_2^ℓ as follows. For every input bid $\mathbf{w} \in \text{supp}\left(\times_{i \in [n]} \left[\mathcal{F}_{z,i}\right]_{\ell,\delta}\right)$, we first find, for each $i \in [n]$, the closest point in ℓ_p norm distance that is in $\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}$; let $\mathbf{w}'_i = \underset{\mathbf{x} \in \text{supp}\left(\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}\right)}{\text{argmin}} \|\mathbf{w} - \mathbf{x}\|_p$ be this point. Notice here, that we assume that we can calculate the above expression exactly. However, according to Definition 7 we can actually compute $\hat{\mathbf{w}}_i = \underset{\mathbf{x} \in \text{supp}\left(\hat{\mathcal{F}}_{z,i}\right)}{\text{argmin}} \|\mathbf{w} - \mathbf{x}\|_p$. For the sake of simplicity, we continue the analysis as if we could calculate the desired expression. However, we can set $\mathbf{w}'_i = r^{(\ell,\delta)}(\hat{\mathbf{w}}_i)$ (as defined in Definition 8) and the following proposition implies that we only lose a small factor of $2\delta k^{\frac{1+p}{p}}\|\mathbf{A}\|_\infty\mathcal{L}$ in the BIC guarantee and a $2n\delta k^{\frac{1+p}{p}}\|\mathbf{A}\|_\infty\mathcal{L}$ factor in the revenue guarantee. In the following proposition and for the rest of our analysis we use the notation $\|\mathbf{x} - \mathcal{A}\|_p$ to denote the distance of a vector \mathbf{x} to the closest vector in a set of vectors \mathcal{A} , i.e., $\|\mathbf{x} - \mathcal{A}\|_p := \min_{\mathbf{y} \in \mathcal{A}} \|\mathbf{x} - \mathbf{y}\|_p$.

Proposition 2. Let \mathcal{B} be a probability distributions supported on $[0, 1]^k$. Then for any \mathbf{x} and $\mathbf{w} = \underset{\mathbf{z} \in \text{supp}(\mathcal{B})}{\text{argmin}} \|\mathbf{x} - \mathbf{z}\|_p$ we have that $\|\mathbf{x} - r^{(\ell,\delta)}(\mathbf{w})\|_p \leq 2k^{1/p}\delta + \|\mathbf{x} - \text{supp}\left(\left[\mathcal{B}\right]_{\ell,\delta}\right)\|_p$.

Proof. For any \mathbf{x} we have that $\|\mathbf{x} - r^{(\ell,\delta)}(\mathbf{x})\|_p \leq k^{1/p}\delta$. Then, by using the triangle inequality for the ℓ_p -norm and chaining the resulting inequalities, the proposition is implied. \square

After finding \mathbf{w}' , we then run \mathcal{M}_1^ℓ on \mathbf{w}' giving a discount to each bidder, and at the same time making sure that the IR constraint is not violated. Let $x(\cdot)$ and $p(\cdot)$ be the allocation and payment rules of \mathcal{M}_1^ℓ . For each $i \in [n]$ we do the following: if $\|w_i - w'_i\|_p \leq (\zeta_p + \delta \cdot k^{1/p})$, bidder i will receive allocation $x_i(\mathbf{w}')$ and pay $\max\{\hat{p}_i(\mathbf{w}'), 0\}$ where $\hat{p}_i(\mathbf{w}') := p_i(\mathbf{w}') - k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L}$. Otherwise, if $\|w_i - w'_i\|_p > (\zeta_p + \delta \cdot k^{1/p})$, she will receive nothing and pays nothing. By construction, and using Proposition 1, mechanism \mathcal{M}_2^ℓ is IR.

Next, we show that \mathcal{M}_2^ℓ is $(O(k\mathcal{L}\|\mathbf{A}\|_\infty\zeta_p + k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L} + \eta), \zeta_p + \delta k^{1/p})$ -BIC w.r.t. $\left[\mathcal{F}_z\right]_{\ell,\delta}$. As a first step, we prove that \mathcal{M}_2^ℓ is $O(k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L} + \eta)$ -BIC w.r.t. $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$.

Recall that \mathcal{M}_1^ℓ is η -BIC w.r.t. to $\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}$, i.e., for all bidders i , latent types $\mathbf{t}_i, \mathbf{t}'_i$ we have,

$$\underset{\mathbf{b}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}}{\mathbb{E}} \left[u_i(\mathbf{t}_i, \mathcal{M}_1^\ell(\mathbf{t}_i, \mathbf{b}_{-i})) \right] \geq \underset{\mathbf{b}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}}{\mathbb{E}} \left[u_i(\mathbf{t}_i, \mathcal{M}_1^\ell(\mathbf{t}'_i, \mathbf{b}_{-i})) \right] - \eta.$$

This further implies that

$$\begin{aligned} & \underset{\mathbf{b}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}}{\mathbb{E}} \left[u_i(\mathbf{t}_i, \mathcal{M}_1^\ell(\mathbf{t}_i, \mathbf{b}_{-i})) + k\left(\zeta_p + \delta \cdot k^{1/p}\right)\|\mathbf{A}\|_\infty\mathcal{L} \right] \\ & \geq \underset{\mathbf{b}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}}{\mathbb{E}} \left[u_i(\mathbf{t}_i, \mathcal{M}_1^\ell(\mathbf{t}'_i, \mathbf{b}_{-i})) + k\left(\zeta_p + \delta \cdot k^{1/p}\right)\|\mathbf{A}\|_\infty\mathcal{L} \right] - \eta \\ & = \underset{\mathbf{b}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}}{\mathbb{E}} \left[v_i^{\mathbf{A}}(\mathbf{t}_i, x_i(\mathbf{t}'_i, \mathbf{b}_{-i})) - p_i(\mathbf{t}'_i, \mathbf{b}_{-i}) + k\left(\zeta_p + \delta \cdot k^{1/p}\right)\|\mathbf{A}\|_\infty\mathcal{L} \right] - \eta \\ & > \underset{\mathbf{b}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}}{\mathbb{E}} \left[v_i^{\mathbf{A}}(\mathbf{t}_i, x_i(\mathbf{t}'_i, \mathbf{b}_{-i})) - \max\{\hat{p}_i(\mathbf{t}'_i, \mathbf{b}_{-i}), 0\} \right] - \eta. \end{aligned} \tag{7}$$

Additionally, using the fact that $x = \max\{x, 0\} + \min\{x, 0\}$ for all $x \in \mathbb{R}$, we can upper bound the left-hand side to obtain,

$$\begin{aligned}
& \mathbb{E}_{\mathbf{b}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} \left[v_i^{\mathbf{A}}(\mathbf{t}_i, x_i(\mathbf{t}_i, \mathbf{b}_{-i})) - p_i(\mathbf{t}_i, \mathbf{b}_{-i}) + k \left(\zeta_p + \delta \cdot k^{1/p} \right) \|\mathbf{A}\|_{\infty} \mathcal{L} \right] \\
&= \mathbb{E}_{\mathbf{b}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} \left[v_i^{\mathbf{A}}(\mathbf{t}_i, x_i(\mathbf{t}_i, \mathbf{b}_{-i})) - \max\{\hat{p}_i(\mathbf{t}'_i, \mathbf{b}_{-i}), 0\} - \min\{\hat{p}_i(\mathbf{t}'_i, \mathbf{b}_{-i}), 0\} \right] \\
&\leq \mathbb{E}_{\mathbf{b}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} \left[v_i^{\mathbf{A}}(\mathbf{t}_i, x_i(\mathbf{t}_i, \mathbf{b}_{-i})) - \max\{\hat{p}_i(\mathbf{t}'_i, \mathbf{b}_{-i}), 0\} \right] + k \left(\zeta_p + \delta \cdot k^{1/p} \right) \|\mathbf{A}\|_{\infty} \mathcal{L}.
\end{aligned} \tag{8}$$

Combining inequalities (7) and (8), gives us the required inequality for the first step, i.e. that \mathcal{M}_2^ℓ is $O(k(\zeta_p + \delta \cdot k^{1/p}) \|\mathbf{A}\|_{\infty} \mathcal{L} + \eta)$ -BIC w.r.t. $\lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}$.

For the second step, we define, for types \mathbf{t}_i , \mathbf{t}_{-i} , and \mathbf{b}_i , the following auxiliary function $u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{t}_{-i}) := u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{t}_i, \mathbf{t}_{-i})) - u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{t}_{-i}))$. In other words $u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{t}_{-i})$ simply represents the difference in utility of i when he reports his true type \mathbf{t}_i as compared to \mathbf{b}_i to the mechanism \mathcal{M}_2^ℓ . Due to the Lipschitz continuity of valuation functions and the fact that \mathcal{M}_2^ℓ is IR, we get that for all choices of $i, \mathbf{t}_i, \mathbf{b}_i, \mathbf{t}_{-i}$ we must have $u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{t}_{-i})) \in [-k\mathcal{L}\|\mathbf{A}\|_{\infty}, k\mathcal{L}\|\mathbf{A}\|_{\infty}]$. Specifically, this follows from the following two observations: first, $v_i^{\mathbf{A}}(\mathbf{0}, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{t}_{-i})) = 0$, thus due to Lipschitz continuity of $v_i^{\mathbf{A}}$ and the fact that $\mathbf{t}_i \in [0, 1]^k$, we have $v_i^{\mathbf{A}}(\mathbf{t}, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{t}_{-i})) \leq k\mathcal{L}\|\mathbf{A}\|_{\infty}$, and second, that the payments in \mathcal{M}_2^ℓ are upper bounded by the maximum utility since \mathcal{M}_2^ℓ is IR. Therefore, $-2k\mathcal{L}\|\mathbf{A}\|_{\infty} \leq u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{t}_{-i}) \leq 2k\mathcal{L}\|\mathbf{A}\|_{\infty}$, or equivalently, for all $\mathbf{x}_{-i}, \mathbf{y}_{-i}$ we have that $\mathbb{E}[u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i})] - \mathbb{E}[u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{y}_{-i})] \leq 4k\mathcal{L}\|\mathbf{A}\|_{\infty} \mathbb{1}\{\mathbf{x}_{-i} \neq \mathbf{y}_{-i}\}$, where the expectation is taken over the randomness of the mechanism.

This implies that for any coupling γ of $\lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}, \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}$ — and in particular for the coupling $\gamma^* = \operatorname{argmin}_{\gamma} \mathbb{E}_{(\mathbf{x}_{-i}, \mathbf{y}_{-i}) \sim \gamma} [\mathbb{1}\{\mathbf{x}_{-i} \neq \mathbf{y}_{-i}\}]$ — we have that,

$$\begin{aligned}
\mathbb{E}_{(\mathbf{x}_{-i}, \mathbf{y}_{-i}) \sim \gamma^*} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i}) - u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{y}_{-i})] &\leq 4k\mathcal{L}\|\mathbf{A}\|_{\infty} \mathbb{E}_{(\mathbf{x}_{-i}, \mathbf{y}_{-i}) \sim \gamma^*} [\mathbb{1}\{\mathbf{x}_{-i} \neq \mathbf{y}_{-i}\}] \\
&\leq 4k\mathcal{L}\|\mathbf{A}\|_{\infty} \left\| \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta} - \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta} \right\|_{TV} \\
&\leq 4k\mathcal{L}\|\mathbf{A}\|_{\infty} \rho^\ell,
\end{aligned}$$

where the final inequality follows from the fact that $\left\| \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta} - \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta} \right\|_{TV} \leq \left\| \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta} - \lfloor \mathcal{F}_z \rfloor_{\ell, \delta} \right\|_{TV} \leq \sum_i \varepsilon_i^\ell = \rho^\ell$. Note that the left hand side can be simplified as $\mathbb{E}_{(\mathbf{x}_{-i}, \mathbf{y}_{-i}) \sim \gamma^*} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i}) - u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{y}_{-i})] = \mathbb{E}_{\mathbf{x}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i})] - \mathbb{E}_{\mathbf{y}_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{y}_{-i})]$, and therefore, rearranging give us,

$$\begin{aligned}
& \mathbb{E}_{\mathbf{y}_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{t}_i, \mathbf{y}_{-i}))] - \mathbb{E}_{\mathbf{y}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{y}_{-i}))] \\
&= \mathbb{E}_{\mathbf{y}_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{y}_{-i})] \\
&\geq \mathbb{E}_{\mathbf{x}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i})] - 4k\mathcal{L}\|\mathbf{A}\|_{\infty} \rho^\ell.
\end{aligned} \tag{9}$$

As the final step of establishing the BIC guarantee of \mathcal{M}_2^ℓ , we lower bound $\mathbb{E}_{\mathbf{x}_{-i} \sim \lfloor (\hat{\mathcal{F}}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i})]$. Note that \mathcal{M}_2^ℓ maps the reported bids to the closest point in $\operatorname{supp} \left(\lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta} \right)$, and hence, without loss of generality we can assume that when bidders misreport they choose a bid $\mathbf{b}_i \in \operatorname{supp} \left(\lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta} \right)$. Letting $\mathbf{t}_i^* = \operatorname{argmin}_{\mathbf{x} \in \operatorname{supp}(\lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta})} \|\mathbf{t}_i - \mathbf{x}\|_p$

and $\lambda = \|\mathbf{t}_i - \text{supp}\left(\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}\right)\|_p$, we have that,

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{x}_{-i})] &= \mathbb{E}_{\mathbf{x}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{t}_i, \mathbf{x}_{-i})) - u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{x}_{-i}))] \\ &= \mathbb{E}_{\mathbf{x}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{t}_i^*, \mathbf{x}_{-i})) - u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{x}_{-i}))] \\ &\geq \mathbb{E}_{\mathbf{x}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i^*, \mathcal{M}_2^\ell(\mathbf{t}_i^*, \mathbf{x}_{-i})) - u_i(\mathbf{t}_i^*, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{x}_{-i}))] - 2k\mathcal{L}\|\mathbf{A}\|_\infty\lambda. \end{aligned}$$

To complete the argument, we use the following two propositions.

Proposition 3. *Let \mathcal{B} and \mathcal{B}' be two probability distributions supported on $[0, 1]^k$ for which $\pi_p(\mathcal{B}, \mathcal{B}') \leq \varepsilon$. Then it must be true that $\pi_p\left(\left[\mathcal{B}\right]_{\ell,\delta}, \left[\mathcal{B}'\right]_{\ell,\delta}\right) \leq \varepsilon + \delta \cdot k^{1/p}$.*

Proof. Intuitively, the proposition follows since two probability masses, when rounded, will move at most by an additive factor of $\delta \cdot k^{1/p}$ in the ℓ_p norm distance compared to each other. Formally, since $\pi_p(\mathcal{B}, \mathcal{B}') \leq \varepsilon$ from Lemma 1 we know that there exist coupling γ of $\mathcal{B}, \mathcal{B}'$ such that $\mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p > \varepsilon] \leq \varepsilon$. From the definition of rounding, we know that $\|r^{(\ell,\delta)}(x) - r^{(\ell,\delta)}(y)\|_p \leq \|x - y + \vec{\delta}\|_p \leq \|x - y\|_p + \delta \|\vec{1}\|_p = \|x - y\|_p + \delta \cdot k^{1/p}$. Therefore, for the same coupling γ we have that,

$$\begin{aligned} \mathbb{P}_{(x,y) \sim \gamma} [\|r^{(\ell,\delta)}(x) - r^{(\ell,\delta)}(y)\|_p > \varepsilon + \delta \cdot k^{1/p}] &\leq \mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p + \delta \cdot k^{1/p} > \varepsilon + \delta \cdot k^{1/p}] \\ &\leq \varepsilon \leq \varepsilon + \delta \cdot k^{1/p}. \end{aligned}$$

□

Proposition 4. *Let \mathcal{B} and \mathcal{B}' be two probability distributions supported on $[0, 1]^k$ for which $\pi_p(\mathcal{B}, \mathcal{B}') \leq \varepsilon$. Then it must be true that $\mathbb{P}_{x \sim \mathcal{B}} [\|x - \text{supp}(\mathcal{B}')\|_p > \varepsilon] \leq \varepsilon$.*

Proof. Using Lemma 1 we know that there exist coupling γ of $\mathcal{B}, \mathcal{B}'$ such that $\mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p > \varepsilon] \leq \varepsilon$. This directly gives us the required inequality, $\mathbb{P}_{x \sim \mathcal{B}} [\|x - \text{supp}(\mathcal{B}')\|_p > \varepsilon] \leq \mathbb{P}_{x \sim \mathcal{B}, y \sim \gamma | x} [\|x - y\|_p > \varepsilon] = \mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p > \varepsilon] \leq \varepsilon$. □

First, as previously shown, our mechanism \mathcal{M}_2^ℓ is a $k(\zeta_p + \delta \cdot k^{1/p}) \|\mathbf{A}\|_\infty \mathcal{L} + \eta$ -BIC w.r.t. $\times_{i \in [n]} \left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}$. Therefore, for every i, \mathbf{t}_i , and \mathbf{b}_i we have $\mathbb{E}_{\mathbf{t}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{t}_{-i})] \geq -k(\zeta_p + \delta \cdot k^{1/p}) \|\mathbf{A}\|_\infty \mathcal{L} - \eta$. Second, we know that $\pi_p\left(\left[\hat{\mathcal{F}}_{z,i}, \mathcal{F}_{z,i}\right]_{\ell,\delta}\right) \leq \zeta_p$, and hence via Proposition 3, we have $\pi_p\left(\left[\hat{\mathcal{F}}_{z,i}\right]_{\ell,\delta}, \left[\mathcal{F}_{z,i}\right]_{\ell,\delta}\right) \leq \zeta_p + \delta \cdot k^{1/p}$. Hence, by invoking Proposition 4, we get that $\lambda = \|\mathbf{t}_i - \text{supp}\left(\left[\hat{\mathcal{F}}_z\right]_{\ell,\delta}\right)\|_p \leq \zeta_p + \delta \cdot k^{1/p}$ with probability at least $1 - \zeta_p - \delta \cdot k^{1/p}$. Using these two observations we get that,

$$\begin{aligned} \mathbb{E}_{\mathbf{t}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i \leftarrow \mathbf{b}_i, \mathbf{t}_{-i})] &\geq -k(\zeta_p + \delta \cdot k^{1/p}) \|\mathbf{A}\|_\infty \mathcal{L} - \eta - 2k\lambda \|\mathbf{A}\|_\infty \mathcal{L} \\ &\geq -3k(\zeta_p + \delta \cdot k^{1/p}) \|\mathbf{A}\|_\infty \mathcal{L} - \eta. \end{aligned}$$

On combining this with Eq. (9), we get that with probability at least $1 - \zeta_p - \delta \cdot k^{1/p}$,

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_{-i} \sim \left[\left(\hat{\mathcal{F}}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{t}_i, \mathbf{x}_{-i}))] &- \mathbb{E}_{\mathbf{x}_{-i} \sim \left[\left(\mathcal{F}_z\right)_{-i}\right]_{\ell,\delta}} [u_i(\mathbf{t}_i, \mathcal{M}_2^\ell(\mathbf{b}_i, \mathbf{x}_{-i}))] \\ &\geq -3k(\zeta_p + \delta \cdot k^{1/p}) \|\mathbf{A}\|_\infty \mathcal{L} - k\mathcal{L}\|\mathbf{A}\|_\infty \rho^\ell - \eta. \end{aligned}$$

That is, \mathcal{M}_2^ℓ is $(O(k\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell + k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L} + \eta), \zeta_p + \delta k^{1/p})$ -BIC w.r.t. $\lfloor \mathcal{F}_z \rfloor_{\ell, \delta}$. Finally, we prove the revenue guarantee of \mathcal{M}_2^ℓ . Since our mechanism operates by offering a discount of $k(\zeta_p + \delta \cdot k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L}$ to each bidder, we have $Rev\left(\mathcal{M}_2^\ell, \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}\right) \geq Rev\left(\mathcal{M}_1^\ell, \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}\right) - nk(\zeta_p + \delta k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L}$. Next, we consider the following two cases. Let $\mathbf{t} \sim \lfloor \mathcal{F}_z \rfloor_{\ell, \delta}$ and $\hat{\mathbf{t}} \sim \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}$. If $\mathbf{t} = \hat{\mathbf{t}}$, mechanism \mathcal{M}_2^ℓ will produce the same revenue. Otherwise, if $\mathbf{t} \neq \hat{\mathbf{t}}$ one instance can extract at most $nk\mathcal{L}\|\mathbf{A}\|_\infty$ more revenue than the other. Since $\|\lfloor \hat{\mathcal{F}} \rfloor_{\ell, \delta} - \lfloor \mathcal{F} \rfloor_{\ell, \delta}\|_{TV} \leq \sum_i \varepsilon_i^\ell = \rho^\ell$, from the characterization of TV distance [LPW09]⁸ we know that there exists a coupling such that $\mathbf{t} \neq \hat{\mathbf{t}}$ with probability less than ρ^ℓ . This gives us the desired bound on revenue:

$$\begin{aligned} Rev\left(\mathcal{M}_2^\ell, \lfloor \mathcal{F}_z \rfloor_{\ell, \delta}\right) &\geq Rev\left(\mathcal{M}_2^\ell, \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}\right) - nk\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell \\ &\geq Rev\left(\mathcal{M}_1^\ell, \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta}\right) - nk\mathcal{L}\|\mathbf{A}\|_\infty\rho^\ell - nk(\zeta_p + \delta k^{1/p})\|\mathbf{A}\|_\infty\mathcal{L}. \quad \square \end{aligned}$$

Lemma 5. *Given a mechanism \mathcal{M}_2^ℓ that is IR and (η, μ) -BIC w.r.t. $\lfloor \mathcal{F}_z \rfloor_{\ell, \delta}$, we can construct a mechanism \mathcal{M}^ℓ that is IR and $(3k\mathcal{L}\|\mathbf{A}\|_\infty\delta + \eta, \mu)$ -BIC w.r.t. \mathcal{F}_z . Moreover, $Rev(\mathcal{M}^\ell, \mathcal{F}_z) \geq Rev(\mathcal{M}_2^\ell, \lfloor \mathcal{F}_z \rfloor_{\ell, \delta}) - nk\mathcal{L}\|\mathbf{A}\|_\infty\delta$.*

Proof. Let x be the allocation rule and p be the payment rule of \mathcal{M}_2^ℓ . The mechanism \mathcal{M}^ℓ operates as follows: given the input bid \mathbf{w}_i of each i , we first construct $\mathbf{w}'_i = r^{(\ell, \delta)}(\mathbf{w}_i)$, and then we allocate to bidder i , the items $x_i(\mathbf{w}')$ and make him pay $\max\{p_i(\mathbf{w}') - k\|\mathbf{A}\|_\infty\mathcal{L}\delta, 0\}$. Note that if $\mathbf{w}_i \sim \mathcal{F}_{z, i}$ then $\mathbf{w}'_i \sim \lfloor \mathcal{F}_{z, i} \rfloor_{\ell, \delta}$.

We first argue that the mechanism is $(3k\mathcal{L}\|\mathbf{A}\|_\infty\delta + \eta, \mu)$ -BIC wrt \mathcal{F}_z .

$$\begin{aligned} \mathbb{E}_{\mathbf{t}_{-i} \sim (\mathcal{F}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}^\ell(\mathbf{w}_i, \mathbf{t}_{-i}))] &\geq \mathbb{E}_{\mathbf{t}'_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{w}_i, \mathcal{M}_2^\ell(\mathbf{w}'_i, \mathbf{t}'_{-i}))] \\ &\geq \mathbb{E}_{\mathbf{t}'_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{w}'_i, \mathcal{M}_2^\ell(\mathbf{w}'_i, \mathbf{t}'_{-i}))] - k\|\mathbf{A}\|_\infty\mathcal{L}\delta. \end{aligned}$$

The first inequality follows from the definition of the mechanism \mathcal{M}_2^ℓ and the last inequality follows from the fact that $\|\mathbf{w}'_i - \mathbf{w}_i\|_\infty \leq \delta$. Towards completing the proof of the BIC guarantee, we will now lower bound the right-hand side. To this end, note that the mechanism \mathcal{M}_2^ℓ is (η, μ) -BIC, i.e., with probability at least $1 - \mu$, for any $\mathbf{b}_i \in \text{supp}(\mathcal{F}_{z, i})$ we have,

$$\begin{aligned} &\mathbb{E}_{\mathbf{t}'_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{w}'_i, \mathcal{M}_2^\ell(\mathbf{w}'_i, \mathbf{t}'_{-i}))] - k\|\mathbf{A}\|_\infty\mathcal{L}\delta \\ &\geq \mathbb{E}_{\mathbf{t}'_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{w}'_i, \mathcal{M}_2^\ell(r^{(\ell, \delta)}(\mathbf{b}_i), \mathbf{t}'_{-i}))] - \eta - k\|\mathbf{A}\|_\infty\mathcal{L}\delta \\ &\geq \mathbb{E}_{\mathbf{t}'_{-i} \sim \lfloor (\mathcal{F}_z)_{-i} \rfloor_{\ell, \delta}} [u_i(\mathbf{w}_i, \mathcal{M}_2^\ell(r^{(\ell, \delta)}(\mathbf{b}_i), \mathbf{t}'_{-i}))] - \eta - 2k\|\mathbf{A}\|_\infty\mathcal{L}\delta \\ &\geq \mathbb{E}_{\mathbf{t}_{-i} \sim (\mathcal{F}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}^\ell(r^{(\ell, \delta)}(\mathbf{b}_i), \mathbf{t}_{-i}))] - \eta - 3k\|\mathbf{A}\|_\infty\mathcal{L}\delta \\ &= \mathbb{E}_{\mathbf{t}_{-i} \sim (\mathcal{F}_z)_{-i}} [u_i(\mathbf{w}_i, \mathcal{M}^\ell(\mathbf{b}_i, \mathbf{t}_{-i}))] - \eta - 3k\|\mathbf{A}\|_\infty\mathcal{L}\delta, \end{aligned}$$

⁸That is, $\|X - Y\|_{TV} = \min_\gamma \mathbb{P}_{(X, Y) \sim \gamma} [X \neq Y]$, where γ is the minimum over all couplings of X and Y .

The second inequality follows from the fact that $\|\mathbf{w}'_i - \mathbf{w}_i\|_\infty \leq \delta$ and the last is due to the definition of the mechanism \mathcal{M}^ℓ .

We now argue that the mechanism \mathcal{M}^ℓ is IR. To establish this, we only need to check instances where $p_i(\mathbf{w}') \neq 0$; if $p_i(\mathbf{w}') = 0$, then the mechanism is trivially IR. Using the fact that \mathcal{M}_2^ℓ is IR and $\|\mathbf{w}_i - \mathbf{w}'_i\|_p \leq \delta$, we have that

$$\begin{aligned} 0 \leq u_i(\mathbf{w}'_i, (\mathcal{M}_2^\ell(\mathbf{w}'_i, \mathbf{w}'_{-i}))) &= v_i^{\mathbf{A}}(\mathbf{w}'_i, (\mathcal{M}_2^\ell(\mathbf{w}'_i, \mathbf{w}'_{-i}))) - p_i(\mathbf{w}'_i, \mathbf{w}'_{-i}) \\ &\leq v_i^{\mathbf{A}}(\mathbf{w}_i, (\mathcal{M}_2^\ell(\mathbf{w}'_i, \mathbf{w}'_{-i}))) - p_i(\mathbf{w}'_i, \mathbf{w}'_{-i}) + k\|\mathbf{A}\|_\infty \mathcal{L}\delta \\ &= u_i(\mathbf{w}_i, (\mathcal{M}^\ell(\mathbf{w}_i, \mathbf{w}'_{-i}))). \end{aligned}$$

Therefore, \mathcal{M}^ℓ is IR. Finally, note that when all bidders bid truthfully \mathcal{M}^ℓ extracts the same revenue as \mathcal{M}_2^ℓ with a cumulative discount of at most $nk\mathcal{L}\|\mathbf{A}\|_\infty\delta$ for all the bidders, this directly implies the revenue bound stated in the lemma statement. \square

Lemma 6. *Let \mathcal{F}_z and $\hat{\mathcal{F}}_z$ be two distributions supported on \mathbb{R}^k such that $\pi_p(\mathcal{F}_z, \hat{\mathcal{F}}_z) \leq \varepsilon$. For any $\delta > 0$, $\mathbb{E}_{\ell \sim \mathcal{U}[0, \delta]^k} \left[\|\lfloor \mathcal{F}_z \rfloor_{\ell, \delta} - \lfloor \hat{\mathcal{F}}_z \rfloor_{\ell, \delta} \|_{TV} \right] \leq \left(1 + \frac{k^{1-1/p}}{\delta}\right) \varepsilon$.*

Proof. Using Lemma 1 we know that there exists a coupling γ of \mathcal{F}_z and $\hat{\mathcal{F}}_z$ so that $\mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p > \varepsilon] \leq \varepsilon$. Thus, we can bound the following probability:

$$\begin{aligned} \mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k, (x,y) \sim \gamma} \left[r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \right] &= \mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k, (x,y) \sim \gamma} \left[r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \wedge \|x - y\|_p > \varepsilon \right] + \\ &\quad \mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k, (x,y) \sim \gamma} \left[r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \wedge \|x - y\|_p \leq \varepsilon \right]. \end{aligned}$$

We can upper bound the first term above by $\mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p > \varepsilon] \leq \varepsilon$, and using Bayes rule, the second term can be written as $\mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k} [r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \mid \|x - y\|_p \leq \varepsilon] \cdot \mathbb{P}_{(x,y) \sim \gamma} [\|x - y\|_p \leq \varepsilon] \leq \mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k} [r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \mid \|x - y\|_p \leq \varepsilon]$. Plugging these upper bounds we get,

$$\begin{aligned} \mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k, (x,y) \sim \gamma} \left[r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \right] &\leq \varepsilon + \mathbb{P}_{\ell \sim \mathcal{U}[0, \delta]^k} \left[r^{(\ell, \delta)}(x) \neq r^{(\ell, \delta)}(y) \mid \|x - y\|_p \leq \varepsilon \right] \\ &\leq \varepsilon + \sum_{i \in [k]} \mathbb{P}_{\ell_i \sim \mathcal{U}[0, \delta]} \left[r_i^{(\ell, \delta)}(x) \neq r_i^{(\ell, \delta)}(y) \mid \|x - y\|_p \leq \varepsilon \right] \\ &\leq \left(1 + \frac{k^{1-1/p}}{\delta}\right) \varepsilon. \end{aligned}$$

The final inequality follows from the fact that $\sum_{i \in [k]} \mathbb{P}_{\ell_i \sim \mathcal{U}[0, \delta]} \left[r_i^{(\ell, \delta)}(x) \neq r_i^{(\ell, \delta)}(y) \right] \leq \sum_{i \in [k]} \frac{|x_i - y_i|}{\delta} = \frac{\|x - y\|_1}{\delta} \leq \frac{k^{1-1/p} \|x - y\|_p}{\delta}$, where in the final inequality here we used Holder's inequality. \square

D Active learning for regression problems via Randomized Linear Algebra: Details

D.1 Our query protocol: details for $p = 1$ and $p = 2$

Our query protocol takes as input the archetype matrix \mathbf{A} and a set of sampling probabilities q_i , $i \in [d]$, summing up to one. It outputs a sampling (and a rescaling) ma-

trix that can be used to select a small subset of types to query bidders' preferences.

Algorithm 1: Sampling & Rescaling Algorithm

Input: $\mathbf{A} \in \mathbb{R}^{d \times k}$, sampling complexity s_p , probabilities $q_i > 0, i \in [d], \sum_{i=1}^d q_i = 1$

Output: sampling matrix $\mathbf{S}_p \in \mathbb{R}^{s_p \times d}$, rescaling matrix $\mathbf{D}_p \in \mathbb{R}^{s_p \times s_p}$

Initialize $\mathbf{S}_p, \mathbf{D}_p$ to be all-zero matrices;

for $t \leftarrow 1$ **to** s_p **do**

Sample index $j \in [d]$ with respect to the probabilities $q_1 \dots q_d$;	
$\mathbf{S}_p(t, j) \leftarrow 1$;	// Set the t th row of \mathbf{S} to \mathbf{e}_j
$\mathbf{D}_p(t, t) \leftarrow \frac{1}{\sqrt{s_p q_j}}$;	// and rescale

end

D.1.1 The $p = 2$ case

We start with the $p = 2$ case and note that ℓ_2 regression is the most studied problem in the RLA literature, including the active learning setting. In this case, the sampling probabilities will be the so-called (row) leverage scores of \mathbf{A} , which can be computed exactly in $O(dk^2)$ time. Leverage scores can be approximated faster, in time that depends basically on the sparsity of the input matrix and we refer the reader to [M. 11, Woo14, DM16] for details on this very well-studied quantities. The main quality-of-approximation result is captured in the following lemma.

Lemma 7. *Let $\mathbf{A} \in \mathbb{R}^{d \times k}$ and $(\mathbf{t} + \epsilon_{nq,2}) \in \mathbb{R}^d$. Assume that the sampling probabilities q_i of Algorithm 1 are the row leverage scores of \mathbf{A} . Let $\tilde{\mathbf{z}} \in \mathbb{R}^k$ be*

$$\tilde{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{D}_2 \mathbf{S}_2 \mathbf{A} \mathbf{z} - \mathbf{D}_2 \mathbf{S}_2 (\mathbf{t} + \epsilon_{nq,2})\|_2.$$

Then, with probability at least 0.99,

$$\|\mathbf{A} \tilde{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,p})\|_p \leq \gamma_2 \text{OPT},$$

where $\text{OPT} = \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{A} \mathbf{z} - (\mathbf{t} + \epsilon_{nq,2})\|_2$. Here $\gamma_2 = 1 + \epsilon_q$, where $\epsilon_q > 0$ and the query complexity s'_2 satisfies

$$s'_2 = O(k \ln k + k/\epsilon_q).$$

The proof of the above lemma follows from Lemmas 4 and 5 of [DMMS11], which each hold with probability at least $1 - \delta/2$, by setting the query complexity to $s_2 = O(k \log k + k/\epsilon_q)$ and using the leverage scores as sampling probabilities. Applying a union bound over the failure probabilities of the two lemmas concludes the proof.

We now use Theorem 3.3 of [MMWY21] to boost the success probability of the above algorithm. Specifically, we repeat our algorithm $O(\ln 1/\delta')$ times to derive multiple candidate solutions for any $\delta' \in (0, 1)$. Then, we can use Algorithm 3 of [MMWY21] to select a solution that satisfies a slightly worse accuracy guarantee with high probability. More precisely, our final solution $\tilde{\mathbf{z}}$ will satisfy

$$\|\mathbf{A} \tilde{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,p})\|_p \leq (3 \cdot (1 + \epsilon_q) + 2) \text{OPT} \tag{10}$$

with probability at least $1 - \delta'$. Fixing $\epsilon_q = 0.5$, we get an overall query complexity equal to

$$s_2 = O(\ln(1/\delta') s'_2) = O(k \ln k \ln(1/\delta')).$$

Eventually, we will need the bound of eqn. (10) to hold for all n bidders via a union bound, so the failure probability must be reduced to $\delta' = \delta/n$. Recall that $\epsilon_q = 0.5$; eqn. (10) becomes

$$\|\mathbf{A} \tilde{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,p})\|_p \leq 6.5 \cdot \text{OPT}. \tag{11}$$

The above bound holds with probability at least $1 - \delta' = 1 - \delta/n$ if

$$s_2 = O(k \ln k \ln(n/\delta)),$$

for any $\delta \in (0, 1)$.

D.1.2 The $p = 1$ case

We now focus on the $p = 1$ case. In this setting, the sampling probabilities will be the Lewis weights, which can be approximated efficiently following the lines of [CP15, CD21]. We do emphasize that approximations to the Lewis weights (or the leverage scores) are sufficient in our setting. We now directly apply Theorem 1.2 of [CD21]). We again need a failure probability that is at most δ/n , since we will need to apply a union bound over all n bidders. In our notation, Theorem 1.2 of [CD21] can be restated as follows.

Lemma 8. *Let $\mathbf{A} \in \mathbb{R}^{d \times k}$ and $(\mathbf{t} + \epsilon_{nq,1}) \in \mathbb{R}^d$. Assume that the sampling probabilities q_i of Algorithm 1 are the row Lewis weights ($p = 1$) of \mathbf{A} . Let $\tilde{\mathbf{z}} \in \mathbb{R}^k$ be*

$$\tilde{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{D}_1 \mathbf{S}_1 \mathbf{A} \mathbf{z} - \mathbf{D}_1 \mathbf{S}_1 (\mathbf{t} + \epsilon_{nq,1})\|_1.$$

Then, with probability at least $1 - (\delta/n)$,

$$\|\mathbf{A} \tilde{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,1})\|_1 \leq \gamma_1 \text{OPT},$$

where $\text{OPT} = \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{A} \mathbf{z} - (\mathbf{t} + \epsilon_{nq,1})\|_1$. Here $\gamma_1 = 1 + \epsilon_q$, where $\epsilon_q > 0$ and the query complexity s_1 satisfies

$$s_1 = O(k/\epsilon_q^2 \ln kn/\epsilon_q \delta).$$

In our proof of Theorem 1, we will set ϵ_q to 0.5 for simplicity.

D.2 Our query protocol: details for $p > 2$

In this section, we show how to adapt the results of [MMWY22, MMWY21] in our setting in order to prove Theorem 1). We restate a sequence of results from [MMWY22, MMWY21], focusing on $p > 2$ and using our notation.

Lemma 9 (Theorem 2.11 in [MMWY21]). *Let $3 \leq p < \infty$ be an integer. There exists a randomized algorithm which constructs a sampling matrix $\mathbf{S}_p \in \mathbb{R}^{s'_p \times d}$ and a rescaling matrix $\mathbf{D}_p \in \mathbb{R}^{s'_p \times s'_p}$ such that, with probability at least 0.99, the ℓ_p subspace embedding property holds:*

$$^{1/2} \|\mathbf{A} \mathbf{x}\|_p \leq \|\mathbf{D}_p \mathbf{S}_p \mathbf{x}\|_p \leq ^{3/2} \|\mathbf{A} \mathbf{x}\|_p, \quad \forall \mathbf{x} \in \mathbb{R}^k.$$

Using Remark 2.21 [MMWY21], we get

$$s'_p = O(k^{p/2} \ln^3 k).$$

Notice that Remark 2.21 [MMWY21] removes the dependency of s'_p on $\ln d$. We also note that we can construct the matrices \mathbf{D}_p and \mathbf{S}_p using Algorithm 1 of [MMWY21]. Then, we solve the sampled ℓ_p regression problem $\hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{D}_p \mathbf{S}_p \mathbf{A} \mathbf{z} - \mathbf{D}_p \mathbf{S}_p (\mathbf{t} + \epsilon_{nq,p})\|_p$, to get guarantees of the form:

$$\|\mathbf{A} \hat{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,p})\|_p \leq \gamma_p \|\mathbf{A} \hat{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,p})\|_p, \quad (12)$$

for some constant $\gamma_p > 1$ that depends on the choice of $3 \leq p < \infty$. In the above,

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{A} \mathbf{z} - (\mathbf{t} + \epsilon_{nq,p})\|_p.$$

Lemma 10 (Theorem 3.2 in [MMWY21]). *Let $\mathbf{A} \in \mathbb{R}^{d \times k}$ and $(\mathbf{t} + \epsilon_{nq,p}) \in \mathbb{R}^d$. Let $3 \leq p < \infty$ be an integer and*

$$\text{OPT} = \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{A} \mathbf{z} - (\mathbf{t} + \epsilon_{nq,p})\|_p.$$

If $\tilde{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{D}_p \mathbf{S}_p \mathbf{A} \mathbf{z} - \mathbf{D}_p \mathbf{S}_p (\mathbf{t} + \epsilon_{nq,p})\|_p$, then, with probability at least 0.99,

$$\|\mathbf{A} \tilde{\mathbf{z}} - (\mathbf{t} + \epsilon_{nq,p})\|_p \leq 6(200)^{1/p} \text{OPT}.$$

The quantity $(200)^{1/p}$ decreases very fast as p increases. To put things into perspective for $p = 3$, $(200)^{1/3} < 6$. Importantly, to boost the success probability, we can compute $O(\ln(1/\delta'))$ candidate solutions by running the algorithm implied by Lemma 10 multiple times, for any $\delta' \in (0, 1)$. Then, we can use Algorithm 3 of [MMWY21] to select a solution that satisfies a slightly worse accuracy guarantee with high probability. More precisely:

Lemma 11 (Theorem 3.3 in [MMWY21]). *Call Algorithm 3 [MMWY21] with inputs $O(\ln(1/\delta))$ candidate solutions computed by the algorithm of Lemma 10. Then, we get a vector $\tilde{\mathbf{z}} \in \mathbb{R}^k$ such that, with probability at least $1 - \delta'$,*

$$\|\mathbf{A}\tilde{\mathbf{z}} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p \leq (18(200)^{1/p} + 2)\text{OPT}.$$

Overall, the query complexity for $3 \leq p < \infty$ is

$$s_p = O(\ln(1/\delta') s'_p) = O(k^{p/2} \ln^3 k \ln(1/\delta')).$$

Recall that we will need our bound to hold for all n bidders via a union bound, so the failure probability must be reduced to $\delta' = \delta/n$. Thus, our final sampling complexity s_p is

$$s_p = O(k^{p/2} \ln^3 k \ln(n/\delta)),$$

for any $\delta \in (0, 1)$.

D.3 The proof of Theorem 1

Proof of Theorem 1. For notational simplicity, in this proof, we use $\tilde{\mathbf{z}}$ instead of $\mathcal{Q}(\mathbf{t})$. Recall from Definition 5 that we can write (dropping the index i for simplicity):

$$\mathbf{t} = \mathbf{A}\mathbf{z} + \boldsymbol{\epsilon}_{\text{md1},p}, \quad (13)$$

where $\|\boldsymbol{\epsilon}_{\text{md1},p}\|_p \leq \varepsilon_{\text{md1},p}$. We now use the active learning protocols of Sections D.1 or D.2 to construct a solution $\tilde{\mathbf{z}} \in \mathbb{R}^k$ such that

$$\|\mathbf{A}\tilde{\mathbf{z}} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p \leq \gamma_p \|\mathbf{A}\hat{\mathbf{z}} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p, \quad (14)$$

for some constant $\gamma_p > 1$ that depends on the choice of $1 \leq p < \infty$. In the above,

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{A}\mathbf{z} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p.$$

For $p = 2$, eqn. (11) implies that $\gamma_2 = 6.5$. For $p = 1$, Lemma 8 implies that $\gamma_1 = 1.5$. For $p \geq 3$, Lemma 11 implies that $\gamma_p = 18(200)^{1/p} + 2$.

We now proceed to bound $\|\mathbf{z} - \tilde{\mathbf{z}}\|_p$. To achieve this, we bound the left- and right-hand-sides of eqn. (14). Substitute $\mathbf{t} = \mathbf{A}\mathbf{z} + \boldsymbol{\epsilon}_{\text{md1},p}$ to the left hand side to get:

$$\begin{aligned} \|\mathbf{A}\tilde{\mathbf{z}} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p &= \|\mathbf{A}(\tilde{\mathbf{z}} - \mathbf{z}) - (\boldsymbol{\epsilon}_{\text{nq},p} + \boldsymbol{\epsilon}_{\text{md1},p})\|_p \\ &\geq \|\mathbf{A}(\mathbf{z} - \tilde{\mathbf{z}})\|_p - \|\boldsymbol{\epsilon}_{\text{nq},p} + \boldsymbol{\epsilon}_{\text{md1},p}\|_p \\ &\geq \|\mathbf{A}(\mathbf{z} - \tilde{\mathbf{z}})\|_p - (\varepsilon_{\text{nq},p} + \varepsilon_{\text{md1},p}). \end{aligned} \quad (15)$$

The first inequality follows from the reverse triangle inequality and the second from the assumptions on $\boldsymbol{\epsilon}_{\text{nq},p}$ and $\boldsymbol{\epsilon}_{\text{md1},p}$. We now manipulate the right hand side of eqn. (14):

$$\begin{aligned} \gamma_p \|\mathbf{A}\hat{\mathbf{z}} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p &\leq \gamma_p \|\mathbf{A}\mathbf{z} - (\mathbf{t} + \boldsymbol{\epsilon}_{\text{nq},p})\|_p \\ &= \gamma_p \|\boldsymbol{\epsilon}_{\text{nq},p} + \boldsymbol{\epsilon}_{\text{md1},p}\|_p \\ &\leq \gamma_p (\varepsilon_{\text{nq},p} + \varepsilon_{\text{md1},p}). \end{aligned} \quad (16)$$

The first inequality follows from the fact that $\hat{\mathbf{z}}$ is the optimum solution of eqn. (14) and the second equality follows from eqn. (13). By combining eqns. (14), (15), and (16), we get

$$\|\mathbf{A}(\mathbf{z} - \tilde{\mathbf{z}})\|_p \leq (\gamma_p + 1)(\varepsilon_{\text{nq},p} + \varepsilon_{\text{md1},p}).$$

Using eqn. (2), we get $\|\mathbf{A}(\mathbf{z} - \tilde{\mathbf{z}})\|_p \geq \sigma_{\min,p} \|\mathbf{z} - \tilde{\mathbf{z}}\|_p$. Let $c_p = \gamma_p + 1$ to conclude:

$$\|\mathbf{z} - \tilde{\mathbf{z}}\|_p \leq \sigma_{\min,p}^{-1}(\mathbf{A}) \cdot c_p (\varepsilon_{\text{nq},p} + \varepsilon_{\text{md1},p}). \quad (17)$$

Using the γ_p values from above, we conclude that for $p = 2$, $c_2 = 7.5$; for $p = 1$, $c_1 = 2.5$; and for $p \geq 3$, $c_p = 18(200)^{1/p} + 3$. (We did not optimize constants, but it is worth noting that reducing c_p below two will need a different approach.) The failure probability of the theorem follows from the failure probability of eqn. (14), which needs to hold for all n bidders with probability at least $1 - \delta$. Applying a union bound and using the failure probabilities presented in Sections D.1 and D.2 concludes the proof of the theorem. \square

D.4 Improving Theorem 2 of [CD22] and comparisons with our work

We revisit the proof of Theorem 2 of [CD22], using our notation. We present a slightly improved analysis for the ℓ_∞ norm case. Let $\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{t}$ and $\mathbf{D}_\infty \mathbf{S}_\infty (\mathbf{t} + \epsilon_{\text{nq},\infty})$ be the bidders' responses to the query protocol \mathcal{Q} with and without the noise of the noisy query model. [CD22] makes the following assumptions, which are completely analogous to our work:

1. $\|\epsilon_{\text{nq},\infty}\|_\infty \leq \epsilon_{\text{nq},\infty}$ (query noise), and
2. $\|\mathbf{t} - \mathbf{A}\mathbf{z}\|_\infty \leq \epsilon_{\text{mdl},\infty}$ (model noise).

[CD22] proceeds by solving the induced least squares (instead of ℓ_∞) regression problem as follows:

$$\begin{aligned} \tilde{\mathbf{z}} &= \arg \min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A} \mathbf{z} - \mathbf{D}_\infty \mathbf{S}_\infty (\mathbf{t} + \epsilon_{\text{nq},\infty})\|_2 \\ &= ((\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^T \mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^{-1} (\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^T \mathbf{D}_\infty \mathbf{S}_\infty (\mathbf{t} + \epsilon_{\text{nq},\infty}) \quad (18) \\ &= (\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger \mathbf{D}_\infty \mathbf{S}_\infty (\mathbf{t} + \epsilon_{\text{nq},\infty}). \quad (19) \end{aligned}$$

In the above we assume that the inverse of the $k \times k$ matrix $(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^T \mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A}$ exists, which is also an implicit assumption in [CD22]. Thus, $\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A} \in \mathbb{R}^{s_\infty \times k}$ has full column rank equal to k . However, unlike [CD22], we will use eqn. (19) instead of eqn. (18) in our analysis. This will result in improved bounds with respect to various quantities that arise in the analysis. We now proceed to bound $\|\tilde{\mathbf{z}} - \mathbf{z}\|_\infty$, as follows:

$$\begin{aligned} \tilde{\mathbf{z}} - \mathbf{z} &= (\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger (\mathbf{D}_\infty \mathbf{S}_\infty (\mathbf{t} + \epsilon_{\text{nq},\infty}) - \mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A} \mathbf{z}) \\ &= (\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger (\mathbf{D}_\infty \mathbf{S}_\infty \epsilon_{\text{nq},\infty}) + (\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger \mathbf{D}_\infty \mathbf{S}_\infty (\mathbf{t} - \mathbf{A} \mathbf{z}). \end{aligned}$$

The first equality follows by the assumption that $\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A} \in \mathbb{R}^{s_\infty \times k}$ has full column rank equal to k , which implies that $(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger \mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A} = \mathbf{I}$. By applying the triangle inequality and using sub-multiplicativity properties of the ℓ_∞ norm, we get

$$\begin{aligned} \|\tilde{\mathbf{z}} - \mathbf{z}\|_\infty &\leq \epsilon_{\text{nq},\infty} \|(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger\|_\infty \|\mathbf{D}_\infty\|_\infty \|\mathbf{S}_\infty\|_\infty \\ &\quad + \|(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger\|_\infty \|\mathbf{D}_\infty\|_\infty \|\mathbf{S}_\infty\|_\infty \|\mathbf{t} - \mathbf{A} \mathbf{z}\|_\infty \\ &\leq \epsilon_{\text{nq},\infty} \|(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger\|_\infty \|\mathbf{D}_\infty\|_\infty + \epsilon_{\text{mdl},\infty} \|(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger\|_\infty \|\mathbf{D}_\infty\|_\infty \\ &= (\epsilon_{\text{nq},\infty} + \epsilon_{\text{mdl},\infty}) \|(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger\|_\infty \|\mathbf{D}_\infty\|_\infty \\ &\leq (\epsilon_{\text{nq},\infty} + \epsilon_{\text{mdl},\infty}) \sqrt{s_\infty} \|(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})^\dagger\|_2 \|\mathbf{D}_\infty\|_\infty \\ &\leq \sqrt{s_\infty} (\epsilon_{\text{nq},\infty} + \epsilon_{\text{mdl},\infty}) \sigma_{\min}^{-1}(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A}) \|\mathbf{D}_\infty\|_\infty, \quad (20) \end{aligned}$$

where for a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ we have $\|\mathbf{X}\|_\infty = \max_{i \in [m]} \sum_{j=1}^n |\mathbf{X}_{ij}|$ and $\|\mathbf{X}\|_2 = \max_{\|\mathbf{y}\|_2=1} \|\mathbf{X}\mathbf{y}\|_2 = \sigma_{\max}(\mathbf{X})$. In the above we used that fact that $\|\mathbf{S}_\infty\|_\infty = 1$, since \mathbf{S}_∞ is a sampling matrix and the property $\|\mathbf{X}\|_\infty \leq \sqrt{n} \|\mathbf{X}\|_2$ for any matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$. The use of eqn. (19) instead of eqn. (18) improves the bound of [CD22] by avoiding an extra \sqrt{k} term and a quadratic dependency on $\sigma_{\min}(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})$.

The bound of eqn. (20) generalizes the bound of [CD22] to allow for sampling and rescaling of the revealed noisy bidder preferences, as well as the corresponding rows of the archetype matrix \mathbf{A} , following the lines of Algorithm 1. It is important to note that in order to get meaningful results using the bound of eqn. (20), we need a lower bound on $\sigma_{\min}(\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A})$ and an upper bound on $\|\mathbf{D}_\infty\|_\infty$. The approach of [CD22] has no rescaling: it sets $\mathbf{D}_\infty = \mathbf{I}$ and, therefore, $\|\mathbf{D}_\infty\|_\infty = 1$. However, this makes it hard to get non-trivial lower bound for the smallest singular value of the matrix $\mathbf{S}_\infty \mathbf{A}$. This matrix is simply a sample of rows of the matrix \mathbf{A} without any rescaling and its smallest singular value could, in general, be arbitrarily close to zero. As a result, only special cases of the archetype matrix \mathbf{A} were analyzed in [CD22]: for those special cases, the smallest singular value of a sample of rows of the input matrix (without rescaling) can be lower bounded. It is a well-known fact that if one were to form the sampling matrix \mathbf{S}_∞ and the rescaling matrix \mathbf{D}_∞ using Algorithm 1 and the row leverage scores of \mathbf{A} as the sampling probabilities, the smallest singular value of $\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A}$ would be close to the smallest singular value of \mathbf{A} . This, however, cannot be done without rescaling, which would necessitate an upper bound on $\|\mathbf{D}_\infty\|_\infty$. As a concrete example, if all the row leverage scores of \mathbf{A} were equal, as would be the case if \mathbf{A} 's columns were a subset of the columns of a

Hadamard matrix, the sampling probabilities would all be equal to $1/d$. In this case, the diagonal rescaling matrix \mathbf{D}_∞ has entries that are proportional to the reciprocal of the sampling probabilities and $\|\mathbf{D}\|_\infty = \sqrt{d/s_\infty}$. This makes the overall error bound depend on the number of bidder types d , since s_∞ (the query complexity) is much smaller than d . This tension between lower-bounding the smallest singular value of $\mathbf{D}_\infty \mathbf{S}_\infty \mathbf{A}$ and upper bounding the infinity norm of \mathbf{D}_∞ does not seem easy to resolve by following the proposed analysis for arbitrary archetype matrices $\mathbf{A} \in \mathbb{R}^{d \times k}$.

D.4.1 Discussing the assumptions of [CD22] for the archetype matrix

We conclude this section with a brief discussion of the three settings of Theorem 2 of [CD22] for the archetype matrix. First, for the deterministic structure, the assumption on the archetype matrix $\mathbf{A} \in \mathbb{R}^{d \times k}$ is that it has to contain a square $k \times k$ submatrix \mathbf{C} , which is both row- and column-diagonally-dominant. It is unclear whether this is a reasonable assumption in the context of archetype matrices; some connections with the non-negative matrix factorization literature are briefly discussed [DS03]. Importantly, it is unclear what the query protocol \mathcal{Q} will be in this setting, since the query matrix \mathbf{Q} has to satisfy $\mathbf{Q} \cdot \mathbf{A} = \mathbf{C}$ and it is not obvious how to recover \mathbf{Q} if \mathbf{C} is unknown. The other two settings are both probabilistic and draw k archetypes either from a d -dimensional Gaussian distribution, or as copies of a d -dimensional random vector. In both cases, the archetypes are drawn from distributions that are designed to construct matrices whose columns are approximately orthogonal, at least in expectation. This observation results in large values for $\sigma_{\min}(\mathbf{A})$, and, under additional technical assumptions stated in Theorem 2 and Proposition 1 of [CD22], large values for the smallest singular value of $\mathbf{S}_\infty \mathbf{A}$. Our work directly characterizes the sample complexity of the protocol in terms of a single parameter of the archetype matrix that, intuitively, measures archetype independence with respect to p -norms. We provide explicit query protocols that leverage information in the archetype matrix and achieve the promised query complexity. We believe that this is a natural way to connect mechanism design with active learning for regression problems.