



An Algorithmic Characterization of Multi-Dimensional Mechanisms

Yang Cai*
EECS, MIT
ycai@csail.mit.edu

Constantinos Daskalakis†
EECS, MIT
costis@mit.edu

S. Matthew Weinberg‡
EECS, MIT
smw79@mit.edu

December 19, 2011

Abstract

We obtain a *characterization* of feasible, Bayesian, multi-item multi-bidder mechanisms with independent, additive bidders as *distributions over hierarchical mechanisms*. Combined with cyclic-monotonicity our results provide a complete characterization of feasible, Bayesian Incentive Compatible mechanisms for this setting.

Our characterization is enabled by a novel, constructive proof of Border’s theorem [5], and a new generalization of this theorem to independent (but not necessarily identically distributed) bidders, improving upon the results of [6, 12]. For a single item and independent (but not necessarily identically distributed) bidders, we show that any feasible reduced form auction can be implemented as a distribution over hierarchical mechanisms. We also give a polynomial-time algorithm for determining feasibility of a reduced form auction, or providing a separation hyperplane from the set of feasible reduced forms. To complete the picture, we provide polynomial-time algorithms to find and exactly sample from a distribution over hierarchical mechanisms consistent with a given feasible reduced form. All these results generalize to multi-item reduced form auctions for independent, additive bidders. Finally, for multiple items, additive bidders with hard demand constraints, and arbitrary value correlation across items or bidders, we give a proper generalization of Border’s Theorem, and characterize feasible reduced form auctions as multi-commodity flows in related multi-commodity flow instances. We also show that our generalization holds for a broader class of feasibility constraints, including the intersection of any two matroids.

As a corollary of our results we obtain revenue-optimal, Bayesian Incentive Compatible (BIC) mechanisms in multi-item multi-bidder settings, when each bidder has arbitrarily correlated values over the items and additive valuations over bundles of items, and the bidders are independent. Our mechanisms run in time polynomial in the total number of bidder types (and *not* type profiles). This running time is polynomial in the number of bidders, but potentially exponential in the number of items. We improve the running time to polynomial in both the number of items and the number of bidders by using recent structural results on optimal BIC auctions in item-symmetric settings [14].

*Supported by NSF Award CCF-0953960 (CAREER) and CCF-1101491.

†Supported by a Sloan Foundation Fellowship and NSF Award CCF-0953960 (CAREER) and CCF-1101491.

‡Supported by a NSF Graduate Research Fellowship and a NPSC Graduate Fellowship.

1 Introduction

In its most general form mechanism design is the task of selecting outcomes and charging prices in a way that maximizes some objective such as revenue or social-welfare. Accordingly, to specify a mechanism one needs to provide two, possibly randomized, functions: one that maps bids to outcomes, and another mapping bids to prices charged to the bidders. Consider, e.g., a setting where n items are (simultaneously) auctioned to m bidders. For every possible collection of bids the mechanism needs to specify (a) a possibly randomized allocation of items to bidders; and (b) a possibly randomized collection of prices charged to them. The latter can be typically summarized by its expectation, which is just a vector of real numbers, recording the price charged to each bidder. But the former is a distribution over allocations and, in principle, would require $(m + 1)^n$ probabilities to be specified, making such explicit specification completely impractical. It turns out that this calculation is too pessimistic, and the Birkhoff-von Neumann decomposition theorem allows one to summarize a randomized allocation by just providing its marginal allocation probabilities $\phi = \{\phi_{ij}\}_{ij}$, where ϕ_{ij} is the probability that item j is given to bidder i [15]. All that these probabilities need to satisfy to be consistent with a joint allocation distribution is that no item is given out more than once in expectation, and no bidder gets more items than she wants (if she has any demand constraints).

Yet, the above description is still too demanding as, even when every bidder can place c possible bids to specify her preferences over the items, there are c^m possible bid vectors, and the mechanism needs to specify what to do for each. To decrease the specification complexity to linear in the number of bidders one can use the *reduced form of the mechanism*, which specifies a pair of possibly randomized allocation and price functions *per bidder*. In particular, if T_i denotes the possible bids that bidder i can submit,¹ the reduced form of the mechanism provides a collection $\mathcal{R} := \{(\pi_i : T_i \rightarrow [0, 1]^n, p_i : T_i \rightarrow \mathbb{R})\}_i$ of functions such that, for all $t_i \in T_i$ and all j , $\pi_{ij}(t_i)$ is the marginal probability (over the randomness in the mechanism and the uncertainty about the bids submitted by the other bidders) that item j is allocated to bidder i when she reports t_i , and $p_i(t_i)$ is the expected price she pays. Indeed, the functions π_i and p_i provide sufficient information for the bidder to decide what bid in T_i optimizes her utility in expectation. The trouble is that, unless the uncertainty over the other bids can be modeled probabilistically, the notion of a marginal allocation probability π_{ij} is ill-defined.

Indeed, Economists go around this obstacle by making average-case assumptions about the preferences of bidders [22]. In particular, it is assumed that every bidder has a *type* $t_i \in T_i$ (forgive the temporary overload of notation), and that the joint type profile of the bidders is sampled from some known, joint distribution \mathcal{D} over $\times_i T_i$. Now, the Revelation Principle implies that every auction is strategically equivalent to one where bidders submit not generic bids, but truthfully report their exact type. Hence w.l.o.g. we can restrict our attention to mechanisms where the bid space of bidder i is just T_i (**EndOf overload of notation**), and where the distribution of bids that bidders submit is exactly \mathcal{D} . For such auctions, the reduced form is perfectly well-defined. But a challenge remains: What conditions are necessary and sufficient for a reduced form $\{\pi_i\}_i$ to be feasible, i.e. correspond to a feasible auction (that never over-allocates items)? Can these conditions be verified efficiently? And, given a feasible reduced form can we efficiently compute an auction implementing it?

Before studying these questions, why do we even care? Even without answers, mechanism design has made great progress in designing auctions, which optimize various objectives in a wide-range

¹We won't make any assumptions about T_i , except that it is measurable. In particular, it won't necessarily be a subset of \mathbb{R} , unless otherwise specified.

of settings. When the objective is social-welfare, there is an optimal auction with a very clean allocation rule, namely the VCG auction [25, 13, 16]. The auction always chooses the welfare-maximizing allocation, and just charges prices to ensure incentive compatibility. When the objective is revenue, things get more murky. Still, when we are auctioning a single item to bidders whose values are independently sampled from known distributions, we can employ Myerson’s celebrated, revenue-optimal auction [22]. Its allocation rule is still rather clean: it always chooses the allocation that maximizes virtual surplus, and just charges prices to ensure incentive compatibility.

While in the above cases we have clean allocation rules, these results remain quite exceptional. In particular, there is no known simple characterization of the set of all feasible single-item mechanisms, and (prior to this work) it is unclear if there is one. Additionally, it seems unlikely that there is a clean allocation rule implementing the revenue-optimal multi-item mechanism. Understanding the set of feasible multi-item mechanisms is therefore a crucial step along the path to understanding revenue-optimal mechanisms in this setting. Indeed, in the past decades a large body of research in Economics has been devoted to finding revenue-optimal mechanisms (extending Myerson’s result) to multi-item settings, but progress has been sporadic (see survey [19] and its references). More recently the problem has entered theory of computation, and there has been a number of papers obtaining constant factor approximations [1, 4, 10]. Our slow progress on this front is intimately related to the lack of a characterization result of feasible mechanisms. Loosely speaking, all known algorithms optimize using relaxed feasibility constraints for the reduced form of the mechanism, such as “every item can be awarded at most once in expectation” [1, 4] or “the expected number of items awarded from every subset must not exceed its rank” [10]. Using such relaxed feasibility constraints, the optimal (possibly infeasible) reduced form is found. To turn this into a feasible reduced form, some “rounding procedure” is applied; this procedure may be just scaling the allocation probabilities down by a constant so that no item is over-allocated [1, 4], or may invoke prophet inequalities [1, 10]. However, rounding comes with a loss in revenue. In view of this computational experience with the problem, it is crucial to try to characterize the feasible reduced forms exactly in a way that makes optimization efficient.

In this paper we carry out this program. We obtain strong characterization results for the set of all feasible single-item mechanisms (for independent bidders), as well as the set of all feasible multiple-item mechanisms when the bidders have arbitrarily correlated additive valuations over bundles of items (but there is no correlation across bidders). Our characterization result (explained in Section 3.2) is based on obtaining novel, constructive proofs of Border’s theorem [5] and recent extensions of this theorem [6, 12], which are overviewed in Section 3.1. Using these constructive proofs and our characterization result, we obtain revenue-optimal, Bayesian Incentive Compatible (BIC) mechanisms in multi-item multi-bidder settings, when each bidder has arbitrarily correlated values over the items and additive valuations over bundles of items, and the bidders are independent. Our mechanisms run in time polynomial in the total number of bidder types, $\sum_i |T_i|$. This running time is polynomial in the number of bidders, but potentially exponential in the number of items. We improve the running time to polynomial in both the number of items and the number of bidders, by using recent structural results on optimal BIC auctions in symmetric settings [14]. These results are overviewed in Section 3.3. In Section 3.4 we present a generalization of Border’s theorem to the multi-item setting where values are arbitrarily correlated across items or bidders, and bidders have additive valuations with hard demand constraints, and provide a multi-commodity flow interpretation of this problem. In addition, we show that our generalization holds for a much broader class of feasibility constraints. Simply put, our generalization holds in all instances where the feasibility constraints can be written as a system of inequalities that are linear in allocation marginals (ϕ_{ij} ’s). This includes feasibility constraints that are the intersection of any two matroids, as well as many other set systems [24].

After a brief enumeration of our theorems and a couple of figures summarizing our results, we fix our notation and proceed to the technical sections. We show the following theorems (and several auxiliary ones along the way): (a) Theorems 2 and 6 state our algorithmic results: If a reduced form auction for independent bidders and a single item is given to us as input, we can efficiently determine if the auction is feasible or not. If it is feasible, we efficiently compute an efficient process that implements it. If it is infeasible, we efficiently provide a hyperplane separating it from the set of feasible mechanisms. These results extend to multi-item settings with additive, independent bidders and arbitrary correlation in the values of each bidder for the items. (b) Theorems 3 and 7 are the core of our characterization result: Every feasible reduced form auction for a single item and independent bidders can be implemented as a distribution over hierarchical mechanisms. Using these theorems we obtain a clean, characterization of the allocation rule of all multidimensional mechanisms for multi-item settings with independent, additive bidders and arbitrary correlation in the values of each bidder for the items. Our characterization result is provided in Section 3.2. (c) Theorem 4 provides our analog of Myerson’s virtual values. It states that we can transform the marginal probabilities π s of a reduced form into *virtual* π s such that the total ordering of the types (across different bidders) induced by the virtual π s allows us to, in linear time, verify the feasibility of a reduced form auction or output a hyperplane separating it from the set of feasible reduced forms. (d) Theorem 12 proves a generalization of Border’s Theorem to a broad class of feasibility constraints that includes the case of many items, many additive bidders with hard demand constraints, and arbitrary value correlation across items or bidders.

Figures 1 and 2 summarize the current state of affairs on the two fronts studied in this paper. The first is on efficiently implementing as well as characterizing the allocation rule of feasible reduced forms, while the second is on efficiently solving the optimal multi-dimensional mechanism design problem. The second only records mechanisms that achieve optimal or near-optimal revenue, i.e. come ϵ -close to optimal revenue for arbitrary accuracy ϵ . In particular, we do not quote constant factor approximations [1, 4, 10], which are not the focus of this work. The “Item-Symmetric” column of the second figure refers to settings that are symmetric with respect to the items (see Section 3.3 for a formal definition of such settings and [14] for a discussion), while the “Bidder-Symmetric” column to settings that are symmetric with respect to the bidders (see Section 3.3 and [14] for a formal definition/discussion). Moreover, the “Many items, Few bidders” row refers to settings where the number of bidders is an absolute constant but the number of items, n , is allowed to scale, and the “Few items, Many bidders” row refers to settings where the number of items is held constant but the number of bidders, m , is allowed to scale. Finally, in the cells of the same table we stress components of the running time/the setting that are the best we could hope for, and do not stress components that could be improved, leaving these improvements to future work.

Parallel Work: We have been informed that, independently from us, [2] also obtained polynomial-time algorithms for checking the feasibility of a given single-item reduced form auction, when the bidders are independent, as well algorithms for efficiently implementing the reduced form, if it is feasible. (See Figure 1 for where their results fit with respect to our results). [2] apply their algorithms to “service constrained” multi-dimensional mechanism design environments, which are orthogonal to the mechanism design settings considered here.

2 Preliminaries and notation

Throughout the paper, we denote the number of bidders by m and the number of items by n . We also use T_i to denote the types of bidder i . We make no assumptions about T_i except that

	Necessary&Sufficient Feasibility Conditions	Efficient Implementation	Characterization of Feasible Mechanisms
Independent Bidders, No Demand Constraints	[5, 6, 12]	[this work] (& independently [2], for single-item auctions)	[this work]
Correlated Bidders, Demand Constraints, and Broader Settings [†]	[this work]	Open Question	Open Question

[†]: See Section 3.4.

Figure 1: Feasibility, Efficient Implementation and Characterization of Reduced Form Multi-Item Auctions. The bidders are additive up to (if stated) some hard demand constraint. Both lines allow correlation of values across different items, but the first row assumes that bidders have independent valuations while the second also allows correlation across bidders.

	Item-Symmetric	Bidder-Symmetric	No Symmetries
Many items, Few bidders	poly(n)[†], demand constraints, item correlation [14]	Same as \rightarrow	poly(n)[§], no demand constraints, MHR independent items [8]
Few items, Many bidders	Same as below [this work]	poly(m)[†], demand constraints, item correlation [14]	Same as below [this work]
Many items Many bidders	poly(n,m)[†], no demand constraints, item correlation [this work]	Same as \rightarrow [this work]	poly($\sum_i T_i$)[‡], no demand constraints, item correlation [this work]

[§]: For all $\epsilon > 0$, a $(1 - \epsilon)$ -fraction of the optimal revenue is achieved in the quoted running time. $1/\epsilon$ appears in the exponent of the running time.

[†]: Exactly optimal revenue is achieved in the quoted running time, when every marginal of the bidders' value distributions has constant size support. The size of the support appears in the exponent of the running time. When the size of the support is non-constant/infinite/unbounded, two types of PTAS's are obtained:

- When every marginal is MHR, a $(1 - \epsilon)$ -fraction of the optimal revenue is achieved in the quoted running time, for all $\epsilon > 0$.
- When the value distributions are bounded and normalized to $[0, 1]^n$, an additive approximation to the optimal revenue is achieved in the quoted running time, for all $\epsilon > 0$. The approximation error is $\epsilon \cdot \min\{n, \sum_i C_i\}$, where C_i is the demand constraint of bidder i (if any).

[‡]: Exactly optimal revenue is achieved in the quoted running time, where recall that T_i is the support of the value distribution of bidder i . When the supports are infinite/unbounded we can alternatively obtain two types of approximation schemes:

- When every marginal of the bidders' value distributions is MHR, a $(1 - \epsilon)$ -fraction of the optimal revenue is achieved in time polynomial in m and $(1/\epsilon)^n$, for all $\epsilon > 0$.
- When the value distributions are bounded and normalized to $[0, 1]^n$, an additive approximation to the optimal revenue is achieved in time polynomial in m and $(1/\epsilon)^n$, for all $\epsilon > 0$. The approximation error is $\epsilon \cdot n$.

Figure 2: Efficient Solution to the Optimal Multi-Dimensional Mechanism Design Problem with Independent Bidders. The bidders are additive up to (if stated) some hard demand constraint.

it is measurable. In particular, it is not assumed that T_i is a subset of \mathbb{R} , and it could well be multidimensional; e.g., when a bidder has additive valuations, the type of the bidder could be just

a vector specifying how much she values each of the items. To obtain computationally meaningful results, we assume in Sections 3.1, 3.2 and 3.4 that the type sets $\{T_i\}_i$ are finite (as opposed to continuous/infinite, but not of constant size), and in this case use c as a shorthand for $\max_i |T_i|$. However, our results in Section 3.3 apply also to continuous type spaces, since when the type distributions are bounded or when they are unbounded but satisfy the *Monotone Hazard Rate* tail condition,² we can truncate and discretize the support of the distribution and still get nearly optimal ϵ -BIC mechanisms (Definition 1) using the techniques of [7] or [14]. We can then employ the ϵ -BIC to BIC reduction of [14] to get nearly optimal BIC mechanisms. See Section 3.3 for details.

To ease notation, we just use A (B , C , etc.) to denote the *type* of a bidder, without emphasizing whether it is a vector or a scalar. The elements of $\times_i T_i$ are called *type profiles*, and specify a type for every bidder. We assume type profiles are sampled from a distribution \mathcal{D} over $\times_i T_i$. We denote by \mathcal{D}_i the marginal of this distribution on bidder i 's type, and use \mathcal{D}_{-i} to denote the marginal of \mathcal{D} 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 $\Pr[t_i = A]$, we mean the probability that bidder i 's type is A . In the i.i.d. case, because $\Pr[t_i = A]$ is the same for all i , we will just write $\Pr[A]$.

The *reduced form* of a mechanism is a vector function $\pi(\cdot)$, specifying values $\pi_{ij}(A)$, for all items j , bidders i and types $A \in T_i$. $\pi_{ij}(A)$ is the probability that bidder i receives item j when reporting type A , where the probability is over the randomness of all other bidders' types and the internal randomness of the mechanism, assuming that the other bidders report their true types to the mechanism. When we only consider a single item (case $n = 1$), we drop the item subscript from the π_{ij} 's, writing $\pi_i(A)$ for the probability that the item is allocated to bidder i . In this case we may think of $\pi(\cdot)$ as a vector in $[0, 1]^{\sum_i |T_i|}$. To emphasize the vector-view of the reduced form we may write $\vec{\pi}$ for the reduced form.

Sometimes we will be considering settings where the bidders are i.i.d., i.e. $T_i = T_{i'} = T$, for all i and i' , and \mathcal{D} is a product distribution over $\times_i T_i$ with the same marginal on every bidder. In such settings, we may be looking at *bidder-symmetric* reduced forms, satisfying $\pi_{ij}(A) = \pi_{i'j}(A)$, for all $j, i, i', A \in T$. In such cases, we will drop the bidder subscript from the π_{ij} 's, writing $\pi_j(A)$, for the probability that a bidder of type $A \in T$ receives item j , over the randomness of the mechanism and the types of the other bidders, assuming that the other bidders report their true types to the mechanism. If there is a single item and the reduced form is bidder-symmetric, we will drop both subscripts writing $\pi(A)$ for the same probability. In this case, we may think of the reduced form as a vector $\vec{\pi}$ in $[0, 1]^{|T|}$.

Given a reduced form π of a mechanism we will be interested in whether the form can be "implemented". By this we mean designing a feasible mechanism M (i.e. one that never over-allocates items) such that the probability $M_{ij}(A)$ that bidder i receives item j when she reports type A to the mechanism is exactly $\pi_{ij}(A)$, where the probability is computed with respect to the randomness in the mechanism and the randomness in the types of the other bidders, assuming that the other bidders report their true types to the mechanism. In fact, we will relax this requirement and say that a mechanism M implements the reduced form π if $M_{ij}(A) \geq \pi_{ij}(A)$, for all i, j, A . This is because we can trivially modify such a mechanism so that all constraints are tight. Indeed, if $M_{ij}(A) = x \cdot \pi_{ij}(A)$ for some $i, j, A, x > 1$, we can rectify the mechanism as follows: whenever M wants to allocate item j to bidder i when he has reported type A , we instead throw the item away with probability $\frac{x-1}{x}$. It's easy to see that this results in $M_{ij}(A) = \pi_{ij}(A)$. When $M_{ij}(A) = \pi_{ij}(A)$ for all i, j, A , we will say that M *exactly* implements π . Finally, if a reduced form can be

²We say that a one-dimensional differentiable distribution F satisfies Monotone Hazard Rate, or succinctly MHR, if $\frac{f(x)}{1-F(x)}$ is monotonically non-decreasing in its domain, where $f = F'$ is the probability density function.

implemented by a mechanism, we say that the reduced form is *feasible*.

When a bidder is additive, we may write \vec{v}_i to denote her type, with the convention that v_{ij} represents her value for item j and that her value for a bundle of items is just the sum of her values for the items in the bundle. Then, to fully specify a (direct-revelation) multi-item mechanism for additive bidders, we need to describe, potentially succinctly, for all type profiles $\vec{v} \in \times_i T_i$, 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} , where $\phi_{ij}(\vec{v})$ is the probability that item j is given to bidder i and $p_i(\vec{v})$ is the price that i pays. The *value* of bidder i for outcome $M_i(\vec{w})$ is just her expected value $\vec{v}_i \cdot \vec{\phi}_i(\vec{w})$ for the bundle allocated to her, while the *utility* of bidder i for the same outcome is $U(\vec{v}_i, M_i(\vec{w})) := \vec{v}_i \cdot \vec{\phi}_i(\vec{w}) - p_i(\vec{w})$. Such bidders subtracting price from expected value are called *quasi-linear*. The relation between ϕ 's and π 's is just the following: for all $i, \vec{v}_i \in T_i$: $\pi_{ij}(\vec{v}_i) = \mathbb{E}_{\vec{v}_{-i} \sim \mathcal{D}_{-i}}[\phi_{ij}(\vec{v}_i; \vec{v}_{-i})]$, where the expectation is computed with respect to the values \vec{v}_{-i} of all bidders except i as these are drawn from \mathcal{D}_{-i} . We conclude by formally defining Bayesian Incentive Compatibility:

Definition 1 ([14]). (*BIC/ ϵ -BIC Mechanism*) A mechanism M is called ϵ -BIC iff the following inequality holds for all i, \vec{v}_i, \vec{w}_i :

$$\mathbb{E}_{\vec{v}_{-i} \sim \mathcal{D}_{-i}} [U(\vec{v}_i, M_i(\vec{v}))] \geq \mathbb{E}_{\vec{v}_{-i} \sim \mathcal{D}_{-i}} [U(\vec{v}_i, M_i(\vec{w}_i; \vec{v}_{-i}))] - \epsilon v_{\max} \cdot \sum_j \pi_{ij}(\vec{w}_i),$$

where v_{\max} is the maximum possible value of any bidder for any item in the support of the value distribution. In other words, M is ϵ -BIC iff when a bidder lies by reporting \vec{w}_i instead of \vec{v}_i , they do not expect to gain more than ϵv_{\max} times the expected number of items that \vec{w}_i receives. A mechanism is called BIC iff it is 0-BIC.³

We conclude our preliminaries section noting that the running times of the algorithms obtained in Sections 3.1.1 and 3.1.2 are quoted without accounting for the number of bits needed to represent a coordinate of $\vec{\pi}$. If ℓ bits are needed to describe a probability in $\vec{\pi}$, then it suffices to multiply all quoted running times by a $\text{poly}(\ell)$ factor.

3 Overview of Results and Techniques

3.1 Feasible Reduced Forms

3.1.1 Single-item, Bidder-Symmetric Reduced Forms, i.i.d. Bidders

In the case of a single item and i.i.d. bidders, Border provided a necessary and sufficient condition for a bidder-symmetric reduced form to be feasible, generalizing prior partial results of Maskin-Riley [20] and Matthews [21]. Let us review Border's theorem.

Theorem 1 ([5]). *Suppose that the bidder's types are i.i.d. distributed according to some measure μ over T . Then a bidder-symmetric reduced form π is feasible if and only if*

$$\forall S \subseteq T : m \cdot \int_S \pi(t) d\mu(t) \leq 1 - (1 - \mu(S))^m. \quad (1)$$

³This definition differs slightly from the more common definition of [3, 17, 18]. We use the definition of [14] to be correct in applying their results.

Simply put, a reduced form is feasible if and only if the probability that the item is awarded to a type in some set S (as computed by the reduced form) is at most the probability that someone with type from S shows up to the auction (as computed by the type distribution), for all subsets of types $S \subseteq T$. We call a set that violates this condition a *constricting set*. Clearly, the existence of a constricting set bears witness that the reduced form is infeasible, as the auctioneer cannot possibly award the item to someone in S if no one in S shows up. Border’s theorem states that this is in fact a sufficient condition.

Border’s original paper considered continuous type spaces (hence the integral in (1)), and the proof was based on measure theory. The following extension of the theorem was also shown: If there exists a constricting set S , then there is also a constricting set of the form S_x , where $S_x = \{A | \pi(A) > x\}$, for some x . In the case of finite type spaces, we can determine the feasibility of a reduced form auction in time $O(c \log c + c \cdot m)$, where $c = |T|$, as after sorting the type space in decreasing π ’s there are only c different subsets of the form S_x , and a dynamic program can find us if any of them violates (1) in time $O(c \cdot m)$. In other words, determining the feasibility of a bidder-symmetric reduced form, for a single item, and many i.i.d. bidders is easy. However, the following important question was left unanswered: Given a feasible reduced form, can we efficiently obtain a mechanism implementing the reduced form? Notice that answering this question in the affirmative is absolutely necessary to be able to run the auction specified by the reduced form. Our first contribution is solving this problem.

Theorem 2. *Under the same assumptions as Theorem 1, given a bidder-symmetric reduced form we can determine if it is feasible, or find a hyperplane separating it from the set of feasible bidder-symmetric reduced forms, in time $O(c \cdot (\log c + m))$, where $c = |T|$. If the reduced form is feasible, we provide a succinct description of a mechanism implementing the reduced form, in time polynomial in $c \cdot m$. The description of the mechanism is just (at most) $c + 1$ probabilities and an equal number of orderings of T .⁴ The mechanism itself runs as follows: given the reported type profile, the mechanism samples a random subset of bidders in time $\text{poly}(m, c)$, and the item is allocated uniformly at random to some bidder in that subset, or the item is thrown away.*

We prove Theorem 2 in Appendix B, as a corollary of Proposition 1 and Theorem 13 of Appendices B and A respectively. In proving our result, we consider the following type of mechanisms:

Definition 2. *A **hierarchical mechanism** consists of a function $H : T \rightarrow [c] \cup \{LOSE\}$; one should interpret $LOSE$ as a value larger than c . On bid vector (A_1, \dots, A_m) , the mechanism has the following behavior: If $H(A_i) = LOSE$ for all i , the mechanism throws the item away. Otherwise, the item is awarded uniformly at random to a bidder in $\text{argmin}_i H(A_i)$.*

In other words, a hierarchical mechanism breaks down the type space into a hierarchy. When the bidders arrive and submit their types, the mechanism finds the highest-priority level of the hierarchy that is populated by the submitted types, and gives the item uniformly at random to a bidder whose type falls in that level of the hierarchy (unless every bidder is a loser, in which case the mechanism throws the item away). We say that a hierarchical mechanism H is *well-ordered* w.r.t. π if: $\pi(A) \geq \pi(A') \Rightarrow H(A) \leq H(A')$. We prove the following characterization result about feasible bidder-symmetric reduced forms:

Theorem 3. *When bidders are i.i.d., every feasible bidder-symmetric reduced form π can be exactly implemented as a distribution over at most $c + 1$ well-ordered with respect to π hierarchical mechanisms.*

⁴An ordering may have equalities, but must be total (compare every pair of elements in T).

Theorem 3 alone is not enough to allow us to implement a given bidder-symmetric reduced form. Indeed, if $\pi(\cdot)$ takes θ (can be as large as c) distinct values, there are 2^θ different well-ordered w.r.t. π hierarchical mechanisms. From here, we switch to our vector-view of reduced forms (as vectors $\vec{\pi}$ in $[0, 1]^{|T|}$) and study the geometry of the space of feasible mechanisms respecting the order on the type-space induced by a given reduced form $\vec{\pi}$, which we will call P . We show that, in fact, P is a θ -dimensional polytope whose corners are *exactly* the 2^θ different well-ordered w.r.t. $\vec{\pi}$ hierarchical mechanisms. We provide a geometric algorithm in Appendix A that in polynomial time outputs a representation of $\vec{\pi}$ as a convex combination of at most $c+1$ corners of P . This convex combination is exactly a distribution over well-ordered w.r.t. $\vec{\pi}$ hierarchical mechanisms that implements $\vec{\pi}$. Our geometric algorithm runs in time $\text{poly}(m, c)$, and sampling from the distribution output by our algorithm also takes time $\text{poly}(m, c)$. We provide the details of our approach and proofs of the relevant claims in Appendices B and A.

3.1.2 Single-item, General Reduced Forms, Non-i.i.d. Bidders

Recently, an alternative proof of Border’s theorem for distributions with finite support was discovered in [6] and again in [12], the latter using a clean network-flow interpretation. These proofs extend Theorem 1 to independent, but not necessarily identical, bidders and non-symmetric reduced forms. In this case, (1) is replaced by the following necessary and sufficient condition:

$$\forall S_1 \subseteq T_1, \dots, S_m \subseteq T_m : \sum_i \sum_{A \in S_i} \pi_i(A) \Pr[t_i = A] \leq 1 - \prod_i (1 - \Pr[t_i \in S_i]). \quad (2)$$

The interpretation of the LHS and RHS of the above inequality is the same as the one given above for (1) except generalized to the non-iid non-symmetric setting. In addition to the above condition, [12] proves a generalization of Border’s extended result: If there is a constricting $S = (S_1, \dots, S_m)$, then there is also a constricting set of the form $S' = (S_{x_1}^{(1)}, \dots, S_{x_m}^{(m)})$, where $S_{x_i}^{(i)} = \{A \in T_i | \pi_i(A) > x_i\}$. In other words, each bidder has a different threshold x_i , and $S_{x_i}^{(i)}$ contains all types of bidder i with π_i above x_i . Unfortunately, despite this simplification, there are still $\prod_i (|T_i| + 1)$ possible constricting sets, and testing each of them would take time exponential in the number of bidders.

One might hope to obtain a stronger theorem that would only require testing a number of sets polynomial in c and m . We prove such a theorem by introducing a notion of a *virtual* π , defined next. We name it such not because the equation involves hazard rates or looks anything like that for virtual valuations [22], but because the spirit of the transformation is the same. Myerson observed that he could make the most revenue not from the bidder with the highest valuation, but from the bidder with the highest virtual valuation. Likewise, in our setting, the most difficult types to satisfy are not the types with the highest π , but the types with the highest virtual π . The definition of virtual π , which we denote $\hat{\pi}$, is actually quite simple.

Definition 3. *If π is a reduced form, we define its corresponding virtual reduced form $\hat{\pi}$ as follows: for all i and type $A \in T_i$, $\hat{\pi}_i(A) := \Pr[\pi_i(t_i) \leq \pi_i(A)]\pi_i(A)$.*

It turns out that this definition exactly captures which types of different bidders are harder to satisfy. In the bidder-symmetric case of Section 3.1.1, we were able to compare a pair of types A and B submitted by bidders $i \neq k$ based only on their corresponding $\pi_i(A)$ and $\pi_k(B)$. This is no longer the case in the non-iid case, resulting in the more complicated constricting sets defined above. Nevertheless, we show that A and B can be compared at face value of $\hat{\pi}_i(A)$ and $\hat{\pi}_k(B)$:

Theorem 4. *Suppose that the bidders are independent and there is a single item for sale. A reduced form π is feasible if and only if: for all x , the sets $S_x^{(i)} = \{A \in T_i | \hat{\pi}_i(A) > x\}$ satisfy:*

$$\sum_i \sum_{A \in S_x^{(i)}} \pi_i(A) \Pr[t_i = A] \leq 1 - \prod_i (1 - \Pr[t_i \in S_x^{(i)}]). \quad (3)$$

In particular, we can test the feasibility of a reduced form, or obtain a hyperplane separating the reduced form from the set of feasible reduced forms, in time linear in $\sum_i |T_i| \cdot (\log(\sum_i |T_i|) + m)$.

Details of the proof can be found in Appendix C. We also prove there two analogues of Theorem 3 in this setting: Theorem 5, which is used for our algorithmic results, and Theorem 7, which provides a much stronger characterization. The analog of Definition 2 is the following:

Definition 4. *A **hierarchical mechanism** consists of a function $H : \bigcup_i (T_i \times \{i\}) \rightarrow [\sum_i |T_i|] \cup \{LOSE\}$; one should interpret $LOSE$ as a value larger than $\sum_i |T_i|$. On bid vector (A_1, \dots, A_m) , if $H(A_i, i) = LOSE$ for all i , the mechanism throws the item away. Otherwise, the item is awarded uniformly at random to a bidder in $\operatorname{argmin}_i H(A_i, i)$.*

We say that a hierarchical mechanism H for non-identical bidders is *partially-ordered* w.r.t. π if for all i and $A, A' \in T_i$, $\pi_i(A) \geq \pi_i(A') \Rightarrow H(A, i) \leq H(A', i)$. We say that a hierarchical mechanism is *strict* if for all bidders i, j and types $A \in T_i, B \in T_j$: $i \neq j \Rightarrow (H(A, i) \neq H(B, j) \vee H(A, i) = H(B, j) = LOSE)$ (i.e. there is always a unique winner in $\operatorname{argmin}_i H(A_i, i)$ if one exists, because each level (except possibly for $LOSE$) contains types from only a single bidder). Our algorithmic extension of Theorem 3 is the following:

Theorem 5. *When bidders are independent, but not necessarily identically distributed, every feasible reduced form π can be exactly implemented as a distribution over at most $\sum_i |T_i| + 1$ strict, partially-ordered w.r.t. π hierarchical mechanisms.*

From here, we take the same geometric approach as in Section 3.1.1 and study the geometry of the set of feasible reduced forms that respect the partial-ordering of types induced by a given reduced-form π . Again we show that this is a $(\sum_i d_i)$ -dimensional polytope, P , where d_i (could be as large as $|T_i|$) is the number of distinct values that $\pi_i(\cdot)$ takes on input from T_i , and that the corners of P are *exactly* the strict, partially-ordered w.r.t. π hierarchical mechanisms. Writing a point in P as a convex combination of $\sum_i |T_i| + 1$ corners is no longer an easy procedure. Not only does P have an exponential number of corners, but there are also exponentially many hyperplanes defining the boundary of P (where there were only $2 \cdot c$ such hyperplanes in the i.i.d. case). Luckily, Theorem 4 provides an efficient separation oracle for membership in P . By making use of this separation oracle instead of checking the exponentially-many boundary equations one by one, the geometric algorithm of Appendix A outputs a representation of a given π as a convex combination of at most $\sum_i |T_i| + 1$ corners of P , which is exactly a distribution over the corresponding $\sum_i |T_i| + 1$ strict, partially-ordered w.r.t. π hierarchical mechanisms. Putting this approach together with Theorems 4 and 5, we obtain in Appendix C the algorithmic result of this section:

Theorem 6. *When bidders are independent, given a reduced form we can determine if it is feasible, or find a hyperplane separating it from the set of feasible reduced forms, in time linear in $\sum_i |T_i| \cdot (\log(\sum_i |T_i|) + m)$. If the reduced form is feasible, we can compute a succinct description of a mechanism implementing the reduced form, in time polynomial in $\sum_i |T_i|$. The description of the mechanism is just (at most) $\sum_i |T_i| + 1$ probabilities and the same number of total orderings of $\bigcup_i (T_i \times \{i\})$. The mechanism itself runs as follows: given the reported type profile, the mechanism samples a random total ordering of all bidders' types in time polynomial in $\sum_i |T_i|$, and allocates the item to the bidder whose reported type is highest in that ordering, or throws the item away.*

While Theorem 5 does provide some structure to the otherwise unknown set of feasible mechanisms for independent bidders, the result is not as compelling as that of Theorem 3. One might have hoped that every feasible reduced form can be implemented as a distribution over *virtually-ordered* hierarchical mechanisms (that is, hierarchical mechanisms such that $\hat{\pi}_i(A) \geq \hat{\pi}_j(B) \Rightarrow H(A, i) \leq H(B, j)$). Unfortunately, this is not true, as is shown in Appendix C. Despite this, we show that a strong generalization of Theorem 3 holds in this setting. Let σ be a total ordering on the elements of $\bigcup_i (T_i \times \{i\})$ (i.e. a mapping $\sigma : \bigcup_i (T_i \times \{i\}) \rightarrow [\sum_i |T_i|]$). We say that σ respects π if $\pi_i(A) > \pi_i(B) \Rightarrow \sigma(A, i) < \sigma(B, i)$. We also say that a hierarchical mechanism H is σ -ordered if $\sigma(A, i) < \sigma(B, j) \Rightarrow H(A, i) \leq H(B, j)$. We prove the following theorem:

Theorem 7. *If a reduced form π is feasible, there exists a total ordering σ on the elements of $\bigcup_i (T_i \times \{i\})$ that respects π such that π can be implemented as a distribution over σ -ordered hierarchical mechanisms.*

We provide the proof of Theorem 7 in Appendix D.

3.2 A Simple Characterization of Feasible Multi-dimensional Mechanisms

The appeal of Myerson’s single-item auction is not only its optimality, but its clean allocation rule. Let us revisit it. In a single-item setting the type-space of each bidder is a subset $T_i \subset \mathbb{R}$, as a bidder’s type is how much she values the item. In brush strokes Myerson’s allocation rule is as follows:

- *The i.i.d. case:* Suppose that the bidders’ values are independently and identically distributed in T . Myerson’s auction in this setting maintains a ranking $H : T \rightarrow [|T|] \cup \{\text{LOSE}\}$ (take $\text{LOSE} := |T| + 1$) satisfying: $A \geq A' \implies H(A) \leq H(A')$; i.e. the ranking respects the order on T (as a subset of \mathbb{R}). If $|H^{-1}(k)| > 1$ for some k , we can think of the types in $H^{-1}(k)$ as having been *ironed* over by H . In fact, let us refer to H as an *ironing* of T . Given the ironing, the allocation rule is simple: on bid vector (A_1, \dots, A_m) , the item is allocated uniformly at random to a bidder in $\arg\min_i \{H(A_i)\}$ (other tie-breaking rules are fine too), or thrown away if $H(A_i) = \text{LOSE}$ for all i .
- *The non-i.i.d. case:* Suppose that the bidders’ values are independent but not necessarily identically distributed. In this case, the allocation rule of Myerson’s auction is a bit more complicated. The auction maintains an ironing $H : \cup_i (T_i \times \{i\}) \rightarrow [\sum_i |T_i|] \cup \{\text{LOSE}\}$ (take $\text{LOSE} := \sum_i |T_i| + 1$) satisfying: for all bidders i , $A_i \geq A'_i \implies H(A_i, i) \leq H(A'_i, i)$; i.e. the ironing now only respects the order between types belonging to the same bidder. Given the ironing, the allocation rule is simple: on bid vector (A_1, \dots, A_m) , the item is allocated uniformly at random to a bidder in $\arg\min_i \{H(A_i, i)\}$, or thrown away if $H(A_i, i) = \text{LOSE}$ for all i .

In multi-item auctions, the type spaces of the bidders are subsets $T_i \subset \mathbb{R}^n$, i.e. are multi-dimensional, and it is a priori not clear whether feasible mechanisms/optimal mechanisms in these settings always have clean allocation rules. Indeed, it is not even clear that all feasible single-item auctions have a clean allocation rule, such as Myerson’s allocation rule in terms of ironings described above. We show a surprising characterization result of feasible multi-item auctions when the bidders are independent, each bidder’s type is sampled from an arbitrary distribution over her type space (i.e. we do not assume that the bidder’s values for the items are necessarily independent), and the bidders have additive valuations over bundles of items. In particular, our results also characterize the set of all feasible single-item mechanisms.

As our characterization aims to characterize a broader set of mechanisms, the allocation rules it contains are necessarily more involved, albeit slightly. All we do is introduce randomization

over ironings. But what are we ironing exactly? We are still ironing the bidders' type spaces. But instead of having a single ironed surface, we maintain, for each item, a distribution over ironed surfaces. Then the allocation rule of the mechanism is rather simple: on a type profile (A_1, \dots, A_m) we allocate each item independently from the other items, by just sampling an ironed surface for this item, looking at the highest-priority plateau in this surface that is populated by at least one reported type, and giving the item to a random bidder whose reported type falls in this plateau.

We proceed to explain how the per-item ironed surfaces look. Given that bidder types are multi-dimensional, ironed surfaces are hard to envision. Interestingly, for each bidder i and item j , the picture is simplified dramatically when we apply $\pi_{ij}(\cdot)$ to embed the type space T_i of bidder i into \mathbb{R} .

Characterization of Multi-Bidder Multi-Item Auctions for Additive Bidders

The allocation rule used by any multi-bidder multi-item auction for additive bidders to allocate item j , for all j , may w.l.o.g. take the following form (the reader should draw immediate connections to Myerson's allocation rule described above):

- *i.i.d. bidders and bidder-symmetric reduced forms:* Suppose that the bidders are i.i.d. and the reduced form π_j on item j is bidder-symmetric. In this case, the auction maintains a distribution \mathcal{H}_j over rankings $H_j : T \rightarrow [|T|] \cup \{\text{LOSE}\}$ (take $\text{LOSE} := |T| + 1$) satisfying: $\pi_j(A) \geq \pi_j(A') \implies H_j(A) \leq H_j(A')$, for all $A, A' \in T$; i.e. each ranking in the support of \mathcal{H}_j respects the order induced on T by its embedding into \mathbb{R} by $\pi_j(\cdot)$. If $|H_j^{-1}(k)| > 1$ for some k and some ranking H_j in the support of \mathcal{H}_j we think of the types in $H_j^{-1}(k)$ as having been *ironed* over by H_j . Given \mathcal{H}_j , the allocation of item j is simple: on bid vector (A_1, \dots, A_m) , a ranking H_j is sampled from \mathcal{H}_j , and the item is allocated uniformly at random to a bidder in $\text{argmin}_i \{H_j(A_i)\}$, or thrown away if $H_j(A_i) = \text{LOSE}$ for all i .
- *Non-i.i.d. bidders or different reduced forms:* When the bidders are independent but non-i.i.d. or the reduced form on item j is not bidder-symmetric, the auction maintains a distribution \mathcal{H}_j over rankings $H_j : \cup_i (T_i \times \{i\}) \rightarrow [\sum_i |T_i|] \cup \{\text{LOSE}\}$ (take $\text{LOSE} := \sum_i |T_i| + 1$) satisfying: for all bidders i , $\pi_{ij}(A) \geq \pi_{ij}(A') \implies H_j(A, i) \leq H_j(A', i)$; i.e. each ranking in the support respects the order between types belonging to the same bidder, for whatever order is induced on T_i by its embedding into \mathbb{R} by $\pi_{ij}(\cdot)$. Additionally, if H_j and H'_j are two rankings in the support of \mathcal{H}_j then, for all $i, i', A_i \in T_i, A_{i'} \in T_{i'}$: $H_j(A_i, i) < H_j(A_{i'}, i') \implies H'_j(A_i, i) \leq H'_j(A_{i'}, i')$, i.e. all rankings in the support of \mathcal{H}_j are consistent with a unique total ordering of $\cup_i (T_i \times \{i\})$. Given \mathcal{H}_j , the allocation of item j is simple: on bid vector (A_1, \dots, A_m) , a ranking H_j is sampled from \mathcal{H}_j , and the item is allocated uniformly at random to a bidder in $\text{argmin}_i \{H_j(A_i, i)\}$, or thrown away if $H_j(A_i, i) = \text{LOSE}$ for all i .

Our characterization result for single-item settings is an immediate corollary of Theorems 3 (for the i.i.d. case) and 7 (for the independent case). That the above characterization holds for multi-item settings and additive bidders follows by noting that, given the lack of demand constraints, we can allocate items sequentially and the resulting allocation will be feasible as long as the supply constraints on the items are met, i.e. every item is allocated exactly once.

3.3 Optimal Multi-Dimensional Mechanism Design

The additivity of the bidders notwithstanding, we know that optimal mechanisms in multi-item settings may be quite involved. For example, in [11] it is shown that even when there are two items, and a single additive bidder whose values over the items are i.i.d. regular, a mechanism that sells lotteries over the items makes strictly more revenue than a mechanism that sells each item separately. This result implies that: we *should not* auction every item separately; and we *should*

use the values of bidders on item j to decide how to allocate item j' . So if bidder i reports higher value on item j' than bidder i' , we should not necessarily expect item j' to go to bidder i more often than bidder i' , even if the bidders are i.i.d.. To rephrase, for allocating item j' , we cannot just rank the bidders according to their bids on that particular item. We need to instead take into account their bids on all the other items.

Despite such complications, it is possible to use Theorems 2 and 6 to solve the following problem:

The BIC many-many problem. Given as input m arbitrarily correlated distributions $\mathcal{F}_1, \dots, \mathcal{F}_m$ over valuation vectors for n items, output a BIC mechanism M whose expected revenue is optimal relative to any other, possibly randomized, BIC mechanism, when played by m additive bidders whose valuation vectors are sampled independently from $\mathcal{F}_1, \dots, \mathcal{F}_m$.

Our approach to solving this problem is to use the separation oracle of Theorems 2 and 6 inside the LPs of [14]. Doing so we obtain the following:

Theorem 8. *There is a solution to the BIC many-many problem with runtime polynomial in n , m , and $\max_{i \in [m]} |\text{supp}(\mathcal{F}_i)|$.*

Theorem 9. *If, for all i , \mathcal{F}_i is item-symmetric (i.e. invariant under permutations of the item names), there exists a solution to the BIC many-many problem with runtime polynomial in m and n^c , where $c = \max_{i,j} |\text{supp}(\mathcal{F}_{ij})|$, where \mathcal{F}_{ij} is the marginal of \mathcal{F}_i on item j .*

Theorem 10. *If, for all i , \mathcal{F}_i is item-symmetric and supported on $[0, 1]^n$ there is an additive PTAS for the BIC many-many problem whose running time is polynomial in m and n (and doesn't depend on the size of the support of $\mathcal{F}_1, \dots, \mathcal{F}_m$). For all $\epsilon > 0$, the PTAS computes a BIC mechanism whose revenue is within an additive $\epsilon \cdot n$ of the optimal revenue achieved by any BIC mechanism.*

Theorem 11. *If, for all i , \mathcal{F}_i is item-symmetric and its n marginals satisfy the MHR condition, there is a multiplicative PTAS for the BIC many-many problem whose running time is polynomial in m and n (and doesn't depend on the size of the support of $\mathcal{F}_1, \dots, \mathcal{F}_m$). For all $\epsilon > 0$, the PTAS computes a BIC mechanism whose revenue is at least a $(1 - \epsilon)$ -fraction of the optimal revenue achieved by any BIC mechanism.*

Further details are provided in Appendix E.

3.4 Non-negative Linear Feasibility Constraints

Aside from the independent results of [2], it is unknown whether Border's theorem applies to other settings, or what the proper statement of Border's Theorem might even be as we depart from the single-item setting. On this front, we prove a generalization of Border's Theorem to settings with *Non-negative Linear Feasibility Constraints*. Throughout this section, P denotes a type profile, and $\vec{\phi}(P)$ denotes an *allocation* in terms of the marginal probability of giving bidder i item j , $\phi_{ij}(P)$, when the reported type profile is P . We are interested in the form of the region F where $\phi(P)$ needs to lie for there to exist a distribution over feasible allocations of items to bidders consistent with the marginal allocation probabilities specified by $\phi(P)$. In the following definition, we represent an inequality h as a vector of coefficients $\vec{c}(h)$, and a right-hand bound $d(h)$. We say that h is non-negative if every coefficient $c_{ij}(h)$ and $d(h)$ are non-negative.

Definition 5. *A set of feasible allocations $F \subseteq [0, 1]^{m \times n}$ has **Non-negative Linear Feasibility Constraints** if there exists a set H of non-negative inequalities such that $\vec{\phi}(P) \in F \Leftrightarrow \vec{c}(h) \cdot \vec{\phi}(P) \leq d(h)$, for all $h \in H$, and $\vec{\phi}(P) \geq \vec{0}$.⁵*

⁵Here, we talk about feasible allocations, not feasible mechanisms (or allocation rules). A feasible mechanism for F maps every profile P to a feasible allocation $\vec{\phi}(P) \in F$.

In a setting where no bidder i should be given more than C_i items and no item j should be given out more than once, an application of the Birkhoff-Von Neumann Theorem implies that F has non-negative linear feasibility constraints [15, 14].⁶ In fact, it is shown in [24] that a much broader range of settings have non-negative linear feasibility constraints, including settings where the constraint on the feasible allocations of items to bidders $\mathcal{S} \subseteq [m] \times [n]$ can be specified as the intersection of two arbitrary matroids.⁷

Below is our generalization of Border's Theorem to this setting, whose proof can be found in Appendix F. In the statement of the theorem, i ranges over all bidders, j ranges over all items, and P ranges over all possible type-profiles, with $\Pr[P]$ denoting the probability that profile P is sampled from \mathcal{D} and P_i denoting the type of bidder i in P .

Theorem 12. *Let F be any set of allocations with non-negative linear feasibility constraints, and \mathcal{D} be any (possibly correlated) distribution over types. Then a reduced form $\vec{\pi}$ is feasible if and only if for all weights $\{W_{ij}(A)\}_{i,j,A} \in [0, 1]^{n \times \sum_i |T_i|}$:*

$$\sum_i \sum_j \sum_{A \in T_i} W_{ij}(A) \cdot \pi_{ij}(A) \cdot \Pr[t_i = A] \leq \sum_P \Pr[P] \max_{\vec{\phi}(P) \in F} \left\{ \sum_i \sum_j W_{ij}(P_i) \cdot \phi_{ij}(P) \right\}.$$

In other words, for fixed weights $\vec{W} = \{W_{ij}(A)\}_{i,j,A}$, consider the allocation rule $M(\vec{W})$ that on profile P picks the max-weight allocation (where the weight of giving item j to bidder i is $W_{ij}(P_i)$). Then a reduced form $\vec{\pi}$ is feasible if and only if for all weights \vec{W} , the LHS is at most the expected weight of items awarded by $M(\vec{W})$.

Like in the single-item setting, it is clear that this is a necessary condition. The expected weight of items awarded by any mechanism that implements $\vec{\pi}$ is exactly the LHS, and this value clearly must be less than the expected weight of items awarded by $M(\vec{W})$, as M maximizes this value over all feasible mechanisms. The content of the theorem is that these conditions are indeed sufficient.

In addition to a proof of Theorem 12, we provide in Appendix F two false generalizations of Border's Theorem that fail even for unit-demand i.i.d. bidders in order to motivate the statement of Theorem 12.

3.4.1 Correlated Bidders, Demand Constraints

As mentioned previously, without demand constraints and without correlation among bidders' valuations, the multi-item allocation problem is no more difficult than the single-item allocation problem, as each item can just be allocated separately and independently from the other items. However, either demand constraints or correlation among bidders negates all known results (including those of Section 3.1) about feasible allocations. On this front, we give a multi-commodity flow interpretation of the problem for arbitrarily correlated bidders with demand constraints, which is similar to the network flow interpretation of [12]. Simply put, the commodities/goods in our multi-commodity flow problem are denoted $G_{ij}(A)$, for all items j , bidders i , and types $A \in T_i$. We have a source node for every item, an intermediate node for every possible profile, and a sink node for every type. We put edges from every item node to every intermediate node and every intermediate node to every sink node. The capacities on the source-intermediate edges are chosen to ensure

⁶In particular, the constraints are: a) $\forall i, \sum_j \phi_{ij}(P) \leq C_i$, b) $\forall j, \sum_i \phi_{ij}(P) \leq 1$, and c) $\forall i, j, \phi_{ij}(P) \geq 0$.

⁷In this case, let $rk_i(S)$ denote the rank of set S in matroid $i = 1, 2$. Then the constraints are: a) $\forall S \subseteq [m] \times [n], \sum_{(i,j) \in S} \phi_{ij}(P) \leq rk_1(S)$, b) $\forall S \subseteq [m] \times [n], \sum_{(i,j) \in S} \phi_{ij}(P) \leq rk_2(S)$, and c) $\forall i, j, \phi_{ij}(P) \geq 0$.

that each item is only given out once. The capacities on the intermediate-sink edges are chosen to ensure that no bidder violates his demand constraint. We then demand that a $\pi_{ij}(A) \cdot \Pr[t_i = A]$ amount of commodity $G_{ij}(A)$ is sent from source node j to sink node (i, A) . We give the details of our flow network and the values of the capacities on its edges in Appendix F, and also explain how to view any solution to the multi-commodity flow problem as a mechanism that implements the reduced form and vice versa.

References

- [1] S. Alaei. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. *Proceedings of FOCS 2010*.
- [2] S. Alaei, H. Fu, N. Haghpanah, J. Hartline, A. Malekian. Personal Communication 2011.
- [3] X. Bei, Z. Huang. Bayesian Incentive Compatibility via Fractional Assignments. *Proceedings of SODA 2011*.
- [4] S. Bhattacharya, G. Goel, S. Gollapudi and K. Munagala. Budget Constrained Auctions with Heterogeneous Items. *Proceedings of STOC 2010*.
- [5] K. C. Border. Implementation of Reduced Form Auctions: A Geometric Approach. *Econometrica*, 59(4): 1175–1187, 1991.
- [6] K. C. Border. Reduced form Auctions Revisited. *Economic Theory*, 31:167–181, 2007.
- [7] Y. Cai and C. Daskalakis. Extreme-Value Theorems for Optimal Multidimensional Pricing. *Proceedings of FOCS 2011*.
- [8] Y. Cai and Z. Huang. Simple and Nearly Optimal Multi-Item Auction. *Manuscript*, 2011.
- [9] S. Chawla, J. D. Hartline and R. D. Kleinberg. Algorithmic Pricing via Virtual Valuations. *Proceedings of EC 2007*.
- [10] S. Chawla, J. D. Hartline, D. Malec and B. Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. *Proceedings of STOC 2010*.
- [11] S. Chawla, D. Malec and B. Sivan. The Power of Randomness in Bayesian Optimal Mechanism Design. *Proceedings of EC 2010*.
- [12] Y. Che, J. Kim and K. Mierendorff. Generalized Reduced-Form Auctions: A Network-Flow Approach. *University of Zürich, ECON-Working Papers*, 2011.
- [13] E. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 18:19–33, 1971.
- [14] C. Daskalakis and S. M. Weinberg. On Optimal Multi-Dimensional Mechanism Design. *arXiv Report*, 2011.
- [15] S. Dobzinski, H. Fu and R. D. Kleinberg. Optimal Auctions with Correlated Bidders are Easy. *Proceedings of STOC 2011*.
- [16] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.

- [17] J. Hartline, R. Kleinberg, A. Malekian. Multi-Parameter Bayesian Algorithmic Mechanism Design. *Proceedings of SODA* 2011.
- [18] J. D. Hartline and B. Lucier. Bayesian Algorithmic Mechanism Design. *Proceedings of STOC* 2010.
- [19] 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.
- [20] E. Maskin and J. Riley. Optimal Auctions with Risk Averse Buyers. *Econometrica*, 52(6): 1473–1518, 1984.
- [21] S. A. Matthews. On the Implementability of Reduced Form Auctions. *Econometrica*, 52(6):1519–1522, 1984.
- [22] R. B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 1981.
- [23] N. Nisan, T. Roughgarden, E. Tardos and V. V. Vazirani (eds.). *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [24] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [25] W. Vickrey. Counterspeculations, Auctions, and Competitive Sealed Tenders. *J. Finance*, 16:15–27, 1961.

A A Geometric Algorithm

Carathéodory’s theorem states that every point \vec{x} inside an n -dimensional polytope P can be written as a convex combination of at most $n + 1$ corners of P . In this section, we provide an efficient algorithm for coming up with such a combination. We will consider polytopes that are described as an intersection of half-spaces. Each half-space is defined by a hyperplane h together with a choice of a side. We use $B(P)$ to denote the set of half-spaces, but overload notation using $B(P)$ to also denote the set of boundary hyperplanes of the polytope P . We reserve the symbol h to denote hyperplanes. In addition, we consider cases where $|B(P)|$ may be exponentially large, and we only have an implicit description of $B(P)$. That is, we have access to a *boundary oracle* BO that outputs yes on input h if $h \in B(P)$, and no otherwise. We also have access to a separation oracle, SO , that outputs yes on input \vec{x} if $\vec{x} \in P$, and outputs some $h \in B(P)$ if \vec{x} is on the wrong side of h (and therefore not in P). We will talk about one more algorithm related to P :

Definition 6. *CO is a **corner oracle** for P if it has the following behavior. Given as input a set of hyperplanes B , CO outputs no if $B \not\subseteq B(P)$, or $(\bigcap_{h \in B} h) \cap P = \emptyset$ (i.e. the hyperplanes are not boundary hyperplanes of P , or they are boundary hyperplanes, but do not intersect inside P). Otherwise, CO outputs a corner of P inside $\bigcap_{h \in B} h$.*

It is clear that CO has well-defined behavior on all inputs. If B contains only boundary hyperplanes of P , and the intersection of these hyperplanes with P is non-empty, this region must contain a corner of P . Now we describe our algorithmic problem in terms of these algorithms:

Question 1. *Given as input a boundary oracle, BO , separation oracle SO , and corner oracle CO for some n -dimensional polytope P , and a point \vec{x} , output no if $\vec{x} \notin P$. Otherwise, output $c_1, \dots, c_{n+1}, \vec{a}_1, \dots, \vec{a}_{n+1}$ such that \vec{a}_i is a corner of P for all i , $\sum_i c_i = 1$, and $\sum_i c_i \vec{a}_i = \vec{x}$.*

It follows from Carathéodory's theorem that such c_i, \vec{a}_i exist whenever $\vec{x} \in P$. We provide an algorithm to find such a solution whenever it exists. At a high level, we begin with the input \vec{x} and maintain at all times two points $\vec{y} \in P, \vec{z} \in P$, such that $\vec{x} = c\vec{y} + (1-c)\vec{z}$, for some $c \in [0, 1]$. After step t of the algorithm is completed, \vec{y} is the convex combination of at most t corners of P , and \vec{z} lies in the $(n-t)$ -dimensional intersection of t hyperplanes of $B(P)$. Hence, after at most n steps, \vec{z} will lie in a 0-dimensional space, and therefore must be a corner, so the algorithm will terminate after at most $n+1$ steps.

To go from step t to step $t+1$, we pick an arbitrary corner, \vec{a}_t , that lies in the intersection of the t hyperplanes where \vec{z} lies. Then, we let c_t be as large as possible without pushing the point $\frac{(1-\sum_{j<t} c_j)\vec{z} - c_t \vec{a}_t}{1-c_t-\sum_{j<t} c_j}$ outside of P . We update \vec{z} to $\vec{z}_{\text{new}} = \frac{(1-\sum_{j<t} c_j)\vec{z}_{\text{old}} - c_t \vec{a}_t}{1-c_t-\sum_{j<t} c_j}$ and update \vec{y} appropriately to include \vec{a}_t in its convex combination of corners. The new \vec{z} must lie at the intersection of the original t hyperplanes where the old \vec{z} lied, as well as a new $h \in B(P)$ that stopped us from further increasing c_t . Below is a formal description of the algorithm. In the description, E denotes the set of hyperplanes whose intersection contains \vec{z} (the empty intersection is the entire space).

Algorithm 1 Algorithm for writing \vec{x} as a convex combination of at most $n+1$ corners

```

1: Initialize:  $i := 1, \vec{y} := \vec{0}, \vec{z} := \vec{x}, E := \emptyset, c_i := 0, \vec{a}_i := \vec{0} \forall i \in [n+1]$ .
2: Invariants:  $c := \sum_i c_i, \vec{y} := \frac{1}{c} \sum_i c_i \vec{a}_i$ , or  $\vec{0}$  if  $c = 0, c\vec{y} + (1-c)\vec{z} = \vec{x}$ .
3: if  $SO(\vec{x}) \neq \text{yes}$  then
4:   Output no.
5: end if
6: while  $c < 1$  do
7:   Set  $\vec{a}_i := CO(E)$ .
8:   if  $\vec{a}_i = \vec{z}$  then
9:     Set  $c_i := 1 - c$ .
10:    Output  $c_1, \dots, c_{n+1}, \vec{a}_1, \dots, \vec{a}_{n+1}$ .
11:   else
12:     Set  $D := \max\{d \mid (1+d)\vec{z} - d\vec{a}_i \in P\}$ .
13:     Set  $E_i = SO((1+D+\epsilon)\vec{z} - (D+\epsilon)\vec{a}_i)$  for sufficiently small  $\epsilon > 0$ . /* the appropriate
        choice of  $\epsilon$  is explained in the proof of Theorem 13*/
14:     Update:  $c_i := (1 - \frac{1}{1+D})(1-c), \vec{z} := \frac{1-c}{1-c-c_i}\vec{z} - \frac{c_i}{1-c-c_i}\vec{a}_i, \vec{y} := \frac{c}{c+c_i}\vec{y} + \frac{c_i}{c+c_i}\vec{a}_i, c := c + c_i,$ 
         $E := E \cup E_i, i := i + 1$ .
15:   end if
16: end while

```

Theorem 13. *Algorithm 1 correctly answers Question 1. Furthermore, if b is the maximum number of bits used in the description of any coefficient of any hyperplane $h \in B(P)$ and any coordinate of \vec{x} , then the runtime of Algorithm 1 is polynomial in n, b and the runtimes of SO and CO .*

Proof. First, we describe how to execute Steps 12 and 13 of the algorithm, as it is clear how to execute every other step. Step 12 can be done by solving a linear program using SO . Explicitly, maximize d subject to $(1+d)\vec{z} - d\vec{a}_i \in P$. For Step 13, we will explain later in the proof how to choose an ϵ small enough so that the following property is satisfied:

(P): for all $h \in B(P)$ and for whatever D is computed in Step 12 of the algorithm, if $(1+D)\vec{z} - D\vec{a}_i$ is not contained in h , then $(1+D+\epsilon)\vec{z} - (D+\epsilon)\vec{a}_i$ is on the same side of h as $(1+D)\vec{z} - D\vec{a}_i$.

We will explain later why (P) suffices for the correctness of the algorithm, how to choose an ϵ so that (P) holds, and why its description complexity is polynomial in n and b .

We start with justifying the algorithm's correctness, assuming that ϵ is chosen so that (P) holds. We observe first that $\sum_i c_i \leq 1$ always. If the algorithm ever increases c , it is because $\vec{z} \neq \vec{a}_i$. If this is the case, then D from Step 12 will have some finite positive value. So $(1 - \frac{1}{1+D})(1 - c) < 1 - c$, and adding c_i to c will not increase c past 1. We also observe that all the invariants declared in Step 2 hold throughout the course of the algorithm. This can be verified by simply checking each update rule in Step 14. Finally, we argue that every time the algorithm updates E , the dimension of $\bigcap_{h \in E} h$ decreases by 1, and $(\bigcap_{h \in E} h) \cap P \neq \emptyset$ is maintained. Because \vec{a}_i and \vec{z} both lie in $\bigcap_{h \in E} h$ when Step 13 is executed, none of the hyperplanes in this intersection can possibly be violated at $(1 + D + \epsilon)\vec{z} - (D + \epsilon)\vec{a}_i$. Therefore, the hyperplane output by $SO((1 + D + \epsilon)\vec{z} - (D + \epsilon)\vec{a}_i)$ must reduce the dimension of $\bigcap_{h \in E} h$ by 1 when added to E at Step 14. Furthermore, because E_i is violated at $(1 + D + \epsilon)\vec{z} - (D + \epsilon)\vec{a}_i$, but not at $(1 + D)\vec{z} - D\vec{a}_i$, it must be the case that $(1 + D)\vec{z} - D\vec{a}_i$ lies in the hyperplane E_i . (This holds because we will guarantee that our ϵ satisfies Property (P), described above.) Because this point is clearly in P , in the hyperplane E_i , and in all of the hyperplanes in E , it bears witness that we maintain $(\bigcap_{h \in E} h) \cap P \neq \emptyset$ always. Hence after at most n iterations of the while loop, the dimension of the remaining space is 0, and we must enter the case where $\vec{a}_i = \vec{z}$. The algorithm then exits outputting a convex combination of corners equaling \vec{x} .

It remains to argue that a choice of ϵ satisfying Property (P) is possible. Assuming the correctness of our algorithm, we show first that all the coefficients c_i computed by the algorithm have low bit complexity. Indeed, let $\vec{b}_i = (\vec{a}_i, 1)$ for all i . Once we know the algorithm is correct, the c_i 's satisfy

$$\sum_i c_i \vec{b}_i = (\vec{x}, 1),$$

where c_i and \vec{a}_i are outputs of our algorithm. We will argue that, for these \vec{a}_i 's, the above system of linear equations has a unique solution. If not, let \vec{c} and \vec{c}' be two different solutions, and $d_i = c_i - c'_i$. We will show by induction on i that $d_i = 0$ for all i . In the base case, consider the hyperplane in E_1 . We can write a corresponding $(n + 1)$ dimensional vector \vec{t}_1 , such that for all $\vec{x}' \in P$, $(\vec{x}', 1) \cdot \vec{t}_1 \leq 0$, and for all $i > 1$, $\vec{b}_i \cdot \vec{t}_1 = 0$. But $\vec{b}_1 \cdot \vec{t}_1 \neq 0$, otherwise, for any D , $(1 + D)\vec{z} - D\vec{a}_1$ does not violate the constraint in E_1 . On the other hand, $\sum_i d_i \vec{b}_i \cdot \vec{t}_1 = 0$, therefore $d_1 = 0$. Now assume when $i < k$, $d_i = 0$, we will argue that $d_k = 0$. Let \vec{t}_k be the corresponding vector for the hyperplane in E_k . For any $j > k$, $\vec{b}_j \cdot \vec{t}_k = 0$, and by the Inductive Hypothesis, for any $i < k$, $d_i = 0$, therefore $d_k \vec{b}_k \cdot \vec{t}_k = 0$. But we know $\vec{b}_k \cdot \vec{t}_k \neq 0$, otherwise, for any D , $(1 + D)\vec{z} - D\vec{a}_k$ does not violate the constraint in E_k . So $d_k = 0$. Thus, $d_i = 0$ for all i .

So we have argued that the c_i 's are in fact the unique solution to the above linear system. We also know that the corners \vec{a}_i (in fact all corners of the polytope) have $\text{poly}(n, b)$ bit complexity. Applying the theory of Gaussian elimination, we deduce that each c_i can be described using no more than $\text{poly}(n, b)$ bits, so the coefficients output by our algorithm have low bit complexity. Hence the \vec{z} maintained by the algorithm has $\text{poly}(n, b)$ bit complexity. So the intersections d_h of the ray $R(d) = \{(1 + d)\vec{z} - d\vec{a}_i\}$ with the hyperplanes $h \in B(P)$ that do not contain both \vec{z} and \vec{a}_i (and hence the whole ray) also have $\text{poly}(n, b)$ bit complexity. This guarantees that we can choose ϵ to be $2^{-\text{poly}(n, b)}$ to satisfy Property (P).

The above reasoning justifies the correctness of the algorithm. It is also now clear that every step runs in time polynomial in b, n , the runtime of SO and the runtime of CO , and each step is executed at most $n + 1$ times. So the entire algorithm runs in polynomial time. \square

B Single Item, I.I.D. Bidders, Bidder-Symmetric Reduced Forms

We provide the details and proofs of the techniques discussed in Section 3.1.1, beginning with some technical lemmas.

Lemma 1. *Every feasible bidder-symmetric reduced form auction π for i.i.d. bidders and a single item can be implemented (not necessarily exactly implemented) as a distribution over well-ordered w.r.t. π hierarchical mechanisms.*

Proof. Our proof is by induction on the number of distinct non-zero values in $\{\pi(A)\}_{A \in T}$.

Base Case: There is a single non-zero value in this set, equal to some $x \leq 1$. The mechanism that gives the item uniformly at random to a bidder whose reported type A satisfies $\pi(A) = x$ (if such bidder shows up) implements this reduced form as long as it is feasible.

Inductive Hypothesis: For all $0 < k < \theta$, every feasible reduced form with k distinct non-zero values in $\{\pi(A)\}_{A \in T}$ can be implemented as a distribution over well-ordered w.r.t. π hierarchical mechanisms.

Inductive Step: We show that the inductive hypothesis extends to the case where there are exactly θ distinct non-zero values in $\{\pi(A)\}_{A \in T}$. Let X denote the set of all distributions over well-ordered w.r.t. π hierarchical mechanisms. Then X can be interpreted as a closed, bounded subset of \mathbb{R}^{2^θ} , where each coordinate denotes the probability of using one of the 2^θ well-ordered w.r.t. π hierarchical mechanisms. Therefore, X is compact. For a distribution over hierarchical mechanisms $M \in X$, denote by $M(A)$ the probability that a bidder reporting type A receives the item under M . Define the function $F : X \rightarrow \mathbb{R}$ as:

$$F(M) = \max_A \{\pi(A) - M(A)\}.$$

Let us use the Euclidean distance in X as a subset of \mathbb{R}^{2^θ} . As X is a compact space, and F is a continuous function, F attains its minimum in X . Let M denote one such minimizer of F . Then if $F(M) \leq 0$, M implements the reduced form. If $F(M) > 0$, we will show a contradiction. Let S denote the subset of types $\operatorname{argmax}_A \{\pi(A) - M(A)\}$, i.e. the subset of types who are the most unsatisfied by M .

We show first that, if S contains every non-zero type, then the reduced form is infeasible. We may, w.l.o.g., assume that M always awards the item to a non-zero type if one shows up, as this will not decrease $M(A)$ for any non-zero A . Therefore, we know that

$$\sum_{A: \pi(A) \neq 0} \Pr[A]M(A) = \Pr[\text{see a non-zero type}].$$

However, if $\pi(A) - M(A) > 0$ for all non-zero types, then we must have $\sum_{A: \pi(A) \neq 0} \Pr[A]\pi(A) > \Pr[\text{see a non-zero type}]$ and the reduced form is infeasible. So if the reduced form is feasible, S must be missing at least one non-zero type.

Now let $s = |\{\pi(A)\}_{A \in S}|$ be the number of distinct non-zero values assigned by π to types in S . We argue that $s < \theta$. To see this, it suffices to observe the following: for all types B and B' , $\pi(B) = \pi(B')$ implies that $M(B) = M(B')$ (this is because, by definition, all hierarchical mechanisms H in the support of M satisfy $H(B) = H(B')$). So in particular either $B, B' \in S$ or $B, B' \notin S$, but it cannot be that one of B, B' is in S and the other is not. And because S is missing at least one non-zero type, $s < \theta$.

To show a contradiction to $F(M) > 0$, let us define a new reduced form π' as follows. For all $A \in S$, set $\pi'(A) = \pi(A)$. For all $A \notin S$, set $\pi'(A) = \max_{B \in S | \pi(B) < \pi(A)} \{\pi(B)\}$, unless $\{B \mid B \in$

$S \wedge \pi(B) < \pi(A)$ is empty in which case we set $\pi'(A) = 0$. Observe that the number of distinct non-zero values in $\{\pi'(A)\}_{A \in T}$ is exactly $s < \theta$. So it follows by our inductive hypothesis that π' can be implemented by a distribution over well-ordered (with respect to π') hierarchical mechanisms. In fact, as $\pi(A) \geq \pi(A') \Rightarrow \pi'(A) \geq \pi'(A')$, every hierarchical mechanism that is well-ordered with respect to π' is also well-ordered with respect to π . Call M' the distribution over well-ordered hierarchical mechanisms implementing π' . Now, set $\epsilon = (F(M) - \operatorname{argmax}_{A \notin S} \{\pi(A) - M(A)\})/2$, and consider the distribution $M'' = (1 - \epsilon)M + \epsilon M'$ (with probability $(1 - \epsilon)$ sample from M , with probability ϵ sample from M').

What is $M''(A)$? If $A \in S$, then $M'(A) = \pi(A)$, so $M''(A) = (1 - \epsilon)M(A) + \epsilon\pi(A)$. So for all $A \in S$, $M''(A) > M(A)$, hence $\pi(A) - M''(A) < F(M)$.

If $A \notin S$, then $M'(A) \geq 0$, so $M''(A) \geq (1 - \epsilon)M(A) \geq M(A) - \epsilon$, so we get $\pi(A) - M''(A) \leq \pi(A) - M(A) + \epsilon < F(M)$. Putting both observations together, we see that $F(M'') < F(M)$, a contradiction.

So we must have $F(M) \leq 0$, meaning M implements the reduced form, completing the inductive step and the proof of the lemma. \square

Corollary 1. *Every feasible bidder-symmetric reduced form π can be exactly implemented as a distribution over well-ordered w.r.t. π hierarchical mechanisms.*

Proof. It follows from Lemma 1 that π can be implemented as a distribution over well-ordered w.r.t. π hierarchical mechanisms. Let then X denote the set of distributions over well-ordered w.r.t. π hierarchical mechanisms that implement π . As in Lemma 1 the set X , viewed as a subset of \mathbb{R}^{2^θ} , where θ is the number of distinct non-zero values in π , is compact. We can also define the function $G : X \rightarrow \mathbb{R}$ as:

$$G(M) = \max_A \{M(A) - \pi(A)\}.$$

Equipping X with the Euclidean distance of \mathbb{R}^{2^θ} , G is a continuous function on X . As X is compact and G continuous, G attains its minimum in X .

We show that the minimum of G is exactly 0 (i.e. that a minimizer of G exactly implements π), following an induction similar to the one used in Lemma 1 in terms of the number of distinct non-zero values in $\{\pi(A)\}_{A \in T}$. We sketch the steps involved for the inductive step. Take any minimizer M of G . If $G(M) \leq 0$, then because M has to implement π , M must exactly implement π . If $G(M) > 0$, then let $S = T - \operatorname{argmax}_A \{M(A) - \pi(A)\}$. Then, for all $A \in S$, define $\pi'(A) = \pi(A)$. For all $A \notin S$, define $\pi'(A) = \max_{B \in S | \pi(B) \leq \pi(A)} \{\pi(B)\}$, unless $\{B : B \in S \wedge \pi(B) \leq \pi(A)\}$ is empty, in which case set $\pi'(A) = 0$. As $\operatorname{argmax}\{M(A) - \pi(A)\}$ can't possibly be empty the number of distinct non-zero values in $\{\pi'(A)\}_{A \in T}$ is smaller than that in $\{\pi(A)\}_{A \in T}$. (We still use the observation that, for all types B and B' , $\pi(B) = \pi(B')$ implies that $M(B) = M(B')$). The rest of the inductive step proceeds identically with that in the proof of Lemma 1, resulting in a contradiction. Hence, it can't be that the minimizer of G satisfies $G(M) > 0$. \square

To proceed we need a definition. Let $> := A_1 > A_2 > \dots > A_k$ be a total ordering of a set $T_k := \{A_1, \dots, A_k\}$ of k types, and consider a setting where we have m bidders whose types are distributed i.i.d. over T_k . We say that a bidder-symmetric reduced form $\pi : T_k \rightarrow [0, 1]$ respects $>$ iff $\pi(A_1) \geq \pi(A_2) \geq \dots \geq \pi(A_k)$. We also say that an hierarchical mechanism H on T_k is well-ordered with respect to $>$ iff $H(A_1) \leq H(A_2) \leq \dots \leq H(A_k)$.⁸ With respect to these definitions we show the following proposition.

⁸Notice that this definition of well-ordered hierarchical mechanism (with respect to $>$) is very similar to its counterpart in the main body (with respect to π), but different. Being well-ordered with respect to π certainly

Proposition 1. Consider m bidders whose types are distributed i.i.d. over $T_k := \{A_1, \dots, A_k\}$ and a total ordering $> := A_1 > A_2 > \dots > A_k$ of T_k . The set of feasible bidder-symmetric reduced forms that respect $>$ is a k -dimensional polytope whose corners are exactly the 2^k hierarchical mechanisms that are well-ordered with respect to $>$.

Proof. As a corollary of Theorem 1, a reduced form respects $>$ and is feasible if and only if

$$\pi(A_i) \geq \pi(A_{i+1}) \quad \forall i \in [k]; \quad (4)$$

$$\sum_{j \leq i} m \cdot \Pr[A_j] \pi(A_j) \leq 1 - \left(1 - \sum_{j \leq i} \Pr[A_j]\right)^m \quad \forall i \in [k]; \quad (5)$$

where for notational convenience we denote $\pi(A_{k+1}) = 0$. We have hence shown that the set of feasible bidder-symmetric reduced forms that respect $>$ is a k -dimensional polytope.

We proceed to show that each well-ordered w.r.t. $>$ hierarchical mechanism is a corner. Let H be such a mechanism and π be the reduced form that it induces. Then, for all i (including $i = k$, denoting $H(A_{k+1}) = \text{LOSE}$) we either have $H(A_i) = H(A_{i+1})$, in which case $\pi(A_i) = \pi(A_{i+1})$, or $H(A_i) < H(A_{i+1})$, in which case $\sum_{j \leq i} m \cdot \Pr[A_j] \pi(A_j) = 1 - (1 - \sum_{j \leq i} \Pr[A_j])^m$, because the item is always awarded to one of the top i types whenever one is present. Therefore, at least k of the inequalities defining the polytope are tight. And it is easy to see that there is a unique reduced form making these inequalities tight. It is also clear that every well-ordered w.r.t. $>$ hierarchical mechanism is inside the polytope. So every well-ordered w.r.t. $>$ hierarchical mechanism is a corner of the polytope.

Finally, we show that there are no other corners. Assume for contradiction that there was a corner π of the polytope that is not a well-ordered w.r.t. $>$ hierarchical mechanism. Then by Corollary 1, we know that π can be written as a convex combination of well-ordered w.r.t. $>$ hierarchical mechanisms, and hence as a convex combination of well-ordered w.r.t. $>$ hierarchical mechanisms. (As π respects $>$ a hierarchical mechanism that is well-ordered w.r.t. $>$ will also be well-ordered w.r.t. $>$). As every well-ordered w.r.t. $>$ hierarchical mechanism is a corner of the polytope, and π is not one of them, this means that π can be written as a convex combination of *other* corners of the polytope, which contradicts that π is itself a corner.

Therefore we have shown that every feasible bidder-symmetric reduced form respecting $>$ lies inside the afore-described polytope, every well-ordered w.r.t. $>$ hierarchical mechanism is a corner of this polytope, and there are no other corners. This establishes the proposition. \square

Now we can put everything together to prove Theorems 3 and 2.

Proof of Theorem 3: Suppose that the bidders' types are sampled i.i.d. from T according to \mathcal{D}_1 , and let π be a feasible bidder-symmetric reduced form. We do the following preprocessing operation on the set of types T :

TYPEMERGE: Find a maximal set of types $A_1, \dots, A_\ell \in T$ such that $\pi(A_1) = \dots = \pi(A_\ell)$. Then remove types A_1, \dots, A_ℓ from T and add super-type $\langle A_1, \dots, A_\ell \rangle$ into T ; change \mathcal{D}_1 to never sample any of A_1, \dots, A_ℓ and sample the super-type $\langle A_1, \dots, A_\ell \rangle$ with probability $\sum_i \Pr[A_i]$; and set $\pi(\langle A_1, \dots, A_\ell \rangle) = \pi(A_1)$. Repeat this procedure until every type in T has different π value, i.e. until the set $\{\pi(A)\}_{A \in T}$ has cardinality $|T|$.

Let T' , π' , \mathcal{D}'_1 be the type-set, reduced-form, type-distribution resulting from the TYPEMERGE operation on input π . We claim that:

imposes the same constraints as being well ordered with respect to any $>$ that π respects. The difference is that being well-ordered with respect to π may also impose some equality constraints, if $\pi(A) = \pi(B)$ for types $A \neq B$.

1. π' is a feasible bidder-symmetric reduced form for bidders sampled i.i.d. from \mathcal{D}'_1 . This follows immediately from the feasibility of π . Indeed it follows from Theorem 1 and our discussion in Section 3.1.1 that a sufficient condition for the feasibility of π' is for it to satisfy Eq. (1) for all subsets of types of the form $\{A \mid A \in T' \wedge \pi'(A) \geq x\}$. On the other hand the feasibility of π implies that π satisfies Eq. (1) for all subsets of types of the form $\{A \mid A \in T \wedge \pi(A) \geq x\}$. This together with the nature of our TYPEMERGE operation implies that π' satisfies the afore-mentioned sufficient conditions for feasibility.
2. A mechanism that exactly implements π' immediately gives a mechanism exactly implementing π . Indeed, to implement π we just run the mechanism implementing π' after replacing in the reported type vector every type A that was removed from T by TYPEMERGE by its corresponding super-type.
3. A hierarchical mechanism H' that is well-ordered w.r.t. π' can be *expanded* into a hierarchical mechanism H that is well-ordered w.r.t. π . Indeed, if super-type $\langle A_1, \dots, A_\ell \rangle$ replaced types A_1, \dots, A_ℓ during the TYPEMERGE operation we set $H(A_i) := H'(\langle A_1, \dots, A_\ell \rangle)$, for all $i = 1, \dots, \ell$. On the other hand, if a type A belongs in both T and T' , we set $H(A) := H'(A)$. Moreover, if $\pi_{H'}$, π_H are respectively the reduced forms induced by H' and H , the following property is satisfied. If super-type $\langle A_1, \dots, A_\ell \rangle$ replaced types A_1, \dots, A_ℓ during the TYPEMERGE operation, then $\pi_H(A_i) := \pi_{H'}(\langle A_1, \dots, A_\ell \rangle)$, for all $i = 1, \dots, \ell$. On the other hand, if a type A belongs in both T and T' , then $\pi_H(A) := \pi_{H'}(A)$.

Now, suppose that the cardinality of T' is k . Given that π' assigns a distinct probability to every type in T' , it induces a total ordering on T' . In particular, suppose that $T' := \{A_1, \dots, A_k\}$, where $\pi'(A_1) > \pi'(A_2) > \dots > \pi'(A_k)$. By Proposition 1, π' lies inside a k -dimensional polytope whose corners are exactly the 2^k hierarchical mechanisms that are well-ordered with respect to the order $> := A_1 > \dots > A_k$. By Carathéodory's Theorem, every point in the polytope can be written as a convex combination of at most $k + 1$ corners. As a convex combination of corners is exactly a distribution over well-ordered hierarchical mechanisms w.r.t. $>$, we get that π' can be written as a distribution over hierarchical mechanisms that are well-ordered w.r.t. $>$, and hence also w.r.t. π' . Now we can expand all hierarchical mechanisms in the support of the distribution, according to the procedure described in Step 3 above, to obtain that π can be written as a distribution over hierarchical mechanisms that are well-ordered w.r.t. π . \square

Proof of Theorem 2: It follows from our discussion in Section 3.1.1 that a bidder-symmetric reduced form $\tilde{\pi}$ is infeasible if and only if it violates Eq. (1) for a subset of types of the form $\{A \mid A \in T \wedge \tilde{\pi}(A) \geq x\}$. Since there are at most $c \equiv |T|$ such sets we can efficiently determine feasibility of a given reduced form $\tilde{\pi}$ or provide a hyperplane separating it from the set of feasible reduced forms.

We now need to describe how to efficiently find a mechanism implementing a reduced form $\tilde{\pi}$ that is feasible. In view of the TYPEMERGE operation defined in the proof of Theorem 3, we can w.l.o.g. assume that $\tilde{\pi}$ assigns a distinct probability to every type in T . (Otherwise we can always run TYPEMERGE to merge types sharing the same $\tilde{\pi}$ -probability to super-types and apply the procedure outlined below to the output of the TYPEMERGE operation, and then go back to the original $\tilde{\pi}$). Under the assumption that $\tilde{\pi}$ assigns distinct probabilities to all types in T , Proposition 1 implies that $\tilde{\pi}$ lies inside a c -dimensional polytope, P , whose corners are the well-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanisms. Therefore we can directly apply Theorem 13 of Appendix A to write $\tilde{\pi}$ as a convex combination of such hierarchical mechanisms, as long as we can describe the boundary oracle BO , corner oracle CO and separation oracle SO that are needed for Theorem 13. BO is trivial to implement, as we just have to include in the set of halfspaces defining the boundary of P those inequalities described in the proof of Proposition 1. For CO , on input B ,

we first check that every hyperplane $h \in B$ satisfies $BO(h) = \text{yes}$. If not, output no. Otherwise we need to check if $\bigcap_{h \in B} h$ contains a corner of P . We know that the corners of P are exactly the well-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanisms. So none of the corners lies in the intersection of the hyperplanes $\pi(A_i) = \pi(A_{i+1})$ and $\sum_{j \leq i} m \cdot \Pr[A_j] \pi(A_j) = 1 - (1 - \sum_{j \leq i} \Pr[A_j])^m$, for any i . (Indeed, for a hierarchical mechanism H and its induced reduced form π , $\pi(A_i) = \pi(A_{i+1})$ implies that $H(A_i) = H(A_{i+1})$, yet $\sum_{j \leq i} m \cdot \Pr[A_j] \pi(A_j) = 1 - (1 - \sum_{j \leq i} \Pr[A_j])^m$ implies $H(A_i) > H(A_{i+1})$). So, if B contains any pair of hyperplanes of this form, output no. Otherwise, for all i such that $\pi(A_i) = \pi(A_{i+1}) \in B$, set $H(A_i) = H(A_{i+1})$, otherwise set $H(A_i) = H(A_{i+1}) - 1$. This defines a well-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanism that is in $\bigcap_{h \in B} h$, so have CO output H . Finally, SO is easy to implement as we can just check each of the $2 \cdot c$ inequalities written in the proof of Proposition 1 one by one.

So because we can implement BO, CO, SO in polynomial time, we can apply Theorem 13 to write $\tilde{\pi}$ as a convex combination of at most $c + 1$ corners, which is exactly a distribution over at most $c + 1$ well-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanisms in polynomial time. \square

C Single Item, Independent Bidders

Here we provide the details of Section 3.1.2, and the proofs of Theorems 4, 5 and 6, postponing the proof of Theorem 7 to Appendix D. Before proving our theorems, we show that the concept of virtual π s is necessary. As in, Theorem 4 would be false if we tried to replace $\hat{\pi}$ with π .

Proposition 2. *There exist reduced forms that are infeasible, yet for all S_x^i of the form $S_x^i = \{A \mid \pi_i(A) > x, \forall i\}$:*

$$\sum_i \sum_{A \in S_x^i} \pi_i(A) \Pr[t_i = A] \leq 1 - \prod_{i=1}^m (1 - \sum_{A \in S_x^i} \Pr[t_i = A]).$$

Proof. Consider the case with two bidders. Bidder 1 has two types, with $\Pr[t_1 = A] = 1/8$, $\Pr[t_1 = B] = 7/8$, $\pi_1(A) = 5/8$, $\pi_1(B) = 0$. Bidder 2 has two types, with $\Pr[t_2 = C] = 1/2$, $\Pr[t_2 = D] = 1/2$, $\pi_2(C) = 1$, $\pi_2(D) = 3/4$.

Then this reduced form is infeasible. Indeed, observe that C must always receive the item whenever $t_2 = C$, which happens with probability $1/2$. So if we have $\pi_2(C) = 1$, we cannot also have $\pi_1(A) > 1/2$. So the set $\{A, C\}$ forms a constricting set. However, the sets of the form S_x^i are $\{C\}, \{C, D\}, \{C, D, A\}, \{C, D, A, B\}$, and they all satisfy the above inequality. \square

Proposition 2 shows us that ordering the types of all bidders by decreasing π doesn't allow us to correctly determine the feasibility of a reduced form. Similarly, a partial ordering of the types that only orders a single bidder's types by decreasing π doesn't give enough structure to efficiently determine the feasibility of the reduced form. What we need is a correct total ordering of the types of all bidders, and we can obtain it using virtual π s. Here is a quick observation about the virtual π s, followed by a proof of Theorem 4.

Observation 1. *For two types $A, B \in T_i$, $\hat{\pi}_i(A) \geq \hat{\pi}_i(B) \Leftrightarrow \pi_i(A) \geq \pi_i(B)$.*

Proof. If $\pi_i(A) \geq \pi_i(B)$, then $\Pr[\pi_i(t_i) \leq \pi_i(A)] \geq \Pr[\pi_i(t_i) \leq \pi_i(B)]$. Therefore, $\hat{\pi}_i(A) \geq \hat{\pi}_i(B)$. The other direction is identical. \square

Proof of Theorem 4: We know from [6, 12], that if a reduced form mechanism is infeasible, then there is some constricting set of the form $S = \bigcup_{i=1}^m S_{x_i}$, where $S_{x_i} = \{A \mid \pi_i(A) \geq x_i, A \in T_i\}$.

(Forgive the abuse of notation here. Formally, S is a collection of m sets of types, one for each bidder. To avoid cumbersome notation and take union casually in this proof, let us assume that a type $A \in T_i$ carries the name of bidder i , for all i .) Now consider any minimal constricting set of this form, i.e. a choice of x_1, \dots, x_m such that replacing S_{x_i} with $S_{x_i} - \{A\}$ ($A \in S_{x_i}$) results in S no longer being a constricting set.⁹ Now let $(i, A) \in \operatorname{argmin}_{i, A \in S_{x_i}} \hat{\pi}_i(A)$. Then by Observation 1 and by our choice of S , $S - \{A\}$ is not a constricting set. Therefore, adding A to $S - \{A\}$ must increase the left-hand bound by more than it increases the right-hand bound:

$$\begin{aligned} \Pr[t_i = A] \pi_i(A) &> \Pr[t_i = A] \prod_{j \neq i} \Pr[\pi_j(t_j) < x_j] \\ \implies \frac{\pi_i(A)}{\prod_{j \neq i} \Pr[\pi_j(t_j) < x_j]} &> 1. \end{aligned}$$

Now consider any other $A' \in T_k$, $A' \notin S$ and $\hat{\pi}_k(A') \geq \hat{\pi}_i(A)$. Observe first that we must have A' from some bidder $k \neq i$, as every $A'' \in T_i$ with $\hat{\pi}_i(A'') \geq \hat{\pi}_i(A)$ has $\pi_i(A'') \geq \pi_i(A) \geq x_i$, so we would have $A'' \in S$. So for this A' , we have:

$$\begin{aligned} \pi_k(A') \Pr[\pi_k(t_k) \leq \pi_k(A')] &\geq \pi_i(A) \Pr[\pi_i(t_i) \leq \pi_i(A)] \\ \implies \pi_k(A') \Pr[\pi_k(t_k) < x_k] &\geq \pi_i(A) \Pr[\pi_i(t_i) < \pi_i(A)] \\ \implies \pi_k(A') \Pr[\pi_k(t_k) < x_k] &\geq \pi_i(A) \Pr[\pi_i(t_i) < x_i] \\ \implies \pi_k(A') \prod_{j \neq i} \Pr[\pi_j(t_j) < x_j] &\geq \pi_i(A) \prod_{j \neq k} \Pr[\pi_j(t_j) < x_j] \\ \implies \frac{\pi_k(A')}{\prod_{j \neq k} \Pr[\pi_j(t_j) < x_j]} &\geq \frac{\pi_i(A)}{\prod_{j \neq i} \Pr[\pi_j(t_j) < x_j]}. \end{aligned}$$

And by our choice of A and the work above, we obtain:

$$\begin{aligned} \frac{\pi_k(A')}{\prod_{j \neq k} \Pr[\pi_j(t_j) < x_j]} &> 1 \\ \implies \Pr[t_k = A'] \pi_k(A') &> \Pr[t_k = A'] \prod_{j \neq k} \Pr[\pi_j(t_j) < x_j]. \end{aligned}$$

This equation tells us directly that we could add A' to S and still get a constricting set. In fact, it tells us something stronger. If $S' = \bigcup_j S'_j$, where $S'_j \subseteq T_j$, is any constricting set containing S , then we could add A' to S' and still have a constricting set. This is because the change to the left-hand side of the inequality is the same, no matter what set we are adding A' to. It is always $\Pr[t_k = A'] \pi_k(A')$. And the change to the right-hand side is exactly $\Pr[t_k = A']$ times the probability that none of the types in $\bigcup_{j \neq k} S'_j$ show up. As we add more types to S , the probability that none of the types in $\bigcup_{j \neq k} S'_j$ show up will never increase. So for any constricting set S' containing S , we can add A' to S'_k and still get a constricting set.

So starting from a constricting set S and a type $A \in T_i$ as above we can add every $B \in T_j$ with $\hat{\pi}_j(B) \geq \hat{\pi}_i(A)$ to S in order to obtain a constricting set of the form $S_x = \{B \mid B \in T_j \wedge \hat{\pi}_j(B) \geq x\}$, where $x = \hat{\pi}_i(A)$. So every infeasible reduced form has a constricting set of this form. Taking the contrapositive proves the theorem. \square

We say that a hierarchical mechanism H is *virtually-ordered* if $\hat{\pi}_i(A) \geq \hat{\pi}_j(A') \Rightarrow H(A, i) \leq H(A', j)$. While the virtual π s give us a nice structural theorem about feasible reduced forms and a

⁹For a minimal set S , there could be many possible choices of x_1, \dots, x_m . We simply use any of them.

linear time separation oracle for determining feasibility (see proof of Theorem 6), the following observation shows that distributions over virtually-ordered hierarchical mechanisms are not sufficient to implement every feasible reduced form when the bidders are non-i.i.d.

Observation 2. *There exist feasible reduced forms that are not implementable as distributions over virtually-ordered hierarchical mechanisms.*

Proof. Consider the following example with two bidders. Bidder one has a single type, A . Bidder two has two types, B and C and is each with probability $1/2$. Then $\pi_1(A) = 1/3$, $\pi_2(B) = 2/3 + \epsilon$, $\pi_2(C) = 2/3 - \epsilon$ is a feasible reduced form. However, $\hat{\pi}_1(A) > \hat{\pi}_2(C)$, so no distribution over virtually-ordered hierarchical mechanisms can possibly have $\pi_2(C) > 1/2$. \square

Now that we have motivated Theorems 5 and 6, we proceed to prove them, after providing the key steps as technical lemmas.

Lemma 2. *Every feasible reduced form π for independent bidders and a single item can be implemented (not necessarily exactly implemented) as a distribution over strict, partially-ordered w.r.t. π hierarchical mechanisms.*

Proof. The proof is almost identical to that of Lemma 1. Here are the main differences: We do induction on $\sum_i d_i$, where d_i is the number of distinct non-zero values in $\{\pi_i(A)\}_{A \in T_i}$. For the inductive step, X is now taken to be the set of distributions over strict, partially-ordered hierarchical mechanisms, and it is still compact, viewed as a subset of the Euclidean space. The function $F : X \rightarrow \mathbb{R}$ is now defined as

$$F(M) = \max_{i, A \in T_i} \{\pi_i(A) - M_i(A)\}.$$

Again, if we use the Euclidean distance on X , as a subset of the Euclidean space, F is continuous. Since F is continuous and X is compact, F achieves its minimum inside X . Let M be a minimizer. For all i , we define $S_i = \{A \in T_i \mid \pi_i(A) - M_i(A) = F(M)\}$. In terms of the sets $\{S_i\}_i$ we can define an alternative reduced form π' as follows. For all i , $A \in T_i$: if $A \in S_i$, then set $\pi'_i(A) = \pi_i(A)$; otherwise, set $\pi'_i(A) = \max_{B \in S_i \mid \pi_i(A) \geq \pi_i(B)} \{\pi(B)\}$, unless $\{B \in S_i \mid \pi_i(A) \geq \pi_i(B)\}$ is empty in which case set $\pi'_i(A)$ equal to 0. With these changes, the proof is truly identical to that of Lemma 1, and we avoid repeating the steps for brevity. \square

Corollary 2. *Every feasible reduced form π for independent bidders and a single item can be exactly implemented as a distribution over strict, partially-ordered w.r.t. π hierarchical mechanisms.*

Proof. The proof is identical to that of Corollary 1 after making the same modifications going from Lemma 1 to Lemma 2. \square

We proceed to prove an analog of Proposition 1 in this setting. We need a definition. For all i , let $>^i := A_{i,1} >^i A_{i,2} >^i \dots >^i A_{i,k_i}$ be a total ordering of the set $T_i := \{A_{i,1}, \dots, A_{i,k_i}\}$ of bidder i 's types. We say that a reduced form π respects $>^i$ iff $\pi_i(A_{i,1}) \geq \pi_i(A_{i,2}) \geq \dots \geq \pi_i(A_{i,k_i})$. We also say that an hierarchical mechanism H is partially-ordered with respect to $>^i$ iff $H(A_{i,1}, i) \leq H(A_{i,2}, i) \leq \dots \leq H(A_{i,k_i}, i)$.¹⁰ With respect to these definitions we show the following proposition.

¹⁰Notice that this definition of partially-ordered hierarchical mechanism (with respect to $\{>^i\}_i$) is similar to its counterpart in the main body (with respect to π), but different. Being partially-ordered with respect to π certainly imposes the same constraints as being partially-ordered with respect to any $\{>^i\}_i$ that π respects. The difference is that being partially-ordered with respect to π may also impose some equality constraints, if $\pi_i(A) = \pi_i(B)$ for types $A \neq B$.

Proposition 3. For all i , let $>^i := A_{i,1} >^i A_{i,2} >^i \dots >^i A_{i,k_i}$ be a total ordering of the set $T_i := \{A_{i,1}, \dots, A_{i,k_i}\}$ of bidder i 's types. The set of feasible reduced forms that respect $>^1, \dots, >^m$ is a $(\sum_i k_i)$ -dimensional polytope whose corners are exactly the strict, partially-ordered w.r.t. $>^1, \dots, >^m$ hierarchical mechanisms.

Proof. We know from [6, 12] that a reduced form π respects $>^1, \dots, >^m$ and is feasible iff

$$\pi_i(A_{i,j}) \geq \pi_i(A_{i,j+1}) \quad \forall i \in [m], j \in [k_i] \quad (6)$$

$$\sum_i \sum_{j \leq x_i} \Pr[t_i = A_{i,j}] \pi_i(A_{i,j}) \leq 1 - \prod_i \left(1 - \sum_{j \leq x_i} \Pr[t_i = A_{i,j}] \right) \quad \forall x_1 \in [k_1], \dots, x_m \in [k_m] \quad (7)$$

where for notational convenience we denote $\pi_i(A_{i,k_i+1}) = 0$. In fact, to make our lives easier in a later proof, we will actually replace (7) with:

$$\sum_i \sum_{j \leq x_i} \Pr[t_i = A_{i,j}] \pi_i(A_{i,j}) \leq 1 - \prod_i \left(1 - \sum_{j \leq x_i} \Pr[t_i = A_{i,j}] \right) \quad \forall x_1 \in [k_1 - 1], \dots, x_m \in [k_m - 1] \quad (8)$$

$$\sum_i \sum_{j \leq k_i} \Pr[t_i = A_{i,j}] \pi_i(A_{i,j}) \leq 1 \quad (9)$$

In the above above replacement, we are basically observing that if (9) holds, then so does (7) for any case where at least one i has $x_i = k_i$. In addition, (8) covers all other cases. We have hence shown that the set of feasible reduced forms that respect $>^1, \dots, >^m$ is a $\sum_i k_i$ -dimensional polytope.

We proceed to show that any strict, partially-ordered w.r.t. $>^1, \dots, >^m$ hierarchical mechanism H whose reduced-form is π is a corner. For convenience in the proof assume that $T_i \cap T_k = \emptyset$, for all $i \neq k$. This is easy to achieve by having every type of bidder i carry the name of bidder i , for all i . Let now $y = \min_x \Pr[\exists i, H(t_i, i) \leq x] = 1$, i.e. the minimum y so that with probability 1 some bidder i will report a type $A_{i,j}$ such that $H(A_{i,j}, i) \leq y$. In terms of this y we define y^* as follows: $y^* := y$, if $y \neq \text{LOSE}$ (case 1), and $y^* := y - 1$, if $y = \text{LOSE}$ (case 2). We observe then that a type $A_{i,j} \in T_i$ with $H(A_{i,j}, i) > y^*$ cannot possibly win the item, as we are either guaranteed to see some bidder whose type lies on a higher level of the hierarchy (case 1) or the type is mapped to LOSE and is hence given no item (case 2). For all such $i, A_{i,j} \in T_i$, we therefore have $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1}) = \dots = \pi_i(A_{i,k_i+1}) = 0$.

We say that a set of types $S = \cup_i \{A_{i,1}, \dots, A_{i,x_i}\}$ for some x_1, \dots, x_m is *near-constricting*, if the corresponding Inequality (7) is tight for x_1, \dots, x_m . Then, for any $i, A_{i,j} \in T_i$ with $H(A_{i,j}, i) \leq y^*$, we know that $H^{-1}([H(A_{i,j}, i)])$ is a near-constricting set, because the item is always awarded to a type in $H^{-1}([H(A_{i,j}, i)])$ whenever at least one type from this set is reported. Moreover, if $H(A_{i,j_1}, i) = \dots = H(A_{i,j_\ell}, i)$, for some types $A_{i,j_1}, \dots, A_{i,j_\ell} \in T_i$, then $\pi_i(A_{i,j_1}) = \dots = \pi_i(A_{i,j_\ell})$. Finally, because H is strict, if some $i, A_{i,j} \in T_i$ satisfy $H(A_{i,j}, i) \leq y^*$, then $H^{-1}(H(A_{i,j}, i)) \cap T_k = \emptyset$, for all $k \neq i$.

Let us now define the following mapping from types to tight inequalities:

- If a type $A_{i,j} \in T_i$ satisfies $H(A_{i,j}, i) > y^*$, then we map $A_{i,j}$ to the constraint $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1})$, i.e. the tightness of inequality $\pi_i(A_{i,j}) \geq \pi_i(A_{i,j+1})$.
- If a type $A_{i,j} \in T_i$ satisfies $H(A_{i,j}, i) \leq y^*$, then:
 - if $H(A_{i,j}, i) = H(A_{i,j+1}, i)$, we map $A_{i,j}$ to the constraint $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1})$, i.e. the tightness of inequality $\pi_i(A_{i,j}) \geq \pi_i(A_{i,j+1})$;

– otherwise, we map $A_{i,j}$ to the tightness of Inequality (7) for the set of types $H^{-1}([H(A_{i,j}, i)])$.

The above discussion ensures that our mapping is injective. Hence π makes at least $\sum_i k_i$ of the inequalities defining our polytope tight. And it is not hard to see that there is a unique feasible reduced form making these inequalities tight. So π is a corner of the polytope. Thus, every strict, partially-ordered w.r.t. $>^1, \dots, >^m$ hierarchical mechanism is a corner of the polytope.

We now make the same observation as in Proposition 1 to argue that there are no other corners. Corollary 2 implies that every point in the polytope can be written as a convex combination of strict, partially-ordered w.r.t. $>^1, \dots, >^m$ hierarchical mechanisms, all of which are corners of the polytope. As no corner of the polytope can be written as a convex combination of *other* corners of the polytope, there must not be any other corners. \square

And now we are ready to prove Theorems 5 and 6.

Proof of Theorem 5: Using a similar argument as in the proof of Theorem 3 we can assume without loss of generality that, for all i , π assigns a distinct probability to every type in T_i . (Otherwise we can define a similar TYPEMERGE operation, like the one defined in the proof of Theorem 3, whereby types in T_i that receive the same π value are merged into super-types.) Under this assumption, Proposition 3 implies that π lies inside a $(\sum_i |T_i|)$ -dimensional polytope, P , whose corners are the strict, partially-ordered w.r.t. π hierarchical mechanisms. By Carathéodory's Theorem, every point in the polytope can be written as a convex combination of at most $(\sum_i |T_i|) + 1$ corners. As a convex combination of corners is exactly a distribution over strict, partially-ordered w.r.t. π hierarchical mechanisms, this proves the theorem. \square

Proof of Theorem 6: The first part of the theorem follows immediately as a corollary of Theorem 4.

We now need to describe how to efficiently find a mechanism implementing a reduced form $\tilde{\pi}$ that is feasible. Using a similar argument as in the proof of Theorem 2 we can assume without loss of generality that, for all i , $\tilde{\pi}$ assigns a distinct probability to every type in T_i . (Otherwise we can merge types in T_i that receive the same $\tilde{\pi}$ value into super-types, like we did in the proof of Theorem 2, then run the procedure outlined below, and finally un-merge types.) Under this assumption, Proposition 3 implies that $\tilde{\pi}$ lies inside a $(\sum_i |T_i|)$ -dimensional polytope, P , whose corners are the strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanisms. Therefore we can directly apply Theorem 13 of Appendix A to write $\tilde{\pi}$ as a convex combination of such hierarchical mechanisms, as long as we can describe the boundary oracle BO , corner oracle CO and separation oracle SO that are needed for Theorem 13. BO is trivial to implement, as we just have to include in the set of halfspaces defining the boundary of P those inequalities described in the proof of Proposition 3. In particular, for convenience in the remaining of the proof, we include the inequalities of the form (6), (8) and (9). For CO , on input B , we first check that every hyperplane $h \in B$ satisfies $BO(h) = \text{yes}$. If not, output no. Otherwise, we need to check if $\bigcap_{h \in B} h$ contains a corner of P . We know that the corners of P are exactly the strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanisms. To check if there is such a corner in $\bigcap_{h \in B} h$ we need to do some work.

First, let us use the same notation as in Proposition 3, denoting $T_i = \{A_{i,1}, \dots, A_{i,k_i}\}$, where $k_i = |T_i|$ and $\tilde{\pi}_i(A_{i,1}) > \dots > \tilde{\pi}_i(A_{i,k_i})$. Also, let us call a set of types S near-constricting if either $S = \cup_i \{A_{i,1}, \dots, A_{i,x_i}\}$ for some $x_1 \leq k_1 - 1, \dots, x_m \leq k_m - 1$ and Inequality (8) is tight for x_1, \dots, x_m , or if $S = \cup_i T_i$ and Inequality (9) is tight.

Now given a set B of hyperplanes, if B contains a near-constricting set hyperplane for the sets of types S_1, \dots, S_k , we check first whether these sets satisfy $S_1 \subset S_2 \dots \subset S_k$ (possibly after renaming). If not, then there are some two near-constricting sets S_i, S_j with $A \in S_i - S_j, B \in S_j - S_i$ for some types $A \neq B$. Because S_i and S_j are different than $\cup_i T_i$ and they are near-constricting they must be of the form making Inequality (8) tight. Hence, both S_i and S_j miss at least one

type of every bidder, so that the right-hand side of the inequality for S_i must be < 1 and similarly the right-hand side of the inequality for S_j must be < 1 . In addition, we cannot have A and B be types of the same bidder, as we only consider near-constricting sets that respect the partial ordering within every bidder's type-set. Therefore, A and B belong to different bidders, and because the probability of seeing a type in S_i is < 1 (and the same holds for S_j), there is a non-zero probability that A and B are both present, but no other type of S_i or S_j is present. Then the near-constricting set equation for set S_i requires that we must give the item to A ,¹¹ and the near-constricting set equation for S_j requires that we must give the item to B , which is impossible. So if we do not have $S_1 \subset S_2 \dots \subset S_k$, the hyperplanes do not intersect in a feasible mechanism, and therefore CO should output no.

Otherwise, CO looks at the other hyperplanes (of the form $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1})$) that belong to B , and chooses an arbitrary strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanism that satisfies the following constraints:

1. $H(A, i) < H(A', i')$ for all $(A, i) \in S_j, (A', i') \in S_{j+1} - S_j$, for all j .
2. $H(A, i) < H(A', i')$ for all $(A, i) \in S_k, (A', i') \notin S_k$.
3. For all $i, A_{i,j} \in T_i$:
 - (a) if the hyperplanes $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1}), \pi_i(A_{i,j+1}) = \pi_i(A_{i,j+2}), \dots, \pi_i(A_{i,k_i}) = 0 (= \pi_i(A_{i,k_i+1}))$ are all in B , then $H(A_{i,j}, i) = \text{LOSE}$;
 - (b) otherwise, if $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1})$ is in B , then either $H(A_{i,j}) = H(A_{i,j+1}) = \text{LOSE}$ or $H(A_{i,j}) \geq H(A_{i,j+1}) - 1$.

We claim that an H satisfying the above constraints exists if and only if $\bigcap_{h \in B} h \cap P \neq \emptyset$. By Proposition 3, we know that if there is a corner π in $\bigcap_{h \in B} h \cap P$, there is a strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanism H that implements it. Without loss of generality we can make two simplifying assumptions about H : (i) For all $t \neq \text{LOSE}$, $|H^{-1}(t)| \leq 1$. Indeed, suppose that $|H^{-1}(t)| > 1$ for some $t \neq \text{LOSE}$. Then because H is strict, location t of the hierarchy defined by H only contains types belonging to the same bidder i . And because H is partially ordered w.r.t. $\tilde{\pi}$ these types are consecutive in T_i . So we can change H to “expand” cell t of the hierarchy into consecutive cells containing a single type each in the right order. This has no effect in the mechanism defined by H . (ii) If H awards the item to bidder i of type $A_{i,j}$ with probability 0, we can assume that $H(A_{i,j}, i) = \text{LOSE}$. Now given (i) and the nature of S_1, \dots, S_k ,¹² for H to have the sets S_1, \dots, S_k near-constricting it must satisfy the first two constraints above. Indeed for these constraints not to be satisfied there must exist $\ell, i, A_{i,j}, i', A_{i',j'}, i \neq i'$ such that $H(A_{i,j}, i) < H(A_{i',j'}, i')$ and $A_{i,j} \notin S_\ell$ while $A_{i',j'} \in S_\ell$. But then it must be $\pi(A_{i,j}) > 0$ (because of assumption (ii)), so there is a positive probability that bidder i of type $A_{i,j}$ is the highest in the hierarchy when every bidder except for i' has declared a type. Still, given that $H(A_{i,j}, i) < H(A_{i',j'}, i')$, even if i' declares $A_{i',j'}$, i will get the item, contradicting that S_ℓ is near-constricting. We argue next that H also needs to satisfy the third constraint above. That constraint 3(a) is satisfied follows immediately from assumption (ii). We argue next that 3(b) needs to be satisfied. Indeed, suppose there exist $A_{i,j}, A_{i,j+1} \in T_i$ and $A_{i',j'} \in T_{i'}$ such that $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1}) > 0$ and $H(A_{i,j}, i) < H(A_{i',j'}, i') < H(A_{i,j+1}, i)$. As $\pi_i(A_{i,j+1}) > 0$, it follows that $H(A_{i,j+1}, i) < \text{LOSE}$, so $H(A_{i',j'}, i') < \text{LOSE}$, which implies that $\pi_{i'}(A_{i',j'}) > 0$ (otherwise it would be that $H(A_{i',j'}, i') = \text{LOSE}$ given (ii)). Now, because bidder i' wins the item with non-zero

¹¹Recall that a way to interpret a near-constricting set equation for set S is the following: whenever at least one bidder reports a type in S to the mechanism, the mechanism gives the item to a bidder who reported a type in S , with probability 1.

¹²in particular, the fact that the sets S_1, \dots, S_k respect the ordering of the type sets as follows: For all $\ell, A_{i,j} \in S_\ell$ implies $A_{i,j'} \in S_\ell$, for all $j' \leq j$.

probability as type $A_{i',j'}$, there is a non-zero probability that $A_{i',j'}$ is on the highest level of the hierarchy after sampling from all bidders except for i . In this case, i will win the item by reporting $A_{i,j}$, and lose by reporting $A_{i,j+1}$. In all other cases, i is at least as likely to win the item by reporting $A_{i,j}$ as $A_{i,j+1}$, and therefore we see that bidder i gets the item strictly more often when reporting $A_{i,j}$ than $A_{i,j+1}$, violating the constraint $\pi_i(A_{i,j}) = \pi_i(A_{i,j+1})$. So if there is a corner π in $\bigcap_{h \in B} h \cap P$ it can be implemented by a strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanism H satisfying the above constraints. The other direction of our claim is even simpler. If a strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanism H satisfies all constraints above, then its induced reduced-form π will immediately satisfy all equalities in B .

Hence to implement CO one just needs to check if there is a strict, partially-ordered w.r.t. $\tilde{\pi}$ hierarchical mechanism H satisfying the four constraints above. This task is easy to do efficiently. If a mechanism is found it satisfies all equalities in B and it is a corner of the polytope by Proposition 3.

SO is also simple to implement. On input $\vec{\pi}$, we first check that Inequalities (6) are satisfied (i.e. that $\vec{\pi}$ respects the total orderings on the bidders' type-spaces induced by $\tilde{\pi}$). Then we use the separation oracle provided by Theorem 4 to verify that $\vec{\pi}$ is feasible. As all three oracles BO, CO, SO run in polynomial time, we can apply Theorem 13 to write $\vec{\pi}$ as a convex combination of at most $\sum_i |T_i| + 1$ corners, which is exactly a distribution over at most $\sum_i |T_i| + 1$ strict, partially-ordered hierarchical mechanisms in polynomial time. \square

D Implementation of General Reduced Forms via Hierarchical Mechanisms

Here we provide the proof of Theorem 7.

Proof of Theorem 7: Let σ be a total ordering over all possible types, $\sigma : \cup_i (T_i \times \{i\}) \rightarrow [\sum_i |T_i|]$. Define the unhappiness $F_\sigma(M)$ of a distribution over σ -ordered hierarchical mechanisms, M , as follows:

$$F_\sigma(M) = \max_{i, A \in T_i} (\pi_i(A) - M_i(A)).$$

As we argued formally in earlier sections F_σ can be viewed as a continuous function over a compact set. Hence it achieves its minimum inside the set. Let then $M^\sigma \in \operatorname{argmin}_M F_\sigma(M)$ (where the minimization is over all distributions over σ -ordered hierarchical mechanisms) and define the set S_σ to be the set of maximally unhappy types under M^σ ; formally, $S_\sigma = \operatorname{argmax}_{i,A} \{\pi_i(A) - M_i^\sigma(A)\}$. If for some σ there are several minimizers M^σ , choose one that minimizes $|S_\sigma|$. Now, let MO be the set of the orderings σ that minimize $F_\sigma(M^\sigma)$. Further refine MO to only contain σ 's minimizing $|S_\sigma|$. Formally, we first set $MO = \operatorname{argmin}_\sigma \{F_\sigma(M^\sigma)\}$ and then refine MO as $MO_{\text{new}} = \operatorname{argmin}_{\sigma \in MO} \{|S_\sigma|\}$. We drop the subscript “new” for the rest of the proof.

From now on, we call a type (A, i) *happy* if $M_i(A) \geq \pi_i(A)$, otherwise we call (A, i) *unhappy*. Intuitively, here is what we have already done: For every ordering σ , we have found a distribution over σ -ordered hierarchical mechanisms M^σ that minimizes the maximal unhappiness and subject to this, the number of maximally unhappy types. We then choose from these (σ, M^σ) 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 σ , such that $F_\sigma(M^\sigma) \leq 0$, and it is natural to start with the ordering that is “closest” to satisfying this property. We are one step away from completing the proof. What we will show next is that, if $\tau \in MO$ does not make every type happy, then we can find some other ordering τ' , such that $F_{\tau'}(M^{\tau'}) = F_\tau(M^\tau)$, $|S_{\tau'}| = |S_\tau|$, and $S_{\tau'} = \{\tau^{-1}(1), \dots, \tau^{-1}(|S_{\tau'}|)\}$. In other words, only the top $|S_{\tau'}|$ types in τ are maximally unhappy. From here, we will show that because $\tau' \in MO$, that $S_{\tau'}$

is a constricting set and get a contradiction.

First, if the maximally unhappy types in S_τ are not the top $|S_\tau|$ ones, let i be the smallest i such that $\tau^{-1}(i+1) \in S_\tau$ but $\tau^{-1}(i) \notin S_\tau$. We proceed to show that by changing either the distribution M or the ordering τ , we can always move $\tau^{-1}(i)$ into S_τ and $\tau^{-1}(i+1)$ out without changing $|S_\tau|$ or the value $F_\tau(M)$. Then by repeating this procedure iteratively, we will get the τ' we want.

Before we describe the procedure, we introduce some terminology. We say there is a *cut* between $\tau^{-1}(i)$ and $\tau^{-1}(i+1)$ in a fixed τ -ordered hierarchical mechanism H if $H(\tau^{-1}(i)) < H(\tau^{-1}(i+1))$, i.e. if $\tau^{-1}(i)$ and $\tau^{-1}(i+1)$ are on different levels of the hierarchy. For the remainder of the proof, we will let l be the level of $\tau^{-1}(i)$ ($H(\tau^{-1}(i))$). When we talk about adding or removing a cut below i , we mean increasing or decreasing $H(\tau^{-1}(j))$ by 1 for all $j > i$. We now proceed with a case analysis, for fixed $\tau^{-1}(i) \notin S_\tau$, $\tau^{-1}(i+1) \in S_\tau$. We let $(A, j) = \tau^{-1}(i)$ and $(B, k) = \tau^{-1}(i+1)$.

• **Case 1:** $j = k$.

Since τ is a linear extension of the bidder's own ordering, then $\pi_j(A) \geq \pi_j(B)$, but we know that

$$\pi_j(A) - M_j^\tau(A) < \pi_j(B) - M_j^\tau(B),$$

thus $M_j^\tau(A) > M_j^\tau(B) \geq 0$. Because A and B are types for the same bidder j , when A and B are in the same level, they get the item with equal probability. Therefore, there must exist some $H \in \text{supp}(M^\tau)$ with a cut below A , and in which A gets the item with non-zero probability. We modify M^τ by modifying the mechanisms H in its support as follows.

Let H be a hierarchical mechanism in the support of M^τ . If there is no cut below A , we do nothing. If all of the types on level l and level $l+1$ are from bidder j , we remove the cut below A . This does not affect $H_q(C)$ (the probability that (C, q) gets the item under H) for any $q, C \in T_q$, because it was impossible for two types in the combined level to show up together anyway. As we have not changed $H_q(C)$ for any q, C in the mechanisms we have touched so far, yet none of these mechanisms has a cut between levels l and $l+1$, there must still be some $H \in \text{supp}(M^\tau)$ with a cut below A and in which A gets the item with non-zero probability (otherwise it couldn't be that $M_j^\tau(A) > M_j^\tau(B) \geq 0$). For such an H , there is at least one type not from bidder j in level l or $l+1$. We distinguish two sub-cases:

- Every bidder has at least one type in level $l+1$ or larger (in other words, every type in level $l+1$ wins the item with non-zero probability). Consider now moving the cut from below i to below $i-1$. Clearly, A will be less happy if we do this. Every type not from bidder j in l will be strictly happier, as now they do not have to share the item with A . Every type not from bidder j in $l+1$ will be strictly happier, as they now get to share the item with A . It is also not hard to see that all types $\neq A$ from bidder j in level l and $l+1$ are not affected by this change, as they never share the item with A in either case. So in particular B is unaffected. Consider instead moving the cut from below i to below $i+1$. Then B is happier, every type not from bidder j in $l+1$ is less happy than before (as they now don't get to share with B), every type not from bidder j in l is also less happy than before (because now they have to share with B), and all types $\neq B$ from bidder j in level l and $l+1$ are not affected by the change (as they never share the item with B in either case). To summarize, we have argued that, when we move the cut to below $i+1$, B becomes strictly happier, and every type that becomes less happy by this change becomes strictly happier if we move the cut to below $i-1$ instead. Also, B is unaffected by moving the cut to $i-1$. So with a tiny probability ϵ , move the cut from below i to below $i-1$, whenever H is sampled from M^τ . This makes all of the types not from bidder j in level l or $l+1$ strictly happier. With a tinier probability δ ,

move the cut from below i to below $i + 1$, whenever H is sampled from M^τ . Choose ϵ to be small enough that we don't make A maximally unhappy, and choose δ to be small enough so that we don't make any types besides A less happy than they were in H . Then we have strictly increased the happiness of B without making A maximally unhappy, or decreasing the happiness of any other types. Therefore, we have reduced $|S_\tau|$, a contradiction.

- If there is a bidder j' whose types are all in levels $1, \dots, l$ (call such bidders *high*), then no type in level $l + 1$ can possibly win the item. We also know that: every high bidder has at least one type in level l by our choice of H (otherwise A would get the item with probability 0); and all high bidders are different than j , since B is in level $l + 1$. Now we can basically use the same argument as above. The only difference is that when we move the cut to below $i - 1$ or the cut to below $i + 1$, types in level $l + 1$ that are different than B will remain unaffected (i.e. the affected types different from B are only those in level l). But since every high bidder has a type in level l , B will be unaffected in the first case but strictly happier in the second, and it is still the case that every type who is made unhappier by moving the cut to below $i + 1$ is made strictly happier by moving the cut to below $i - 1$. So we can carry over the same proof as above, and get a contradiction.

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

- **Case 2:** $j \neq k$ and there is never a cut below A .

This case is easy. If we switch (A, j) and (B, k) in τ , then the set S_τ is exactly the same, and the distribution M^τ is exactly the same. However, we have now relabeled the types in S_τ so that $\tau^{-1}(i) \in S_\tau$ and $\tau^{-1}(i + 1) \notin S_\tau$.

- **Case 3:** $j \neq k$ and there is sometimes a cut below A .

Pick a mechanism H in the support of M^τ that has a cut between A and B and in which A gets the item with positive probability. (If such a mechanism doesn't exist we can remove the cut between i and $i + 1$ in all mechanisms in the support without changing the allocation probabilities and return to Case 2). Let now $i^* = \max_{i' < i} \{i' \mid \tau^{-1}(i') \in S_\tau\}$. By our choice of i (specifically, that it is the smallest i such that $\tau^{-1}(i + 1) \in S_\tau$ but $\tau^{-1}(i) \notin S_\tau$), we see that $\tau^{-1}(i') \in S_\tau$ for all $i' \leq i^*$, and $\tau^{-1}(i') \notin S_\tau$ for all $i^* < i' \leq i$. There are again two sub-cases:

- $H(\tau^{-1}(i^*)) < l$. By our choice of i^* , this means that everyone in level l is *not* maximally unhappy. By our choice of H , everyone in level l receives the item with non-zero probability, so there is at least one type from each bidder in level l or larger. If we pick a tiny ϵ , and with probability ϵ remove the cut from below i (whenever H is sampled from M^τ), then everyone in level $l + 1$ is happier, everyone in level l is unhappier, and everyone else is unaffected. In particular, B will be strictly happier with this change, as he now gets to share with A (and possibly others). If we choose a sufficiently small ϵ , no one in level l will be made maximally unhappy, and (B, k) will be removed from S_τ , a contradiction.
- $H(\tau^{-1}(i^*)) = l$. In this case, introduce a cut below i^* with some probability ϵ whenever H is sampled from M^τ . The only types who may become happier by this change are those in level l with $\tau(C, q) \leq i^*$. The only types who may become unhappier by this change are those in level l with $\tau(C, q) > i^*$. Everyone else is unaffected by this change. But, if we can make any type happier, then we can choose ϵ small enough, so that we remove this type from S_τ (this type must be in S_τ as all types in level l with $\tau(C, q) \leq i^*$ are) without making any new type maximally unhappy (as all types that can possibly become unhappier with this change are *not* in S_τ). Again, we obtain a contradiction.

because this would decrease $|S_\tau|$ without increasing $F_\tau(M^\tau)$. Thus, this change cannot make anyone happier, and therefore cannot make anyone unhappier. So we may modify M^σ by introducing a cut below i^* with probability 1 whenever M^τ samples H , thereby removing H from the support of M^τ (without making anyone happier or unhappier) and replacing it with H' satisfying: $H'(\tau^{-1}(i^*)) < H'(\tau^{-1}(i)) < H'(\tau^{-1}(i+1))$ and H' awards the item to $\tau^{-1}(i)$ with non-zero probability. After this modification, we may return to the previous sub-case to obtain a contradiction.

Hence, it can not be the case that $j \neq k$ with sometimes a cut below A .

At the end of all three cases, we see that if we ever have $\tau^{-1}(i) \notin S_\tau$ and $\tau^{-1}(i+1) \in S_\tau$, then these types must belong to different bidders, and no mechanism in the support of M^τ ever places a cut between these types. Hence, we can simply swap these types in τ (as we described in Case 2 above), and we do that repeatedly until we have $S_\tau = \{\tau^{-1}(1), \dots, \tau^{-1}(|S_\tau|)\}$. Once such a τ has been found, let $k = |S_\tau|$. Now consider a mechanism in the support of M^τ that has no cut below k , and consider putting a cut there with some tiny probability ϵ whenever this mechanism is sampled. The only effect this *might* have is that when the item went to a type outside S_τ , it now goes with some probability to a type inside S_τ . Therefore, if anyone gets happier, it is someone in S_τ . However, if we make anyone in S_τ happier and choose ϵ small enough so that we don't make anyone outside of S_τ maximally unhappy, we decrease $|S_\tau|$, getting a contradiction. Therefore, putting a cut below 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 k in all mechanisms in the support of M^τ . But now we get that the item always goes to someone in S_τ whenever someone in S_τ shows up, yet everyone in this set is unhappy. Therefore, S_τ is a constricting set, certifying that the given π is infeasible.

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

E Optimal multi-dimensional mechanisms

We sketch how to use our results from the previous sections and recent results of [14] to obtain the proofs of Theorems 8, 9, 10 and 11. Simply put, [14] provides a LP formulation for the problem of finding the optimal BIC mechanism for value distributions of finite support. In addition, [14] exploits any symmetry that may exist in the value distribution to reduce the size of the resulting LPs. However, without a separation oracle to determine the feasibility of the π s, these LP formulations will have size exponential in the number of bidders in the settings we consider, even when the bidders are i.i.d., because the LPs need to maintain variables for the ex-post allocation rule described by the ϕ 's (See Section 2). Using the separation oracles of Theorems 2 and 6, we can eliminate the need for ϕ 's in these LPs, turning the number of variables to polynomial in the number of bidders, and solving them in time polynomial in m , n and $\max_i |supp(\mathcal{F}_i)|$, thereby establishing Theorem 8. Theorem 9 follows immediately by combining our separation oracles with the symmetric LPs of Section 5 of [14]. Theorem 10 follows by discretizing the value distributions, applying the aforementioned symmetric LPs and the ϵ -BIC to BIC reduction of Section 6 of [14]. Finally, Theorem 11 is obtained by truncating and discretizing the value distributions, running the symmetric LP on the resulting distributions, and using the ϵ -BIC to BIC reduction of Section 6 of [14]. (This procedure is described in detail in Appendix J of [14].)

For completeness, in Figure 3 we describe how to modify the LP of [14] to exploit our separation oracles. We only describe the LP for the item-symmetric case to illustrate some extra work that

is needed to handle symmetries. It is clear how to “un-symmetrize” the LP to treat the non-symmetric problem. The LP shown in Figure 3 is identical to the item-symmetric LP of [14] except that we have replaced the naïve feasibility constraints for the reduced form used in [14] (these were explicitly finding consistent $\vec{\phi}(P)$ ’s for every profile) with a separation oracle SO that takes as input a reduced form and decides whether it is feasible or not, in which case it also provides a separation hyperplane from the feasible set. The reader should refer to [14] for the correctness of the unmodified LP. So we only need to justify here that our separation oracle works properly. While this is obvious in the non-symmetric case (as it follows directly from Theorem 6), it is slightly trickier for the symmetric LP of Figure 3, as the LP does not explicitly maintain the whole reduced form. In particular, for every bidder i , we only maintain the reduced form for all types of bidder i in the set $E_i = \{\vec{v}_i \mid v_{i1} \geq \dots \geq v_{in}\}$, and it is assumed that (because of symmetries) the reduced form of a type $\vec{w}_i \notin E_i$ is just $\{\pi_{ij}(\vec{w}_i) = \pi_{i\sigma^{-1}(j)}(\vec{v}_i)\}_j$ for whatever $\vec{v}_i \in E_i$ and $\sigma \in S$ satisfy $\sigma(\vec{v}_i) = \vec{w}_i$, where S is the set of all item permutations. (For why it is OK to assume that such symmetries are satisfied by the optimal reduced form, the reader is referred to [14].)

We proceed to describe how our separation oracle works in the item-symmetric case. The input to SO is the succinct description $\{\pi_{ij}(\vec{v}_i)\}_{i,j,\vec{v}_i \in E_i}$ of a reduced form and the value distribution \mathcal{D} . If we define $\Pi := \cup_{i,j,\vec{v}_i \in E_i} \{\hat{\pi}_{ij}(\vec{v}_i)\}$, where the $\hat{\pi}$ ’s are as in Definition 3, it follows from Theorem 4 and the symmetries satisfied by the complete reduced form (which is implied by the succinct one provided in the input to SO) that, for each item j , we only need to check $|\Pi|$ constricting sets for feasibility. (In fact, because the mechanism is item-symmetric, it suffices to only check these constraints for a single item.) Checking a particular set can be done in time polynomial in $|\Pi|$, i.e. without explicitly listing all types that belong to the set, due to the symmetries. If none of the sets is found constricting, output “yes.” Otherwise, output the linear constraint violated by the constricting set. This is a hyperplane certifying that the $\pi_{ij}(\vec{v}_i)$ ’s are infeasible. If SO outputs “yes,” this means that it is possible to allocate each item individually and satisfy the $\pi_{ij}(\vec{v}_i)$ ’s. Because bidders have no demand constraints, this also means that we can simultaneously satisfy the $\pi_{ij}(\vec{v}_i)$ ’s of all items by simply allocating each item independently. It is now clear that SO acts correctly as a separation oracle determining whether or not the $\pi_{ij}(\vec{v}_i)$ ’s are feasible. Moreover, notice that $|E_i| = \text{poly}(n^c)$, for all i , where $c = |\text{supp}(\mathcal{F}_{ij})|$. Hence, $|\Pi|$ is polynomial in m and n^c . So our SO runs in time polynomial in m and n^c .

Our LP has size polynomial in m , n^c and our separation oracle runs in time polynomial in m , n^c . So we can solve the LP to compute the optimal reduced form, $\{\pi_{ij}(\vec{v}_i)\}_{i,j,\vec{v}_i \in E_i}$. To execute the mechanism corresponding to the computed reduced form, we use Theorem 6, treating each item separately and allocating each item independently of the others.

F Generalization of Border’s Theorem: Non-Negative Linear Feasibility Constraints

We provide details from Section 3.4 and prove Theorem 12. Throughout the section, we let P denote a type-profile, $\Pr[P]$ the probability of seeing profile P , and P_i the type of bidder i in P .

Proof of Theorem 12: It is clear that this is a necessary condition for a reduced form to be feasible. Indeed, if M implements the reduced form, then we have:

$$\pi_{ij}(A) \Pr[t_i = A] \leq M_{ij}(A) \Pr[t_i = A] \quad \forall i, j, A \in T_i;$$

Variables:

- $\pi_{ij}(\vec{v}_i)$, for all bidders i , items j , and $\vec{v}_i \in E_i$ ($mn|E_i|$), the ex-interim probability that bidder i gets item j when reporting type \vec{v}_i .
- $q_i(\vec{v}_i)$, for all bidders i , $\vec{v}_i \in E_i$ ($m|E_i|$), the ex-interim expected price that bidder i pays when reporting type \vec{v}_i .

Constraints:

- $0 \leq \pi_{ij}(\vec{v}_i) \leq 1$, for all $i, j, \vec{v}_i \in E_i$, guaranteeing that each $\pi_{ij}(\vec{v}_i)$ is a probability ($mn|E_i|$).
- $\sum_j v_{ij} \pi_{ij}(\vec{v}_i) - q_i(\vec{v}_i) \geq 0$, for all $i, \vec{v}_i \in E_i$, guaranteeing that the mechanism is ex-interim Individually Rational ($m|E_i|$).
- $\sum_j v_{ij} \pi_{ij}(\vec{v}_i) - q_i(\vec{v}_i) \geq \sum_j v_{ij} \pi_{ij}(\vec{v}'_i) - q_i(\vec{v}'_i)$, for all $i, \vec{v}_i \in E_i, \vec{v}'_i \in E_i$, guaranteeing that the mechanism is BIC ($m|E_i|^2$).
- $SO(\{\pi_{ij}(\vec{v}_i)\}_{i,j,\vec{v}_i}, \mathcal{D}) = \text{“yes”}$, guaranteeing that there is a feasible mechanism implementing the $\pi_{ij}(\vec{v}_i)$ s.

Maximizing:

$\sum_{i,\vec{v}_i \in E_i} q_i(\vec{v}_i) \Pr[\vec{w}_i \leftarrow \mathcal{D}_i, \vec{w}_i \in \cup_{\sigma \in S} \{\sigma(\vec{v}_i)\}]$, the expected revenue.

Figure 3: LP of [14] for the item-symmetric setting, where we use a separation oracle to determine feasibility of the reduced form. In parentheses at the end of each line is the number of such variables/constraints .

and therefore for any weights $W_{ij}(A)$:

$$\begin{aligned} \sum_{i,j} \sum_{A \in T_i} W_{ij}(A) \pi_{ij}(A) \Pr[t_i = A] &\leq \sum_{i,j} \sum_{A \in T_i} W_{ij}(A) M_{ij}(A) \Pr[t_i = A] \\ &\leq \sum_P \Pr[P] \cdot \max_{\vec{\phi}(P) \in F} \left\{ \sum_i \sum_j W_{ij}(P_i) \cdot \phi_{ij}(P) \right\}, \end{aligned}$$

where the last inequality comes from observing that M defines a feasible mechanism, and therefore for every profile provides a valid choice of $\vec{\phi}(P) \in F$ for the maximization.

Now, we argue that if the reduced form is infeasible, we can find *constricting weights*, i.e. weights making the inequality in the statement of the theorem fail. For any mechanism M , define the value of M , $V(M)$ as follows:

$$V(M) = \sum_i \sum_j \sum_{A \in T_i} \min\{M_{ij}(A) \Pr[t_i = A], \pi_{ij}(A) \Pr[t_i = A]\}.$$

We can now say that M implements the reduced form if and only if:

$$V(M) = \sum_i \sum_j \sum_{A \in T_i} \pi_{ij}(A) \Pr[t_i = A].$$

By our assumption on the feasibility constraints, there exist some set of linear constraints H , with all $c_{ij}(h) \geq 0$ and $d(h) \geq 0$, for all $h \in H$, such that a mechanism is feasible if and only if its induced $\{\phi(P)\}_P$ satisfies the following for all profiles P :

$$\vec{\phi}(P) \cdot \vec{c}(h) \leq d(h) \quad \forall h \in H \quad \wedge \quad \vec{\phi}(P) \geq \vec{0}.$$

Therefore, the mechanism M maximizing $V(M)$ is a solution to the following LP:

Primal LP 1

Variables:

- $\phi_{ij}(P)$, for all bidders i , items j and profiles P .

Constraints:

- $\vec{\phi}(P) \cdot \vec{c}(h) \leq d(h)$ for all $P, h \in H$. Rewritten as:

$$\Pr[P] \cdot \vec{\phi}(P) \cdot \vec{c}(h) \leq \Pr[P] \cdot d(h).$$

- $\phi_{ij}(P) \geq 0$, for all i, j, P .
- $\sum_{P|P_i=A} \phi_{ij}(P) \cdot \Pr[P] \leq \pi_{ij}(A) \cdot \Pr[t_i = A]$, for all $i, j, A \in T_i$.

Maximizing:

- $\sum_{i,j,P} \phi_{ij}(P) \cdot \Pr[P]$.

In other words, for any feasible mechanism M , Primal LP 1 modifies M to sometimes throw away an item rather than have $M_{ij}(A) > \pi_{ij}(A)$, then counts the expected number of items that M awards. Primal LP 1 maximizes the expected number of items awarded by a feasible mechanism subject to these constraints. Consider now the dual to this LP:

Dual LP 1

Variables:

- $x_h(P)$, for all $h \in H$, and profiles P .
- $y_{ij}(A)$, for all i, j and $A \in T_i$.

Constraints:

- $\sum_{h \in H} \Pr[P] \cdot c_{ij}(h) \cdot x_h(P) + \Pr[P] \cdot y_{ij}(P_i) \geq \Pr[P]$, for all i, j, P . Rewritten as:

$$\sum_{h \in H} c_{ij}(h) \cdot x_h(P) + y_{ij}(P_i) \geq 1.$$

- $x_h(P) \geq 0$, for all $P, h \in H$.
- $y_{ij}(A) \geq 0$, for all i, j and $A \in T_i$.

Minimizing:

- $\sum_{i,j} \sum_{A \in T_i} y_{ij}(A) \pi_{ij}(A) \Pr[t_i = A] + \sum_P \Pr[P] \cdot \sum_{h \in H} d(h) \cdot x_h(P)$.

By setting all $\phi_{ij}(P)$ to 0, we know Primal LP 1 is feasible. Also, Dual LP 1 is feasible, as we can set all $x_h(P)$ to 0 and $y_{ij}(A)$ to 1. Hence, by strong LP duality, we know the values of the two LPs are the same, so the value of Dual LP 1 is strictly smaller than $\sum_{i,j} \sum_{A \in T_i} \pi_{ij}(A) \Pr[t_i = A]$ if and only if the reduced form is infeasible. Now let us consider Dual LP 1 and try to find the constricting weights.

Observe first, that there is always an optimal solution to Dual LP 1 with $y_{ij}(A) \leq 1$ for all i, j, A . This is because each $y_{ij}(A)$ has a non-negative coefficient in the objective function, each $c_{ij}(h) \geq 0$, and each $x_h(P) \geq 0$. So decreasing $y_{ij}(A)$ from above 1 to 1 does not violate any feasibility constraints, or increase the value of the objective function.

Next, we claim that in Dual LP 1 the optimal feasible choice of $x_h(P)$ to minimize the objective function depends only on the values $y_{ij}(P_i)$ and $x_{h'}(P)$ for other $h' \in H$, *but not on $x_{h'}(P')$ for any other profile P'* . This is easy to see, as there are no inequalities with non-zero coefficients for both $x_h(P)$ and $x_{h'}(P')$ for any inequalities h, h' and profiles $P \neq P'$. Therefore, if we are given a choice of $y_{ij}(A)$ for all i, j, A , we can find the optimal choices of $x_h(P)$ for all h, P *independently for*

each profile P (but still considering all $h \in H$ at once). In particular, once the $y_{ij}(A)$ s are fixed, the optimal choice of $x_h(P)$ for a fixed profile P is just a solution to the following LP:

Primal LP 2

Variables:

- $x_h(P)$ for all $h \in H$ (P is fixed for the purposes of this LP).

Constraints:

- $\sum_{h \in H} c_{ij}(h) \cdot x_h(P) \geq 1 - y_{ij}(P_i)$, for all i, j .
- $x_h(P) \geq 0$ for all $h \in H$.

Minimizing:

- $\Pr[P] \cdot \sum_{h \in H} d(h) \cdot x_h(P)$.

We can now take the dual of Primal LP 2, for a fixed profile P :

Dual LP 2

Variables:

- $\gamma_{ij}(P)$, for all i, j .
- $\phi_{ij}(P)$, for all i, j . Extra variables satisfying $\Pr[P] \cdot \phi_{ij}(P) = \gamma_{ij}(P)$, used for convenience.

Constraints:

- $\sum_{i,j} \gamma_{ij}(P) \cdot c_{ij}(h) \leq \Pr[P] \cdot d(h)$. Replacing $\gamma_{ij}(P)$ with $\Pr[P] \cdot \phi_{ij}(P)$ yields:

$$\vec{\phi}(P) \cdot \vec{c}(h) \leq d(h).$$

- $\Pr[P] \cdot \phi_{ij}(P) = \gamma_{ij}(P)$ for all i, j .
- $\phi_{ij}(P) \geq 0$ for all i, j .

Maximizing:

- $\sum_{i,j} \gamma_{ij}(P) \cdot (1 - y_{ij}(P_i))$. Replacing $\gamma_{ij}(P)$ with $\Pr[P] \cdot \phi_{ij}(P)$ yields:

$$\Pr[P] \cdot \sum_{i,j} \phi_{ij}(P) \cdot (1 - y_{ij}(P_i)).$$

We now make our key observation. For a fixed profile P , and for a fixed choice of $y_{ij}(A)$ for all i, j, A , Dual LP 2 is *exactly* trying to maximize the weighted sum $\sum_{i,j} (1 - y_{ij}(P_i)) \cdot \phi_{ij}(P)$, subject to the same constraints guaranteeing that $\vec{\phi}(P) \in F$. By the strong LP duality (it is again easy to verify that both Primal LP 2 and Dual LP 2 are feasible), we know that the value of this maximum weighted feasible allocation is exactly the value attained by the optimal solution to Primal LP 2. Therefore, we may now claim that, for any choice of $y_{ij}(A)$ s in Dual LP 1, the minimum value of Dual LP 1 subject to this choice is exactly:

$$\sum_{i,j} \sum_{A \in T_i} y_{ij}(A) \cdot \pi_{ij}(A) \cdot \Pr[t_i = A] + \sum_P \Pr[P] \cdot \max_{\vec{\phi}(P) \in F} \sum_{i,j} (1 - y_{ij}(P_i)) \cdot \phi_{ij}(P_i).$$

Finally, as the reduced form is infeasible, there is some optimal solution $\{x_h^*(P), y_{ij}^*(A)\}$ to Dual LP 1 with all $y_{ij}^*(A) \in [0, 1]$ whose value is strictly less than $\sum_{i,j} \sum_{A \in T_i} \pi_{ij}(A) \cdot \Pr[t_i = A]$. Therefore, if we take this choice of $y_{ij}^*(A)$ and set $W_{ij}(A) = 1 - y_{ij}^*(A)$ we obtain:

$$\begin{aligned}
& \sum_{i,j} \sum_{A \in T_i} y_{ij}^*(A) \cdot \pi_{ij}(A) \cdot \Pr[t_i = A] + \sum_P \Pr[P] \cdot \max_{\vec{\phi}(P) \in F} \sum_{i,j} (1 - y_{ij}^*(P_i)) \cdot \phi_{ij}(P_i) \\
&= \sum_{i,j} \sum_{A \in T_i} y_{ij}^*(A) \pi_{ij}(A) \Pr[t_i = A] + \sum_P \Pr[P] \cdot \sum_{h \in H} d(h) \cdot x_h^*(P) \\
&< \sum_{i,j} \sum_{A \in T_i} \pi_{ij}(A) \cdot \Pr[t_i = A].
\end{aligned}$$

By subtracting $\sum_{i,j} \sum_{A \in T_i} y_{ij}^*(A) \cdot \pi_{ij}(A) \cdot \Pr[t_i = A]$ from both sides, we obtain a proof of the theorem:

$$\sum_P \Pr[P] \cdot \max_{\vec{\phi}(P) \in F} \sum_{i,j} W_{ij}(P_i) \cdot \phi_{ij}(P_i) < \sum_{i,j} \sum_{A \in T_i} W_{ij}(A) \cdot \pi_{ij}(A) \cdot \Pr[t_i = A].$$

□

F.1 A Special Case: Many Items, Demand Constraints, Correlated Bidders

Recall the definition of our multi-commodity flow instance outlined in Section 3.4.1. Let us make it more precise:

Nodes:

- $S(j)$, for all items j , these are sources.
- $I(P)$, for all possible profiles.
- $T(i, A)$ for all i and $A \in T_i$, these are sinks.

Edges:

- From $S(j)$ to $I(P)$ of capacity $\Pr[P]$, for all items j and profiles P .
- From $I(P)$ to $T(i, P_i)$ of capacity $C_i \Pr[P]$, where C_i is bidder i 's demand, for all bidders i and profiles P .

Commodities:

- $G_{ij}(A)$ for all bidders i , items j and all types $A \in T_i$ with a demand of $\Pr[t_i = A] \pi_{ij}(A)$.

We observe that the feasibility of the above multi-commodity flow instance is equivalent to the feasibility of the reduced form π .

Proposition 4. *The above multi-commodity flow problem has a feasible solution satisfying every demand if and only if the reduced form $\vec{\pi}$ is feasible.*

Proof. Say the reduced form is feasible. Then there exist variables $\phi_{ij}(P)$ such that:

1. $\sum_j \phi_{ij}(P) \leq C_i$, for all i, P .
2. $\sum_i \phi_{ij}(P) \leq 1$, for all j, P .
3. $\sum_{P|P_i=A} \phi_{ij}(P) \cdot \Pr[P] = \pi_{ij}(A) \cdot \Pr[t_i = A]$, for all i, j, A .

So along the edges from $S(j)$ to $I(P)$ and $I(P)$ to $T(i, P_i)$, we will send $\Pr[P] \cdot \phi_{ij}(P)$ of commodity $G_{ij}(P_i)$. It is clear that by the conditions on the ϕ s that all of the edge capacities are respected, and that the total amount of commodity $G_{ij}(A)$ received by node $T(i, A)$ is exactly $\pi_{ij}(A) \cdot \Pr[t_i = A]$, i.e. matching its demand.

Taking the inverse of this procedure (by interpreting as $\Pr[P] \cdot \phi_{ij}(P)$ the amount of commodity $G_{ij}(A)$ that is sent through the edge $(S(j), I(P))$) shows how to take any feasible flow satisfying

every demand and capacity constraint and turn it into a set of ϕ s that implement the reduced form (by just reading out the $\phi_{ij}(P)$'s from the edges of the flow network, using the afore-mentioned interpretation). \square

F.2 False Extensions

A short discussion is warranted on why this is the right extension of Border's Theorem. There are stronger, more natural extensions that one might hope to be true. For instance, below are two natural extensions that are **not** true, *even in the special case that the feasibility constraints are only demand constraints*.

1. **False Extension 1.** If the demand constraints are met for every type in expectation, and it is feasible to allocate each item individually, then the reduced form is feasible. Formally,

$$\sum_j \pi_{ij}(A) \leq C_i \quad \forall i, A \in T_i$$

$$\sum_i \sum_{A \in S_i} \pi_{ij}(A) \leq 1 - \prod_i (1 - \Pr[t_i \in S_i]) \quad \forall j, \forall i, S_i \subseteq T_i$$

2. **False Extension 2.** Theorem 12 holds when we only consider $\{W_{ij}(A)\}$ such that $W_{ij}(A) = W_{ij'}(A)$ for all bidders i , items j, j' , and $A \in T_i$. I.E. the weights for a type of bidder i does not depend on the item.

Proposition 5. *Both False Extensions are indeed false, even when all bidders are i.i.d. and when the feasibility constraints are just demand constraints.*

Proof. First, consider the case with three unit-demand i.i.d. bidders and two items. Every bidder has two possible types A, B , and $\Pr[A] = \epsilon$, $\Pr[B] = 1 - \epsilon$. We use $\pi_j(t)$ to denote the probability that a bidder with type t receives item j . Let $\pi_1(A) = \pi_2(A) = 1/2$. $\pi_1(B) = \pi_2(B) = 0$. Then clearly the hypotheses of False Extension 1 are met. No type receives more than one item in expectation, and it is feasible to allocate each item individually. However, as each bidder is unit-demand, and type A wants exactly one item in expectation, A must always leave with an item every time they show up to the auction to possibly satisfy the π s. Clearly this is not possible, as it is possible for all three A 's to show up at once, but there are only two items. So we cannot possibly give A an item every time they show up. Therefore, False Extension 1 is indeed false.

To get a counterexample to the second false extension, modify the above example by adding a third item, with $\pi_3(A) = \pi_3(B) = 0$. Once we have $W_i(A) = W_{i1}(A) = W_{i2}(A) = W_{i3}(A)$ for all i , the LHS of the constricting-weights inequality in the statement of Theorem 12 becomes:

$$\sum_{i,j} \sum_{t \in T_i} W_i(t) \pi_j(t) \Pr[t_i = t] = \epsilon(W_1(A) + W_2(A) + W_3(A)).$$

This is because $\pi_{ij}(B) = 0$ for all i, j , and $\Pr[t_i = A] = \epsilon$ for all i .

Now consider setting $\phi_{ii}(P) = 1$ for all P , and $\phi_{ij}(P) = 0$ for all P and $i \neq j$. For any choice of weights $W_i(A), W_i(B)$, the RHS of the inequality (the expected weights induced by the ϕ s) will be:

$$\epsilon(W_1(A) + W_2(A) + W_3(A)) + (1 - \epsilon)(W_1(B) + W_2(B) + W_3(B)).$$

This is because we always give exactly one item to each bidder. When the type of bidder i is A , we get weight $W_i(A)$, when it is B we get weight $W_i(B)$. Because $t_i = A$ with probability ϵ , we get the above calculation.

Obviously, for any choice of weights, the following holds.

$$\epsilon(W_1(A) + W_2(A) + W_3(A)) + (1 - \epsilon)(W_1(B) + W_2(B) + W_3(B)) \geq \epsilon(W_1(A) + W_2(A) + W_3(A)).$$

Then if we choose the ϕ s in a way that it maximizes the expected weights, LHS will still be less than RHS. This shows that there are no constricting weights satisfying $W_{ij}(A) = W_{ij'}(A)$ for all i, j, j' . However, we still know that this reduced form is infeasible. The key is that we need to let $W_{i3}(A) = 0$ in order to find constricting weights, without having W_{i1} and $W_{i2} = 0$. So we need different weights for different items. \square