

The Sample Complexity of Up-to- ε Multi-Dimensional Revenue Maximization

Yannai A. Gonczarowski
*The Hebrew University of Jerusalem and
 Microsoft Research
 Israel*
Email: yannai@gonch.name

S. Matthew Weinberg
*Princeton University
 NJ, USA*
Email: smweinberg@princeton.edu

Abstract—We consider the sample complexity of revenue maximization for multiple bidders in unrestricted multi-dimensional settings. Specifically, we study the standard model of n additive bidders whose values for m heterogeneous items are drawn independently. For any such instance and any $\varepsilon > 0$, we show that it is possible to learn an ε -Bayesian Incentive Compatible auction whose expected revenue is within ε of the optimal ε -BIC auction from only polynomially many samples.

Our approach is based on ideas that hold quite generally, and completely sidestep the difficulty of characterizing optimal (or near-optimal) auctions for these settings. Therefore, our results easily extend to general multi-dimensional settings, including valuations that aren't necessarily even *subadditive*, and arbitrary allocation constraints. For the cases of a single bidder and many goods, or a single parameter (good) and many bidders, our analysis yields exact incentive compatibility (and for the latter also computational efficiency). Although the single-parameter case is already well-understood, our corollary for this case extends slightly the state-of-the-art.

Keywords—algorithmic game theory; algorithmic mechanism design; auctions; sample complexity; generalization bounds; PAC learning; approximate revenue maximization; multi-dimensional auctions;

I. INTRODUCTION

A fundamental question at the heart of the literature on mechanism design is that of *revenue maximization* by a single seller who is offering for sale any number of goods to any number of (potential) bidders. In the classic economic literature, this problem is studied in a *Bayesian* setting: the seller has prior knowledge of (often, independent) distributions from which the valuation of each bidder for each good is drawn, and wishes to devise a truthful mechanism that maximizes her revenue in expectation over these prior distributions. Over the past few years, numerous works at the interface of economics and computation are now studying a more demanding model: that of *mechanism design from samples*. In this model, rather than possessing complete knowledge of the distributions from which the bidders' values for the various items are drawn, the seller more realistically only has access to samples from these distributions (e.g., past data). The goal in this setting is to learn with high probability

A full version of this paper with complete proofs is available on arXiv.org.

an auction with good revenue guarantees given polynomially many (in the parameters of the problem) samples.

Revenue maximization from samples is somewhat ubiquitously seen as a “next step” beyond Bayesian revenue maximization. That is, existing works so far in this context take settings for which simple auctions in the related Bayesian problem are already well-understood and prove that these simple auctions can be learned efficiently via samples (up to an ε loss, which will always be lost when optimizing from samples). For example: in single-parameter settings, seminal work of Myerson [1] completely characterizes a simple and optimal auction in the Bayesian setting, and works such as Cole and Roughgarden [2], Morgenstern and Roughgarden [3], Devanur et al. [4], Hartline and Taggart [5], Roughgarden and Schrijvers [6], Gonczarowski and Nisan [7] prove that these simple mechanisms or variants thereof can be learned with polynomially many samples. Similarly, in multi-parameter settings with independent items, works of Chawla et al. [8]–[10], Hart and Nisan [11], Babaioff et al. [12], Rubinstein and Weinberg [13], Yao [14], Cai et al. [15], Chawla and Miller [16], Cai and Zhao [17] prove that simple mechanisms achieve constant-factor approximations in rich multi-dimensional settings, and works of Morgenstern and Roughgarden [18], Balcan et al. [19], [20], Cai and Daskalakis [21], Syrgkanis [22] prove that simple mechanisms with these guarantees can be learned with polynomially many samples. These analyses rely on a delicate understanding of the structure and/or inherent dimensionality of auctions that give such revenue guarantees to show how to learn such an auction without overfitting the samples.

It is therefore unsurprising that the problem of learning an up-to- ε revenue-maximizing multi-item auction from samples has not been previously studied, since the structure/dimensionality of optimal (precisely or up-to- ε) multi-item auctions is not understood even when there is only one bidder, and even with independent items. Such auctions are known to be extremely complex, suffering from properties such as randomization [23], uncountable menu complexity [24], and non-monotonicity [25]. Such domains provably lack the natural starting point of all previous works: a structured/low-dimensional mechanism in the Bayesian setting to learn via samples.

In this paper we show that despite these challenges, up-to- ε optimal multi-item auctions can be learned from polynomially many samples from the underlying bidder-item distributions. More formally, in a setting with n bidders and m items where the value of each bidder i for each item j is drawn independently from a distribution $V_{i,j}$ supported on $[0, H]$ for some H that is known to the seller, we show that polynomially many samples suffice for learning, with probability at least $1 - \delta$, an m -item almost-truthful auction that maximizes the expected revenue among all possible m -item almost-truthful auctions up to an additive ε . Below, BIC refers to *Bayesian Incentive Compatible*: an auction for which it is in every bidders’ interest to bid truthfully, given that all other bidders do so as well.

Theorem 1 (Main Result — informal version of Theorem 4). *For n bidders with independent values for m items supported on $[0, H]$, for every $\varepsilon, \delta > 0$ and for every $\eta \leq \text{poly}(n, m, H, \varepsilon)$, the sample complexity of learning, w.p. $1 - \delta$, an η -BIC auction that maximizes revenue (among all η -BIC auctions) up to an additive ε is $\text{poly}(n, m, H, 1/\varepsilon, 1/\eta, \log 1/\delta)$.*

The above theorem is informal mostly because we have not specified exactly how bidders value bundles of items. Essentially the bidders may have arbitrary (i.e., not necessarily additive, not even necessarily subadditive) valuations subject to some Lipschitz condition (i.e., changing the value of bidder i for item j by ε only changes the bidder’s value for any outcome by at most $L\varepsilon$ for some absolute constant L).¹ We defer a formal definition to Section II, but only note here that commonly studied classes of valuations such as additive, unit-demand, or “additive subject to constraints” with independent items (as well as several natural subadditive and superadditive valuation classes) all satisfy our definition with Lipschitz constant 1.

The main challenge in proving our result for $m > 1$ items is noted above: the structure of (up-to- ε) optimal mechanisms for such settings is not understood, even for additive valuations. In particular, there is no known low-dimensional class of mechanisms that is guaranteed to contain an (up-to- ε) optimal mechanism for any product distribution, thus barring the use of many learning-theoretic arguments. Our result relies on a succinct structured argument, allowing to reduce revenue maximization from samples to related problems of revenue maximization from given discrete distributions.

As the corresponding Bayesian question remains open (i.e., whether one can find, given the distributions explicitly, an up-to- ε optimal mechanism in poly-time), our result is of course information-theoretic: it shows that polynomially many samples suffice for a computationally unbounded seller, but provides no computationally efficient learning

¹Our results in fact hold even more generally: to arbitrary outcomes that do not even correspond to bundles of items. See Section II for the full details.

algorithm. Concretely, the algorithm that we give uses as a black box an oracle that can perform (optimal or almost-optimal) multi-item Bayesian revenue maximization given (the full description of) finite prior distributions.²

A. Brief Overview of Techniques

Most prior works (for single- as well as multi-dimensional settings) take the following approach: first, define a class \mathcal{C}_ε of auctions as a function of ε . Second, prove that, for all possible distributions \mathcal{D} , the class \mathcal{C}_ε contains an up-to- ε optimal mechanism for \mathcal{D} . Finally, prove that the best-in-class (up to ε) for \mathcal{C}_ε can be learned with polynomially many samples. In prior works, ingenuity is required for both steps: \mathcal{C}_ε is explicitly defined, proved to contain up-to- ε optimal auctions, and proved to have some low-dimensional structure allowing efficient learnability.

Our approach indeed follows this rough outline, with two notable simplifying exceptions. First is our approach to defining \mathcal{C}_ε . Here, we first define $\mathcal{C}_{\varepsilon,S}$ be the space of all auctions that are optimal for an empirical distribution over S -many $\text{poly}(\varepsilon)$ -rounded samples (that is, optimal for any discrete product distribution where each marginal is: a) only supported on multiples of $\text{poly}(\varepsilon)$ and b) uniform over a multiset of size S). While, unlike popular existing approaches, the set $\mathcal{C}_{\varepsilon,S}$ grows with the number of samples S , we show that the rate of its growth is moderate enough so that there exists a “sweet-spot” number of samples $S = \text{poly}(1/\varepsilon)$ such that on the one hand $\mathcal{C}_{\varepsilon,S}$ contains an auction that is up-to- ε optimal for the “true distribution” \mathcal{D} and on the other hand, the best-in-class from $\mathcal{C}_{\varepsilon,S}$ can be learned from S samples. So in the language of prior work, one could say that we set $\mathcal{C}_\varepsilon = \mathcal{C}_{\varepsilon,S}$ for this $S = \text{poly}(1/\varepsilon)$.

To show that \mathcal{C}_ε does in fact contain, for all distributions \mathcal{D} , an auction that is up-to- ε optimal for \mathcal{D} , we simply take enough samples to guarantee uniform convergence (of the revenue) over \mathcal{C}_ε and additionally the optimal auction for \mathcal{D} . It’s far from obvious why this should suffice, as the optimal auction for \mathcal{D} is not an element of \mathcal{C}_ε , nor even of the same format.³ Still existing tools (namely, the ε -BIC-to-BIC reduction of Daskalakis and Weinberg [26], Rubinstein and Weinberg [13]), when applied correctly, suffice to complete the argument. This part of our proof is conceptually much simpler than prior works (despite making use of a big technical hammer), as this approach holds quite generally and is robust due to not requiring the analysis of any specific class of mechanisms.

Second, our argument that the best-in-class can be in fact

²Note however that if computationally efficient algorithms were to be developed for up-to- ε optimal mechanisms given an explicit prior, then our approach would immediately become computationally efficient as well.

³That is, the optimal auction for \mathcal{D} is a mapping from the support of \mathcal{D} to outcomes, whereas the elements of \mathcal{C}_ε are mappings from a finite space to outcomes. Furthermore, notions of Bayesian incentive compatibility for \mathcal{D} do not imply nor are implied by these notions for the various discrete distributions defining \mathcal{C}_ε .

learned (up to ε) with $S = \text{poly}(\varepsilon)$ samples is simply a counting argument, and does not require any notions of a learning dimension. This is indeed in the spirit of some recent *single-dimensional results*, however in those results the counting argument is highly dependent on the structure of auctions in C_ε . As discussed above, such dependence is damning for multi-dimensional settings where such structure provably doesn't exist. Again, the proof does require some hammers (notably, arguments originally developed for *reduced forms* via samples in Cai et al. [27], and a concentration inequality of Babichenko et al. [28], Devanur et al. [4]), but they are applied in a fairly transparent manner.

The above approach should help explain how we are able to extend far beyond prior works, which relied on a detailed analysis of specific structured mechanisms: The key tools we use are applicable quite generally, whereas the specific mechanisms analyzed in prior work are only known to maintain guarantees only in restricted settings. For example, Theorem 1 already constitutes the first up-to- ε optimal-mechanism learning result for any multi-parameter setting *even if it held only for additive valuations (and one bidder)*. But the approach is so general that extending it to arbitrary Lipschitz valuations with independent items is simply a matter of updating notation.

B. Applications and Extensions

Specialized to a single-bidder setting, our construction in fact yields exact truthfulness (more on that in Section VI), showing that an ε -optimal mechanism can be found for a single bidder with independent item values (with Lipschitz valuations) using only polynomially many samples. This should be contrasted with a result of Dughmi et al. [29], which shows that achieving this is *not* possible for correlated distributions, even for a buyer with additive valuations.

Corollary 1 (Single Bidder — informal version of Theorem 5). *For one bidder with independent values for m items supported on $[0, H]$, for every $\varepsilon, \delta > 0$, the sample complexity of learning, w.p. $1 - \delta$, an IC auction that maximizes revenue (among all IC auctions) up to an additive ε is $\text{poly}(m, H, 1/\varepsilon, \log 1/\delta)$.*

Specialized to single-dimensional settings, our analysis once again yields a strengthened result, both in giving exact Dominant Strategy Incentive Compatibility (DSIC)⁴, and in providing a computationally efficient algorithm (due to known efficient solutions [1], [30] to single-parameter revenue maximization from given discrete distributions):

Corollary 2 (Single-Parameter — informal version of Theorem 6). *For n single-parameter bidders with independent values in $[0, H]$, for every $\varepsilon, \delta > 0$, the sample complexity of*

⁴A mechanism is DSIC if it is a dominant strategy for each bidder to bid truthfully. For a single good, Myerson [1] shows that the maximal revenue attainable by a BIC mechanism and by a DSIC mechanism is the same.

efficiently learning, w.p. $1 - \delta$, a DSIC auction that maximizes revenue (among all DSIC auctions) up to an additive ε is $\text{poly}(n, H, 1/\varepsilon, \log 1/\delta)$.

Corollary 2 nicely complements the existing literature on single-parameter sample complexity in the following ways. First, our algorithm/analysis immediately follows as a special case of Theorem 1 (without referencing structural results about optimal single-parameter auctions), so it is in some sense more principled. Second, our analysis holds even for arbitrary constraints on the allocations (putting it in the same class as the state-of-the-art⁵ single-parameter results [5], [7], and even slightly beyond⁶).

Finally, portions of our approach are specific to Bayesian Incentive Compatible auctions (versus Dominant Strategy Incentive Compatible auctions), but portions are not. We're therefore able to use the same techniques to conclude similar, albeit qualitatively weaker, results for ε -DSIC auctions in the full paper.

C. Related Work and Brief Discussion

Two active lines of work are directly related to the present paper. First are papers that study rich multi-dimensional settings, and aim to show that mechanisms with good approximation guarantees can be learned with few samples, such as Morgenstern and Roughgarden [18], Balcan et al. [19], [20], Cai and Daskalakis [21], Syrgkanis [22]. The main approach in each of these works is to show that specific classes of structured mechanisms (e.g., classes that are known to allow for constant-factor revenue maximization) are inherently low-dimensional with respect to some notion of dimensionality. Our results are stronger than these in some regards and weaker in others. More specifically, our results are stronger in the sense that with comparably many samples, our mechanisms guarantee an up-to- ε approximation to the optimal mechanism instead of a constant-factor. Our results are weaker in the sense that our learning algorithms are information-theoretic (do not run in poly-time), and our mechanisms are not "as simple." As discussed earlier, both weaknesses are necessary in order to possibly surpass the constant-factor barrier (at least, barring the resolution of major open questions, such as a computationally efficient up-to- ε approximation even when all distributions are explicitly known. Again, note that should this question be resolved affirmatively, our results would immediately become computationally efficient as well).

⁵Here and throughout the paper when we refer to "state-of-the-art" for single-parameter settings, we are specifically referring to allocation constraints that can be accommodated.

⁶Note that for the single-parameter setting, our algorithm in fact coincides with that of Devanur et al. [4]. However our analysis, unlike theirs, extends to arbitrary allocation constraints. Our approach also transparently handles mild extensions of constraints *beyond* those considered in Hartline and Taggart [5], Gonczarowski and Nisan [7]. Gonczarowski and Nisan [7] explicitly state that their techniques cannot handle such extensions and leave this question (which we successfully resolve) open.

Most related to our work, at least in terms of techniques, is the rich line of works on *single-dimensional settings* [2]–[7], [31], [32]. These works show that up-to- ε optimal mechanisms can be learned in richer and richer settings. In comparison to these works, our single-dimensional results slightly extend the state-of-the-art [5], [7] *as a corollary* of a more general theorem that applies to multi-dimensional settings. Even restricted to single-dimensional settings, our proof is perhaps more transparent.

We conclude with a brief discussion and an open problem. Corollaries 1 and 2 are both deduced from Theorem 1 by use of an argument as to why the resulting ε -BIC auction is in fact BIC, or by using an ε -BIC to BIC reduction that loses negligible revenue. Given that we have explicitly referenced the existence of a quite general ε -BIC-to-BIC reduction, the reader may be wondering why this reduction does not in fact allow our general results to be exactly BIC as well.

The main barrier is the following: in order to actually *run* the ε -BIC-to-BIC reduction as part of our auction for $n > 1$ bidders, one must take samples *exponential in the number of items* from each bidders’ value distribution. This means that even though we can learn an ε -BIC mechanism with few samples, plugging it through the reduction to remove the ε would cost us exponentially many samples in addition. Note that our current use of these theorems is non-constructive: we only use them to claim that the revenues achievable by the optimal BIC and ε -BIC mechanisms are not far off. This conclusion does not actually require *running* the reduction, but rather simply observing that it could be run (more details in the full paper).

When bidder valuations are drawn from a product distribution, it seems conceivable (especially given our results), that sample complexity polynomial in the number of items should suffice. Indeed, if each bidders’ values are drawn i.i.d., this is known due to exploitations of symmetry [26]. But subexponential sample complexity is not known to suffice for any other restricted class of distributions, despite remarkable recent progress in developing connections to combinatorial Bernoulli factories [33]. We state below what we consider to be the main open problem left by our work in the context of this paper, but readers familiar with black-box reductions in Bayesian mechanism design will immediately recognize a corresponding open problem for the original welfare-maximization setting studied in Hartline et al. [34], Bei and Huang [35] that is equally enticing.

Open Problem 1. *Given an ε -BIC auction for some product distribution, even in an additive multi-item setting, is it possible to transform it into a (precisely) BIC auction with negligible ($\text{poly}(\varepsilon) \cdot \text{poly}(n, m, H)$) revenue loss using polynomially many samples from this product distribution?*

The remainder of this paper is structured as follows. In Section II, we formally present the model and setting. In Section III, we formally state our results, which are

informally stated above as Theorem 1 and Corollaries 1 and 2. In Section IV, we overview the main ideas behind the proof of Theorem 1. Our learning algorithm is presented in Section V. In Section VI, we derive Corollaries 1 and 2. We present some extensions in Section VII. In the full paper, we also state and prove a result analogous to Theorem 1 for DSIC auctions, using similar proof techniques. Full proofs are given in the full paper.

II. MODEL AND PRELIMINARIES

The Decision Maker (Seller), Bidders, and Outcomes:

A single decision maker has the power to choose a social outcome, such as who gets which good that is for sale, or such as which pastime activities are offered in which of the weekends of the upcoming year. There are n bidders who have stakes in this outcome. (The decision maker will be able to charge the bidders and will wish to maximize her revenue.) The possible set of allowed outcomes is denoted by X and can be completely arbitrary. A central example is that of an m -parameter auction: the decision maker is a seller who has m items for sale, and the set of outcomes/allocations is $X \subseteq [0, 1]^{n \cdot m}$, where an allocation $\vec{x} = (x_{i,j})_{i \in [n], j \in [m]} \in X$ specifies for each bidder i and good j the amount of good j that bidder i wins. The traditional multi-item setting is the special case with $X = X_{\text{multi-item}} \triangleq \{(x_{i,j})_{i,j} \in \{0, 1\}^{n \cdot m} \mid \forall j : \sum_{i=1}^n x_{i,j} \leq 1\}$, while outcomes with fractional coordinates occur for example in the canonical model of position auctions, where smaller coordinates denote smaller click-through rates.

Values: Bidder $i \in [n]$ has a valuation function $v_i(\cdot)$ over the set of possible outcomes X . This function is parametrized by m values $v_{i,1}, \dots, v_{i,m}$ (we will not explicitly write $v_{i,\vec{v}_i}(\cdot)$, but refer to the parameters implicitly for ease of notation. Moreover, as $v_i(\cdot)$ is completely determined by $v_{i,1}, \dots, v_{i,m}$, we will sometimes simply refer to \vec{v}_i as bidder i ’s value, and to $v_{i,j}$ as bidder i ’s value for parameter j) and drawn from a given distribution such that:

- (Independent items) The $v_{i,j}$ s are independent random variables, drawn from distributions $V_{i,j}$ which are all supported in $[0, H]$.
- (Lipschitz) There exists an absolute constant L , such that if $v'_{i,1}, \dots, v'_{i,m}$ is obtained from $v_{i,1}, \dots, v_{i,m}$ by modifying one of the $v_{i,j}$ s by at most an additive ε , then $|v_i(x) - v'_i(x)| \leq L\varepsilon$ for all $x \in X$.

For example, in the multi-item setting described above, $v_i(x) = \sum_{j=1}^m x_{i,j} \cdot v_{i,j}$ (and $L = 1$).⁷ An additive-up-

⁷This case is actually Lipschitz in a stronger sense: if $v'_{i,1}, \dots, v'_{i,m}$ is obtained from $v_{i,1}, \dots, v_{i,m}$ by modifying one of the $v_{i,j}$ s by at most an additive ε , then $|v_i(\vec{x}) - v'_i(\vec{x})| \leq |\vec{x}|_1 \cdot L\varepsilon$ for $L = 1$. We note that using this stronger property (as well as other properties of the multi-item setting such as monotonicity), our analysis (mutatis mutandis) can be used to quantitatively improve the polynomial dependency of our sample complexity on the parameters of the problems, however we do not follow this direction in this paper. In general, in this paper we always choose generality of results over tighter polynomials.

to- k -items setting may also be easily captured using this setting (again with $L = 1$), and so can even settings with complementarities, such as a setting in which good $2j$ is worth $v_{i,j}$ to bidder i iff bidder i also gets good $2j + 1$ (and is otherwise worthless to bidder i). Since X can be completely arbitrary (in particular does not have to be a subset of $[0, 1]^{n-m}$), we can most generally think of the $v_{i,j}$ s as parameters that capture the “relevant attributes” of each bidder, such as affinity to action films, affinity to winter sports, willingness to spend a lot of time in a pastime activity, etc. The only requirement is that these attributes are a “natural” parametrization in the sense that the utility of the bidder from any given outcome $x \in X$ smoothly depends on (i.e., is Lipschitz in) each of them,⁸ and that they are independently drawn.

We note that both properties above (independent items and Lipschitz) together imply that the valuation of each bidder for each outcome is bounded in $[0, mLH]$.

Payments, Priced Outcomes, and Mechanisms: A payment specification $p = (p_i)_{i \in [n]}$ specifies for each bidder i to be charged p_i . A priced outcome is a pair (x, p) of allocation and payment specification. The utility of bidder i with value $v_i(\cdot)$ from priced outcome (x, p) is $u_i(v_i, (x, p)) = v_i(x) - p_i$. An auction/mechanism is a function that maps each valuation profile $(v_{i,j})_{i \in [n], j \in [m]}$ to a distribution over priced outcomes. The seller’s expected revenue from a mechanism μ is $\mathbb{E}_{v \sim \times_{i,j} V_{i,j}} \left[\sum_{i \in [n]} p_i(v) \right]$, where $p(v)$ is the payment specification chosen by the mechanism for the valuation profile v .

Truthfulness: An auction μ is individually rational (IR) if the expected utility of a truthful bidder is nonnegative at any valuation profile, i.e.: $\mathbb{E} [u_k(v_k, \mu(v))] \geq 0$ for every $k \in [n]$ and $v \in [0, H]^{n-m}$, where the expectation is over the randomness of the auction. For $\varepsilon > 0$, an auction μ is ε -dominant-strategy incentive compatible (ε -DSIC) if truthful bidding maximizes a bidder’s expected utility at any valuation profile up to an additive ε , i.e.: $\mathbb{E} [u_k(v_k, \mu(v))] \geq \mathbb{E} [u_k(v_k, \mu(v'_k, v_{-k}))] - \varepsilon$ for every $k \in [n]$, $v \in [0, H]^{n-m}$, and $v'_k \in [0, H]^m$, where the expectation is once again over the randomness of the auction. An auction is DSIC if it is 0-DSIC. An auction μ is ε -Bayesian incentive compatible (ε -BIC) if truthful bidding maximizes a bidder’s utility in expectation over all valuations of the other bidders, up to an additive ε , i.e.: $\mathbb{E}_{v_{-k} \sim \times_{i,j} V_{i,j}} [u_k(v_k, \mu(v_k, v_{-k}))] \geq \mathbb{E}_{v_{-k} \sim \times_{i,j} V_{i,j}} [u_k(v_k, \mu(v'_k, v_{-k}))] - \varepsilon$ for every $k \in [n]$ and $v_k, v'_k \in [0, H]^m$, where the expectation is both over the valuations of the bidders other than k and over the randomness of the auction. An auction is BIC if it is 0-BIC.

⁸This rules out such “tricks” as using bit-interleaving to condense the m parameters into a single parameter.

Additional Notation: We will use the following additional notation in our analysis, where $\varepsilon > 0$:

- For $v \in [0, H]$, we denote by $\lfloor v \rfloor_\varepsilon$ the value of v , rounded down to the nearest integer multiple of ε .
- We use $[0, H]_\varepsilon = \{\lfloor v \rfloor_\varepsilon \mid v \in [0, H]\}$ to denote the set of integer multiples of ε in $[0, H]$.
- For every i, j , we denote by $\lfloor V_{i,j} \rfloor_\varepsilon$ the distribution of $\lfloor v_{i,j} \rfloor_\varepsilon$ for $v_{i,j} \sim V_{i,j}$.

Existing Tools: In our analysis, we will make use of the following two theorems, which we state below in a way that is adapted to the notation of our paper. The first shows the optimal revenue over all ε -BIC auctions and the optimal revenue over all BIC auctions are close (while this is stated in Rubinstein and Weinberg [13] with respect to multi-parameter settings with allocations in $\{0, 1\}^{n-m}$, the same proof holds verbatim for arbitrary outcome sets X):

Theorem 2 (Rubinstein and Weinberg [13],⁹ see also Daskalakis and Weinberg [26]). *Let \mathcal{D} be any joint distribution over arbitrary valuations, where the valuations of different bidders are independent. The maximum revenue attainable by any IR and ε -BIC auction for a given product distribution is at most $2n\sqrt{mLH\varepsilon}$ greater than the maximum revenue attainable by any IR and BIC auction for that distribution.*

The second is a Chernoff-style concentration inequality for product distributions:

Theorem 3 (Babichenko et al. [28]; see also Devanur et al. [4]). *Let W_1, \dots, W_ℓ be discrete distributions. Let $S \in \mathbb{N}$. For every i , draw S independent samples from W_i , and let $W_i^{(S)}$ be the uniform distribution over these samples. For every $\varepsilon > 0$ and $f : \prod_{i=1}^\ell \text{supp } W_i \rightarrow [0, H]$, we have that $\Pr \left(\left| \mathbb{E}_{\prod_{i=1}^\ell W_i^{(S)}} [f] - \mathbb{E}_{\prod_{i=1}^\ell W_i} [f] \right| > \varepsilon \right) \leq \frac{4H}{\varepsilon} \exp\left(-\frac{\varepsilon^2 S}{8H^2}\right)$.*

III. MAIN RESULTS

In this section, we formally state our main results, which were informally presented as Theorem 1 and Corollaries 1 and 2 in the introduction. We start with our main result.

Theorem 4 (Main Result). *For every $\varepsilon, \delta > 0$ and for every $\eta \leq \text{poly}(n, m, L, H, \varepsilon)$, the sample complexity of learning an up-to- ε optimal IR and η -BIC auction is $\text{poly}(n, m, L, H, 1/\varepsilon, 1/\eta, \log 1/\delta)$. That*

⁹If we denote by R_ε the maximum expected revenue attainable by any IR and ε -BIC auction for the bidders’ product distribution, and by R the maximum revenue attainable by any IR and BIC auction for the same distribution, then the result of Rubinstein and Weinberg [13] is that for any $\eta > 0$, it is the case that $R \geq (1 - \eta) \cdot (R_\varepsilon - \frac{n\varepsilon}{\eta})$. Choosing $\eta = \sqrt{\frac{\varepsilon}{mLH}}$ yields Theorem 2 as stated above, since R_ε is trivially bounded from above by the maximum possible sum of valuations, i.e., by $nmLH$.

is, there exists a deterministic algorithm¹⁰ that given $\text{poly}(n, m, L, H, 1/\varepsilon, 1/\eta, \log 1/\delta)$ samples from each $V_{i,j}$, with probability $1 - \delta$ outputs an IR and η -BIC auction that attains from $\times_{i,j} V_{i,j}$ expected revenue at most an additive ε smaller than any IR and η -BIC auction.

The following corollary of our main result should be contrasted with a result of Dughmi et al. [29], which shows that finding an ε -optimal mechanism for a single additive bidder with *correlated* item distributions requires exponentially many samples.

Theorem 5 (Single Bidder). *When there is $n = 1$ bidder, for every $\varepsilon, \delta > 0$, the sample complexity of learning an up-to- ε optimal IR and IC¹¹ auction is $\text{poly}(n, m, L, H, 1/\varepsilon, \log 1/\delta)$. That is, there exists a deterministic algorithm¹² that given $\text{poly}(n, m, L, H, 1/\varepsilon, \log 1/\delta)$ samples from each $V_{i,j}$, with probability $1 - \delta$ outputs an IR and IC auction that attains from $\times_{i,j} V_{i,j}$ expected revenue at most an additive ε smaller than any IR and IC auction.*

The following corollary of our main result unifies and even somewhat extends the state-of-the-art results for single-parameter ($m = 1$) revenue maximization. To state it we restrict ourselves to the setting where revenue maximization has been solved by Myerson [1]: assume that $X \subseteq [0, 1]^n$ and that $v_i(x) = v_i \cdot x_i$.¹³

Theorem 6 (Single-Parameter). *In an $m = 1$ -parameter setting with $X \subseteq [0, 1]^n$ and $v_i(x) = v_i \cdot x_i$, for every $\varepsilon, \delta > 0$, the sample complexity of **efficiently** learning an up-to- ε IR and DSIC auction is $\text{poly}(n, m, L, H, 1/\varepsilon, \log 1/\delta)$. That is, there exists a deterministic algorithm with running time $\text{poly}(n, m, L, H, 1/\varepsilon, \log 1/\delta)$ that given $\text{poly}(n, m, L, H, 1/\varepsilon, \log 1/\delta)$ samples from each $V_{i,j}$, with probability $1 - \delta$ outputs an IR and DSIC auction that attains from $\times_{i,j} V_{i,j}$ expected revenue at most an additive ε smaller than any IR and BIC auction.*

IV. MAIN RESULT PROOF OVERVIEW

In this section we roughly sketch our learning algorithm and present each of the main ideas behind its analysis, by presenting a proof overview structured to present each of these ideas separately. The proof overview is given in this section only for an additive multi-item setting, and

¹⁰Recall that this result is information-theoretic and not computationally efficient (by necessity, without resolving major open problems), so our decision maker (seller) is computationally unbounded, and we allow the algorithm to make calls to any deterministic oracle that has no access to any $V_{i,j}$. In particular, we assume access to an oracle that can solve the revenue maximization problem on any precisely given $V'_{i,j}$ of finite support.

¹¹Recall that for a single bidder, the notions of BIC and DSIC coincide.

¹²See Footnote 10.

¹³This setting is slightly more general than the state-of-the-art single-parameter results, that assume $X \subseteq \{0, 1\}^n$ and/or the special setting of position auctions [5], [7]. As noted in the introduction, Goncezarowski and Nisan [7] explicitly state that their techniques cannot accommodate arbitrary $X \subseteq [0, 1]^n$ and leave this question (which we successfully resolve) open.

some elements of the proof are omitted or glossed over for readability. The full proof, which contains all omitted details and applies to a general arbitrary Lipschitz setting, and in which the main ideas that are surveyed in this section separately are quite intermingled, is given in the full paper.¹⁴

Our learning algorithm is similar in nature to the one presented in Devanur et al. [4] for certain single-parameter environments, however the analysis that we will use to show that it does not overfit the samples is completely different (even for single-parameter environments, where our analysis holds for arbitrary allocation constraints). Recall that our result is (necessarily) information-theoretic and not computationally efficient. Therefore, some steps in the algorithm perform operations that are not known to be performable in poly-time (but can certainly be performed without access to any $V_{i,j}$). In particular, our algorithm will solve an instance of a Bayesian revenue maximization problem for a precisely given input of finite support (step 2).

Algorithm: We start with S (to be determined later) independent samples from each $V_{i,j}$. Our algorithm roughly proceeds as follows:

- 1) For each item i and good j , round all samples from $V_{i,j}$ down to the nearest multiple of ε . Denote the uniform distribution over these rounded samples by $W_{i,j}$.
- 2) Find an IR and $O(\varepsilon)$ -IC (see below) multi-item auction that maximizes the revenue from the product of the rounded empirical distributions $W_{i,j}$. Denote this auction by μ .
- 3) Return the auction μ^ε , which on input \vec{v} , rounds down all actual bids to the nearest multiple of ε , $\lfloor \vec{v} \rfloor_\varepsilon$, and allocates and charges payments according to the output of $\mu(\lfloor \vec{v} \rfloor_\varepsilon)$ when run on these rounded bids.

We start by showing that if in step 2 of our algorithm we take an IR and $O(\varepsilon)$ -DSIC auction that maximizes the revenue from the product of the rounded empirical distributions, then there exists $S = \text{poly}(n, m, H, 1/\varepsilon, \log 1/\delta)$ such that the auction μ^ε output by our algorithm is $O(\varepsilon)$ -DSIC and its revenue from $\times_{i,j} V_{i,j}$ is, with probability at least $1 - \delta$, up-to- $O(\varepsilon)$ -close to the maximum revenue attainable from $\times_{i,j} V_{i,j}$ by any DSIC auction. (The formal statement and full proof are given in the full paper.) We note that the auction output by the algorithm is indeed $O(\varepsilon)$ -DSIC, since the output μ in step 2 is $O(\varepsilon)$ -DSIC, and the rounding of the actual bids as defined in step 3 only loses another $m\varepsilon$.¹⁵

¹⁴The somewhat less involved proof for DSIC auctions that is sketched in this section as an intermediary proof is also given in the full paper, both since we find that result interesting in its own right, and to allow interested readers to familiarize themselves with that proof before diving into the more involved proof for BIC auctions.

¹⁵The astute reader will notice that $m\varepsilon \notin O(\varepsilon)$. As all our bounds are polynomial in $m, 1/\varepsilon$ anyway, this is immaterial, and one example of a detail that we glossed over in this section in the interest of cleanliness, as promised.

Uniform Convergence of the Revenue of all Possible Output Mechanisms: Note that for every i, j , each rounded sample from step 1 of the algorithm is independently distributed according to $[V_{i,j}]_\varepsilon$. The main challenge is in showing that the resulting auction gives up-to- $O(\varepsilon)$ optimal revenue not only on the rounded empirical distributions $\times_{i,j} W_{i,j}$, but also on the rounded true distributions $\times_{i,j} [V_{i,j}]_\varepsilon$. That is, the main challenge is in showing that no overfitting occurs, in the absence of any structural properties that we can exploit for the mechanisms that are optimal (or up-to- $O(\varepsilon)$ optimal) for $\times_{i,j} W_{i,j}$.

This is the point where our approach makes a sharp departure from prior works. Prior work deems this task to be hopeless, and proceeds by proving structural results on optimal mechanisms for restricted domains. We circumvent this by instead simply counting the number possible inputs we will ever query in step 2, and observing that the number of mechanisms over which we have to obtain uniform convergence is at most this number. A crucial observation is that while we do have to consider more and more mechanisms as the number of samples S grows, the number of mechanisms that we have to consider grows moderately enough so as to not eclipse our gains from increasing the number of samples that we take. For this argument to hold, it is *essential* that our distributions are product distributions.

Let \mathcal{V} be the set of all product distributions $\times_{i,j} W'_{i,j}$ where each $W'_{i,j}$ is the uniform distribution over some multiset of S values from $[0, H]_\varepsilon$. Let M be the set of all mechanisms returned by step 2 of the algorithm for some distribution $\times_{i,j} W'_{i,j} \in \mathcal{V}$. **At the heart of our analysis, and of this part of our analysis in particular, is the observation that $|\mathcal{V}| \leq (S + 1)^{n \cdot m \cdot \lceil H/\varepsilon \rceil}$. Crucially, this expression has S only in the base and not in the exponent.** Indeed, for every $\times_{i,j} W'_{i,j} \in \mathcal{V}$, for every i, j , and for every integer multiple of ε in $[0, H]$ (there are $\lceil H/\varepsilon \rceil$ many such values), the probability of this value in $W'_{i,j}$ can be any of the $S + 1$ values $0, 1/S, \dots, 1$. Therefore, $|M| \leq (S + 1)^{n \cdot m \cdot \lceil H/\varepsilon \rceil}$.

We will choose S so that with probability at least $1 - \delta$, it simultaneously holds for all mechanisms $\mu \in M$ that

$$|\text{Rev}_{\times_{i,j} W_{i,j}}(\mu) - \text{Rev}_{\times_{i,j} [V_{i,j}]_\varepsilon}(\mu)| \leq \varepsilon. \quad (1)$$

To this end, we will use a Chernoff-style concentration bound (Theorem 3) for product distributions, which when applied to our setting shows that for each auction separately Equation (1) is violated with probability exponentially small in $\frac{\varepsilon^2 S}{m^2 H^2}$. So, to have Equation (1) hold with probability at most $1 - \delta$ for all auctions in M simultaneously, we choose S so that the violation probability for each auction separately is at most $\frac{\delta}{|M|}$, and use a union bound. Since $|M| \leq (S + 1)^{n \cdot m \cdot \lceil H/\varepsilon \rceil}$, we have that it is enough to take S such that $\frac{\varepsilon^2 S}{m^2 H^2}$ is of order of magnitude at least $\log |M|/\delta = \log \frac{(S+1)^{n \cdot m \cdot \lceil H/\varepsilon \rceil}}{\delta} \approx \log 1/\delta + n \cdot m \cdot H/\varepsilon \log S$,

which is clearly possible by taking a suitable S that is polynomial in $n, m, H, 1/\varepsilon$, and $\log 1/\delta$. So, taking a number of sample of this magnitude gives that with probability at least $1 - \delta$, Equation (1) simultaneously holds for all mechanisms in M and so the mechanism output by step 2 of the algorithm gets up-to- $O(\varepsilon)$ the same revenue on the product of the rounded empirical distributions $\times_{i,j} W_{i,j}$ as it does on the product of the rounded true distributions $\times_{i,j} [V_{i,j}]_\varepsilon$. So, the revenue that the mechanism μ^ε output by (step 3 of) the algorithm attains from $\times_{i,j} V_{i,j}$ is identical to the revenue that the mechanism μ output by step 2 of the algorithm attains from $\times_{i,j} [V_{i,j}]_\varepsilon$, which is up-to- $O(\varepsilon)$ the optimal revenue attainable from $\times_{i,j} W_{i,j}$.

Revenue Close to Optimal: Our next task is to show that with high probability the optimal revenue attainable from $\times_{i,j} W_{i,j}$ by any $O(\varepsilon)$ -DSIC auction is up-to- $O(\varepsilon)$ the same as the optimal revenue attainable from $\times_{i,j} V_{i,j}$ by any DSIC auction, which would imply that the revenue that μ^ε attains from $\times_{i,j} V_{i,j}$ is close to optimal, as required. Let OPT be the DSIC auction that maximizes the revenue (among such auctions) in expectation over $\times_{i,j} V_{i,j}$. **At the heart of this part of our analysis is the fact that while our algorithm cannot hope to find OPT, we can nonetheless carefully reason about it in our analysis, as it is nonetheless fixed and well-defined (in particular, it does not depend on the drawn samples).** Let OPT_ε be the mechanism defined over $\times_{i,j} [V_{i,j}]_\varepsilon$ as follows: for each bidder i and item j , let $w_{i,j}$ be the input bid of bidder i for item j (a multiple of ε), and replace it by a bid $v_{i,j}$ independently drawn from the distribution $V_{i,j}$ conditioned upon being in the interval $[w_{i,j}, w_{i,j} + \varepsilon)$; the auction OPT_ε allocates and charges payments according to the output of OPT when run on these drawn replacement bids. Obviously, the auction OPT_ε is an $O(\varepsilon)$ -DSIC auction whose revenue from $\times_{i,j} [V_{i,j}]_\varepsilon$ is identical to that of the auction OPT from $\times_{i,j} V_{i,j}$, i.e., to the optimal revenue from $\times_{i,j} V_{i,j}$, so it is enough to show that the revenue of the auction OPT_ε from $\times_{i,j} [V_{i,j}]_\varepsilon$ and from $\times_{i,j} W_{i,j}$ is the same up-to- $O(\varepsilon)$ with high probability, that is, that Equation (1) also holds for the mechanism OPT_ε with high probability. To do so, we modify the definition of the set M to also include the (well-defined even prior to sampling, despite being unknown to our algorithm) mechanism OPT_ε — since the order of magnitude of $|M|$ does not change, the order of magnitude of the number of samples required to guarantee that Equation (1) holds for all auctions in M (including OPT_ε) does not change.

Bayesian Incentive Compatibility: We conclude our proof overview by adapting the proof to the more delicate BIC notion of incentive compatibility, thus showing that if in step 2 of our algorithm we take an $O(\varepsilon)$ -BIC (rather than $O(\varepsilon)$ -DSIC) and IR auction that maximizes the revenue from the product of the rounded empirical distributions,

then there exists $S = \text{poly}(n, m, H, 1/\varepsilon, \log 1/\delta)$ such that the auction μ^ε output by our algorithm is, with probability at least $1 - \delta$, an $O(\varepsilon)$ -BIC auction whose revenue from $\times_{i,j} V_{i,j}$ is up-to- $O(\varepsilon)$ -close to the maximum revenue attainable from $\times_{i,j} V_{i,j}$ by any BIC auction (and therefore, by Theorem 2, up-to- $O(\sqrt{\varepsilon})$ -close to the revenue attainable from this distribution by any $O(\varepsilon)$ -BIC auction)¹⁶. The challenge here is that (approximate) BIC is a distribution-dependent property of a mechanism (as opposed to DSIC, which is a distribution-agnostic incentive compatibility notion). Indeed, examining our analysis above with (ε) -DSIC replaced by (ε) -BIC, we note that the resulting analysis falls short of carrying through in two points: it is unclear why OPT_ε is $O(\varepsilon)$ -BIC not only with respect to $\times_{i,j} V_{i,j}$ but also with respect to $\times_{i,j} W_{i,j}$, and it is unclear why any mechanism that can be output by step 2 of our algorithm is $O(\varepsilon)$ -BIC not only with respect to $\times_{i,j} W_{i,j}$ but only with respect to $\times_{i,j} V_{i,j}$. **At the heart of this part of our analysis is the observation that the set of all interim expected utilities, of all bidders' possible types, from all possible reported types, in all mechanisms¹⁷ on the one hand is comprised of a small-enough number of random variables to still enable uniform convergence, and on the other hand contains sufficient information to show that incentive constraints do not deteriorate much.** Concretely, we will choose S so that with probability at least $1 - \delta$, simultaneously for all mechanisms in M (including OPT_ε) not only does Equation (1) hold, but also the following holds for every bidder $k \in [n]$ and values $v_k, v'_k \in [0, H]_\varepsilon^m$:

$$\begin{aligned} & \left| \mathbb{E}_{v_{-k} \sim \times_{i,j \neq k} W_{i,j}} [u_k(v_k, \mu(v'_k, v_{-k}))] - \right. \\ & \left. - \mathbb{E}_{v_{-k} \sim \times_{i,j \neq k} [V_{i,j}]_\varepsilon} [u_k(v_k, \mu(v'_k, v_{-k}))] \right| \leq \varepsilon. \quad (2) \end{aligned}$$

We note that for every mechanism μ , we require that Equation (2) hold for $n \cdot \lceil H/\varepsilon \rceil^{2m}$ distinct combinations of $k \in [n]$ and $v_k, v'_k \in [0, H]_\varepsilon^m$. Crucially, this number does not depend on S . So, the number of instances of Equation (2) that we would like to hold simultaneously with high probability is $|M| \cdot n \cdot \lceil H/\varepsilon \rceil^{2m}$, and so we have $|M| \cdot (1 + n \cdot \lceil H/\varepsilon \rceil^{2m}) \leq (1 + n \cdot \lceil H/\varepsilon \rceil^{2m}) \cdot (S + 1)^{n \cdot \lceil H/\varepsilon \rceil^{2m}}$ instances of either Equation (1) or Equation (2) that we would like to hold simultaneously with high probability.¹⁸ As this number still has S only in the base and not in the exponent, we can proceed as above to guarantee this with high probability using only a polynomial number of samples.

¹⁶Theorem 2 in fact allows us to also reduce to a cleaner oracle, which finds an optimal BIC auction rather than an optimal $O(\varepsilon)$ -BIC auction, in Section V.

¹⁷The set of all such interim expected utilities for a single mechanism is sometimes referred to as the *reduced form* of the mechanism.

¹⁸A somewhat similar idea appeared in Cai et al. [27], albeit without exploiting independence across items.

V. LEARNING ALGORITHM

In this section, we give the fully detailed description of the learning algorithm from our main result, Theorem 4. The algorithm is described in Figure 1, and the proof of the correctness of the claims in that figure, whose overview was given in Section IV, is given in the full paper.

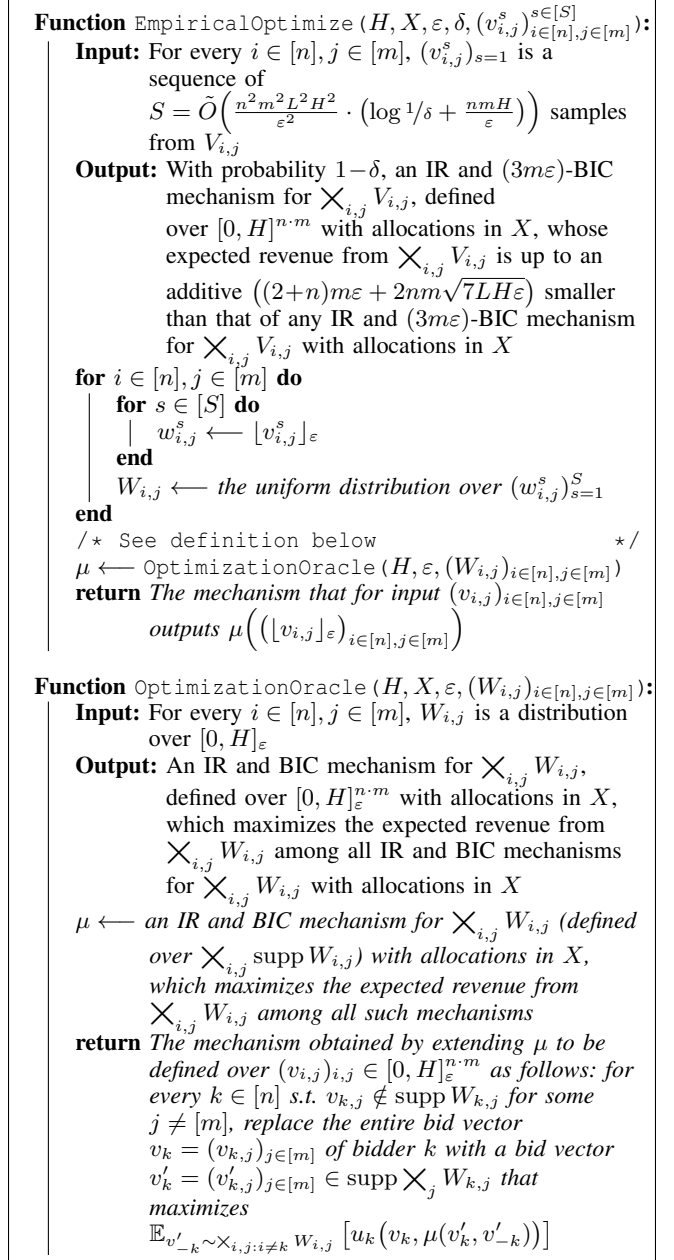


Figure 1. Empirical Multi-Parameter Up-to- ε BIC Revenue Maximization.

VI. FROM APPROXIMATE TO EXACT INCENTIVE COMPATIBILITY

In this section, we derive sample complexity results for exact incentive compatibility for the special cases of a single

bidder (Theorem 5) or a single good / single-parameter setting (Theorem 6). As mentioned in the introduction, whether this can also be done for more general settings remains an open question.

A. One Bidder

In this section, we will prove Theorem 5. For a single bidder, the following theorem, which to the best of our knowledge first implicitly appeared in Balcan et al. [36], where it is attributed to Nisan,¹⁹ provides an ε -IC to IC reduction with negligible revenue loss.

Theorem 7 (Nisan, circa 2005). *Let μ be an IR and ε -IC²⁰ mechanism for a single bidder. Modifying each possible priced outcome by multiplying the payment in that priced outcome by $1 - \sqrt{\varepsilon}$ and letting the bidder choose the (modified) priced outcome that maximizes her utility yields an IR and IC mechanism μ' with expected revenue at least $(1 - \sqrt{\varepsilon}) \cdot (\text{Rev}(\mu) - \sqrt{\varepsilon})$.*

For completeness, we provide a proof of this theorem in the full paper. The idea is that discounting more expensive priced outcomes more heavily makes sure that incentives do not drive the buyer toward a much cheaper priced outcome. More concretely, due to the auction being only ε -IC, the utility of a buyer from choosing a cheaper priced outcome can be higher by at most ε . Since for any priced outcome whose price is cheaper by more than a $\sqrt{\varepsilon}$ compared to the buyer's original priced outcome, the given discount is smaller by more than $\sqrt{\varepsilon}^2 = \varepsilon$, this smaller discount more than eliminates any potential utility gain due to choosing the cheaper priced outcome, so such a cheaper priced outcome would not become the most-preferred one. Applying Theorem 7 to the auction output by the algorithm of Figure 1 yields Theorem 5.

B. Single-Parameter Settings

In this section, we will prove Theorem 6. The algorithm presented in Section V constitutes a black-box reductions from ε -BIC revenue maximization from samples to BIC revenue maximization from given distributions. As noted in the introduction, the latter are mostly unsolved for more than one good. For a single good, however, the problem of DSIC/BIC revenue maximization was completely resolved by the seminal work of Myerson [1] (who, in particular, showed that the optimal BIC mechanism is DSIC), and

¹⁹ It appears there and in following papers [8], [29], [37]–[39] as part of a two-step reduction sometimes called “nudge and round” (this is the “nudge” part), which reduces the menu size of a single-bidder auction with negligible revenue loss. To the best of our knowledge, the first reference to this argument as a general ε -IC to IC reduction rather than as part of a “nudge and round” operation (where it fixes IC issues resulting from rounding) is in Daskalakis and Weinberg [26], who also attribute it to Nisan following Chawla et al. [8], who in turn attribute it to Nisan following Balcan et al. [36].

²⁰Recall once again that for a single bidder, the notions of BIC and DSIC coincide.

the computation complexity of the solution for discrete distributions was shown by Elkind [30] to be polynomial.

Definition (Myersonian Auction [1]). An n -bidder Myersonian auction (for valuations in $[0, H]$) is a tuple $(\phi_i)_{i \in N}$, where for every $i \in [n]$, $\phi_i : [0, H] \rightarrow \mathbb{R}$ is a nondecreasing function called the *ironed virtual valuation* of bidder i . The chosen outcome is $x \in X$ that maximizes $\sum_{i \in [n]} x_i \cdot \phi_i(v_i)$, with ties broken in a consistent manner. The payment is defined by the payment identity of Myerson [1], which guarantees that the auction is IR and DSIC.

Theorem 8 ([1]). *For every product distribution $\times_{i=1}^n W_i$, there exists a (DSIC) Myersonian auction $(\phi_i)_{i=1}^n$ that attains maximum revenue from $\times_{i=1}^n W_i$ among all IR and BIC auctions. Moreover, for every $i \in [n]$, the ironed virtual valuation ϕ_i depends only on W_i .*

Theorem 9 ([30]). *Let $S \in \mathbb{N}$. There exists an algorithm that runs in time $\text{poly}(S)$, such that given a discrete distribution W with support of size at most S , outputs a nondecreasing function $\phi : \text{supp } W \rightarrow \mathbb{R}$, such that for every product $\times_{i=1}^n W_i$ of discrete distributions each having support of size at most S , the (DSIC) Myersonian auction $(\phi_i)_{i \in N}$ (where ϕ_i is the output of the algorithm given W_i) maximizes the expected revenue from $\times_{i=1}^n W_i$ among all IR and BIC auctions.*

Plugging²¹ Theorem 9 into the algorithm of Figure 1 brings us closer (by making the algorithm efficient) to proving Theorem 6, however seems to still result in an ε -DSIC (rather than precisely DSIC) auction. Indeed, while the auction output by `OptimizationOracle` is DSIC, it seems that the auction output by `EmpiricalOptimize` is only $O(m\varepsilon)$ -DSIC. To complete the proof of Theorem 6, we note that the latter is in fact exactly DSIC in this case. Indeed, its allocation rule is monotone, and it satisfies the payment identity of Myerson [1] for every bidder.

VII. EXTENSIONS

A. On Computational Efficiency

As mentioned above, it is currently not known how to efficiently implement the optimization oracle, outputting an up-to- $O(\varepsilon)$ -optimal auction, as used in our algorithm. Nonetheless, there has been quite some work on efficiently finding auctions with weaker revenue guarantees with respect to optimal, such as guaranteeing some constant fraction of the optimal revenue. As the structured argument of our analysis provides a black-box reduction from BIC-revenue-maximization given a full distributions description to ε -BIC-revenue-maximization from samples, we can plug in any

²¹Extending the mechanism μ returned by the optimizer in `OptimizationOracle` to be defined over all of $[0, H]_\varepsilon^{n \cdot m}$ as follows, though (following Gonczarowski and Nisan [7]): for every $i \in [n]$ and $j \in [m]$ s.t. $v_{i,j} \notin \text{supp } W_{i,j}$, replace $v_{i,j}$ with $\max\{w_{i,j} \in \text{supp } W_{i,j} \mid w_{i,j} \leq v_{i,j}\}$.

such algorithm into our analysis (in lieu of the optimization oracle) to obtain a learning algorithm with matching computational complexity and revenue degradation:

Theorem 10 (“Meta Theorem”: Black-Box Reduction for Efficient Up-to-Constant Guarantees). *If there exists a polynomial-time algorithm for Bayesian revenue-maximization up to a constant factor C given an explicitly specified finite product distribution, then for every $\varepsilon, \delta > 0$ and for every $\eta \leq \text{poly}(n, m, L, H, \varepsilon)$, there exists a polynomial-time algorithm that given $\text{poly}(n, m, L, H, 1/\varepsilon, 1/\eta, \log 1/\delta)$ samples from each $V_{i,j}$, with probability at least $1 - \delta$ outputs an IR and η -BIC auction that attains from $\times_{i,j} V_{i,j}$ expected revenue at most an additive ε smaller than a C fraction of that of any IR and η -BIC mechanism.*

B. Partial Correlations

In some settings, there could be partial correlations between the distributions of the values of each bidder for the various goods. Our construction and analysis can also be modified to analyze such settings, to obtain sample bounds that are polynomial in the independent dimensions. To give a few examples:

- If there are correlations across values of different goods for the same bidder, but different bidders’ valuations are independent, then our upper bound for $|M|$ becomes $(S + 1)^{m \cdot (\lceil H/\varepsilon \rceil)^n}$, and so our analysis would yield sample complexity that is polynomial in the number of bidders (as our upper bound of $|M|$ in this case is still singly exponential in the number of bidders) and exponential in the number of goods (as our upper bound of $|M|$ in this case is doubly exponential in the number of goods).
- If for each bidder i the values of every two goods $2j, 2j+1 \in [m]$ are correlated, but are independent of the values of these goods for any other bidder, and of the values of any other good for any bidder, then our upper bound for $|M|$ becomes $(S + 1)^{n \cdot m/2 \cdot (\lceil H/\varepsilon \rceil)^2}$, and so our analysis still yields sample complexity that is polynomial in the both the number of bidders and the number of goods (as our upper bound of $|M|$ in this case is still singly exponential in both parameters). This is an example for a weaker form of correlation for which our analysis can still yield sample complexity that is polynomial in all parameters.

ACKNOWLEDGMENT

Yannai Gonczarowski is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. His work is supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389. This project has received funding from

the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 740282). Matt Weinberg is supported by NSF CCF-1717899. We thank Noam Nisan and Costis Daskalakis for helpful conversations.

REFERENCES

- [1] R. Myerson, “Optimal auction design,” *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58–73, 1981.
- [2] R. Cole and T. Roughgarden, “The sample complexity of revenue maximization,” in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, 2014, pp. 243–252.
- [3] J. Morgenstern and T. Roughgarden, “On the pseudo-dimension of nearly optimal auctions,” in *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, 2015, pp. 136–144.
- [4] N. R. Devanur, Z. Huang, and C.-A. Psomas, “The sample complexity of auctions with side information,” in *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016, pp. 426–439.
- [5] J. Hartline and S. Taggart, “Non-revelation mechanism design,” 2016, unpublished manuscript available on arXiv.org.
- [6] T. Roughgarden and O. Schrijvers, “Ironing in the dark,” in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 2016, pp. 1–18.
- [7] Y. A. Gonczarowski and N. Nisan, “Efficient empirical revenue maximization in single-parameter auction environments,” in *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017, pp. 856–868.
- [8] S. Chawla, J. D. Hartline, and R. Kleinberg, “Algorithmic pricing via virtual valuations,” in *Proceedings of the 8th ACM Conference on Electronic Commerce (EC)*, 2007, pp. 243–251.
- [9] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, “Multi-parameter mechanism design and sequential posted pricing,” in *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 2010, pp. 311–320.
- [10] S. Chawla, D. L. Malec, and B. Sivan, “The power of randomness in Bayesian optimal mechanism design,” *Games and Economic Behavior*, vol. 91, pp. 297–317, 2015, preliminary version in Proceedings of the 11th ACM Conference on Electronic Commerce (EC), 2010.
- [11] S. Hart and N. Nisan, “Approximate revenue maximization with multiple items,” in *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 2012, p. 656.
- [12] M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg, “A simple and approximately optimal mechanism for an additive buyer,” in *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014, pp. 21–30.

- [13] A. Rubinstein and S. M. Weinberg, “Simple mechanisms for a combinatorial buyer and applications to revenue monotonicity,” in *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, 2015, pp. 377–394.
- [14] A. C.-C. Yao, “An n -to-1 bidder reduction for multi-item auctions and its applications,” in *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015, pp. 92–109.
- [15] Y. Cai, N. Devanur, and S. M. Weinberg, “A duality based unified approach to bayesian mechanism design,” in *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016, pp. 926–939.
- [16] S. Chawla and J. B. Miller, “Mechanism design for subadditive agents via an ex ante relaxation,” in *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 2016, pp. 579–596.
- [17] Y. Cai and M. Zhao, “Simple mechanisms for subadditive buyers via duality,” in *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017, pp. 170–183.
- [18] J. Morgenstern and T. Roughgarden, “Learning simple auctions,” in *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*, 2016, pp. 1298–1318.
- [19] M.-F. Balcan, T. Sandholm, and E. Vitercik, “Sample complexity of automated mechanism design,” in *Proceedings of the 30th Conference on Neural Information Processing Systems (NIPS)*, 2016, pp. 2083–2091.
- [20] —, “A general theory of sample complexity for multi-item profit maximization,” in *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 2018, pp. 173–174.
- [21] Y. Cai and C. Daskalakis, “Learning multi-item auctions with (or without) samples,” in *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2017, pp. 516–527.
- [22] V. Syrgkanis, “A sample complexity measure with applications to learning optimal auctions,” in *Proceedings of the 31th Annual Conference on Neural Information Processing Systems (NIPS)*, 2017, pp. 5352–5359.
- [23] J. Thanassoulis, “Haggling over substitutes,” *Journal of Economic Theory*, vol. 117, no. 2, pp. 217–245, 2004.
- [24] C. Daskalakis, A. Deckelbaum, and C. Tzamos, “Mechanism design via optimal transport,” in *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, 2013, pp. 269–286.
- [25] S. Hart and P. J. Reny, “Maximal revenue with multiple goods: nonmonotonicity and other observations,” *Theoretical Economics*, vol. 10, no. 3, pp. 893–922, 2015.
- [26] C. Daskalakis and S. M. Weinberg, “Symmetries and optimal multi-dimensional mechanism design,” in *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 2012, pp. 370–387.
- [27] Y. Cai, C. Daskalakis, and S. M. Weinberg, “Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization,” in *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012, pp. 130–139.
- [28] Y. Babichenko, S. Barman, and R. Peretz, “Empirical distribution of equilibrium play and its testing application,” *Mathematics of Operations Research*, vol. 42, no. 1, pp. 15–29, 2017, preliminary version (“Simple approximate equilibria in large games”) in *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.
- [29] S. Dughmi, L. Han, and N. Nisan, “Sampling and representation complexity of revenue maximization,” in *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, 2014, pp. 277–291.
- [30] E. Elkind, “Designing and learning optimal finite support auctions,” in *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007, pp. 736–745.
- [31] P. Dhangwatnotai, T. Roughgarden, and Q. Yan, “Revenue maximization with a single sample,” *Games and Economic Behavior*, vol. 91, pp. 318–333, 2015.
- [32] Z. Huang, Y. Mansour, and T. Roughgarden, “Making the most of your samples,” in *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, 2015, pp. 45–60.
- [33] S. Dughmi, J. D. Hartline, R. Kleinberg, and R. Niazadeh, “Bernoulli factories and black-box reductions in mechanism design,” in *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017, pp. 158–169.
- [34] J. D. Hartline, R. Kleinberg, and A. Malekian, “Bayesian incentive compatibility via matchings,” in *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011, pp. 734–747.
- [35] X. Bei and Z. Huang, “Bayesian incentive compatibility via fractional assignments,” in *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011, pp. 720–733.
- [36] M.-F. Balcan, A. Blum, J. D. Hartline, and Y. Mansour, “Mechanism design via machine learning,” in *Proceedings of the 46th Annual IEEE Symposium on the Foundations of Computer Science (FOCS)*, 2005, pp. 426–439, full version: Technical Report CMU-CS-05-143.
- [37] S. Hart and N. Nisan, “The menu-size complexity of auctions,” in *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, 2013, p. 565.
- [38] M. Babaioff, Y. A. Gonczarowski, and N. Nisan, “The menu-size complexity of revenue approximation,” in *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017, pp. 869–877.
- [39] Y. A. Gonczarowski, “Bounding the menu-size of approximately optimal auctions via optimal-transport duality,” in *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC)*, 2018, pp. 123–131.