arXiv:1112.4572v2 [cs.GT] 19 Nov 2017

A Constructive Approach to Reduced-Form Auctions with Applications to Multi-Item Mechanism Design

Yang Cai*
School of Computer Science, McGill
cai@cs.mcgill.ca

Constantinos Daskalakis†
EECS, MIT
costis@mit.edu

S. Matthew Weinberg‡
Computer Science, Princeton
smweinberg@princeton.edu

November 21, 2017

Abstract

We provide a constructive proof of Border's theorem [Bor91, HR15a] and its generalization to reduced-form auctions with asymmetric bidders [Bor07, MV10, CKM13]. Given a reduced form, we identify a subset of Border constraints that are necessary and sufficient to determine its feasibility. Importantly, the number of these constraints is linear in the total number of bidder types. In addition, we provide a characterization result showing that every feasible reduced form can be induced by an ex-post allocation rule that is a distribution over ironings of the same total ordering of the union of all bidders' types.

We show how to leverage our results for single-item reduced forms to design auctions with heterogeneous items and asymmetric bidders with valuations that are additive over

*Supported by NSERC Discovery RGPIN-2015-06127, FRQNT 2017-NC-198956 and NSF Awards CCF-0953960 (CAREER), CCF-1101491, and CCF-1617730. Work done in part while the author was a Research Fellow at the Simons Institute for the Theory of Computing.

†Supported by a Sloan Foundation Fellowship, a Microsoft Research Faculty Fellowship, and NSF Awards CCF-0953960 (CAREER), CCF-1101491, and CCF-1617730. Work done in part while the author was a Research Fellow at the Simons Institute for the Theory of Computing.

‡Research completed in part while the author was supported by a NSF Graduate Research Fellowship, and in part while the author was a Microsoft Research Fellow at the Simons Institute for the Theory of Computing.

items. Appealing to our constructive Border's theorem, we obtain polynomial-time algorithms for computing the revenue-optimal auction. Appealing to our characterization of feasible reduced forms, we characterize feasible multi-item allocation rules.

Keywords: Reduced forms, multi-dimensional mechanism design, revenue maximization

1 Introduction

Consider a mechanism design setting with n bidders whose types lie in some finite set T and one copy of a single indivisible item. In this setting, explicitly describing an ex-post allocation rule requires $|T|^n$ probability distributions. In particular, for every type profile one needs to specify the probability that the item is allocated to each bidder. In applications where a succinct characterization of optimal mechanisms is lacking[1] (c.f. revenue optimal auctions with independent quasi-linear and risk-neutral bidders [Mye81]), it is desirable to take as a first step an optimization approach. However, optimizing over ex-post allocation rules is too expensive computationally and provides little structural insight into the optimal mechanism.

The above considerations motivate the study of *interim allocation rules*, also called *reduced-form auctions*, or simply *reduced forms*. Formally, suppose that bidder i's type is distributed according to some distribution $\mathcal{D}_i$ over T, and that bidders' types are independent. The reduced form of an allocation rule is a collection of functions $\mathcal{R} := \{\pi_i : T \to [0,1]\}_{i\in[n]}$. For all bidders i and types $\tau \in T$, $\pi_i(\tau)$ is the probability that the item is allocated to bidder i conditioning on her report being τ. The conditional probability is defined with respect to the reports of the other bidders, assumed to be drawn from the product distribution $\times_{j\neq i}\mathcal{D}_j$, and any randomization used by the allocation rule itself.

A key question surrounding reduced forms is: *under what conditions is a reduced form feasible?* More specifically, given a reduced form does there exist an ex-post allocation rule inducing it? This question was studied by Matthews [Mat84] and Maskin and Riley [MR84], and Border [Bor91] provided a collection of linear inequalities that are necessary and sufficient for the feasibility of a symmetric reduced form, namely when $\mathcal{D}_i = \mathcal{D}_j$ and $\pi_i(\cdot) = \pi_j(\cdot)$ for all bidders i and j. In more recent work, Border [Bor07], Manelli and Vincent [MV10], and Che et al. [CKM13] extend Border's conditions to the general (asymmetric) case. It follows

[1] Such applications include, for example, settings where bidders are risk-averse [MR84], budget-constrained [LR96], or have multi-dimensional preferences [RC98].

from these works that a reduced form $\mathcal{R}$ is feasible if and only if

$$\forall x_1, \ldots, x_n : \sum_i \sum_{\tau_i : \pi_i(\tau_i) \geq x_i} \pi_i(\tau_i) \Pr_{t_i \sim \mathcal{D}_i}[t_i = \tau_i] \leq 1 - \prod_i \left(1 - \Pr_{t_i \sim \mathcal{D}_i}[\pi_i(t_i) \geq x_i]\right). \quad (1)$$

Intuitively, the left hand side of (1) represents the probability that the item is allocated to some bidder i whose realized type τ_i satisfies $\pi_i(\tau_i) \geq x_i$, as computed by the reduced form $\mathcal{R}$. The right hand side represents the probability that there exists some bidder i whose realized type τ_i satisfies $\pi_i(\tau_i) \geq x_i$. Clearly, if $\mathcal{R}$ is feasible, it outght to satisfy (1) for any choice of thresholds $x_1, \ldots, x_n$. What is less clear is that if the inequality is satisfied for all thresholds, then $\mathcal{R}$ is feasible (but indeed this is the case).

From an optimization standpoint, one drawback of the afore-described conditions is that there are about $|T|^n$ inequalities that one needs to verify. Our first main result is that in fact it suffices to only check a subset of $n|T|$ constraints to verify the feasibility of a given reduced form.

Theorem 1. *A reduced form $\mathcal{R}$ is feasible if and only if*

$$\forall x : \sum_i \sum_{\tau_i \in S_x^{(i)}} \pi_i(\tau_i) \Pr_{t_i \sim \mathcal{D}_i}[t_i = \tau_i] \leq 1 - \prod_i \left(1 - \Pr_{t_i \sim \mathcal{D}_i}\left[t_i \in S_x^{(i)}\right]\right), \quad (2)$$

where $S_x^{(i)} = \{\tau_i \in T \mid \pi_i(\tau_i) \cdot \Pr_{t_i \sim \mathcal{D}_i}[\pi_i(\tau_i) \geq \pi_i(t_i)] > x\}$. In particular, one can test the feasibility of a reduced form or obtain a hyperplane separating $\mathcal{R}$ from the set of feasible reduced forms in time $O(|T|n \cdot \log(|T|n))$.

Note that the collection of inequalities (2) are a subset of inequalities (1). In particular, we have only kept $n|T|$ out of about $|T|^n$ inequalities. One way to interpret our theorem is that it coordinates which combinations of thresholds $x_1, \ldots, x_n$ it suffices to check simultaneously in (1). A priori, the simplest approach that could conceivably work would be to set all the x_i's equal, but we show that this does not suffice; see Example Three in Section 4.1.1. Instead of comparing different bidders' interim probabilities of allocation at face value, our theorem describes a way to shade these probabilities depending on each bidder's type distribution. The resulting "shaded interim probabilities" *can* be compared at face value. We provide several examples showing how Theorem 1 can be used to turn the task of verifying the feasibility of interim allocation rules analytically tractable in Section 4.1.1.

In addition to determining the feasibility of reduced forms, it is also important to understand the structure of ex-post allocation rules that induce them. To this end, Manelli and Vincent [MV10] provide an interesting characterization result, using the notion of a

hierarchical allocation rule [Bor91]. A hierarchical allocation rule maintains a weak total ordering $\succeq$ over the elements of $[n] \times T \cup \{(0, \perp = \tau_0)\}$. On input $(\tau_1, \ldots, \tau_n)$, the allocation rule computes the subset of indices $\mathcal{W} = \{i \mid (i, \tau_i) \succeq (j, \tau_j), \forall j\}$, then selects a uniformly random index i in $\mathcal{W}$. If $i > 0$, the item is allocated to bidder i. If $i = 0$, the item is not allocated. Manelli and Vincent show that, if a reduced form $\mathcal{R}$ is feasible, then there exists a distribution over hierarchical allocation rules inducing it. Moreover, each hierarchical allocation rule in the support of the distribution uses a weak total ordering $\succeq$ satisfying $\pi_i(\tau_i') \geq \pi_i(\tau_i) \implies (i, \tau_i') \succeq (i, \tau_i)$, i.e. "stronger types" of bidder i are ranked higher than "weaker types" of bidder i in every hierarchical allocation rule in the support. We strengthen this characterization as follows.

Theorem 2. *If a reduced form $\mathcal{R}$ is feasible then there exists a strict total ordering $\succ$ over the elements of $[n] \times T \cup \{(0, \perp)\}$ such that $\mathcal{R}$ can be induced by a distribution over hierarchical allocation rules, each using a weak ordering that irons $\succ$.*[2]

In addition to the *per-bidder* notion of "strength" of types guaranteed by Manelli and Vincent's characterization, Theorem 2 guarantees the existence of a *global* notion of strength of types in the sense that for every profile $(\tau_1, \ldots, \tau_n)$ with $\tau_{i_1} \succ \tau_{i_2} \succ \ldots \succ \tau_{i_n}$, the ex-post allocation probabilities $p_1, \ldots, p_n$ of the item to the bidders satisfy $p_{i_1} \geq p_{i_2} \geq \ldots \geq p_{i_n}$.

1.1 Multi-Dimensional Mechanism Design

Designing revenue-optimal auctions in multi-item settings has been a challenging application domain in mechanism design. A characterization theorem à la Myerson [Mye81] is unknown, and it is well-understood that optimal multi-item mechanisms exhibit much richer structure compared to optimal single-item auctions, involving bundling and randomization even in the case of a single additive bidder;[3] see discussion in Section 1.2. In light of this, it is valuable to develop optimization tools to compute optimal multi-item mechanisms.

Towards this end we adopt a linear programming approach. It is easy to write a linear program optimizing expected revenue over ex-post allocation and price rules of feasible, Bayesian incentive compatible mechanisms. However, this approach has two drawbacks. First, describing ex-post allocation and price rules requires $(m+1)n|T|^n$ numbers (for each of $|T|^n$ type profiles, and each of n bidders, one must list an allocation probability for each of m items along with a price paid). The exponential dependence on n makes this approach

[2] We say that a weak ordering $\succeq$ over some set S "irons" a strict ordering $\succ'$ over the same set S iff, for all $i, j \in S$, $i \succ' j \implies i \succeq j$.

[3] A bidder is additive if her valuation for a set of items is equal to the sum of her values for each item in that set.

computationally intractable. And, even if we could solve this linear program, it would be hard to discern useful structural insights from a "laundry list" of allocation probabilities and prices for every type profile.

The afore-described difficulties motivate a linear programming formulation with respect to the interim description of a mechanism instead. To write such a linear program, we need to identify linear constraints guaranteeing that an interim allocation/price rule pair is both feasible and Bayesian Incentive Compatible. When bidders are additive, Bayesian Incentive Compatibility is easy to express directly in terms of the interim allocation and price rule with a short list of $n|T|^2$ linear constraints. What is not clear is how to concisely express the feasibility of a multi-item interim allocation rule.

Our key observation is the following: a multi-item interim allocation rule is feasible if and only if the single-item interim rules that it projects onto each item are all feasible. Therefore, it suffices to invoke our single-item results above to resolve this problem. To be absolutely clear, even when the bidders are additive, it is folklore knowledge and well-understood that in the revenue-optimal auction **the interim probability that a bidder receives some item must in principle depend on her values for the other items.** It is exactly this property that makes multi-dimensional mechanism design notoriously difficult and not just a product of tractable single-item problems, and we are not claiming otherwise. However, the very specific subproblem of **determining whether a multi-item interim allocation rule is feasible can be solved separately across items**. Making use of Theorem 1 as a subroutine inside a linear program solver, we obtain the following computational result:

Theorem 3. *There is a polynomial-time algorithm that finds a revenue-optimal, BIC mechanism in multi-item settings with additive bidders. The algorithm takes as input the type distributions $\mathcal{D}_1, \ldots, \mathcal{D}_n$ of the bidders, and outputs a concise description of an optimal mechanism in time polynomial in the number of bidders n, the number of items m and the size of the type-space, $|T|$. The bidders are assumed independent, but each $\mathcal{D}_i$ may be an arbitrarily correlated distribution over item values.*

Besides Theorem 3, our key observation stated above, combined with Theorem 2, directly implies a characterization of feasible multi-item interim allocation rules, as follows.

Characterization of Feasible Multi-Item Interim Allocation Rules
Every feasible multi-item interim allocation rule can be implemented as follows:

- Every item is allocated independently of the other items.
- The allocation rule of every item ℓ maintains:
  - a strict ordering $\succ_\ell$ over the elements of $[n] \times T \cup \{(0, \perp = \tau_0)\}$; and
  - a distribution over ironings of $\succ_\ell$.
- Each item is then allocated as follows. First, a random ironing $\succeq'_\ell$ is sampled. On an input of reported types $(\tau_1, \ldots, \tau_n)$, the allocation rule computes the subset of indices $\mathcal{W}_\ell = \{i \mid (i, \tau_i) \succeq'_\ell (j, \tau_j), \forall j\}$, then selects a uniformly random index i_ℓ in $\mathcal{W}_\ell$. If $i_\ell > 0$, item ℓ is allocated to bidder i_ℓ. If $i_\ell = 0$, item ℓ is not allocated.

Recall that the set T in the characterization above is the set of types a bidder may have. In particular, each element $\tau \in T$ determines the values of a bidder of type τ for each bundle of items. Hence, the content of the first bullet is that **while the ex-post allocation rule for item ℓ indeed must depend on bidders' values for items $\neq \ell$, it need not depend on how items $\neq \ell$ are themselves allocated**.

1.2 Related Work

Reduced Forms. A necessary and sufficient condition for the feasibility of a bidder-symmetric reduced form was provided by Border [Bor91], building on prior work by Maskin and Riley [MR84] and Matthews [Mat84]. A simpler proof of Border's theorem and alternative criteria for feasibility were also provided by Hart and Reny [HR15a]. For all these works, the necessary and sufficient conditions take the form of $|T|$ linear inequalities. Border's conditions for the symmetric setting were generalized to the asymmetric setting by Border [Bor07], Manelli and Vincent [MV10], and Che et al. [CKM13]. For these works, the necessary and sufficient conditions take the form of $|T|^n$ linear inequalities.[4] In comparison, Theorem 1 shows that $|T|n$ linear inequalities suffice.

Let us review in more detail some of the most related works on reduced forms. Manelli and Vincent characterize the extreme points of the space of feasible, monotone reduced forms as monotone hierarchical allocation rules in both the bidder-symmeteric and asymmetric case when bidders have continuous type spaces. This implies that every monotone reduced form has an ex-post allocation rule inducing it that is also ex-post monotone. In

[4] To be more precise, Border's [Bor07] conditions took the form of $2^{|T|n}$ linear inequalities, and Che et al. [CKM13] identified a necessary and sufficient subset of $|T|^n$ linear inequalities.

this work, we offer alternative proofs of these results (Theorems 5 and 6) when type spaces are finite. Owing to the simpler nature of finite versus infinite dimensional geometry (in particular that counting arguments over finite sets are considerably simpler than counting arguments over infinite sets), our proofs more clearly isolate the key insights and are considerably shorter. Therefore, we include these proofs both for the sake of completeness and intuition. As detailed previously, in comparison to these characterization results, Theorem 2 provides a stronger characterization of reduced forms as implementable via randomizations over hierarchical allocation rules that iron the same *global* total ordering of all bidders' types.

Che et al. provide a clean network-flow interpretation of Border's theorem for asymmetric bidders, and show how to also accommodate bidders' capacity constraints in multi-unit generalizations (for example, that the set G of bidders must never receive more than $C(G)$ units or less than $L(G)$ units on any profiles). Their necessary and sufficient conditions remain in the form of $|T|^n$ linear inequalities (same as for a single item with asymmetric bidders) despite the significant increase in generality where their results apply. Independently from our work, Alaei et al. [AFH+12] also provide a computationally efficient algorithm to determine the feasibility of reduced forms via a "token-passing game." Their work shows in fact that there exists a collection of roughly $(n|T|)^2$ inequalities that define the space of feasible reduced forms. Their characterization is what is called an "extended formulation" - they introduce an additional $(n|T|)^2$ variables and their inequalities are not of the form (1). The results of Alaei et al. further extend to settings where the allocation is subject to matroid constraints. We address neither capacity constraints nor matroid constraints in this work. One high-level distinction between the main contributions of these works and our Theorems 1 and 2 is that their results extend Border's theorem to more general settings, whereas our work provides deeper insight into the core single-item setting (and by extension, as observed earlier, the multi-item setting where each item can be feasibly allocated to any bidder, regardless of other items she is allocated).

Finally, several recent papers have provided polynomial-time algorithms for determining whether a reduced form is *approximately* feasible in multi-item settings with more complex allocation constraints [CDW12b, CDW13a, CDW13b]. On the other hand, Gopalan et al. show essentially that approximation is the best one can hope for: their work identifies a formal barrier to the existence of exact and "computationally useful" Border-like theorems beyond single-item settings [GNR15].

Multi-Item Auctions. Prior work on multi-dimensional mechanism design is extensive (see e.g. survey [MV07]), driven by the scarcity of settings where the optimal mechanism has a clean allocation rule (such as Myerson's revenue-optimal auction for single-

dimensional settings [Mye81], or the welfare-optimal VCG auction in quite general settings [Vic61, Cla71, Gro73]). Indeed, numerous formal barriers have been identified to the existence of clean, revenue-optimal multi-item mechanisms, such as the necessity of randomization [RC98, Tha04, Pav11], large menu complexity [BCKW15, HN13, DDT13, DDT17], large description complexity [DDT14], and revenue non-monotonicity [HR15b, RW15]. Daskalakis et al. [DDT17] have recently provided a characterization of single-bidder revenue-optimal mechanisms using optimal transport theory. Older work of Rochet and Choné [RC98] had provided a characterization of optimal single-bidder mechanisms in the related setting where there is no bound on the number of units per item but the seller has a strictly convex production cost for generating more units. Finally, the problem has recently entered the Theory of Computation, where the emphasis has mostly been in deriving computationally efficient algorithms for computing optimal mechanisms. A number of results have emerged obtaining constant-factor approximations in polynomial time [CD11, Ala11, BGGM10, CHMS10, CMS15, KW12, CH13, BILW14, Yao15, CM16, CZ17].

In comparison to these works, ours is the first to provide a poly-time algorithm and corresponding characterization of revenue-optimal multi-item auctions without any distributional assumptions (such as a hazard rate condition or item-value independence). Indeed, following the announcement of portions of this work [CDW12a], several works (including some by the authors) provided computationally efficient algorithms to find approximately-optimal mechanisms in increasingly general multi-item settings [CDW12b, CDW13a, CDW13b, BGM13, DDW15, CDW16]. In comparison to these works, the present paper remains unique in containing a computationally efficient algorithm to find the exact optimal mechanism in a multi-item setting without any approximation error.

1.3 Roadmap of Remaining Sections

Section 2 below makes clear the notation we use and formal questions we study with respect to reduced forms. Section 3 studies bidder-symmetric reduced forms as a warm-up for Section 4, which studies asymmetric reduced forms. Section 5 provides our results on multi-item auctions. Appendix A contains some omitted proofs.

2 Preliminaries and notation

Throughout the paper, we denote the number of bidders by n. We also use T to denote the possible types of a bidder. In order to obtain computationally meaningful results, we assume that T is finite and use c as shorthand for $|T|$, but make no other assumptions on

T. In particular, it is not necessary to assume that T is a subset of $\mathbb{R}$.

We use τ to denote the *type* of a bidder, without emphasizing whether it is a vector or a scalar (or otherwise). The elements of T^n are called *type profiles*, and specify a type for every bidder. We assume type profiles are sampled from a distribution $\mathcal{D} = \prod_{i=1}^{n} \mathcal{D}_i$ over T^n, where $\mathcal{D}_i$ the marginal of this distribution on bidder i's type, and use $\mathcal{D}_{-i}$ to denote the marginal distribution over the types of all bidders, except bidder i. We use t_i for the random variable representing the type of bidder i. So when we write $\Pr[t_i = \tau]$, we mean the probability that bidder i's type is τ. If bidders are i.i.d., because $\Pr[t_i = \tau]$ is the same for all i, we will just write $\Pr[\tau]$.

The *reduced form* $\mathcal{R}$ of an allocation rule specifies a vector function $\pi(\cdot)$, specifying values $\pi_i(\tau)$, for all bidders i and types $\tau \in T$. $\pi_i(\tau)$ is the probability that bidder i receives the item when reporting type τ, where the probability is over the randomness of all other bidders' types and the internal randomness of the allocation rule, assuming that the other bidders report their true types. We may think of $\mathcal{R}$ as a vector in $[0,1]^{nc}$, by simply listing $\pi_i(\tau)$ for all i, τ, and will sometimes write $\vec{\pi}$ to emphasize this view.

In Section 3, we consider settings where the bidders are i.i.d., i.e. $\mathcal{D}_i = \mathcal{D}_j$ for all i, j, and the reduced forms are bidder-symmetric, which satisfy $\pi_i(\tau) = \pi_j(\tau)$, for all i, j, τ. In such cases, we will drop the subscript i, writing just $\pi(\tau)$, for the probability that a bidder of type $\tau \in T$ receives the item, over the randomness of the allocation rule and the types of the other bidders, assuming that the other bidders report their true types. Given a reduced form $\mathcal{R}$, we will be interested in whether it is "feasible." By this, we mean "does there exist an ex-post allocation rule that never over-allocates the item whose reduced form is $\mathcal{R}$?" If the answer to this question is "yes" then we will also say that the ex-post allocation rule whose reduced form is $\mathcal{R}$, "induces $\mathcal{R}$" or "implements $\mathcal{R}$." Note that there is some subtlety in defining feasibility of a reduced form if $\Pr[t_i = \tau] = 0$ for some $i \in [n], \tau \in T$. There are a couple reasonable choices that are qualitatively the same - we choose to define a reduced form to be feasible only if $\Pr[t_i = \tau] = 0 \Rightarrow \pi_i(\tau) = 0$.

We also note that the running times of the algorithms obtained in Sections 3 and 4 are quoted without accounting for the *bit complexity* of numbers involved. The bit complexity of a rational number x is the number of bits b required so that x can be expressed as the ratio of two binary numbers with b bits of precision. If b upper bounds the bit complexity of $\Pr[t_i = \tau]$ for all i, τ, and all coordinates of the input reduced form $\vec{\pi}$, then it suffices to multiply all quoted running times by a factor that is polynomial in b.

Order Notation. Throughout the text we use the $O(\cdot)$ notation. Let $f(x)$, $g(x)$ be two positive functions defined on some infinite subset of $\mathbb{R}_+$. Then we write $f(x) = O(g(x))$ iff

there exist some positive reals α and x_0 such that $f(x) \leq \alpha g(x)$, for all $x > x_0$. We also write $f(x) = \text{poly}(x)$ iff there exist positive reals α and x_0 such that $f(x) \leq x^\alpha$, for all $x > x_0$.

Finally, we provide some brief geometric preliminaries.

Definition 1. *(**Corner**) Let P be a closed, convex subset of Euclidean space defined as the intersection of finitely many halfspaces. Namely, $P = \cap_{i \in \mathcal{I}} \{\vec{x} | \vec{a}_i \cdot \vec{x} \leq b_i\}$, for some finite index set $\mathcal{I}$. We say that $\vec{x}^*$ is a* corner *of P if $\vec{x}^* \in P$, and the set of equations $\{\vec{a}_i \cdot \vec{x} = b_i\}_{i \in S}$ has as a unique solution the point $\vec{x}^*$, where $S = \{i \in \mathcal{I} | \vec{a}_i \cdot \vec{x}^* = b_i\}$.*

Definition 2. *(**Separation Oracle**) Let P be a closed, convex subset of Euclidean space. Then a* Separation Oracle *for P is an algorithm that takes as input a point $\vec{x}$ and outputs "*Yes*" if $\vec{x} \in P$, or a hyperplane $(\vec{w}, c)$ such that $\vec{y} \cdot \vec{w} \leq c$ for all $\vec{y} \in P$, but $\vec{x} \cdot \vec{w} > c$. Note that because P is closed and convex, such a hyperplane always exists whenever $\vec{x} \notin P$.*

A separation oracle is poly-time *if on inputs of bit complexity b, it terminates in time poly(b, x), where x is the maximum bit complexity of any coordinate in any halfspace defining P.*

We will also make use of the following theorem, reworded from [Kha79, GLS81, KP82].

Theorem 4. *([Kha79, GLS81, KP82]) Let P be a d-dimensional closed, convex subset of $\mathbb{R}^d$ defined as the intersection of finitely many halfspaces, and SO be a poly-time separation oracle for P. Then it is possible to do the following:*

- *Find an element in $\text{argmax}_{\vec{x} \in P}\{\vec{c} \cdot \vec{x}\}$ for any $\vec{c} \in \mathbb{Q}^d$ (i.e. solve linear programs) in time polynomial in d, and b, where b upper bounds the bit complexity of all coordinates of the vector $\vec{c}$, and all coordinates of the halfspaces defining P.*

- *Decompose any $\vec{x} \in P$ into a convex combination of at most $d+1$ corners of P in time polynomial in d, and b, where b upper bounds the bit complexity of all coordinates of the vector $\vec{x}$, and all coordinates of the halfspaces defining P.*

3 Warm-Up: Bidder-Symmetric Reduced Forms

This section serves as a warm-up for our main results by viewing bidder-symmetric reduced forms through a computational lens. Some of the key ideas for our main results in Section 4 can be more cleanly illustrated for symmetric bidders below.

Let us begin by reviewing Border's theorem [Bor91], which specializes (1) to the case of i.i.d. bidders and bidder-symmetric reduced forms. A bidder-symmetric reduced form is

feasible if and only if:

$$\forall x: \quad n \cdot \sum_{\tau:\pi(\tau)\geq x} \pi(\tau)\Pr[\tau] \leq 1 - \left(1 - \Pr_{t\sim\mathcal{D}_1}[\pi(t) \geq x]\right)^n. \tag{3}$$

The semantic meaning of the above equations are the same as those in Equation (1): the left-hand side denotes the probability that some bidder whose type τ satisfies $\pi(\tau) \geq x$ receives the item, as promised by the reduced form, and the right-hand side denotes the probability that some bidder i has $\pi(t_i) \geq x$ as computed by the probability distribution. Note that there are drastically fewer inequalities to check of form (3): only $|T|$ instead of $|T|^n$. This is essentially because if there exist thresholds $x_1, \ldots, x_n$ for which an equation of form (1) is violated, there is also a single x such that Equation (3) is violated at x. As there are only $|T|$ inequalities to check, Equation (3) directly implies Corollary 1 below: one can determine the feasibility of a reduced form in time $O(c(\log c + \log n))$ via a routine computation. A proof is included in Appendix A.

Corollary 1 (of [Bor91])**.** *The feasibility of a given bidder-symmetric reduced form can be determined in time $O(c(\log c + \log n))$. If it is infeasible, a violated inequality of form (3) can be determined in the same time.*

Now that it is easy to determine the feasibility of a reduced form, we wish to understand how to find an ex-post allocation rule inducing a given feasible reduced form in poly-time. To this end, let us formally define hierarchical allocation rules, again specialized to the bidder-symmetric case.

Definition 3. *([Bor91]) A **hierarchical allocation rule** consists of a weak total ordering $\succeq$ over $T \cup \{\perp\}$. On reports $(\tau_1, \ldots, \tau_n)$, the allocation rule computes the subset of indices $\mathcal{W} = \{i \geq 1 \mid \tau_i \succeq \perp \text{ and } \tau_i \succeq \tau_j, \forall j\}$, then selects a uniformly random bidder i in $\mathcal{W}$, if $\mathcal{W}$ is non-empty. If $\mathcal{W}$ is empty, the item is unallocated.*

We say that a hierarchical allocation rule induced by $\succeq$ is well-ordered *with respect to a reduced form $\mathcal{R}$ if $\pi(\tau) \geq \pi(\tau') \Rightarrow \tau \succeq \tau'$.*

For every feasible reduced form $\mathcal{R}$, Theorem 5 below characterizes the corners of a convex region containing it - which is intimately connected to ex-post allocation rules inducing $\mathcal{R}$.

Theorem 5. *(implied by [MV10]) Every feasible reduced form $\mathcal{R}$ lies inside a c-dimensional polytope P whose corners are all reduced forms of hierarchical allocation rules that are well-ordered w.r.t. $\mathcal{R}$. Furthermore, there is a distribution over at most $c+1$ hierarchical allocation rules, all well-ordered w.r.t. $\mathcal{R}$, that induces $\mathcal{R}$.*

Proof. For ease of notation, first relabel all types in T so that $\pi(\tau_1) \geq \pi(\tau_2) \ldots \geq \pi(\tau_c)$. Let $S = \{i | \pi(\tau_i) = \pi(\tau_{i+1})\}$ (for notational convenience, denote by $\pi(\tau_{c+1}) = 0$). Consider the convex polytope $P \subseteq [0,1]^c$ specified by the following constraints.

$$\tilde{\pi}(\tau_i) = \tilde{\pi}(\tau_{i+1}) \qquad \forall i \in S; \tag{4}$$

$$\tilde{\pi}(\tau_i) \geq \tilde{\pi}(\tau_{i+1}) \qquad \forall i \in [c] - S; \tag{5}$$

$$\sum_{j \leq i} n \cdot \Pr[t_j] \tilde{\pi}(\tau_j) \leq 1 - \left(1 - \sum_{j \leq i} \Pr[t_j]\right)^n \qquad \forall i \in [c]; \tag{6}$$

where for notational convenience we denote $\tilde{\pi}(\tau_{c+1}) = 0$ (so in particular $\tilde{\pi}(\tau_1), \cdots, \tilde{\pi}(\tau_c)$ are the free variables and $\tilde{\pi}(\tau_{c+1})$ is not, and the afore-desribed polytope is a subset of $[0,1]^c$). By (3), $\mathcal{R}$ is feasible if and only if $\pi \in P$. Consider the corners of this polytope. As there are c variables, every corner must satisfy at least c of the above inequalities with equality. Refer to the constraint $\tilde{\pi}(\tau_i) \geq \tilde{\pi}(\tau_{i+1})$ or $\tilde{\pi}(\tau_i) = \tilde{\pi}(\tau_{i+1})$ (whichever is included in the definition of P) as the ith monotonicity constraint, and the constraint $\sum_{j \leq i} n \cdot \Pr[t_j] \tilde{\pi}(\tau_j) \leq 1 - \left(1 - \sum_{j \leq i} \Pr[t_j]\right)^n$ as the ith Border constraint.

We first show that no feasible reduced form satisfies both the ith monotonicity constraint *and* the ith Border constraint with equality. Consider then a feasible reduced form $\tilde{\pi}$, and an ex-post allocation rule $\tilde{M}$ inducing $\tilde{\pi}$. Recall that if the ith Border constraint is tight, then the probability that a type in $\{\tau_1, \ldots, \tau_i\}$ receives the item is exactly the probability that such a type is reported to $\tilde{M}$. Therefore, whenever one or more types from this set are reported to $\tilde{M}$, a bidder with a type from this set necessarily wins the item. In particular, this means that we must have $\tilde{\pi}(\tau_i) \geq (1 - \sum_{j \leq i} \Pr[t_j])^{n-1}$, as τ_i must certainly win whenever all other reported types have index strictly larger than i. In fact, we must have $\tilde{\pi}(\tau_i) > (1 - \sum_{j \leq i} \Pr[t_j])^{n-1}$, as τ_i must also win with non-zero probability in the disjoint event that there is at least one other reported type equal to τ_i, and the remaining types have indicies strictly larger than i (because $\mathcal{R}$ is bidder-symmetric). Similarly, we necessarily have $\tilde{\pi}(\tau_{i+1}) \leq (1 - \sum_{j \leq i} \Pr[t_j])^{n-1}$, as τ_{i+1} can *only* win in the event that all other types have index strictly larger than i. Therefore, if the ith Border constraint is tight, the ith monotonicity constraint is not.

Now let us consider a corner $\tilde{\pi}$ of P. Since P lies in $\mathbb{R}^c$, there must be at least c constraints that $\tilde{\pi}$ satisfies with equality. As there are $2c$ constraints in total, and no feasible reduced form can satisfy both the ith Border constraint and the ith monotonicity constraint, we see that every corner must satisfy either the ith Border constraint or the ith monotonicity constraint with equality. Such a reduced form corresponds to a hierarchical allocation rule with $\tau_i \succeq \tau_j$ for all $i < j$, and $\tau_j \succeq \tau_i$ for some $i < j$ if and only if the kth monotonicity

constraints are tight for all $k \in \{i, \ldots, j-1\}$. Additionally, if $\tilde{\pi}(\tau_i) = 0$, then $\perp \succeq \tau_i$ and $\tau_i \not\succeq \perp$. If $\tilde{\pi}(\tau_i) > 0$, then $\tau_i \succeq \perp$ and $\perp \not\succeq \tau_i$. It is easy to see that this hierarchical allocation rule has a feasible reduced form that satisfies the desired equalities (hence it equals $\tilde{\pi}$) and is well-ordered w.r.t. $\mathcal{R}$.

By Carathéodory's theorem, we can write any feasible reduced form π as $\pi = \sum_{j=1}^{c+1} w_j \tilde{\pi}_j$, where $\sum_{j=1}^{c+1} w_j = 1$, and for all j, $w_j \geq 0$ and $\tilde{\pi}_j$ is a corner of P. The last step of the proof is an immediate consequence of the following observation.

Observation 1. If a reduced form $\mathcal{R}$ can be written as $\pi = \sum_j w_j \tilde{\pi}_j$, where each $w_j \geq 0$ and $\sum_j w_j = 1$, and each $\tilde{\pi}_j$ is induced by the ex-post allocation rule M_j, then $\mathcal{R}$ is induced by the ex-post allocation rule $\sum_j w_j M_j$ (sample j with probability w_j, then use M_j). □

Corollary 2. *Given a bidder-symmetric reduced form $\mathcal{R}$ we can determine if it is feasible, or find a hyperplane separating it from the set of feasible bidder-symmetric reduced forms in time $O\left(c(\log c + \log n)\right)$. If $\mathcal{R}$ is feasible, we provide a succinct description of an ex-post allocation rule inducing it, in time polynomial in c and $\log n$. In particular, the allocation rule is a distribution of at most $c+1$ hierarchical allocation rules, all well-ordered w.r.t. $\mathcal{R}$.*

Proof. The first sentence immediately follows from Corollary 1, and the observation that any violated Border inequality is exactly a hyperplane separating $\mathcal{R}$ from the space of feasible bidder-symmetric reduced forms.

We now need to describe how to computationally efficiently find an ex-post allocation rule implementing a reduced form π that is feasible. By Theorem 5, any feasible $\mathcal{R}$ lies inside a c-dimensional polytope P whose corners are all reduced forms of hierarchical allocation rules that are well-ordered w.r.t. $\mathcal{R}$. We now observe that we have defined a separation oracle for P in the first paragraph that runs in time $O(c(\log c + \log n))$. So Theorem 4 implies that we may decompose $\mathcal{R}$ into a convex combination of corners of P in time polynomial in c and $c(\log c + \log n)$ (resulting in a runtime polynomial in both c and $\log n$). Observation 1 completes the proof.

Notice in particular that in the proof of Theorem 5, we described an easy procedure to define a hierarchical allocation rule that implements any corner $\tilde{\pi}$ of P in terms of the inequalities of the polytope P that are tight at $\tilde{\pi}$. □

4 Asymmetric Reduced Forms

Let's begin this section by recapping the major components leading to Corollary 2 for symmetric reduced forms:

1. First, we need a computationally-efficient algorithm that takes as input a prospective reduced form and finds a violated Border constraint, if it exists (and otherwise claims that all Border constraints are satisfied). For symmetric reduced forms, Border's Theorem [Bor91] gives us this for free as there are only $|T|$ constraints to check. For asymmetric reduced forms, it will take exponential time to check all $|T|^n$ constraints of form (1), so we show in Section 4.1 that in fact it suffices to check only $n|T|$ constraints by properly shading the interim probabilities.

2. Next, to implement feasible reduced forms, we need to understand the corners of a convex region containing all feasible reduced forms. Work of Manelli and Vincent accomplishes this for the symmetric case (Theorem 5) and the asymmetric case (Theorem 6) for continuous type spaces. Again, we include a proof for the asymmetric case in Section 4.2 for finite type spaces that follows the same intuitive approach as our proof of Theorem 5.

3. Finally, we want to gain more insights into the structure of the space of feasible reduced forms than only understanding the corners of the feasible region. In the symmetric case, it is unclear what one might hope for beyond Theorem 5. In the asymmetric case, however, Theorem 6 doesn't tell the whole story. More specifically, Theorem 6 proves that every feasible reduced form can be induced by a distribution over hierarchical allocation rules *all of which respect the same partial ordering within a single bidder's types, but may not respect any global ordering across all bidders' types.*

   Indeed, a stronger charcaterization is possible in the form of Theorem 2: for every feasible reduced form $\mathcal{R}$ there exists a *global* total ordering of (bidder, type) pairs $\succ$, such that $\mathcal{R}$ can be induced by a distribution of hierarchical allocation rules that all respect $\succ$. We prove Theorem 2 in Section 4.3 using a combinatorial approach, and also show that the algebraic approach of Sections 4.1 and 4.2 cannot possibly yield such a theorem.

Let us now proceed by first recalling Equation (1), which states that a reduced form $\mathcal{R}$ is feasible if and only if for all thresholds $x_1, \ldots, x_n$, the probability that $\mathcal{R}$ awards the item to a bidder i whose reported type τ_i satisfies $\pi_i(\tau_i) \geq x_i$ is at most the probability that such a type is reported. When (1) is violated at some choice of thresholds $\vec{x}$, we call these thresholds *constricting*. Unfortunately, there are roughly c^n relevant $\vec{x}$ to test, so testing each of them separately is computationally/analytically intractable.

What we would really like is a more structured subset of Border constraints that are sufficient to check. To this end, we show that properly *shading* interim allocation probabilities

according to a bidder's type distribution allows interim probabilities to be compared across bidders at face value.

4.1 The Shaded Reduced Form

In this section, we define our notion of a *shaded reduced form*. We first provide the definition and main proposition, followed by some illustrative examples where analysis is greatly simplified by our approach.

Definition 4. *For any $s_i(\cdot)$ such that $s_i(\tau) \in [\Pr[\pi_i(t_i) < \pi_i(\tau)], \Pr[\pi_i(t_i) \leq \pi_i(\tau)]]$ for all i, τ, and any reduced form π, we define the corresponding* shaded reduced form *$\hat{\pi}$ as follows: for all i and types $\tau \in T$, $\hat{\pi}_i(\tau) := s_i(\tau) \cdot \pi_i(\tau)$.*

Observation 2. For all bidders i and any types $\tau, \tau' \in T$, $\hat{\pi}_i(\tau) \geq \hat{\pi}_i(\tau') \implies \pi_i(\tau) \geq \pi_i(\tau')$.

Proof. If $\hat{\pi}_i(\tau) \geq \hat{\pi}_i(\tau')$ but $\pi_i(\tau) < \pi_i(\tau')$, clearly $s_i(\tau) > s_i(\tau')$. On the other hand, $s_i(\tau) \leq \Pr[\pi_i(t_i) \leq \pi_i(\tau)] \leq \Pr[\pi_i(t_i) < \pi_i(\tau')] \leq s_i(\tau')$. Contradiction. □

Proposition 1. *Let π be an infeasible reduced form, and $s_i(\cdot)$ be any shading satisfying $s_i(\tau) \in [\Pr[\pi_i(t_i) < \pi_i(\tau)], \Pr[\pi_i(t_i) \leq \pi_i(\tau)]]$ for all i, τ. Then there exists a single threshold x such that: $\sum_i \sum_{\tau_i | \hat{\pi}_i(\tau_i) \geq x} \pi_i(\tau_i) \cdot \Pr[t_i = \tau_i] > 1 - \prod_i (1 - \Pr[\hat{\pi}_i(t_i) \geq x])$.*

In other words, for any valid shading of the reduced form, one can determine feasibility of a reduced form by checking Border's constraints where the threshold for the shaded reduced form is constant across all bidders.

Proof. If a reduced form $\mathcal{R}$ is infeasible, consider any maximal choice of constricting thresholds $x_1, \ldots, x_n$, i.e. a choice of $x_1, \ldots, x_n$ such that $(x_i + \delta, x_{-i})$ is not constricting for any $i, \delta > 0$. Now let $(i, \tau) \in \operatorname{argmin}_{j, \mu : \pi_j(\mu) \geq x_j} \{\hat{\pi}_j(\mu)\}$. Then by the maximality of $x_1, \ldots, x_n$, it must be the case that decreasing from $x_i + \delta$ to x_i causes us to go from satisfying (1) to violating it, and therefore $x_i = \pi_i(\tau)$ and this change must increase the LHS by more than it increases the RHS. The change in the LHS is easy to compute: Observation 2 implies that we are simply including additional (bidder, type) pairs in our calculations, namely (i, τ) (and all (i, τ') with $\pi_i(\tau') = \pi_i(\tau)$). The change in the RHS is also easy to compute: we have increased the probability that some bidder k exists with $\pi_k(t_k) \geq x_k$ by exactly the probability that all bidders $j \neq i$ have $\pi_j(t_j) < x_j$ and $\pi_i(t_i) = \pi_i(\tau)$. This therefore implies:

$$\Pr[\pi_i(t_i) = \pi_i(\tau)]\pi_i(\tau) > \Pr[\pi_i(t_i) = \pi_i(\tau)] \prod_{j \neq i} \Pr[\pi_j(t_j) < x_j] \iff \frac{\pi_i(\tau)}{\prod_{j \neq i} \Pr[\pi_j(t_j) < x_j]} > 1.$$

Now consider any other τ', k, $\pi_k(\tau') < x_k$ and $\hat{\pi}_k(\tau') \geq \hat{\pi}_i(\tau)$. Observe first that we must have $k \neq i$, as Observation 2 would otherwise imply $\pi_i(\tau') \geq x_i$. So we must have:

$$\begin{aligned}
&\pi_k(\tau') \cdot \Pr[\pi_k(t_k) \leq \pi_k(\tau')] \geq \hat{\pi}_k(\tau') \geq \hat{\pi}_i(\tau) \geq \pi_i(\tau) \cdot \Pr[\pi_i(t_i) < \pi_i(\tau)] \quad \text{(valid shading)} \\
\Longrightarrow &\pi_k(\tau') \cdot \Pr[\pi_k(t_k) < x_k] \geq \pi_i(\tau) \cdot \Pr[\pi_i(t_i) < \pi_i(\tau)] \quad \text{(because } \pi_k(\tau') < x_k) \\
\Longleftrightarrow &\pi_k(\tau') \cdot \Pr[\pi_k(t_k) < x_k] \geq \pi_i(\tau) \cdot \Pr[\pi_i(t_i) < x_i] \quad \text{(because } \pi_i(\tau) = x_i) \\
\Longleftrightarrow &\frac{\pi_k(\tau')}{\prod_{j \neq k} \Pr[\pi_j(t_j) < x_j]} \geq \frac{\pi_i(\tau)}{\prod_{j \neq i} \Pr[\pi_j(t_j) < x_j]}.
\end{aligned}$$

So by our choice of (i, τ) and the work above, we obtain:

$$\frac{\pi_k(\tau')}{\prod_{j \neq k} \Pr[\pi_j(t_j) < x_j]} > 1 \iff \Pr[t_k = \tau'] \pi_k(\tau') > \Pr[t_k = \tau'] \prod_{j \neq k} \Pr[\pi_j(t_j) < x_j].$$

This inequality (combined with Observation 2) tells us that we could lower x_k to $\pi_k(\tau')$ and still have constricting thresholds. In fact, it tells us something even stronger. If $y_1, \ldots, y_n$ are constricting thresholds with $y_j \leq x_j$ for all j, then we could decrease y_k to $\pi_k(\tau')$ and still have constricting thresholds. This is because lowering y_k to $\pi_k(\tau')$, causes the LHS of (1) to increase by $\sum_{\tau_k : \pi_k(\tau_k) \in [\pi_k(\tau'), y_k)} \Pr[t_k = \tau_k] \pi_k(\tau_k) \geq \Pr[\pi_k(t_k) \in [\pi_k(\tau'), y_k)] \, \pi_k(\tau')$. And the RHS increases by exactly $\Pr[\pi_k(t_k) \in [\pi_k(\tau'), y_k)]$ times the probability that $\pi_j(t_j) < x_j$ (respectively, y_j) for all $j \neq k$. As we decrease from x_j to y_j, this probability will clearly never increase.

So starting from any constricting thresholds $x_1, \ldots, x_n$ and a bidder type pair (i, τ) as above, we can lower each bidder k's threshold x_k to the lowest $\pi_k(\tau')$ such that $\pi_k(\tau') \geq \hat{\pi}_i(\tau)$. The resulting thresholds remain constricting and have the desired form.

□

The proof of Theorem 1 now readily follows from Proposition 1 and similar routine computation to Corollary 1. A complete proof of Theorem 1 appears in Appendix A. We now proceed with a few examples to clearly illustrate the benefits of an improved characterization.

4.1.1 Some Illustrative Examples

Example One: Monomials. Consider a reduced form where $T = [0, 1]$ and each $\mathcal{D}_i = U(T)$ (the uniform distribution on $[0, 1]$). Moreover, for all i let $\pi_i(\tau) = \tau^{\alpha_i}$, for some constants (but not necessarily identical) α_i. Using prior work, one could check all constraints of the form (1), which is a multi-variate optimization problem: maximize $\sum_{i=1}^{n} \int_{y_i}^{1} \tau^{\alpha_i} d\tau - (1 - \prod_{i=1}^{n} y_i)$, over all $\vec{y} \in [0, 1]^n$ (to map onto the variables used in Equation (1), set

$x_i = y_i^{\alpha_i}$). If the maximum happens to yield a value ≤ 0, then all constraints are satisfied. If the maximum yields a value > 0, then we have explicitly found a violated constraint. In turn, identifying the maximum requires considering second-order conditions at all local optima in addition to all points on the boundary, and is tedious.[5]

With Theorem 1 in hand, we observe that we need not optimize over all settings of thresholds $\vec{y} \in [0,1]^n$. Rather, it suffices to only consider thresholds that are jointly parametrized via a single threshold $x \in [0,1]$ on the shaded reduced form. For all $x \in [0,1]$, the corresponding cutoffs $\vec{y}$ satisfy $y_i^{\alpha_i+1} = x$, for all i: This is because for any $\tau \in [0,1]$, $\Pr[\pi_i(t_i) \leq \pi_i(\tau)] = \tau$, and $\pi_i(\tau) = \tau^{\alpha_i}$, so $\Pr[\pi_i(t_i) \leq \pi_i(\tau)] \cdot \pi_i(\tau) = \tau^{\alpha_i+1}$. So we see that our problem now reduces to a single-variate optimization: maximize $\sum_{i=1}^n \int_{x^{1/(\alpha_i+1)}}^1 \tau^{\alpha_i} d\tau - (1 - \prod_{i=1}^n x^{1/(\alpha_i+1)})$, over all $x \in [0,1]$.

This integral happens to be extremely simple to evaluate, and our objective function is just $\sum_{i=1}^n \frac{1-x}{\alpha_i+1} - \left(1 - x^{(\sum_{i=1}^n 1/(\alpha_i+1))}\right)$. We can then take a derivative with respect to x, which is $\left(\sum_{i=1}^n \frac{1}{\alpha_i+1}\right) \cdot \left(x^{(\sum_{i=1}^n 1/(\alpha_i+1)-1)} - 1\right)$. At this point, we can observe that if $\sum_{i=1}^n 1/(\alpha_i+1) > 1$, then the derivative is negative on the entire interval $[0,1)$, and therefore the maximum occurs at $x = 0$. At $x = 0$, the objective function evaluates to $\sum_i 1/(\alpha_i+1) - 1 > 0$, and therefore such reduced forms are infeasible. If instead, $\sum_i 1/(\alpha_i+1) \leq 1$, then the derivative is non-negative on the entire interval $[0,1)$, and therefore a maximum occurs at $x = 1$. At $x = 1$, the objective function evaluates to 0, and therefore such reduced forms are feasible.

In conclusion, Theorem 1 combined with the single-variable optimization above provides a complete proof that reduced forms of the above form are feasible if and only if they promise at most one item in expectation ex ante, which occurs if and only if $\sum_{i=1}^n 1/(\alpha_i+1) \leq 1$.

Example Two: Uniform Distributions of Support Two. Now, consider a reduced form where $T = \{H, L\}$, and each $\mathcal{D}_i = U(\{H, L\})$. Then if $\pi_i(H) \geq \pi_j(H)/2$ for all i, j, a valid shading sets $\hat{\pi}_i(H) = \min_j\{\pi_j(H)\}$, and $\hat{\pi}_i(L) = 0$ for all i. Proposition 1 then guarantees that we only need to check two constraints: $\sum_{i=1}^n \pi_i(H)/2 \leq 1 - 1/2^n$, and $\sum_{i=1}^n \pi_i(H)/2 + \pi_i(L)/2 \leq 1$. On the other hand, using prior work would require checking 2^n equations of the form (1) (with some extra thought, this can be reduced to n by appealing to the fact that the type distributions are iid. But getting all the way down to 2 requires reasoning à la Proposition 1).

Similarly, if each $\mathcal{D}_i$ assigns probability p_i to L, and $\pi_i(H) \geq p_j \pi_j(H)$ for all i, j, a valid shading sets $\hat{\pi}_i(H) = \min_j\{\pi_j(H)\}$, and $\hat{\pi}_i(L) = 0$ for all i. Proposition 1 again

[5]Note that Che et al. extend Border's Theorem to continuous type spaces (the theorem statement is identical to that for discrete type spaces) [CKM13]. Given this, it is easy to see that our proof of Proposition 1 extends to continuous type spaces as well as-is (subject to some change in notation).

guarantees that we only need to check two constraints: $\sum_{i=1}^{n}(1-p_i)\cdot\pi_i(H) \leq 1-\prod_{i=1}^{n} p_i$, and $\sum_{i=1}^{n}(1-p_i)\cdot\pi_i(H)+p_i\cdot\pi_i(L) \leq 1$. On the other hand, using prior work would again require checking 2^n equations of the form (1) (and this time it is not obvious how to check any fewer).

Example Three: A Correct Ordering is Necessary. The above two examples clearly illustrate why a tighter characterization is helpful, but it is not a priori clear that the even simpler approach of just considering thresholds where $x_i = x_j$ for all i, j fails. To see that indeed this approach fails, consider the following example with two bidders, and two types per bidder, H and L. For bidder 1, we set $\Pr[t_1 = H] = 1/8$, $\Pr[t_1 = L] = 7/8$, $\pi_1(H) = 5/8$, $\pi_1(L) = 0$. For bidder 2, we set $\Pr[t_2 = H] = 1/2$, $\Pr[t_2 = L] = 1/2$, $\pi_2(H) = 1$, $\pi_2(L) = 3/4$.

This reduced form is infeasible. Indeed, observe that bidder 2 must always receive the item whenever $\tau_2 = H$, which happens with probability 1/2. So if we have $\pi_2(H) = 1$, we cannot also have $\pi_1(H) > 1/2$. So (1) is violated when $x_1 = 5/8$ and $x_2 = 1$.

However, one can check that Border's conditions are satisfied whenever $x_1 = x_2$.[6] Essentially the problem is that to find constricting thresholds, we need x_1 to be small enough to include $(1, H)$, yet x_2 big enough to exclude $(2, L)$, which is not possible when $x_1 = x_2$.

Examples: Summary. Examples One and Two demonstrate the usefulness of a tighter characterization. On the other hand, Example Three shows that such a tighter characterization must be constructed carefully. Indeed, the shaded reduced form is exactly what is necessary to determine whether a reduced form is feasible or not.

4.2 Asymmetric Hierarchical Allocation Rules

Theorem 1 tightens the necessary and sufficient conditions of Border's theorem in a way that allows for computationally efficient determination of the feasibility of reduced forms. We proceed by examining which ex-post allocation rules are neccessary to induce all feasible reduced forms, similar to Theorem 5, first formally stating the definition of a hierarchical allocation rule in asymmetric settings.

Definition 5. *A **hierarchical allocation rule** consists of a weak total ordering $\succeq$ on $T \times [n] \cup \{(0, \perp)\}$. On reported types $(\tau_1, \ldots, \tau_n)$, the allocation rule computes the subset of*

[6] There are four inequalities of this form: $\Pr[t_2 = H]\cdot\pi_2(H) = 1/2 \leq 1-\pi_2(L)$, $\Pr[t_2 = H]\cdot\pi_2(H)+\Pr[t_2 = L]\cdot\pi_2(L) = 7/8 \leq 1$, $\Pr[t_2 = H]\cdot\pi_2(H)+\Pr[t_2 = L]\cdot\pi_2(L)+\Pr[t_1 = H]\cdot\pi_1(H) = 61/64 \leq 1$, and $\Pr[t_2 = H]\cdot\pi_2(H)+\Pr[t_2 = L]\cdot\pi_2(L)+\Pr[t_1 = H]\cdot\pi_1(H)+\Pr[t_1 = L]\cdot\pi_1(L) = 61/64 \leq 1$, all of which are satisfied.

indices $\{i|(i,\tau_i) \succeq (j,\tau_j), \forall j\}$*, then selects a uniformly random index* $i \in \mathcal{W}$*. If* $i > 0$*, the item is allocated to bidder* i*. If* $i = 0$*, the item is not allocated.*

We say that a hierarchical allocation rule $\succeq$ *for non-identical bidders is* partially-ordered *with respect to* $\mathcal{R}$ *if for all* i *and* $\tau, \tau' \in T$*,* $\pi_i(\tau) \geq \pi_i(\tau') \Rightarrow (i,\tau) \succeq (i,\tau')$*. We say that a hierarchical allocation rule is* strict *if for all bidders* $i \neq j$ *and types* $\tau, \tau' \in T$*:* $(i,\tau) \succeq (j,\tau') \wedge (i,\tau) \succeq (0,\perp) \Rightarrow (j,\tau') \not\succeq (i,\tau)$*, and* $(i,\tau) \succeq (0,\perp) \Rightarrow (0,\perp) \not\succeq (i,\tau)$ *(i.e.* $|\mathcal{W}| = 1$ *on all inputs).*

Similar to the symmetric case, a simple counting argument shows that every feasible reduced form can be implemented as a distribution over strict, partially-ordered hierarchical allocation rules. The proof for the asymmetric case follows the same outline, but requires one additional technical lemma whose proof is deferred to the Appendix A.

Theorem 6. *(implied by [MV10]) Every feasible reduced form* $\mathcal{R}$ *lies in a* cn*-dimensional polytope whose corners are all strict, partially-ordered w.r.t.* $\mathcal{R}$ *hierarchical allocation rules. Furthermore, there is a distribution over at most* $cn+1$ *hierarchical allocation rules, all strict and partially-ordered w.r.t.* $\mathcal{R}$*, that induces* $\mathcal{R}$*.*

Proof. For ease of notation, relabel all types in T (differently for each bidder) so that $\pi_i(\tau_{i,1}) \geq \ldots \geq \pi_i(\tau_{i,c})$ for all i, and so that $\Pr[t_i = \tau_{i,j}] = 0 \Rightarrow \Pr[t_i = \tau_{i,k}] = 0$ for all $k > j$ (i.e. since $\pi_i(\tau) = 0$ for all τ such that $\Pr[t_i = \tau] = 0$, we are free to put them at the end of the list). Let $c_i \leq c$ denote the number of types $\tau \in T$ such that $\Pr[t_i = \tau] > 0$. Again for notational convenience, denote by $\pi_i(\tau_{i,c+1}) = 0$. Let S_i denote the set of j such that $\pi_i(\tau_{i,j}) = \pi_i(\tau_{i,j+1})$. Consider the closed, convex polytope $P \subseteq [0,1]^{cn}$ specified by the following constraints.

$$\tilde{\pi}_i(\tau_{i,j}) = \tilde{\pi}_i(\tau_{i,j+1}) \qquad \forall i \in [n], j \in S_i \tag{7}$$

$$\tilde{\pi}_i(\tau_{i,j}) \geq \tilde{\pi}_i(\tau_{i,j+1}) \qquad \forall i \in [n], j \in [c] - S_i \tag{8}$$

$$\sum_i \sum_{j<z_i} \Pr[t_i = \tau_{i,j}]\tilde{\pi}_i(\tau_{i,j}) \leq 1 - \prod_i \left(1 - \sum_{j<z_i} \Pr[t_i = \tau_{i,j}]\right) \qquad \forall z_1, \ldots, z_n \in [c+1] \tag{9}$$

where for notational convenience we denote $\tilde{\pi}_i(\tau_{i,c+1}) = 0$ (again, $\tilde{\pi}_i(\tau_{i,c+1})$ is not a free

variable). In fact, we can also replace (9) with:

$$\sum_i \sum_{j<z_i} \Pr[t_i = \tau_{i,j}]\tilde{\pi}_i(\tau_{i,j}) \leq 1 - \prod_i \left(1 - \sum_{j<z_i} \Pr[t_i = \tau_{i,j}]\right) \quad \forall z_1 \in [c_1], \ldots, z_n \in [c_n] \tag{10}$$

$$\sum_i \sum_{j<c_i+1} \Pr[t_i = \tau_{i,j}]\tilde{\pi}_i(\tau_{i,j}) \leq 1 \tag{11}$$

$$\tilde{\pi}_i(\tau_{i,j}) = 0 \qquad \forall i \in [n], j > c_i \tag{12}$$

In the above replacement, we are observing that if (11) holds, then so does (9) for any case where at least one i has $z_i = c_i + 1$, as the left-hand side of all such inequalities is upper bounded by the left-hand side of Equation (11), and the right-hand side of all such inequalities is also 1. In addition, (10) covers all other cases, and Equation (12) just states that $\tilde{\pi}_i(\tau) = 0$ whenever $\Pr[t_i = \tau] = 0$. We know from [Bor07, CKM13] that $\mathcal{R}$ is feasible if and only if $\pi \in P$.

We proceed to show that any corner of this polytope is the reduced form of a strict hierarchical allocation rule. Note that any corner corresponds to a set of at least cn tight inequalities from (7), (8), (10), (11), or (12) whose tightness determine a unique solution. Similarly to Theorem 5, we examine the structure of what constraints can be simultaneously tight. We again refer to the inequalities in (7)/(8)/(12) as the $(i, j)^{th}$ monotonicity constraints, and those in (10)/(11) as the $\vec{z}$ Border constraints (where $\vec{z}$ indexes the constraint as above).

We first argue that the tight Border constraints must be *nested*. That is, if the Border constraint is tight for both $\vec{z}$ and $\vec{w}$, then we must either have $z_i \leq w_i$ for all i, or $w_i \leq z_i$ for all i. Assume for contradiction that this is not the case. Then there exists some bidders j, k such that $w_j > z_j$ and $z_k > w_k$. Clearly, we have $z_i < c_i + 1$ (respectively, $w_i < c_i + 1$) for all i, because otherwise we would necessarily have $z_i = c_i + 1$ for all i (respectively, $w_i = c_i + 1$ for all i) - the only such constraint with any $z_i = c_i + 1$ is (11). Consider now the type profile where bidder j has type τ_{j,z_j}, bidder k has type τ_{k,w_k}, and every other bidder i has type τ_{i,c_i} (note that this profile indeed arises with non-zero probability by definition of c_i). Now, any ex-post allocation rule that induces a feasible reduced form whose Border constraint at $\vec{w}$ is tight *must* award the item to bidder j on this profile, as every other bidder i's type has index at least w_i. Similarly, any ex-post allocation rule that induces a feasible reduced form whose Border constraint at $\vec{z}$ is tight *must* award the item to bidder k on this same profile, as every other bidder i's type has index at least z_i. Clearly, no feasible ex-post allocation rule can award the item to both bidders, so no feasible reduced form can have both Border constraints be tight.

Now that we know that for any corner $\tilde{\pi}$ of P the tight Border constraints are nested, we (suggestively) define the relation:

$$(i, \tau_{i,j}) \succeq (k, \tau_{k,\ell}) \Leftrightarrow |\{\vec{z} \,|z_i > j \wedge \text{ Border}(\vec{z}) \text{ is tight}\}| \geq |\{\vec{z} \,|z_k > \ell \wedge \text{ Border}(\vec{z}) \text{ is tight}\}|\,,$$

$$(i, \tau_{i,j}) \succeq (0, \perp) \Leftrightarrow \tilde{\pi}_i(\tau_{i,j}) > 0, \quad (0, \perp) \succeq (i, \tau_{i,j}) \Leftrightarrow \tilde{\pi}_i(\tau_{i,j}) = 0.$$

Next, we reason about what monotonicity constraints can possibly be tight simultaneously with a nested set of Border constraints that define $\succeq$ with the following lemma (whose proof appears in Appendix A).

Lemma 1. *For $\succeq$ defined as above for some $\tilde{\pi} \in P$, the $(i, j)^{th}$ monotonicity constraint can be tight and linearly independent of the tight Border constraints only if $(i, \tau_{i,j+1}) \succeq (i, \tau_{i,j})$.*

The remainder of the proof is just a counting argument. Define an equivalence relation $(i, \tau) \sim (j, \tau') \Leftrightarrow (i, \tau) \succeq (j, \tau') \succeq (i, \tau)$. Then because the tight Border constraints are nested, the number of tight Border constraints is exactly the number of equivalence classes under $\sim$ among (type, bidder) pairs $\succeq (0, \perp)$. It is also now clear, from Lemma 1, that a monotonicity constraint can be tight and linearly independent of the tight Border constraints only if it is between two types $(i, \tau_{i,j})$ and $(i, \tau_{i,j+1})$ of the same bidder i and $(i, \tau_{i,j}) \sim (i, \tau_{i,j+1})$. Therefore, there are no tight monotonicity constraints across equivalence classes, and the number of tight monotonicity constraints in each equivalence class is at most the number of types in that class minus one. Moreover, the number of tight monotonicity constraints in each equivalence class can only be equal to the number of types in that class minus one if all types in that class are from the same bidder. We simply observe that if $(i, \tau) \sim (j, \tau') \succeq (0, \perp)$ for any $i \neq j$, that the above counting shows we can't possibly have cn tight linearly independent constraints. So we may conclude that $(i, \tau) \not\sim (j, \tau')$ for any $(i, \tau), (j, \tau') \succeq (0, \perp)$.

Finally, we now want to conclude that the hierarchical allocation rule defined by $\succeq$ induces the proposed corner $\tilde{\pi}$. We make one slight modification to $\succeq$ to fit exactly Definition 5, and merge adjacent equivalence classes that contain types from the same bidder. More formally, if there exist two types $(i, \tau_{i,j})$, $(i, \tau_{i,j+1}) \succeq (0, \perp)$ that do not lie in the same equivalence class and no bidder $k \neq i$ has a type τ with $(i, \tau_{i,j}) \succeq (k, \tau) \succeq (i, \tau_{i,j+1})$, we merge the two equivalence classes by modifying $\succeq$ so that $(i, \tau_{i,j+1}) \succeq (i, \tau_{i,j})$ (but keeping $\succeq$ otherwise the same).[7] Now, it is clear that $\succeq$ is a weak total-ordering that is strict (because $(i, \tau) \not\sim (j, \tau')$ for any $(i, \tau), (j, \tau') \succeq (0, \perp)$) and partially-ordered w.r.t. to $\tilde{\pi}$, and the

[7] This adjustment is technically necessary to claim that when $\tilde{\pi}_i(\tau_{i,j}) = \tilde{\pi}_i(\tau_{i,j+1})$, we have $\tau_{i,j+1} \succeq \tau_{i,j}$. This adjustment doesn't affect the implementation of the hierarchical allocation rule according to $\succeq$ at all, since there is only ever one type per bidder present at the auction.

hierarchical allocation rule corresponding to $\succeq$ uniquely implements the corner $\tilde{\pi}$ (because the tight Border constraints uniquely determine a winner on every possible type profile, exactly the strongest type according to $\succeq$).

The final sentence of the theorem statement is again a consequence of Carathéodory's Theorem and Observation 1. $\square$

Now that we know that every corner of the polytope can be implemented as a strict, partially-ordered w.r.t. $\mathcal{R}$ hierarchical allocation rule, we want to ensure that the hierarchy $\succeq$ can be found computationally efficiently.

Lemma 2. *Let $\tilde{\pi}$ be any corner of P and $\succeq$ be the strict and partially-ordered w.r.t. $\mathcal{R}$ hierarchical allocation rule that implements $\tilde{\pi}$. Then the shaded reduced form defined as $\hat{\tilde{\pi}}_i(\tau) = \tilde{\pi}_i(\tau) \cdot \Pr\left[\tilde{\pi}_i(t_i) \leq \tilde{\pi}_i(\tau)\right]$ respects $\succeq$. Specifically, for any two types (i, τ) and (j, τ'), $(i, \tau) \succeq (j, \tau') \iff \hat{\pi}_i(\tau) \geq \hat{\pi}_j(\tau')$. Therefore, given $\tilde{\pi}$ we can construct the ordering $\succeq$ in time $O(cn \log(cn))$.*

Proof. First, observe that for any type (i, τ), $\tilde{\pi}_i(\tau) = \prod_{k \neq i} \Pr[(i, \tau) \succeq (k, t_k)]$ and $\Pr\left[\tilde{\pi}_i(t_i) \leq \tilde{\pi}_i(\tau)\right] = \Pr[(i, \tau) \succeq (i, t_i)]$. Therefore, $\hat{\tilde{\pi}}_i(\tau) = \prod_{k=1}^{n} \Pr[(i, \tau) \succeq (k, t_k)]$. We may therefore immediately conclude that $(i, \tau) \succeq (j, \tau') \Leftrightarrow \hat{\tilde{\pi}}_i(\tau) \geq \hat{\tilde{\pi}}_j(\tau')$. So in order to find the ordering $\succeq$, we only need to compute the shaded reduced form and sort its components, which can clearly be done in time $O(cn \log(cn))$. $\square$

And now, we can draw the main conclusion of this section: given as input any reduced form, we can computationally efficiently determine whether or not it is feasible. If it is feasible, we can computationally efficiently output an implementation.

Corollary 3. *Given an asymmetric reduced form $\mathcal{R}$, one can determine if it is feasible or find a hyperplane separating it from the set of feasible reduced forms, in time $O(cn \log(cn))$. If $\mathcal{R}$ is feasible, a succinct description of an allocation rule implementing $\mathcal{R}$ can be found in time polynomial in c and n. The output allocation rule is a distribution over at most $cn + 1$ hierarchical allocation rules, all strict and partially-ordered w.r.t. $\mathcal{R}$.*

Proof. We first observe that Theorem 1 provides an algorithm that determines if $\mathcal{R}$ is feasible, or provides a hyperplane separating it from the space of feasible reduced forms (the violated Border constraint) that runs in time $O(cn \log(cn))$. We now have to describe how to efficiently find an ex-post allocation rule implementing a reduced form $\mathcal{R}$ that is feasible. Theorem 6 implies that π lies inside a cn-dimensional polytope, P, whose corners are the reduced forms of the strict, partially-ordered w.r.t. $\mathcal{R}$ hierarchical allocation rules. Theorem 1 provides a separation oracle for P, so Theorem 4 implies that we can decompose π

into a convex combination of corners of P in time polynomial in cn and $cn\log(cn)$ (resulting in a runtime polynomial in both c and n). Lemma 2 shows how to implement a hierarchical allocation rule in time $O(cn\log(cn))$ given a corner. Observation 1 completes the proof. $\square$

4.3 A Tighter Characterization Result

Let's first briefly recall the goal of this section. Theorem 6 provides a nice characterization result: every feasible reduced form can be induced by a distribution over hierarchical allocation rules, all of which respect the same partial ordering within a single bidder's types, but may not respect any global ordering across all bidders' types. The purpose of this section is to prove Theorem 2 and show that in fact the distribution over hierarchical allocation rules may be taken to respect the same global ordering over all bidders' types.

At this point, we note that it would be great if the shaded reduced form provided an easy way to extend Theorem 5 to the asymmetric setting. Specifically, we can say that a hierarchical allocation rule $\succeq$ is *shaded-ordered* if $\hat{\pi}_i(\tau) \geq \hat{\pi}_j(\tau') \Rightarrow (i,\tau) \succeq (j,\tau')$, and hope it is the case that every feasible reduced form can be induced by a distribution over shaded-ordered hierarchical allocation rules. Unfortunately, although the shaded reduced form provides a nice structural theorem about feasible reduced forms and a near-linear time algorithm for determining feasibility, the following example shows that distributions over shaded-ordered hierarchical allocation rules are not sufficient to implement every feasible reduced form when the bidders are non-i.i.d. For completeness, we rule out all possible shadings.

Proposition 2. *There exist feasible reduced forms that cannot be induced by distributions over shaded-ordered hierarchical allocation rules.*

Proof. Consider the following example with two bidders and two types, with $\epsilon < 1/4$. Bidder one has $\Pr[t_1 = H] = 1-\epsilon^2, \Pr[t_1 = L] = \epsilon^2$, $\pi_1(H) = 1-\epsilon^2$ and $\pi_1(L) = 1-2\epsilon^2$. Bidder two has $\Pr[t_2 = H] = \epsilon, \Pr[t_2 = L] = 1-\epsilon$, $\pi_2(H) = \epsilon$ and $\pi_2(L) = 0$.

Then for any shading, we have $\hat{\pi}_1(H) \geq (1-\epsilon^2)^2 > \epsilon \geq \hat{\pi}_2(H) \geq \epsilon\cdot(1-\epsilon) > \epsilon^2\cdot(1-2\epsilon^2) \geq \hat{\pi}_1(L)$. So any shaded-ordered reduced form necessarily has $(2,H) \succeq (1,L)$, and therefore cannot possibly have $\pi_1(L) > 1-\epsilon/2 \geq 1-2\epsilon^2$.

Observe also that this reduced form is clearly feasible: consider the allocation rule that awards the item to bidder one whenever $t_2 = L$, awards the item to bidder one with probability $1-\epsilon$ when $t_1 = t_2 = H$ (and bidder two otherwise), and awards the item to bidder one with probability $1-2\epsilon$, to bidder two with probability ϵ when $t_1 = L, t_2 = H$ (and throws the item away otherwise). Then bidder one receives the item with probability $1-\epsilon^2$

when $t_1 = H$, with probability $1 - 2\epsilon^2$ when $t_1 = L$, and bidder two receives the item with probability ϵ when her type is H, and 0 otherwise. $\square$

In light of Proposition 2, it seems that geometric techniques will not get us all the way to a proof of Theorem 2 (which provides a global ordering instead of just a partial ordering as in Theorem 6), so our proof below appeals more to analytical tools. Throughout the proof, we will use the term $\succ$-ordered hierarchical allocation rule to denote a hierarchical allocation rule corresponding to some weak total ordering $\succeq$ that irons the strict total ordering $\succ$.

Proof of Theorem 2: The high-level approach is to find a strict total ordering, $\succ$, and an allocation rule that is a distribution over $\succ$-ordered hierarchical allocation rules, M, that is "closest" to inducing $\mathcal{R}$ by some measure (over all $\succ, M$). We will then argue that if M does not already induce $\mathcal{R}$, then the fact that we cannot improve M witnesses a violated Border constraint.

We first formally introduce a dummy bidder 0 with $\Pr[t_0 = \perp] = 1$ and $\pi_0(\perp) = 1 - \sum_{i>0} \sum_{\tau_i} \Pr[t_i = \tau_i]\pi_i(\tau_i)$. With this addition, we now have $\sum_{i\geq 0} \sum_{\tau_i} \Pr[t_i = \tau_i]\pi_i(\tau_i) = 1$. So if we find a feasible allocation rule M whose reduced form π^M satisfies $\pi_i^M(\tau) \geq \pi_i(\tau)$ for all $i \geq 0, \tau$, then we must have $\pi^M = \pi$, and M induces $\mathcal{R}$.

Let $\succ$ be a strict total ordering over all possible types that respects all the per-bidder partial orderings induced by π. Namely, for all i, if $(i, \tau) \succ (i, \tau')$, then $\pi_i(\tau) \geq \pi_i(\tau')$. Define the unhappiness $F_\succ(M)$ of a distribution over $\succ$-ordered hierarchical allocation rules, M (with reduced form π^M), as follows:

$$F_\succ(M) = \max_{i\geq 0, \tau\in T} (\pi_i(\tau) - \pi_i^M(\tau)).$$

$F_\succ$ can be viewed as a continuous function over a compact set: There are finitely many $\succ$-ordered hierarchical allocation rules, so their convex hull is a compact set (and exactly the space of distributions over $\succ$-ordered hierarchical allocation rules). Each function $\pi_i(\tau) - \pi_i^M(\tau)$ is linear in this space (and therefore continuous), and the maximum of continuous functions is continuous. Hence, $F_\succ$ achieves its infimum. Let then $M^\succ \in \operatorname{argmin}_M F_\succ(M)$ (where the minimization is over all distributions over $\succ$-ordered hierarchical allocation rules) and define the set $S_\succ$ to be the set of maximally unhappy types under $M^\succ$; formally, $S_\succ = \operatorname{argmax}_{i,\tau}\{\pi_i(\tau) - \pi_i^{M^\succ}(\tau)\}$. If for some $\succ$ there are several minimizers $M^\succ$, choose one that minimizes $|S_\succ|$. Now, let MO (stands for Minimal Orderings) be the set of the orderings $\succ$ that minimize $F_\succ(M^\succ)$, further refined to only contain $\succ \in MO$ that also minimizing $|S_\succ|$. Formally, first set $MO = \operatorname{argmin}_\succ\{F_\succ(M^\succ)\}$ and then refine MO as $MO_{\text{new}} = \operatorname{argmin}_{\succ\in MO}\{|S_\succ|\}$. We drop the subscript "new" for the rest of the proof.

From now on, we call a (bidder, type) pair (i, τ) *happy* if $\pi_i^M(\tau) \geq \pi_i(\tau)$, otherwise we

call (i, τ) *unhappy*. Intuitively, here is what we have already done: For every ordering $\succ$, we have found a distribution over $\succ$-ordered hierarchical allocation rules $M^{\succ}$ that minimizes the maximal unhappiness and subject to this, the number of maximally unhappy types. We then choose from these $(\succ, M^{\succ})$ pairs those that minimize the maximal unhappiness, and subject to this, the number of maximally unhappy types. We have made these definitions because we want to eventually show that there is an ordering $\succ$, such that $F_{\succ}(M^{\succ}) = 0$, and it is natural to start with the ordering that is "closest" to satisfying this property.

What we will show next is that, if $\exists \succ \in MO$ that does not make every (bidder, type) pair happy, then there also exists some $\succ' \in MO$, such that $F_{\succ'}(M^{\succ'}) = F_{\succ}(M^{\succ})$, $|S_{\succ'}| = |S_{\succ}|$, and $(i, \tau) \succ' (j, \tau')$ for all $(i, \tau) \in S_{\succ'}, (j, \tau') \notin S_{\succ'}$. In other words, only the top $|S_{\succ'}|$ types in $\succ'$ are maximally unhappy. From here, we will show that because $\succ' \in MO$, that $S_{\succ'}$ is a constricting set for $\mathcal{R}$, contradicting its feasibility. We begin by showing the existence of $\succ'$, beginning with an arbitrary $\succ \in MO$.

Before we begin, we introduce some terminology. We say that two (bidder, type) pairs $(i, \tau), (j, \tau')$ are *adjacent* if $(i, \tau) \succ (k, \tau'') \Leftrightarrow (j, \tau') \succ (k, \tau'')$ for all $(k, \tau'') \notin \{(i, \tau), (j, \tau')\}$. For any $\succeq$, we also define an equivalence relation $\sim_{\succeq}$ with $(i, \tau) \sim_{\succeq} (j, \tau') \Leftrightarrow (i, \tau) \succeq (j, \tau') \wedge (j, \tau') \succeq (i, \tau)$. Finally, we say that there is a *cut* between two adjacent types (i, τ) and (j, τ') in $\succeq$ if $(i, \tau) \succ (j, \tau')$ and $(j, \tau') \not\succeq (i, \tau)$. When we talk about adding a cut below (i, τ), we mean modifying $\succeq$ so that $(j, \tau') \not\succeq (i, \tau)$ for all $(i, \tau) \succ (j, \tau')$ (but otherwise keeping $\succeq$ the same). When we talk about removing a cut between two equivalence classes A and B, we mean modifying $\succeq$ so that $(i, \tau) \sim_{\succeq} (j, \tau')$ for all $(i, \tau), (j, \tau') \in A \cup B$ (but otherwise keeping $\succeq$ the same).

Now, if $S_{\succ}$ is not the highest $|S_{\succ}|$ (bidder, type) pairs, let (i, τ) be the maximal element under $\succ$ in $S_{\succ}$ such that there exists some $(k, \tau'') \notin S_{\succ}$ with $(k, \tau'') \succ (i, \tau)$. Then the adjacent (bidder, type) pair (j, τ') with $(j, \tau') \succ (i, \tau)$ is necessarily $\notin S_{\succ}$. We proceed to show that we can change $\succ$ to swap $(i, \tau) \succ (j, \tau')$ (keeping M, $S_{\succ}$ and $F_{\succ}(M)$ as-is). We can repeat these swaps iteratively and they will terminate in the $\succ'$ we want (with $S_{\succ'}$ equal to the first $|S_{\succ'}|$ (bidder, type) pairs).

We now proceed with a case analysis, for fixed $(j, \tau') \notin S_{\succ}$, $(i, \tau) \in S_{\succ}$, $(j, \tau') \succ (i, \tau)$ and $(j, \tau'), (i, \tau)$ adjacent.

- **Case 1:** $i = j$.

  Since $\succ$ is a linear extension of the bidder's own ordering, we must have $\pi_i(\tau') \geq \pi_i(\tau)$, but we know that

  $$\pi_i(\tau') - \pi_i^{M^{\succ}}(\tau') < \pi_i(\tau) - \pi_i^{M^{\succ}}(\tau),$$

  thus $\pi_i^{M^{\succ}}(\tau') > \pi_i^{M^{\succ}}(\tau) \geq 0$. In any hierarchical mechanism $\succeq$, if there is no cut

between (i, τ') and (i, τ), then they would receive the item with the same probability. Therefore, there must exist some $\succeq$ in the support of $M^{\succ}$ with a cut below (i, τ'), and in which (i, τ') gets the item with non-zero probability. We modify $M^{\succ}$ by modifying the hierarchical allocation rules $\succeq$ in its support as follows.

Let $\succeq$ be a hierarchical allocation rule in the support of $M^{\succ}$. If there is no cut below (i, τ'), we do nothing. If all (bidder, type) pairs equivalent to (i, τ') *and* those equivalent to (i, τ) are of bidder i, we remove the cut below (i, τ'). This does not affect the allocation probabilities at all, because it was impossible for two types equivalent to either (i, τ') or (i, τ) to show up together anyway. So after this "modification," we haven't changed $M^{\succ}$ at all, meaning that there must still exist some $\succeq$ in the support of $M^{\succ}$ with a cut below (i, τ'), and in which (i, τ') gets the item with non-zero probability, and clearly it is not one of the allocation rules we just modified by removing the cut below (i, τ'). For such an $\succeq$, there is at least one (bidder, type) pair with bidder $\neq i$ equivalent to (i, τ') or (i, τ). We distinguish two sub-cases:

- Every bidder $k \neq i$ has at least one type τ_k such that $(i, \tau) \succeq (k, \tau_k)$ (in other words, every (bidder, type) pair equivalent to (i, τ) wins the item with non-zero probability). Consider now moving the cut from below (i, τ') to right above (i, τ'). Clearly, (i, τ') will be less happy if we do this. Every (bidder, type) pair with bidder $\neq i$ that was formerly equivalent to (i, τ') will be strictly happier, as now they do not have to share the item with (i, τ'), whereas previously they did with positive probability. Every (bidder, type) pair with bidder $\neq i$ that is equivalent to (i, τ) will be strictly happier, as they now get to share the item with (i, τ'), whereas previously they always lost to (i, τ'). It is also clear to see that all (i, τ''), for $\tau'' \neq \tau'$ are unaffected by this change. So in particular (i, τ) is unaffected.

  Consider instead moving the cut from below (i, τ') to right below (i, τ). Then (i, τ) is clearly strictly happier, every (bidder, type) pair with bidder $\neq i$ that was formerly equivalent to (i, τ) is less happy than before (as they now don't get to share with (i, τ)), every (bidder, type) pair with bidder $\neq i$ that is equivalent to (i, τ') is also less happy than before (because now they have to share with (i, τ')), and all (i, τ''), for $\tau'' \neq \tau$ are not affected by the change.

  To summarize, we have argued that when we move the cut from below (i, τ') to just below (i, τ), (i, τ) becomes strictly happier, and every (bidder, type) pair that becomes less happy by this change becomes strictly happier if we instead move the cut to just above (i, τ') instead. Also, (i, τ) is unaffected by moving the cut to just above (i, τ). So with a tiny probability ϵ, move the cut from below

(i, τ') to just above (i, τ'), whenever $\succeq$ is sampled from $M^{\succ}$. This makes all of the (bidder, type) pairs with bidder $\neq i$ that were equivalent to either (i, τ') or (i, τ) strictly happier. With a tinier probability δ, move the cut from below (i, τ') to below (i, τ), whenever $\succeq$ is sampled from $M^{\succ}$. Choose ϵ to be small enough that we don't make (i, τ') maximally unhappy, and choose δ to be small enough so that we don't make any (type, bidder) pairs besides (i, τ') less happy than they were in $\succeq$. Then we have strictly increased the happiness of (i, τ) without making (i, τ') maximally unhappy, or decreasing the happiness of any other (bidder, type) pairs. Therefore, we have reduced $|S_{\succ}|$, contradicting the choice of $M^{\succ}$.

- If there is a bidder k such that $(i, \tau) \not\succeq (k, \tau_k)$ for all τ_k, (call such bidders *high*), then no (bidder, type) pair equivalent to (i, τ) can possibly win the item. We also know that every high bidder k has at least one type τ_k such that $(k, \tau_k) \sim_{\succeq} (i, \tau')$ by our choice of $\succeq$ (otherwise (i, τ') would get the item with probability 0). Now we can basically use the same argument as above. The only difference is that when we move the cut to just above (i, τ') or just below (i, τ), (bidder, type) pairs formerly equivalent to (i, τ) (other than (i, τ) itself) will remain unaffected. But since every high bidder k has a type τ_k with $(k, \tau_k) \sim_{\succeq} (i, \tau')$, (i, τ) will be strictly happier if we move the cut to just below (i, τ). Therefore, it is still the case that every (bidder, type) pair who is made unhappier by moving the cut to just below (i, τ) is made strictly happier by moving the cut to just above (i, τ'). So we can carry over the same reasoning as above (choosing ϵ, δ sufficiently small), and again contradict the choice of $M^{\succ}$.

Therefore, it can not be the case that $i = j$.

- **Case 2:** $i \neq j$ and there is never a cut below (j, τ').

  This case is easy. If we switch (j, τ') and (i, τ) in $\succ$, then the set $S_{\succ}$ is exactly the same, and the distribution $M^{\succ}$ is exactly the same. However, we have now relabeled the types in $S_{\succ}$ to get closer to the top $|S_{\succ}|$ elements being in $S_{\succ}$. Note that all $\succeq$ with no cut below (j, τ') are all $\succ$-ordered for the new $\succ$ as well, so this is a valid swap.

- **Case 3:** $i \neq j$ and there is sometimes a cut below (j, τ').

  Pick a $\succeq$ in the support of $M^{\succ}$ that has a cut between (j, τ') and (i, τ) and in which (j, τ') gets the item with positive probability. Note that if such a $\succeq$ doesn't exist, we can remove the cut below (j, τ') in all $\succeq$ in the support of $M^{\succ}$ without changing the allocation probabilities and return to Case 2. From here, there are again two subcases:

- $(k, \tau'') \notin S_{\succ}$ for all $(k, \tau'') \sim_{\succeq} (j, \tau')$. This means that *all* (bidder, type) pairs equivalent to (j, τ') are *not* maximally unhappy. Therefore, if we pick a tiny ϵ and remove the cut below (j, τ') with probability ϵ, only the types equivalent to (j, τ') will become unhappier. So there is a sufficiently small $\epsilon > 0$ for which this operation does not create any new maximally unhappy types. At the same time, because $i \neq j$ and (j, τ') receives the item with non-zero probability under $\succeq$, this operation makes (i, τ) strictly happier, as she now sometimes shares the item with (j, τ') (whereas previously she always lost). So this operation will create no new maximally unhappy (bidder, type) pairs, while making (i, τ) strictly happier, decreasing the size of $|S_{\succ}|$ and contradicting the choice of $M^{\succ}$.
- There exists a $(k, \tau'') \in S_{\succ}$ with $(k, \tau'') \sim_{\succeq} (j, \tau')$ (and $(k, \tau'') \neq (j, \tau')$). Let (k, τ'') be the minimal such (bidder, type) pair under $\succ$. Note that by our choice of (i, τ), that all $(\ell, \tau''') \succ (k, \tau'')$ are also in $S_{\succ}$ (maximally unhappy). Now consider introducing a cut below (k, τ'') with some tiny probability ϵ. Then the only (bidder, type) pairs who may become unhappier with this change are those that are still equivalent to (j, τ'), and all such types are *not* maximally unhappy. The only (bidder, type) pairs who may become happier with this change are those that are (k, τ'') or those that are $\succ (k, \tau'')$, all of which are in $S_{\succ}$. So if *any* of these types become happier at all with this change, there is a sufficiently small probability ϵ with which we can make this change without introducing any new maximally unhappy types and therefore decreasing $|S_{\succ}|$, a contradiction. So we must not make any (bidder, type) pairs happier with this change, and therefore we must also not make any (bidder, type) pairs unhappier (note that it is impossible to make any (bidder, type) pair unhappier without making some other (bidder, type) pair happier, since we are treating $(0, \perp)$ as a regular type). So we may in fact introduce a cut below (k, τ'') with probability 1 whenever $M^{\succ}$ samples $\succeq$ without affecting $\pi^{M^{\succ}}$ at all, but removing the original $\succeq$ from the support of $M^{\succ}$, and replacing it with a $\succeq$ in which all (bidder, type) pairs equivalent to (j, τ') are *not* maximally unhappy. After doing this for all such $\succeq$, we must return to the previous sub-case, after which we again obtain a contradiction.

Hence, it can not be the case that $i \neq j$ with a cut sometimes below (j, τ').

At the end of all three cases, we see that if we ever have $(j, \tau') \notin S_{\succ}$ and $(i, \tau) \in S_{\succ}$, with $(j, \tau') \succ (i, \tau)$ and $(j, \tau'), (i, \tau')$ adjacent, then we must have $i \neq j$, and no $\succeq$ in the support of $M^{\succ}$ ever places a cut directly below (j, τ'). Hence, we can simply swap the order

of these types in $\succ$ without affecting $S_\succ$ or $F_\succ(M)$ (as we described in Case 2 above), and we do that repeatedly until $S_\succ$ is equal to the top $|S_\succ|$ (bidder, type) pairs according to $\succ$.

Now that we have shown the existence of such a $\succ$, we show that it implies a constricting set. Label the elements in $S_\succ$ as $(i_1, \tau_1) \succ \ldots \succ (i_k, \tau_k)$ $(k = |S_\succ|)$. Now consider a $\succeq$ in the support of $M^\succ$ that has no cut below (i_k, τ_k), and consider putting a cut there with some tiny probability ϵ whenever $\succeq$ is sampled. The only effect this *might* have is that when the item is awarded to a (bidder, type) pair outside $S_\succ$, it is now awarded to a (bidder, type) pair inside $S_\succ$ instead with some probability. Therefore, if anyone gets happier, it is someone in $S_\succ$. However, if we make anyone in $S_\succ$ happier and choose ϵ small enough so that we don't make anyone outside of $S_\succ$ maximally unhappy, we decrease $|S_\succ|$, contradicting the choice of $M^\succ$. Therefore, putting a cut below (i_k, τ_k) cannot possibly make anyone happier, and therefore cannot make anyone unhappier. So we may w.l.o.g. assume that there is a cut below (i_k, τ_k) in all $\succeq$ in the support of $M^\succ$. But now we get that the item always goes to someone in $S_\succ$ whenever a (bidder, type) pair in $S_\succ$ is reported, yet all (bidder, type) pairs in this set are unhappy. Therefore, $S_\succ$ is a constricting set, certifying that the given $\mathcal{R}$ is infeasible.

Putting everything together, we have shown that if there is no $\succ$ with $F_\succ(M^\succ) = 0$ then the reduced form is infeasible. So there must be some $\succ$ with $F_\succ(M^\succ) = 0$, and such an $M^\succ$ induces the reduced form by sampling only $\succ$-ordered hierarchical allocation rules, completing the proof. $\square$

5 Multi-Item Mechanism Design

In this section, we show how our results above on reduced forms can be useful for multi-item mechanism design as well. Essentially, our key observation is that an m-item interim allocation rule is feasible if and only if the m projected single-item reduced forms are feasible, so the question is simply whether or not an m-item reduced form contains enough useful information for mechanism design. When buyers are additive, this information indeed suffices to guarantee that a mechanism inducing to this interim allocation rule is *Bayesian Incentive Compatible*, which allows us to formulate and solve an optimization problem. We make these statements more precise shortly, but first provide some notation specific to multi-item mechanism design not covered in Section 2.

5.1 Notation

For Section 5, there are n bidders and m items. All bidders' valuation functions are additive. We write $\vec{v}_i$ to denote the type of bidder i, with the convention that v_{ij} represents her value for item j and that her value for a bundle S of items is simply $\sum_{j \in S} v_{ij}$. We still let T denote the space of possible types, which is now a subset of $\mathbb{R}^n$.

To fully specify a (direct-revelation) multi-item mechanism for additive bidders, we need to describe, potentially succinctly, for all type profiles $\vec{v} \in T^n$, and for every bidder i, the outcome $M_i(\vec{v}) = (\vec{\phi}_i(\vec{v}), p_i(\vec{v}))$ given by M to bidder i when the reported bidder types are $\vec{v}$. Here, $\phi_{ij}(\vec{v})$ is the ex-post probability that item j is given to bidder i when the reported types are $\vec{v}$, and $p_i(\vec{v})$ is the ex-post price that i pays. The *value* of bidder i for outcome $M_i(\vec{w})$ is just her expected value $\vec{v}_i \cdot \vec{\phi}_i(\vec{w})$, while her utility is *quasi-linear,* meaning that bidder i's utility for the same outcome is $U(\vec{v}_i, M_i(\vec{w})) := \vec{v}_i \cdot \vec{\phi}_i(\vec{w}) - p_i(\vec{w})$. The relation between the ex-post probabilities ϕ and interim probabilities π is just the following: for all i,j, $\vec{v}_i \in T$: $\pi_{ij}(\vec{v}_i) = \mathbb{E}_{\vec{v}_{-i} \sim \mathcal{D}_{-i}}[\phi_{ij}(\vec{v}_i \ ; \ \vec{v}_{-i})]$. We now formally define Bayesian Incentive Compatibility (BIC):

Definition 6. *(Bayesian Incentive Compatible Mechanism) A mechanism M is called BIC iff the following inequality holds for all $i \in [n], \vec{v}_i, \vec{w}_i \in T$:*

$$\mathbb{E}_{\vec{v}_{-i} \sim \mathcal{D}_{-i}} \left[U(\vec{v}_i, M_i(\vec{v})) \right] \geq \mathbb{E}_{\vec{v}_{-i} \sim \mathcal{D}_{-i}} \left[U(\vec{v}_i, M_i(\vec{w}_i \ ; \ \vec{v}_{-i})) \right].$$

5.2 Optimal Multi-Item Mechanism Design

We begin this section with our key observation, essentially stating that some single-item results (specifically, those in Section 4) can be extended for free to some multi-item settings. Let us begin by being clear what we mean by an interim allocation rule projecting a reduced form onto item j.

Definition 7. *(Projected reduced form) Let there be m heterogeneous items, T be some arbitrary type space, and $\mathcal{R} = \{\pi_{ij}(\cdot)\}_{i \in [n], j \in [m]}$ be some interim allocation rule of a mechanism. Then the projected reduced form of $\mathcal{R}$ onto item j, $\mathcal{R}_j$, is just $\mathcal{R}_j = \{\pi_{ij}(\cdot)\}_{i \in [n]}$. Note that $\mathcal{R}_j$ still takes as input types in the original type space T.*

Observation 3. An m-item interim allocation rule $\mathcal{R}$ is feasible if and only if for all j, the projected reduced form $\mathcal{R}_j$ onto item j is feasible. Furthermore, if $\mathcal{R}$ is feasible, $\mathcal{R}$ is induced by the ex-post allocation rule that allocates each item j according to $\mathcal{R}_j$ independently of the others.

Proof. First, assume that $\mathcal{R}$ is feasible, and let M be an ex-post allocation rule that induces $\mathcal{R}$. We wish to come up with an allocation rule for item j that induces $\mathcal{R}_j$. Define M_j[8] to be the single-item allocation rule that runs M and awards the single item to whichever bidder was awarded item j under M. Clearly, M_j implements $\mathcal{R}_j$, so if $\mathcal{R}$ is feasible, so is each projection $\mathcal{R}_j$.

Next, assume that each $\mathcal{R}_j$ is feasible, and let M_j be an ex-post allocation rule that induces $\mathcal{R}_j$. Then let M be the allocation rule that on every input type profile, runs M_j on that type profile for all j and awards item j to whoever receives the single item under M_j. Clearly, the projection of the reduced form of M onto item j will be exactly $\mathcal{R}_j$, so the reduced form of M is exactly $\mathcal{R}$. Therefore, M induces $\mathcal{R}$, and $\mathcal{R}$ is feasible.

So $\mathcal{R}$ is feasible if and only if each projection $\mathcal{R}_j$ is feasible. Furthermore, the above argument shows that when $\mathcal{R}$ is feasible, $\mathcal{R}$ can be induced by an ex-post allocation rule that allocates each item separately. □

Observation 3 combined with Theorem 2 immediately yields our characterization of feasible multi-item interim allocation rules (from Section 1). Replacing Theorem 2 with Theorem 5 provides a tighter characterization in the symmetric case. Note that at this point we have absolutely not addressed the issue of when this characterization is useful for multi-item mechanism design - all we have done is observed (somewhat trivially) that the structure of reduced forms is preserved under concatenation.

So now, let's address this issue and discuss the multi-item settings in which Observation 3 is useful for mechanism design. Essentially, we observe that the interim allocation rule as we have defined it provides sufficient information to determine whether or not a mechanism is BIC if and only if bidders' valuations are additive. We first present an example illustrating that this fails, for instance, when bidders are instead unit-demand.[9]

Example 1. *Consider a setting with a single unit-demand bidder, two items, and one possible type,* $(1, 1)$ *(value one for each item). Consider the following two ex-post allocation rules:*

- *Pick* j *uniformly random from* $\{1, 2\}$ *and award item* j*.*
- *Award the set of items* $\{1, 2\}$ *with probability* $1/2$*, and the set* $\emptyset$ *with probability* $1/2$*.*

Then these two ex-post allocation rules have the same interim allocation rule: $\pi_{11}(1, 1) = \pi_{12}(1, 1) = 1/2$*. But the bidder's expected value under the first ex-post allocation rule is* 1*, whereas under the second it is* $1/2$*. Therefore, the interim allocation rule simply does not*

[8] Note that M_j takes the *whole type* $\vec{v_i}$ of each bidder i as input, and not just v_{ij}.

[9] A bidder is unit-demand if whenever they have value v_j for item j, their value for set S is equal to $\max_{j \in S}\{v_j\}$.

contain enough information for the bidder to compute her expected value for reporting a given type to an ex-post allocation rule inducing it - because it depends on which ex-post allocation rule is chosen.

Note that if instead the bidder were additive, she would have expected value 1 *under both ex-post allocation rules, and this would not be an issue.*

Observation 4. When bidders are additive, the per-item interim allocation rule $\mathcal{R}$ contains enough information for a bidder to compute her expected value when reporting type $\vec{w}_i$ to *any* ex-post allocation rule inducing $\mathcal{R}$. It is exactly $\sum_j v_{ij} \cdot \pi_{ij}(\vec{w}_i)$.

Proof. Let M be any ex-post allocation rule with ex-post allocation probabilities $\phi_{ij}(\cdot)$. Then we can write the expected value of a buyer with type $\vec{v}_i$ for reporting $\vec{w}_i$ to M as:

$$\begin{aligned}\mathbb{E}_{\vec{v}_{-i}\sim\mathcal{D}_{-i}}[\sum_j v_{ij} \cdot \phi_{ij}(\vec{w}_i \ ; \ \vec{v}_{-i})] &= \sum_j \mathbb{E}_{\vec{v}_{-i}\sim\mathcal{D}_{-i}}[v_{ij} \cdot \phi_{ij}(\vec{w}_i \ ; \ \vec{v}_{-i})] \\ &= \sum_j v_{ij} \cdot \mathbb{E}_{\vec{v}_{-i}\sim\mathcal{D}_{-i}}[\phi_{ij}(\vec{w}_i \ ; \ \vec{v}_{-i})] = \sum_j v_{ij} \cdot \pi_{ij}(\vec{w}_i).\end{aligned}$$

□

Essentially what makes additive buyers unique in comparison to other multi-dimensional valuation functions is that the marginal value for item j is completely independent of the set it is being added to. Let us again emphasize that an ex-post allocation rule M implementing an interim allocation rule $\mathcal{R}$ absolutely takes into consideration the *value* of bidders for items $\ell \neq j$ when determining how to allocate item j, as this information is stored in their types. However, Observations 3 shows that it need not consider how items $\ell \neq j$ are themselves *allocated* when determining how to allocate item j. It is well-known that even when there is just a single additive bidder and two items, and even if the bidder's values for these items are distributed i.i.d., that the allocation rule of the revenue-optimal mechanism necessarily considers values for items $\ell \neq j$ when deciding the allocation of item j. Below is a folklore example (that appears concretely, for instance, in [DDT14]). This shows that even if we are willing to restrict to a characterization only of interim allocation rules that are revenue-optimal for simple multi-item instances, we should not hope for a stronger characterization that (say) allocates each item j independent of bidders' values for items $\ell \neq j$.

Example 2. *There is a single additive bidder and two items. Each v_j is drawn independently and uniformly from the set $\{1, 2\}$. Then the revenue optimal mechanism awards both items to the bidder whenever $v_1 + v_2 \geq 3$, and charges price* 3.

Note that when $v_1 = 1$, whether or not the bidder receives item 1 *depends on v_2 (namely, she receives item* 1 *iff $v_2 = 2$).*

With all this in mind, let us now formally state the multi-item auction problem we solve:

BIC Multi-item auction. Given as input n distributions $\mathcal{D}_1, \ldots, \mathcal{D}_n$ over valuation vectors for m items, output a BIC mechanism M whose expected revenue is optimal relative to any other, possibly randomized, BIC mechanism, when played by n additive bidders whose valuation vectors are sampled independently from $\mathcal{D}_1, \ldots, \mathcal{D}_n$. Note that each $\mathcal{D}_i$ need not be a product distribution, a single bidder's values for different items may be arbitrarily correlated.

Our approach to solving proving Theorem 3 is to use the separation oracle for checking the feasibility of a reduced form developed in Corollary 3 inside a linear program that optimizes over all feasible interim allocation rules and interim price rules. We begin with the LP formulation in Figure 1 below, followed by a proof that the LP is correct.

Variables:

- $\pi_{ij}(\vec{v}_i)$, for all bidders $i \in [n]$, items $j \in [m]$, and $\vec{v}_i \in T$, the interim probability that bidder i gets item j when reporting type $\vec{v}_i$ (mnc variables).
- $q_i(\vec{v}_i)$, for all bidders $i \in [n]$, $\vec{v}_i \in T$, the interim expected price that bidder i pays when reporting type $\vec{v}_i$ (nc variables).

Constraints:

- $0 \leq \pi_{ij}(\vec{v}_i) \leq 1$, for all $i \in [n]$, $j \in [m]$, $\vec{v}_i \in T$, guaranteeing that each $\pi_{ij}(\vec{v}_i)$ is a probability (mnc constraints).
- $\sum_{j\in[m]} v_{ij}\pi_{ij}(\vec{v}_i) - q_i(\vec{v}_i) \geq 0$, for all $i \in [n]$, $\vec{v}_i \in T$, guaranteeing that the mechanism is interim Individually Rational (interim IR) (nc constraints).
- $\sum_{j\in[m]} v_{ij}\pi_{ij}(\vec{v}_i) - q_i(\vec{v}_i) \geq \sum_{j\in[m]} v_{ij}\pi_{ij}(\vec{v}'_i) - q_i(\vec{v}'_i)$, for all $i \in [n], \vec{v}_i, \vec{v}'_i \in T$, guaranteeing that the mechanism is Bayesian Incentive Compatible (BIC) (nc^2 constraints).
- $SO(\vec{\pi}, \mathcal{D}) =$"YES", guaranteeing that there is an ex-post allocation rule inducing π;

Maximizing:

- $\sum_{i\in[n],\vec{v}_i\in T} q_i(\vec{v}_i) \Pr[\vec{v}_i \leftarrow \mathcal{D}_i]$, the expected revenue.

Figure 1: A folklore LP (that appears concretely, e.g., in [DW12]), where we use a separation oracle to determine feasibility of the interim allocation rule. In parentheses at the end of each line is the number of such variables/constraints .

Proposition 3. *Provided that SO acts as a valid separation oracle for the space of feasible interim allocation rules, the LP of Figure 1 outputs the revenue optimal interim allocation*

rule and interim price rule for BIC multi-item auction in time polynomial in n, m, c and the runtime of SO.

Proof. First, it is clear that any output of the LP of Figure 1 (henceforth, just "the LP") is interim IR, BIC, and feasible, as long as SO is correct. This is because the constraints on an interim allocation rule/price rule pair to be interim IR are linear and explicitly included in the LP. The same holds for BIC. Therefore, as long as SO is correct, any interim allocation rule/price rule pair accepted by the LP must be interim IR, BIC, and feasible. It is also clear that the objective function correctly computes the expected revenue of any interim price rule considered. So the LP outputs exactly the feasible, interim IR, BIC interim allocation rule/price rule pair that maximizes expected revenue with respect to all feasible, interim IR, BIC interim allocation rule/price rule pairs.

Second, it is clear that every feasible, IR, BIC mechanism has an interim allocation rule/price rule pair, and that this pair is also interim IR, feasible, and BIC. So the revenue-optimal interim allocation/price rule pair output by the LP is indeed optimal with respect to all feasible, IR, BIC mechanisms.

Finally, it is clear that the number of variables and constraints (excluding SO) is polynomial in nmc. Therefore, the LP can be solved in time polynomial in n, m, c and the runtime of SO via a direct application of Theorem 4. □

Proof of Theorem 3: Theorem 3 now follows immediately from Proposition 3, Observation 3, and Corollary 3. Corollary 3 combined with Observation 3 guarantees that we can design the desired separation oracle, which terminates in time $\text{poly}(n, m, c)$, and Proposition 3 guarantees that we can use this to solve the LP in time $\text{poly}(n, m, c)$. After finding the optimal interim allocation/price rule pair $(\vec{\pi}^*, \vec{q}^*)$, Corollary 3 shows how to find in time $\text{poly}(n, m, c)$ a succinct description of an ex-post allocation rule implementing $\vec{\pi}^*$ (as a distribution over at most $cn + 1$ strict hierarchical allocation rules). So one can implement the optimal mechanism by using this ex-post allocation rule, and charging bidder i price $q_i^*(\vec{v}_i)$ when her bid is $\vec{v}_i$. □

We also note that the LP of Figure 1 only requires that the mechanism be interim individually rational. A well-known simple trick (shown e.g. in [DW12]) converts any BIC, interim IR mechanism into one that is BIC and ex-post IR with no loss in revenue, so this is w.l.o.g. For completeness, we describe the trick below. Note that the trick does not require any assumptions on the valuation functions of bidders except that they are quasi-linear and risk-neutral.

Proposition 4. *Let M be any BIC and interim IR mechanism for quasi-linear and risk-neutral bidders. There exists another mechanism M' that uses exactly the same ex-post allocation rule and interim price rule as M that is also ex-post IR.*

Proof. For any bidder i and valuation function/type $v_i(\cdot) : 2^{[m]} \to \mathbb{R}$ (v_i takes as input a set of items and outputs a value), let $V_i^M(v_i)$ denote the expected value of bidder i when truthfully reporting type v_i to M (over any randomness in M, and the randomness in other bidders' types). Let also $q_i^M(v_i)$ denote the interim price paid by i when reporting type v_i to M (over the same randomness). As M is interim IR, we must have $q_i^M(v_i) \leq V_i^M(v_i)$.

Now define the following ex-post price rule for M': whenever bidder i receives items S after reporting type v_i, charge her $\frac{q_i^M(v_i)}{V_i^M(v_i)} \cdot v_i(S)$. Observe that we have not changed the ex-post allocation rule, so clearly M and M' have the same ex-post allocation rule. As $\frac{q_i^M(v_i)}{V_i^M(v_i)} \leq 1$, M' is clearly ex-post IR. All that remains to verify is that M and M' have the same interim price rule. But this is also clear: if we define $\phi_S(\vec{v})$ to be the probability that M allocates set S to bidder i on input $\vec{v}$, then the interim price paid by bidder i when reporting type v_i to M' is exactly:

$$\begin{aligned}\mathbb{E}_{\vec{v}_{-i}\sim\mathcal{D}_{-i}}[\sum_{S\subseteq[m]} \phi_S(\vec{v}) \cdot \frac{q_i^M(v_i)}{V_i^M(v_i)} \cdot v_i(S)] &= \frac{q_i^M(v_i)}{V_i^M(v_i)} \cdot \mathbb{E}_{\vec{v}_{-i}\sim\mathcal{D}_{-i}}[\sum_{S\subseteq[m]} \phi_S(\vec{v}) \cdot v_i(S)] \\ &= \frac{q_i^M(v_i)}{V_i^M(v_i)} \cdot V_i^M(v_i) = q_i^M(v_i).\end{aligned}$$

□

6 Conclusions and Discussion

Motivated by settings where an optimization approach is necessary to develop optimal auctions, we study single-item reduced-form auctions for asymmetric bidders. We provide a linear-sized subset of necessary and sufficient Border conditions (down from exponential-sized of prior work) that can be checked in nearly-linear time, and also show that every feasible reduced form can be implemented as a distribution of hierarchical allocation rules that all respect the same total ordering over all bidders' types. We further show that our results imply both polynomial-time algorithms and structural characterizations for multi-item auctions with additive bidders.

Our work demonstrates how a better understanding of reduced-form auctions in the core single-item setting can be utilized for much more general settings, and also that a computational lens can lead not only to tractable optimization algorithms, but also improved

structural understanding.

We conclude with a very brief discussion on how one could extend our approach for multi-item auction design beyond additive buyers. Indeed, the core issue is that the interim allocation probabilities no longer contain enough information to verify that a mechanism is BIC (Example 1), so in order to get mileage out of our approach one first needs to come up with a new interim description of auctions. It is not hard to come up with definitions that work for, say, unit-demand buyers: first observe that it is w.l.o.g. to only consider auctions which award each bidder at most one item.[10] Then simply define $\pi_{ij}(\vec{v}_i)$ to be the interim probablity that bidder i receives item j when submitting valuation vector $\vec{v}_i$, and exactly the same LP as written in Figure 1 finds the revenue-optimal BIC mechanism, provided we can design a computationally efficient separation oracle for this new space of feasible interim allocation rules (which never allocate the same item to multiple bidders *or* the same bidder multiple items). However, obtaining such a separation oracle is no longer simply the product of n single-item problems, and in fact Gopalan et. al. prove that, under well-believed complexity-theoretic assumptions, no computationally efficient separation oracle for this space exists [GNR15].

Still, our approach has proven useful for the design of *nearly*-optimal multi-item auctions in settings with unit-demand buyers (and in fact significantly more general settings as well) in follow-up works by the authors and others [CDW12b, CDW13a, CDW13b, BGM13, DDW15, CDW16], essentially by making use of computationally-efficient but *approximate* characterization of the feasible interim allocation probabilities in these settings. The present paper remains unique in this line of works for finding the optimal mechanism without approximation error.

References

[AFH+12] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012. 1.2

[Ala11] Saeed Alaei. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011. 1

[10] To see this, take any auction that awards multiple items to bidder i with reported type $\vec{v}_i$, and modify it to throw away all items except the highest value item (according to $\vec{v}_i$). Then if bidder i was telling the truth, their utility doesn't change at all. If she was lying, her utility might go down. So the mechanism is still truthful, and the revenue is the same because the payments are the same.

[BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing Lotteries. *Journal of Economic Theory*, 156:144–174, 2015. 1

[BGGM10] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget Constrained Auctions with Heterogeneous Items. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010. 1

[BGM13] Anand Bhalgat, Sreenivas Gollapudi, and Kamesh Munagala. Optimal Auctions via the Multiplicative Weight Method. In *the 14th ACM Conference on Economics and Computation (EC)*, 2013. 1, 6

[BILW14] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 21–30, 2014. 1

[Bor91] Kim C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4):1175–1187, 1991. (document), 1, 1, 1.2, 3, 1, 3, 1

[Bor07] Kim C. Border. Reduced Form Auctions Revisited. *Economic Theory*, 31:167–181, 2007. (document), 1, 1.2, 4, 4.2

[CD11] Yang Cai and Constantinos Daskalakis. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011. 1

[CDW12a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012. 1

[CDW12b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012. 1, 1, 6

[CDW13a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing Revenue to Welfare Maximization: Approximation Algorithms and other Generalizations. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013. 1, 1, 6

[CDW13b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013. 1, 1, 6

[CDW16] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th ACM Conference on Theory of Computation (STOC)*, 2016. 1, 6

[CH13] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013. 1

[CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010. 1

[CKM13] Yeon-Koo Che, Jinwoo Kim, and Konrad Mierendorff. Generalized Reduced-Form Auctions: A Network-Flow Approach. *Econometrica*, 2013. (document), 1, 1.2, 4, 5, 4.2

[Cla71] Edward H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11(1):17–33, 1971. 1

[CM16] Shuchi Chawla and J. Benjamin Miller. Mechanism design for subadditive agents via an ex-ante relaxation. In *Proceedings of the seventeenth ACM Conference on Economics and Computation (EC)*, pages 579–596, 2016. 1

[CMS15] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015. 1

[CZ17] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017. 1

[DDT13] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism Design via Optimal Transport. In *The 14th ACM Conference on Electronic Commerce (EC)*, 2013. 1

[DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014. 1, 5.2

[DDT17] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767, 2017. 1

[DDW15] Constantinos Daskalakis, Nikhil Devanur, and S. Matthew Weinberg. Revenue maximization and ex-post budget constraints. In *the 16th ACM Conference on Economics and Computation (EC)*, 2015. 1, 6

[DW12] Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012. 1, 5.2

[GLS81] Martin Grötschel, László Lovász, and Alexander Schrijver. The Ellipsoid Method and its Consequences in Combinatorial Optimization. *Combinatorica*, 1(2):169–197, 1981. 2, 4

[GNR15] Parikshit Gopalan, Noam Nisan, and Tim Roughgarden. Public projects, boolean functions and the borders of border's theorem. In *the 16th ACM Conference on Economics and Computation (EC)*, 2015. 1, 6

[Gro73] Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973. 1

[HN13] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*, 2013. 1

[HR15a] Sergiu Hart and Philip J. Reny. Implementation of Reduced Form Mechanisms: A Simple Approach and a New Characterization. *Economic Theory Bulletin*, 3(1):1–8, 2015. (document), 1.2

[HR15b] Sergiu Hart and Philip J. Reny. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics*, 10(3):893–922, 2015. 1

[Kha79] Leonid G. Khachiyan. A Polynomial Algorithm in Linear Programming. *Soviet Mathematics Doklady*, 20(1):191–194, 1979. 2, 4

[KP82] Richard M. Karp and Christos H. Papadimitriou. On linear characterizations of combinatorial optimization problems. *SIAM J. Comput.*, 11(4):620–632, 1982. 2, 4

[KW12] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012. 1

[LR96] Jean-Jacques Laffont and Jacques Robert. Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186, 1996. 1

[Mat84] Steven Matthews. On the Implementability of Reduced Form Auctions. *Econometrica*, 52(6):1519–1522, 1984. 1, 1.2

[MR84] Eric Maskin and John Riley. Optimal Auctions with Risk Averse Buyers. *Econometrica*, 52(6):1473–1518, 1984. 1, 1, 1.2

[MV07] A. M. Manelli and D. R. Vincent. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007. 1

[MV10] A. M. Manelli and D. R. Vincent. Bayesian and Dominant-Strrategy Implementation in the Independent Private-Values Model. *Econometrica*, 78(6):1905–1938, 2010. (document), 1, 1, 1.2, 5, 6

[Mye81] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981. 1, 1.1, 1

[Pav11] Gregory Pavlov. Optimal mechanism for selling two goods. *The B.E. Journal of Theoretical Economics*, 11(3), 2011. 1

[RC98] Jean-Charles Rochet and Philippe Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, 1998. 1, 1

[RW15] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the 16th ACM Conference on Electronic Commerce*, 2015. 1

[Tha04] John Thanassoulis. Haggling over substitutes. *Journal of Economic Theory*, 117:217–245, 2004. 1

[Vic61] William Vickrey. Counterspeculations, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961. 1

[Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015. 1

A Omitted Proofs

The first omitted proof is of Corollary 1, which simply claimed that we can find a violated Border constraint in the symmetric case in time $O(c(\log c + \log n))$.

Proof of Corollary 1. First, sort the types τ in T in decreasing order according to $\pi(\tau)$ in time $O(c \log c)$, and label the distinct values of $\pi(\tau)$ as $x_1, \ldots, x_k$ ($k \leq c$), and define $X_i = \{\tau \mid \pi(\tau) = x_i\}$. Next, compute $\sum_{\tau:\pi(\tau)\geq x_i} \pi(\tau) \Pr[\tau]$ for all i in the following way: First, compute $\sum_{\tau \in X_i} \pi(\tau) \Pr[\tau]$ for all i. This can clearly be done in time $O(|X_i|)$ for all i, and therefore the total computation takes time $O(c)$. Next, observe that $\sum_{\tau:\pi(\tau)\geq x_{i+1}} \pi(\tau) \Pr[\tau] = \sum_{\tau:\pi(\tau)\geq x_i} \pi(\tau) \Pr[\tau] + \sum_{\tau \in X_{i+1}} \pi(\tau) \Pr[\tau]$. So we can compute $\sum_{\tau:\pi(\tau)\geq x_{i+1}} \pi(\tau) \Pr[\tau]$ from $\sum_{\tau:\pi(\tau)\geq x_i} \pi(\tau) \Pr[\tau]$ using just $O(1)$ additional computation. Therefore, we can compute $\sum_{\tau:\pi(\tau)\geq x_i} \pi(\tau) \Pr[\tau]$ for all i in total time $O(c)$. This immediately gives us the LHS for all inequalities (3) in total time $O(c \log c)$. Similarly, we can compute $\sum_{\tau \in X_i} \Pr[\tau]$ for all i in total time $O(c)$, and use these in the same way to compute $\Pr_{t\sim D_1}[\pi(t) \geq x_i] = \sum_{\tau:\pi(\tau)\geq x_i} \Pr[\tau]$ for all i in total time $O(c \log c)$. With this, we can then compute the RHS of each inequality (3) using repeated squaring in time $O(\log n)$ per inequality.

So all together, we have a pre-processing stage which takes time $O(c \log c)$, and an additional $O(c \log n)$ to compute each RHS, resulting in a total runtime of $O(c(\log c + \log n))$. □

The next omitted proof is of (the algorithmic portion of) Theorem 1, which claims that we can find a violated Border constraint in the asymmetric case in time $O(cn \cdot \log(cn))$.

Proof of Theorem 1: Proposition 1 immediately yields the first part of Theorem 1 (that a reduced form is feasible if and only if it satisfies Equation 2 for all x). The only remaining detail to check is that a violated Border constraint (if it exists) can be found in the desired runtime, which follows from a similar routine computation as in Corollary 1. For the sake of completeness, we include with a full proof of this.

First, sort all (bidder, type) pairs $(i, \tau) \in [n] \times T$ in decreasing order according to $\hat{\pi}_i(\tau)$ in time $O(nc \log(nc))$ ($c = |T|$), label the distinct values of $\hat{\pi}_i(\tau)$ as $x_1, \ldots, x_k$ ($k \leq nc$), and store $X_j = \{(i, \tau) \mid \hat{\pi}_i(\tau) = x_j\}$ (i.e. X_j is the set of (bidder, type) pairs for which $\hat{\pi}_i(\tau) = x_j$). Next, to compute all left-hand sides of Equation (2), begin by computing $\sum_{(i,\tau)\in X_j} \Pr[t_i = \tau] \cdot \pi_i(\tau)$ for all j. This can clearly be done in time $O(|X_j|)$ for all j, and therefore the total computation takes time $O(nc)$. Next, observe that $\sum_{(i,\tau):\hat{\pi}_i(\tau)\geq x_{j+1}} \Pr[t_i = \tau] \cdot \pi_i(\tau) = \sum_{(i,\tau):\hat{\pi}_i(\tau)\geq x_j} \Pr[t_i = \tau] \cdot \pi_i(\tau) + \sum_{(i,\tau)\in X_{j+1}} \Pr[t_i = \tau] \cdot \pi_i(\tau)$. So we can compute $\sum_{(i,\tau):\hat{\pi}_i(\tau)\geq x_{j+1}} \Pr[t_i = \tau] \cdot \pi_i(\tau)$ from $\sum_{(i,\tau):\hat{\pi}_i(\tau)\geq x_j} \Pr[t_i = \tau] \cdot \pi_i(\tau)$ with $O(1)$ additional computation (by making use of our pre-processed values). So starting with x_1, we can

compute all left-hand sides of Equation (2) with $O(1)$ additional computation per constraint, for a total computation time of $O(nc)$.

Computing all right-hand sides is a touch trickier, because we don't actually want to multiply n numbers together for each of nc inequalities. So we will make use of the fact that most of the terms going into the product on the RHS don't change between successive inequalities. To compute all right-hand sides, additionally for all i, sort all types $\tau \in T$ in decreasing order of $\hat{\pi}_i(\tau)$ in time $O(c \log c)$ per bidder, or $O(nc \log c)$ in total, label the distinct values of $\hat{\pi}_i(\tau)$ as $x_{i1}, \ldots, x_{ik_i}$ ($k_i \leq c$), and store $X_{ij} = \{\tau \mid \hat{\pi}_i(\tau) = x_{ij}\}$. Compute $\sum_{\tau \in X_{ij}} \Pr[t_i = \tau]$ for all i, j in total time $O(nc)$. Recall $S_x^{(i)} = \{\tau \mid \hat{\pi}_i(\tau) \geq x\}$, and observe that $\Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{ij}}^{(i)}] = \sum_{\tau : \hat{\pi}_i(\tau) \geq x_{ij}} \Pr[t_i = \tau]$ can be computed for fixed i and all j in time $O(c)$, and therefore for all i and all j in total time $O(nc)$. With this, finally compute $\frac{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{i(j+1)}}^{(i)}]}{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{ij}}^{(i)}]}$ for all i and $j \in \{0, \ldots, k_i - 1\}$ (denoting $S_{x_i 0}^{(i)} = \emptyset$), again in time $O(c)$ per bidder, or $O(nc)$ in total (it is just at most one division per $(i, \tau) \in [n] \times T$). Now, starting from $j = 1$, we will inductively compute $\prod_i (1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}])$ for all j, using the result for $j - 1$ as well as our precomputed terms. Observe that $\prod_i (1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}]) = \prod_i (1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{j-1}}^{(i)}]) \cdot \prod_i \frac{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}]}{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{j-1}}^{(i)}]}$. For many i, it might be that $\frac{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}]}{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{j-1}}^{(i)}]} = 1$ (namely, all i such that $(i, \tau) \notin X_j$ for all $\tau \in T$). So to compute $\prod_i \frac{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}]}{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{j-1}}^{(i)}]}$ simply multiply together for all i such that $\exists \tau \in T, (i, \tau) \in X_j$ the corresponding $\frac{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}]}{1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_{j-1}}^{(i)}]}$, which we have precomputed. This can clearly be done in time $O(|X_j|)$. So the total time to compute $\prod_i (1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}])$ for all j is $O(\sum_j |X_j|) = O(nc)$. From here, we just take $1 - \prod_i (1 - \Pr_{t_i \sim \mathcal{D}_i}[t_i \in S_{x_j}^{(i)}])$ to get the corresponding right-hand sides for all x_j. So all inequalities of the form (2) can be checked in total time $O(nc \log(nc))$. □

The final omitted proof is a technical lemma used inside the proof of Theorem 6. The lemma states that monotonicity constraints can be tight (and linearly independent) only if they are between two types in the same equivalence class.

Proof of Lemma 1: There are six cases to consider. Below we are studying which monotonicity constraints might potentially be tight in *any* possible ex-post allocation rule that induces $\tilde{\pi}$.

- What if $(i, \tau_{i,j+1}) \succeq (i, \tau_{i,j})$? Then $(i, j)^{th}$ monotonicity constraint might be tight.
- What if $(i, \tau_{i,j}) \succeq (0, \perp) \succeq (i, \tau_{i,j+1})$? Then by definition of $\succeq$, $\tilde{\pi}_i(\tau_{i,j}) > 0 = \tilde{\pi}_i(\tau_{i,j+1})$, so the $(i, j)^{th}$ monotonicity constraint can't be tight.
- What if $(0, \perp) \not\succeq (i, \tau_{i,j+1}) \not\succeq (i, \tau_{i,j})$ and there exists a bidder $k \neq i$ and type τ' with $(i, \tau_{i,j+1}) \not\succeq (k, \tau') \not\succeq (i, \tau_{i,j})$? First, it is clear that bidder i with type $\tau_{i,j+1}$

can *only* win when $(\tau_{i,j+1}) \succeq (\ell, t_\ell)$ for all $\ell \neq i$ (because on any other profile, there exists a (bidder, type) pair present in a tight Border constraint without $(i, \tau_{i,j+1})$), and that bidder i with type $\tau_{i,j}$ *must* win in all these cases (exactly because on all such profiles, $(i, \tau_{i,j})$ is present in a tight Border constraint without (ℓ, t_ℓ) for all ℓ). So for every opposing profile where $\tau_{i,j+1}$ has a chance of winning, $\tau_{i,j}$ certainly wins. Now we show that an opposing profile exists where bidder i loses with type $\tau_{i,j+1}$ but wins with type $\tau_{i,j}$, meaning that the $(i,j)^{th}$ monotonicity constraint can't be tight. Because $(i, \tau_{i,j+1}) \succeq (0, \perp)$, we necessarily have $(i, \tau_{i,j+1}) \succeq (\ell, \tau_{\ell,c_\ell})$ for all bidders ℓ. So consider the opposing profile where bidder k has type τ', and all other bidders $\ell \notin \{i, k\}$ have type τ_{ℓ,c_ℓ}, which arises with non-zero probability by our choice of c_ℓ. Clearly, bidder i with type $\tau_{i,j}$ must win the item against these opponents, while bidder i with type $\tau_{i,j+1}$ must lose. Therefore, the $(i,j)^{th}$ monotonicity constraint can't be tight.

- What if $(0, \perp) \not\succeq (i, \tau_{i,j+1}) \not\succeq (i, \tau_{i,j})$ and there exists a bidder $k \neq i$ and type τ' with $(i, \tau_{i,j}) \succeq (k, \tau') \succeq (i, \tau_{i,j})$? We know that bidder i with type $\tau_{i,j+1}$ *must* lose whenever there exists a bidder ℓ with $(\ell, \tau_\ell) \succeq (i, \tau_{i,j})$. However, if bidder i also loses with type $\tau_{i,j}$ against *all* such profiles, then the Border constraint for the set of all (type, bidder) pairs $(\ell, \tau) \neq (i, \tau_{i,j})$ with $(\ell, \tau) \succeq (i, \tau_{i,j})$ would necessarily be tight as well, meaning that the type (k, τ') is present in a tight Border constraint that $(i, \tau_{i,j})$ is not, contradicting $(i, \tau_{i,j}) \succeq (k, \tau')$.[11] So bidder i with type $\tau_{i,j}$ must win with non-zero probability against *some* such opposing profile, meaning again that the $(i,j)^{th}$ monotonicity constraint can't be tight.

- What if $(0, \perp) \not\succeq (i, \tau_{i,j+1}) \not\succeq (i, \tau_{i,j})$ and there exists a bidder $k \neq i$ and type τ' with $(i, \tau_{i,j+1}) \succeq (k, \tau') \succeq (i, \tau_{i,j+1})$? The reasoning is symmetric to the above case. We know that bidder i with type $\tau_{i,j}$ *must* win whenever all other bidders ℓ have $(i, \tau_{i,j+1}) \succeq (\ell, t_\ell)$. However, if bidder i also wins with type $\tau_{i,j+1}$ against *all* such profiles, then the Border constraint for the set $(i, \tau_{i,j+1})$ and all (type, bidder) pairs (ℓ, τ) with $(i, \tau_{i,j+1}) \not\succeq (\ell, \tau)$ (i.e. the set $\{(\ell, \tau) | (i, \tau_{i,j+1}) \not\succeq (\ell, \tau)\} \cup \{(i, \tau_{i,j+1})\}$) would necessarily be tight as well, meaning that the type $(i, \tau_{i,j+1})$ is in a tight Border constraint that (k, τ') is not, contradicting $(k, \tau') \succeq (i, \tau_{i,j+1})$.[12] So bidder i with type $\tau_{i,j+1}$ must lose with non-zero probability against *some* such opposing profile, meaning again that the $(i,j)^{th}$ monotonicity constraint can't be tight.

[11] Observe that all (ℓ, τ) with $(\ell, \tau) \succeq (i, \tau_{i,j})$ are necessarily have $\tau \neq \tau_{\ell,c}$, as such types are present in at most one tight Border constraint (namely (11)), in which $(i, \tau_{i,j+1})$ is itself present. So the specified Border constraint is indeed one that remains in Equation (10).

[12] Again, observe that the specified Border constraint is indeed one that remains in Equation (10), because if $j + 1 = c$, then $\tilde{\pi}_k(\tau')$ will be 0 contradicting the assumption that $(k, \tau') \succeq (i, \tau_{i,j+1}) \succeq (0, \perp)$.

- What if $(0, \perp) \not\succeq (i, \tau_{i,j+1}) \not\succeq (i, \tau_{i,j})$ and for all bidders $k \neq i$ and types τ', $(i, \tau_{i,j}) \not\succeq (k, \tau')$ or $(k, \tau') \not\succeq (i, \tau_{i,j+1})$? In this case, the $(i, j)^{th}$ monotonicity constraint is guaranteed to be tight, but linearly dependent on the set of tight Border constraints. We first argue that the Border constraint is tight for the set $A = \{(k, \tau') | (k, \tau') \succeq (i, \tau_{i,j})\} - (i, \tau_{i,j})$: on any profile $(t_1, \ldots, t_n)$ where $t_i \neq \tau_{i,j}$, but some $(k, t_k) \succeq (i, \tau_{i,j})$ (possibly $k = i$), clearly some bidder k with $(k, t_k) \in A$ must win (as the Border constraint for $B = \{(k, \tau') | (k, \tau') \succeq (i, \tau_{i,j})\}$ is tight, by definition of $\succeq$). Moreover, when $t_i = \tau_{i,j}$ and the profile $(t_1, \ldots, t_n)$ has some type $(k, t_k) \in A$ then $k \neq i$, and we must have $(i, \tau_{i,j}) \not\succeq (k, t_k)$ (by hypothesis defining this case). Therefore, the winner must still be some $(k, t_k) \in A$, and we conclude that the Border constraint for A must be tight.

  Similar reasoning implies that the Border constraint for set $C = \{(k, \tau') | (k, \tau') \succeq (i, \tau_{i,j})\} \cup \{(i, \tau_{i,j+1})\}$ is tight. Finally, one can take a linear combination of the tight Border constraints for the three sets $\{(k, \tau') | (k, \tau') \succeq (i, \tau_{i,j})\} - (i, \tau_{i,j})$, $\{(k, \tau') | (k, \tau') \succeq (i, \tau_{i,j})\}$ and $\{(k, \tau') | (k, \tau') \succeq (i, \tau_{i,j})\} \cup \{(i, \tau_{i,j+1})\}$ to recover the $(i, j)^{th}$ monotonicity constraint.[13] To be thorough, we do the calculation below: we first list the three tight Border's constraints.

$$\sum_k \sum_{j|(k,\tau_{k,j}) \in A} \Pr[t_k = \tau_{k,j}] \tilde{\pi}_k(\tau_{k,j}) = 1 - \prod_k \left(1 - \sum_{j|(k,\tau_{k,j}) \in A} \Pr[t_k = \tau_{k,j}] \right) \tag{13}$$

$$\sum_k \sum_{j|(k,\tau_{k,j}) \in B} \Pr[t_k = \tau_{k,j}] \tilde{\pi}_k(\tau_{k,j}) = 1 - \prod_k \left(1 - \sum_{j|(k,\tau_{k,j}) \in B} \Pr[t_k = \tau_{k,j}] \right) \tag{14}$$

$$\sum_k \sum_{j|(k,\tau_{k,j}) \in C} \Pr[t_k = \tau_{k,j}] \tilde{\pi}_k(\tau_{k,j}) = 1 - \prod_k \left(1 - \sum_{j|(k,\tau_{k,j}) \in C} \Pr[t_k = \tau_{k,j}] \right) \tag{15}$$

Clearly, Equation (13) and (14) imply that

$$\Pr[t_i = \tau_{i,j}] \cdot \tilde{\pi}_i(\tau_{i,j}) = \prod_{k \neq i} \left(1 - \sum_{j|(k,\tau_{k,j}) \in A} \Pr[t_k = \tau_{k,j}] \right) \cdot \Pr[t_i = \tau_{i,j}]. \tag{16}$$

[13] Again, all three Border constraints remain in Equation (10).

Also, Equation (14) and (15) imply that

$$\begin{aligned}\Pr[t_i = \tau_{i,j+1}] \cdot \tilde{\pi}_i(\tau_{i,j+1}) &= \prod_{k \neq i} \left(1 - \sum_{j|(k,\tau_{k,j}) \in B} \Pr[t_k = \tau_{k,j}]\right) \cdot \Pr[t_i = \tau_{i,j+1}] \\ &= \prod_{k \neq i} \left(1 - \sum_{j|(k,\tau_{k,j}) \in A} \Pr[t_k = \tau_{k,j}]\right) \cdot \Pr[t_i = \tau_{i,j+1}]. \qquad (17)\end{aligned}$$

Hence,

$$\tilde{\pi}_{i,j}(\tau_{i,j}) = \tilde{\pi}_{i,j+1}(\tau_{i,j+1}) = \prod_{k \neq i} \left(1 - \sum_{j|(k,\tau_{k,j}) \in A} \Pr[t_k = \tau_{k,j}]\right).$$

From this case analysis, we now conclude that the $(i,j)^{th}$ monotonicity constraint can be tight and linearly independent of the tight Border constraints *only if* $(i, \tau_{i,j+1}) \succeq (i, \tau_{i,j})$. $\square$