

On Symmetries and Fairness in Multi-Dimensional Mechanism Design

SUBMISSION 150

We consider a revenue-maximizing seller with multiple items for sale to a single population of buyers. In ad auction domains the items correspond to views from particular demographics, and recent works have therefore identified a novel *fairness constraint*: equally-qualified users from different demographics should be shown the same desired ad at equal rates. Prior work abstracts this to the following fairness guarantee: if an advertiser places an identical bid on two users, those two users should view the ad with the same probability [27, 35].

We first propose a relaxation of this guarantee from worst-case to Bayesian settings, which circumvents strong impossibility results from these works. We then study this guarantee through the lens of *symmetries*, as any item-symmetric auction is also fair (by this definition). Our main result shows that for a single population of additive buyers with independent (but not necessarily identically distributed) item values, bundling all items together achieves a constant-factor approximation to the revenue-optimal item-symmetric mechanism. Observe that in this setting, bundling all items together corresponds to concealing all demographic data and treating all users the same [22].

1 INTRODUCTION

Ad auctions are a significant source of revenue for numerous firms, causing their theoretical study to be a mainstay in both the Economics and Computer Science communities. Classical works typically design and analyze auctions that optimize the participants' collective utility (i.e. the sum of bidders' values for items they receive, also called the *welfare*) [19, 30, 50], or perhaps just the auctioneer's utility (i.e., her *revenue*) [42]. Recently, the ubiquity of ad auctions in domains where fairness constraints are first-order concerns has motivated a new desideratum for consideration: the *items'* utility for the outcome selected.

While it makes little sense to consider the utility of an apple or orange, recall that the items in ad auction domains are in fact users. That is, when an advertiser (bidder) wins an item (impression on a particular user), that item (the user) also enjoys some utility.

In practice, these utilities are hard to quantify (even moreso than typical values for an item), and this side of the market is typically not monetized. As a result, there are no 'bids' or 'utilities' of the users (items) to consider. However, some examples of high-utility ads include those for desirable jobs, low-interest loans, etc. and are subject to anti-discrimination laws. Specifically, it is considered unfair for users (items) who are equally qualified for jobs/loans/etc. to view protected ads at different rates. Therefore, recent works have proposed considering the utility of users (items) through the lens of *fairness* [27, 35]. That is, these works propose to still consider the utility of the auctioneer and bidders (ads) in the classical sense, but to additionally ensure that the outcomes are fair to the users (items).

In practice, a non-discriminatory advertiser (bidder) might submit identical bids for equally-qualified users (items) of different demographics. But [27] observes that this *fair behavior* is insufficient to achieve a *fair outcome*. Indeed, protected ads are bidding against non-protected ads (e.g. Men's Shoes, Maternity Clothes, etc.) which legally place discriminatory bids. If discriminatory advertisers (bidders) place higher bids on demographic *A* than *B* (items), then the price of impressions for demographic *A* (items) will be higher, and then *even a non-discriminatory ad will be displayed in a discriminatory manner*.

To have a simple example in mind (taken from [27]), consider the case that the auctioneer runs a second-price auction on each of two items. This auction format is ostensibly fair: there is nothing in its description that seems to bias it against any item (user) or bidder (ad). But consider when Bidder One (a protected ad) submits a bid of 1 for both items (users), and Bidder Two (a non-protected ad) submits a bid of 2 for item one and 0 for item two. Then Bidder One wins item two, and Bidder Two wins item one. As a result, the (user) demographic corresponding to item one views no protected ads, while the (user) demographic corresponding to item two views protected ads with probability one. That is, despite the fact that the protected advertiser bids in a non-discriminatory manner, and that the auction is ostensibly fair, the result is an unfair outcome (assuming that users prefer to see the protected ad).

While the above example is clearly stylized, this phenomenon is not just of theoretical concern. Indeed, automated systems are constantly making decisions that affect our daily lives. These systems rely on advanced algorithms and big data to make decisions which, in theory, have the potential to be better informed and more equitable. However, studies have shown that they may instead internalize and perpetuate societal biases [1, 2, 18, 36, 39, 44]. Efforts on mitigating bias have ranged from examinations of the data [8, 49, 51], to machine learning algorithmic contributions [7, 47, 54], to theoretical analyses [15, 27, 35]. In our domain of study, works indeed find that impressions for female users are more expensive than impressions for male users, and that female users are less

likely to see ads for high paying jobs [24] and STEM jobs [38], even when advertisers for such jobs are unbiased [3].

[27, 35] posit that it is the job of the auction designer to guarantee fair outcomes, and propose formal fairness definitions motivated by individual fairness [26]. Requiring an auction to satisfy these constraints of course limits the auctioneer’s ability to optimize its revenue and/or the bidders’ welfare, so these works study the tradeoff between fairness and optimality. Our work follows this same paradigm, but deviates from prior work in a few fundamental ways, highlighted below.¹

Bayesian vs. Worst-Case. [15, 27, 35] consider worst-case definitions of fairness. For example, they might demand that for all possible bids of non-protected advertisers, if a protected advertiser submits an identical bid for two demographics (items), those demographics view the protected ad with (almost) equal probability. These definitions arise naturally from the fairness literature upon which they build, but unfortunately also lead to strong impossibility results.

Ad auctions, however, are executed millions of times daily, and auctioneers have quite extensive Bayesian priors. Indeed, revenue optimization is typically studied in Bayesian settings, where the designer seeks to maximize their expected revenue. *We propose to also consider a Bayesian, rather than worst-case, notion of fairness.* Indeed, unfairness is undesirable exactly when it is systemic, and Bayesian notions are best suited to capture systemic phenomena. A formal definition appears in Section 2, and a discussion appears in Section A.1. By considering a Bayesian notion of fairness, we’re able to circumvent the impossibility results proved in [35].

Revenue vs. Welfare. [15, 27, 35] consider auctions that attempt to maximize welfare (more specifically: they attempt to maximize the “declared welfare,” but don’t assume that the declared bids correspond to the bidders’ actual values). In the absence of fairness constraints, such auctions are extremely well-understood, and are particularly simple in the settings considered (e.g. if the ads/bidders are additive², welfare is maximized by awarding the item/user to the highest bidder).

However, if an auctioneer truly wishes to optimize their expected revenue, it is well-understood that revenue-optimal auctions are *significantly* more complex [17, 20, 21, 33, 41, 43, 48]. *We propose to consider an auctioneer who wishes to optimize their expected revenue in a multi-dimensional Bayesian setting,* rather than one who wishes to optimize the declared welfare. Our model is the standard setup for multi-dimensional mechanism design (formally defined in Section 2): the auctioneer has multiple items for sale, and bidders’ values for the items (users) are drawn independently.

Direct vs. Indirect Competition. [15, 27, 35] consider advertisers (bidders) who directly compete to display their ads to a limited supply of users (items). In such settings, even a benevolent platform must fail to show some ads to some users. While our model is rich enough to capture this setting, our analysis isolates a different source of competition.

Specifically, even when there is unlimited supply, a revenue-maximizing designer may choose not to show every ad to every user. Indeed, if they cannot offer different prices to different advertisers, they may achieve greater revenue by setting a high price that excludes some advertisers from purchasing. In this sense, advertisers indirectly compete with each other: one advertiser’s bids affect the impressions sold to another due to the fact that the seller wishes to optimize their revenue, rather than due to limited supply. *We study the seller’s revenue objective (rather than the limited supply of users) as a driving source of unfairness.*

¹The below paragraphs distinguish our model from [15, 27, 35]. We discuss other related works such as [12] in Section 1.2.

²A valuation function $v(\cdot)$ is additive if for all sets S of items, $v(S) = \sum_{j \in S} v(\{j\})$.

Connection to Symmetries. We adopt (a Bayesian version of) the individual fairness notion proposed in [27]: an auction is fair if whenever a bidder (ad) submits an identical bid for two items (users), they receive those items with the same probability (in expectation over other bidders' bids). We further observe that this notion is implied by the stronger definition of *item-symmetric* [23]. Specifically, an auction is item-symmetric if whenever a bidder (ad) swaps their bids for two items (users), this swaps the probabilities with which they receive those items (in expectation over other bidders' bids).

Item-symmetric auctions have been studied in the multi-dimensional mechanism design literature for their own sake as a tool to optimize revenue in a computationally-efficient manner [23, 37], but we use them here as a tool to guarantee fair outcomes. Specifically, *we target the design of item-symmetric auctions.*

1.1 Results and Technical Highlights

The previous paragraphs motivate our modeling decisions within the multi-dimensional mechanism design domain: we consider a single seller with m items (users) for sale, and the buyers' (ads') values are drawn from a distribution known to the seller (Bayesian vs. Worst-Case). The seller's goal is to design a truthful auction that optimizes their expected revenue (Revenue vs. Welfare). We focus on the case where there is unlimited supply, which can alternatively be represented by a single population of potential bidders (Direct vs. Indirect Competition).³

From here, we wish to study item-symmetric auctions, and do so through the lens of simplicity vs. optimality [13, 14, 32, 34]: *is there a simple, approximately-optimal item-symmetric auction?* For example, one particularly simple auction is to *bundle the items together* (that is, pick a price p and allow the buyer to receive all items for price p , or no items). It is also easy to see that this auction is item-symmetric (and therefore fair). Indeed, in the language of ad auctions, it corresponds to an auction which does not use personalized data at all, and chooses to display an ad to whatever user shows up independently of their demographics [22].

Another particularly simple auction is to *sell the items separately* (that is, for each item j , pick a price p_j , and allow the buyer to pick any set of items S to purchase at price $\sum_{j \in S} p_j$). Auctions of this format, while simple, are not necessarily item-symmetric (nor fair): if $p_1 > p_2$, then an advertiser could submit an identical bid of p_2 for both items yet receive only item two. A proper subclass of such auctions, which is item-symmetric (and therefore fair), is to *sell the items separately and symmetrically* (that is, set a single price p , and allow the buyer to pick any set of items and pay p per item). In the language of ad auctions, selling separately and symmetrically corresponds to setting a price of p to display an ad (independently of any data), but letting the advertisers choose any subsets of demographics to display their ads.

Our main result is that bundling the items together achieves a constant-factor approximation to the revenue-optimal item-symmetric mechanism, and that this factor can be improved by considering the better of bundling together and selling separately and symmetrically.

Main Result (See Theorems 3.1 and 3.2): For a single additive buyer, and any number of independent items, bundling together achieves a $O(1)$ -approximation to the revenue-optimal item-symmetric mechanism. The maximum between bundling together and selling separately and symmetrically improves this constant.

³To quickly see why the unlimited supply setting is equivalent to the single bidder setting: Because there are no supply constraints, it is feasible to pick any single-bidder mechanism and just use it for every bidder.

We also provide several auxiliary results in Appendix F, which study the relationships between simple mechanisms such as bundling together, selling separately, and selling separately and symmetrically.

1.2 Related Work

Multi-dimensional Mechanism Design. At a technical level, the most closely-related field to our work is that of simple vs. optimal multi-dimensional mechanism design. For example, our proof outline is reminiscent of [5, 40], and we provide an alternative proof outline reminiscent of [9]. Note also that our main result (that bundling together is a constant-factor approximation of the optimal revenue from any symmetric auction) implies one result from [40] (that bundling together is a constant-factor approximation of the optimal revenue when all items are i.i.d.).⁴ [22] were the first to explicitly note the connection between the sale of an “uncertain item” (e.g. an impression to a user whose demographic is known only to the designer) and the classic multi-dimensional mechanism design setting.

Individual Fairness in Auction Design. To our knowledge, [27] were the first to consider fairness in auction design from a theoretical perspective. Their work provides fairness definitions based on individual fairness and motivating examples demonstrating that unfair outcomes can arise from ostensibly fair auctions and non-discriminatory behavior of auctioneers. Follow-up works such as [15, 35] proceed in similar models. These works provide strong impossibility results in the worst-case, but also provide matching positive results (and improve the guarantees of these positive results under restrictions on the otherwise worst-case input). Section 1 discusses extensively several ways in which our work contributes to this line of work, along with the technical differences.

The other related work in this direction is [12], who also consider the theoretical design of fair auctions from the perspective of a revenue-maximizing seller. The biggest difference between their work and ours is that they essentially consider a single-dimensional setting (that is, they seek to optimize the Myersonian virtual value of the winning bidder (ad) for each item (user)), but place fairness constraints across auctions for different items. They formulate a linear program in their setting and optimize their problem (exactly) computationally efficiently. Put another way, their work exclusively considers auctions which “sell items separately” from a revenue perspective, but with cross-item constraints concerning fairness.

One simple way to compare our work to these is that we focus on a simple notion of fairness, but in the sophisticated multi-dimensional mechanism design setting, whereas these prior works consider more sophisticated/quantitative notions of fairness but in simpler auction settings (either welfare-maximization or single-dimensional revenue-maximization).

Empirical Studies of Fairness in Auction Design. Empirical studies on the rate at which ads are displayed to different demographics motivate the line of work to which we contribute [3, 24, 38]. For example, [24] finds that ads for high-paying jobs are shown to more men than women, and [38] draws the same conclusion for STEM jobs. The empirical studies in [3] support the conjecture that this is due to “spillover effects” caused by higher competition for female views.

2 PRELIMINARIES

Our main results consider a single seller (the advertising platform) with m items for sale (each item corresponds to an impression for a different demographic of user) to a single buyer (representing the

⁴This follows as the optimal auction when all items are i.i.d. is in fact item-symmetric.

population of advertisers, who do not directly compete for limited supply). However, we instantiate our model with $n \geq 1$ buyers (directly competing advertisers), to best compare with prior work.

Each buyer i has a value v_{ij} for each item j , and has value $\sum_{j \in S} v_{ij}$ for set S (that is, the bidders are additive). Each v_{ij} is drawn independently from a distribution D_{ij} , and we define $D_i := \times_j D_{ij}$ to represent the i^{th} population of advertisers,⁵ and a particular \vec{v}_i represents a particular advertiser. We denote by $D := \times_i D_i$ as the entire population, and \vec{v} as a particular profile of advertisers. When there is just a single bidder, we abuse notation and let $D := D_1$, and $\vec{v} := \vec{v}_1$. For discrete distributions, we let $f^D(\vec{v}) := \Pr_{\vec{w} \leftarrow D}[\vec{w} = \vec{v}]$. For a single variable discrete distribution D , we let $F^D(\cdot)$ denote the CDF, so $F^D(v) := \sum_{w \in [0, v]} f^D(w)$.

The seller's goal is to design a truthful mechanism that maximizes their expected revenue. Specifically, a mechanism consists of a mapping from valuation profiles \vec{v} to ex-post allocation probabilities $x_{ij}(\vec{v})$ for all bidders i and items j , and an ex-post price $p_i(\vec{v})$ for all bidders i . This denotes the probability that bidder i gets item j when the full vector of bids is \vec{v} , and the price that bidder i pays (respectively). The interim allocation rule is a mapping from valuation vector \vec{v}_i to the interim allocation probability $\pi_{ij}(\vec{v}_i) := \mathbb{E}_{\vec{v}_{-i} \leftarrow D_{-i}}[x_{ij}(\vec{v}_i; \vec{v}_{-i})]$. The interim price paid is $q_i(\vec{v}_i) := \mathbb{E}_{\vec{v}_{-i} \leftarrow D_{-i}}[p_i(\vec{v}_i; \vec{v}_{-i})]$.⁶ These quantities denote the probability that bidder i receives item j and the price bidder i pays (respectively), conditioned on reporting \vec{v}_i and in expectation over the remaining bids being drawn from D_{-i} . If we wish to emphasize that these terms come from a specific mechanism M , we will write $x_{ij}^M(\cdot)$, $p_i^M(\cdot)$, etc. We say that a mechanism is truthful if it is Bayesian individually rational and Bayesian incentive compatible. That is:

$$\sum_j v_{ij} \cdot \pi_{ij}(\vec{v}_i) - q_i(\vec{v}_i) \geq 0, \quad \forall i, \vec{v}_i. \quad (\text{Bayesian Individually Rational})$$

$$\sum_j v_{ij} \cdot \pi_{ij}(\vec{v}_i) - q_i(\vec{v}_i) \geq \sum_j v_{ij} \cdot \pi_{ij}(\vec{v}'_i) - q_i(\vec{v}'_i), \quad \forall i, \vec{v}_i, \vec{v}'_i. \quad (\text{Bayesian Incentive Compatible})$$

The seller's goal is to find, over all truthful mechanisms, the one maximizing her expected revenue (which can be written either as $\mathbb{E}_{\vec{v} \leftarrow D}[\sum_i p_i(\vec{v})]$ or $\sum_i \mathbb{E}_{\vec{v}_i \leftarrow D_i}[q_i(\vec{v}_i)]$).

Fairness and Symmetries. Motivated by the discussion in Section 1, we define an auction to be *fair* if whenever an advertiser places the same bid for an impression for two different demographics, those two demographics view the ad with the same probability. After mapping from advertiser to buyer, and demographics to items, this yields the following two definitions, depending on whether we seek a guarantee ex-post or in the interim.

Definition 2.1 (Fair Auction). An auction is *ex-post fair* with respect to bidder i if for all valuation profiles \vec{v} , and items j, k :

$$v_{ij} = v_{ik} \implies x_{ij}(\vec{v}) = x_{ik}(\vec{v}) \quad \forall i.$$

An auction is *interim fair* if for all bidders i , valuation vectors \vec{v}_i , and items j, k :

$$v_{ij} = v_{ik} \implies \pi_{ij}(\vec{v}_i) = \pi_{ik}(\vec{v}_i).$$

Intuitively, an auction is ex-post fair with respect to bidder i if no matter the bids of the other bidders (advertisers), when bidder i places an identical bid for two items (views from particular

⁵For example, it could be that $D_i = D_{i'}$ for all i, i' , and each bidder is drawn from the same population. This represents settings where the platform cannot price-discriminate based on properties of the advertiser. It could also be that $D_i \neq D_{i'}$. In such settings, perhaps D_i is the population of 'big' advertisers, and $D_{i'}$ is the population of 'small' advertisers, and the platform knows from which population each individual advertiser is drawn.

⁶We use the standard notation \vec{v}_{-i} to refer to the vector of bids excluding bidder i , and D_{-i} to refer to the distribution over valuation profiles, excluding bidder i .

demographics), they receive those items with the same probability (those demographics view the ad with the same probability). An auction is interim fair with respect to bidder i if when bidder i places an identical bid for two items, they receive those items with the same probability in expectation over the other bidders' bids (assuming they are drawn from D_{-i}).

Both definitions are implied by the following stronger definitions (respectively), which require that the auction be invariant under relabeling of items/demographics. Below, the notation $\sigma(\vec{v})$ refers to a vector satisfying $(\sigma(\vec{v}))_{i\sigma(j)} = \vec{v}_{ij}$ for all i, j (that is, the items/demographics have been relabeled according to σ).

Definition 2.2 (Symmetric Auction). An auction is *ex-post symmetric* with respect to bidder i if for all permutations σ on $[m]$, valuation vectors \vec{v}_i , and partial valuation profiles \vec{v}_{-i} :

$$\sigma(\vec{x}_i(\vec{v}_i; \vec{v}_{-i})) = \vec{x}_i(\sigma(\vec{v}_i), \vec{v}_{-i}).$$

An auction is *interim symmetric* if for all permutations σ , bidders i , and valuation vectors \vec{v}_i :

$$\sigma(\vec{\pi}_i(\vec{v}_i)) = \vec{\pi}_i(\sigma(\vec{v}_i)).$$

Intuitively, an auction is symmetric if permuting a valuation vector by σ permutes the allocation vector by σ as well. We briefly observe that symmetry implies fairness.

OBSERVATION 1. *If an auction is ex-post (respectively, interim) symmetric, it is also ex-post (respectively, interim) fair.*

PROOF. Let $v_{ij} = v_{ik}$, and consider the permutation σ which swaps j and k . Then $\sigma(\vec{v}_i) = \vec{v}_i$. Symmetry therefore implies⁷ that $x_{ij}(\vec{v}_i; \vec{v}_{-i}) = x_{i\sigma(j)}(\sigma(\vec{v}_i); \vec{v}_{-i}) = x_{ik}(\vec{v}_i; \vec{v}_{-i})$. This completes the proof for ex-post.

Similarly by symmetry: $\pi_{ij}(\vec{v}_i) = \pi_{i\sigma(j)}(\sigma(\vec{v}_i)) = \pi_{ik}(\vec{v}_i)$. This completes the proof for interim. \square

A Stronger Fairness Guarantee via Symmetry. The fairness guarantees above (and those in prior work) demand that equally-valued users are shown an ad with the same probability. A stronger fairness guarantee might instead demand that if demographic (item) i is *valued higher* by an advertiser (bidder) than demographic (item) j , then users from demographic i are shown that ad at least as often as those from demographic j . We term this property *strong monotonicity in fairness*, defined below (after mapping from advertiser to buyer, and demographics to items).

Definition 2.3 (Strong Monotonicity in Fairness). An auction satisfies *ex-post strong monotonicity in fairness* with respect to bidder i , if for all valuation profiles \vec{v} , and items j, k :

$$v_{ij} \geq v_{ik} \Rightarrow x_{ij}(\vec{v}) \geq x_{ik}(\vec{v}) \forall i.$$

An auction is *interim strong monotonicity in fairness* if for all bidders i , valuation vectors \vec{v}_i , and items j, k :

$$v_{ij} \geq v_{ik} \Rightarrow \pi_{ij}(\vec{v}_i) \geq \pi_{ik}(\vec{v}_i).$$

Note that ex-post (resp. interim) fairness does not imply ex-post (resp. interim) strong monotonicity in fairness. Interestingly, however, [23] has already previously studied interim strong monotonicity in fairness (under the name strong monotonicity), and shown that it is implied by interim symmetry! That is, while previously studied notions of fairness alone do not imply this stronger fairness notion, symmetry does. Below we briefly repeat their observation (and it's short proof, for completeness).

⁷To see this, recall that $\sigma(\vec{x}_i(\vec{v}_i; \vec{v}_{-i}))$ is a vector that puts $x_{ij}(\vec{v}_i; \vec{v}_{-i})$ in the $i, \sigma(j)$ coordinate.

OBSERVATION 2 ([23]). *If an auction is interim symmetric and Bayesian Incentive Compatible, then it satisfies interim strong monotonicity in fairness.*

PROOF. By Observation 1, any auction that is interim symmetric is also interim fair. This means that if $v_{ij} = v_{ik}$, then $\pi_{ij}(\vec{v}_i) \geq \pi_{ik}(\vec{v}_i)$, so the conditions for interim strong monotonicity in fairness hold whenever $v_{ij} = v_{ik}$, and we need only consider the case when $v_{ij} > v_{ik}$.

Assume for contradiction that $v_{ij} > v_{ik}$ but $\pi_{ij}(\vec{v}) < \pi_{ik}(\vec{v})$. Advertiser i could lie and swap v_{ij} and v_{ik} . By symmetry, this swaps $\pi_{ij}(\vec{v})$ and $\pi_{ik}(\vec{v})$, and strictly increases the advertiser's interim expected value (by $(v_{ij} - v_{ik}) \cdot (\pi_{ik}(\vec{v}) - \pi_{ij}(\vec{v}))$). The auctioneer still charges advertiser i the same interim expected price (also by symmetry), giving them strictly more expected utility, and contradicting that the auction is Bayesian Incentive Compatible. \square

We briefly note that ex-post symmetry and ex-post incentive compatibility also imply ex-post strong monotonicity in fairness, and the proof outline is identical (but we will not formally state/prove this, as we did not formally define ex-post incentive compatibility).

OBSERVATION 3. *If an auction satisfies ex-post (respectively, interim) strong monotonicity in fairness, then it is also ex-post (respectively, interim) fair.*

PROOF. Let $v_{ij} = v_{ik}$, then $v_{ij} \geq v_{ik}$ and $v_{ij} \leq v_{ik}$. By ex-post strong monotonicity in fairness, we get $x_{ij}(\vec{v}) \geq x_{ik}(\vec{v})$ and $x_{ij}(\vec{v}) \leq x_{ik}(\vec{v})$. Therefore, $x_{ij}(\vec{v}) = x_{ik}(\vec{v})$.

Similarly, by interim strong monotonicity in fairness, we get $\pi_{ij}(\vec{v}) \geq \pi_{ik}(\vec{v})$ and $\pi_{ij}(\vec{v}) \leq \pi_{ik}(\vec{v})$. Therefore, $\pi_{ij}(\vec{v}) = \pi_{ik}(\vec{v})$. This completes the proof for ex-post. \square

Selling Separately and Bundling Together. For a single bidder distribution D , we use the following notation:

- $\text{REV}_M(D)$: the revenue of a particular mechanism M for distribution D .
- $\text{REV}(D)$: the optimal revenue achieved by any truthful mechanism for D (formally, this is: $\sup_{M, M \text{ is truthful}} \{\text{REV}_M(D)\}$). Observe that mechanisms achieving $\text{REV}(D)$ are not necessarily fair nor symmetric.
- $\text{SYMREV}(D)$: the optimal revenue achieved by any truthful and interim symmetric mechanism for D . By definition, the mechanism witnessing $\text{SYMREV}(D)$ is symmetric.
- $\text{SREV}(D)$: the optimal revenue achieved by *selling separately* to a bidder drawn from D . That is, the seller sets a price $p_j := \arg \max_p \{p \cdot \Pr_{v \leftarrow D_j} [v \geq p]\}$ on item j , and the buyer purchases all items for which $v_j \geq p_j$. Observe that such mechanisms are not necessarily fair nor symmetric. In the context of our running example, this corresponds to the platform setting a different price to display an ad to each demographic, and allowing each advertiser to choose which demographic views to purchase.
- $\text{SSREV}(D)$: the optimal revenue achieved by *symmetrically selling separately* to a bidder drawn from D . That is, the seller sets the same price $p := \arg \max_q \{q \cdot \sum_j \Pr_{v \leftarrow D_j} [v \geq q]\}$ on all items, and the buyer purchases all items for which $v_j \geq p$. Observe that such mechanisms are both fair and symmetric. In the context of our running example, this corresponds to the platform setting the same price to display an ad to each demographic, and allowing each advertiser to choose which demographic views to purchase.
- $\text{BREV}(D)$: the optimal revenue achieved by *bundling together*. That is, the seller sets a price $p := \arg \max_q \{q \cdot \Pr_{\vec{v} \leftarrow D} [\sum_j v_j \geq q]\}$ on the grand bundle of all items, and the buyer either purchases all items at total price p (when $\sum_j v_j \geq p$), or nothing. Observe that such mechanisms are both fair and symmetric. In the context of our running example, this corresponds to the platform ignoring all demographic information, and allowing advertisers to show their ads to all users or none.

Mapping between ad auctions and multi-dimensional mechanism design. We briefly repeat the connection between ad auctions and the classical multi-item auction setup formally identified by [22, Theorem 2]. An item j in the classic setting corresponds to a demographic j in the ad auction domain. Moreover, awarding the buyer item j with probability x_j corresponds to showing a user with type j their ad with probability x_j . Therefore, if advertisers have a value of v_j per click from demographic j , and demographic j represents a d_j fraction of the population, then the advertiser’s value for an allocation \vec{x} is $\sum_j v_j \cdot d_j x_j$. Observe that when each demographic represents the same fraction of the population (e.g. male/female) that the advertiser’s valuation is simply additive. Therefore, our main results on item-symmetric mechanisms with an additive buyer directly have bite in the ad auction domain when each demographic is equally likely.⁸

We also remind the reader that [22, Theorem 2] observes that bundling items together in the classic setting corresponds to concealing demographic data in the ad auction setting. Similarly, selling separately in the classic setting corresponds to setting a price p_j to display an ad to demographic j , and letting advertisers choose which subset of demographics to target. Selling separately and symmetrically further enforces that $p_i = p_j$ for all i, j .

3 MAIN RESULT: BREV IS A CONSTANT-FACTOR APPROXIMATION TO SYMREV.

In this section, we prove our main result: BREV is a constant factor approximation to SYMREV. Recall that in our setting, BREV corresponds to the optimal revenue achieved by a mechanism which ignores demographic data entirely.

THEOREM 3.1. *Let D be any additive single-bidder distribution over any number of independent items. Then $204\text{BREV}(D) \geq \text{SYMREV}(D)$.*

We prove Theorem 3.1 in two steps. The first step is the main step, and proves a theorem reminiscent of the main result of [5], establishing that either $\text{BREV}(D)$ or $\text{SSREV}(D)$ is a constant-factor approximation to $\text{SYMREV}(D)$ (Theorem 3.2). The second step argues that in fact $\text{BREV}(D)$ is a constant factor approximation to $\text{SSREV}(D)$ (Proposition 3.3).

THEOREM 3.2. *Let D be any additive single-bidder distribution over any number of independent items. Then $24\text{BREV}(D) + 20\text{SSREV}(D) \geq \text{SYMREV}(D)$.*

PROPOSITION 3.3. *Let D be any additive single-bidder distribution over any number of independent items. Then $9\text{BREV}(D) \geq \text{SSREV}(D)$.*

We defer the proof of Theorem 3.1 to Appendix B. In Appendix B.1 we show how to upper bound $\text{SYMREV}(D)$ with $\text{REV}(D')$ for a modified distribution D' . To prove Theorem 3.2, we will first provide a modified distribution D' , show that its revenue is close to that of D , and then design a flow for D' . In Appendix B.2 we upper bound $\text{REV}(D')$ with $24\text{BREV}(D) + 20\text{SSREV}(D)$, and we provide a proof based on tools used in [32]. In Section 3.1 we prove Proposition 3.3.

In Appendix C, we provide an alternative proof based on the [9] duality framework (for the case when the distribution of the bidder’s maximum value for the items is regular). In Appendix C.2 we also overview a naive attempt at applying their framework (using their “canonical flow”), which helps provide intuition for the need to go through D' .

3.1 Comparing BREV to SSREV

In this section, we prove Proposition 3.3: BREV is a constant factor approximation to SSREV. Recall that in our setting, BREV corresponds to the optimal revenue achieved by a mechanism which

⁸We also note that it is an interesting open direction to extend our main results from an additive bidder to a ‘scaled additive’ bidder so that this connection holds even for non-uniform demographic distributions.

ignores demographic data entirely, while SSREV corresponds to the optimal revenue achieved by a mechanism which sets the same price to display an ad to each demographic, and allows each advertiser to choose which demographic views to purchase.

To prove Proposition 3.3, consider any mechanism that sets price p on each item separately, and let $q_j(p)$ denote the probability that the bidder purchases item j (that is, that $v_j \geq p$), and let $q(p) := \sum_j q_j(p)$ denote the expected number of items purchased at price p (and therefore, $\text{SSREV}(D) := \sup_p \{p \cdot q(p)\}$). We will show that there is always a price p' for the grand bundle that collects a constant fraction of $p \cdot q(p)$.

We will first consider the case where $q(p) \leq 8$ (that is, at most 8 items are sold at price p in expectation). Unsurprisingly, in this case it suffices to set a price $p' := p$ on the grand bundle.

LEMMA 3.4. *Let $q(p) \leq 8$. Then selling the grand bundle at price p generates expected revenue at least $p \cdot q(p)/9$.*

PROOF. The proof follows from straight-forward calculations:

$$\begin{aligned} \Pr_{\vec{v} \leftarrow D} \left[\sum_j v_j \geq p \right] &\geq \Pr_{\vec{v} \leftarrow D} \left[\max_j \{v_j\} \geq p \right] = 1 - \Pr_{\vec{v} \leftarrow D} \left[\forall j, v_j < p \right] \\ &= 1 - \prod_j \Pr_{v_j \leftarrow D_j} \left[v_j < p \right] = 1 - \prod_j (1 - q_j(p)) \\ &\geq 1 - \prod_i e^{-q_j(p)} = 1 - e^{-q(p)} \\ &\geq q(p)/9. \end{aligned}$$

Above, the first line holds since the distribution for $\sum_j v_j$ stochastically dominates $\max_j v_j$. The second line follows as values are independent. The third line holds as $e^{-q_j(p)} \geq 1 - q_j(p)$, and the last line holds for all values $q(p) \leq 8$. In particular, this means that we can set price p on the grand bundle, and it will sell with probability at least $q(p)$, completing the proof. \square

This completes our analysis of the first case. We now consider the case where $q(p) > 8$, and we set the grand bundle price to be $p' = q(p)p/2$. We will also use the notation $\sigma^2(p)$ to denote the variance of the random variable $\sum_j \mathbb{I}(v_j \geq p)$. We quickly observe a bound on $\sigma^2(p)$, which follows as all values are independent.

OBSERVATION 4. $\sigma^2(p) = \sum_j q_j(p)(1 - q_j(p)) \leq \sum_j q_j(p) = q(p)$.

LEMMA 3.5. *Let $q(p) > 8$. Then selling the grand bundle at price $p \cdot q(p)/2$ generates expected revenue at least $p \cdot q(p)/9$.*

PROOF. Observe that certainly $\sum_j v_j \geq p \cdot q(p)/2$ when there are at least $q(p)/2$ items with value greater than p . We lower bound the probability of this event using Chebyshev's inequality:

$$\begin{aligned} \Pr_{\vec{v} \leftarrow D} \left[\sum_j v_j \geq p \cdot q(p)/2 \right] &\geq \Pr_{\vec{v} \leftarrow D} \left[\sum_j \mathbb{I}(v_j \geq p) \geq \frac{1}{2}x(p) \right] \\ &\geq \Pr_{\vec{v} \leftarrow D} \left[\left| \sum_j \mathbb{I}(v_j \geq p) - q(p) \right| \leq \frac{1}{2}x(p) \right] \\ &\geq 1 - \frac{\sigma^2(p)}{q(p)^2/4} \\ &\geq 1 - 4/q(p) \geq 1/2 \end{aligned}$$

\square

Proposition 3.3 now follows from Lemma 3.4 and Lemma 3.5.

While not related to our main result, we also explore the relationship between $SSREV$ and $BREV$ in the other direction, and include a proof in Appendix E. The outline is similar to a related claim in [5].

THEOREM 3.6 (SSREV IS A LOG APPROXIMATION OF BREV). *For any distribution D for a single additive buyer and m not necessarily independent items, $BREV(D) \leq 5 \log(m)SSREV(D)$.*

In Appendix F, we analyze several examples demonstrating the relationship between $BREV$ and $SSREV$.

4 CONCLUSIONS

Motivated by recent works which consider fairness constraints in welfare-maximizing or single-dimensional auctions [12, 15, 27, 35], we introduce fairness considerations in multi-dimensional mechanism design. We study interim (rather than worst-case) notions of fairness, and use this to motivate the study of simple item-symmetric auctions. Our main technical result is that the simple auction which ignores demographic information entirely is a constant-factor approximation to the optimal item-symmetric auction.

REFERENCES

- [1] 2013. Google Translate’s Gender Problem (And Bing Translate’s, And Systran’s). <https://www.fastcompany.com/3010223/google-translates-gender-problem-and-bing-translates-and-systrans>.
- [2] 2015. Google Photos labeled black people ‘gorillas’. USA Today. <https://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/>.
- [3] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through Optimization: How Facebook’s Ad Delivery Can Lead to Biased Outcomes. *Proc. ACM Hum. Comput. Interact.* 3, CSCW (2019), 199:1–199:30. <https://doi.org/10.1145/3359301>
- [4] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. 2017. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 869–877. <https://doi.org/10.1145/3055399.3055426>
- [5] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*. 21–30. <https://doi.org/10.1109/FOCS.2014.11>
- [6] Hedyeh Beyhaghi and S. Matthew Weinberg. 2019. Optimal (and Benchmark-Optimal) Competition Complexity for Additive Buyers over Independent Items. In *Proceedings of the 51st ACM Symposium on Theory of Computing Conference (STOC)*.
- [7] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems*. 4349–4357.
- [8] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *ACM Conference on Fairness, Accountability, Transparency (FAccT)* (2018).
- [9] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. 2016. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the 48th ACM Conference on Theory of Computation (STOC)*.
- [10] Yang Cai and Zhiyi Huang. 2013. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [11] Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. 170–183. <https://doi.org/10.1145/3055399.3055465>
- [12] L. Elisa Celis, Anay Mehrotra, and Nisheeth K. Vishnoi. 2019. Toward Controlling Discrimination in Online Ad Auctions. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA (Proceedings of Machine Learning Research)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.), Vol. 97. PMLR, 4456–4465. <http://proceedings.mlr.press/v97/mehrotra19a.html>
- [13] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*.

- [14] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.
- [15] Shuchi Chawla and Meena Jagadeesan. 2020. Fairness in ad auctions through inverse proportionality. *CoRR* abs/2003.13966 (2020). arXiv:2003.13966 <https://arxiv.org/abs/2003.13966>
- [16] Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*. 579–596. <https://doi.org/10.1145/2940716.2940756>
- [17] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Pappas, Xiaorui Sun, and Mihalis Yannakakis. 2015. On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*. 1464–1479. <https://doi.org/10.1109/FOCS.2015.93>
- [18] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *arXiv preprint arXiv:1703.00056* (2017).
- [19] Edward H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11, 1 (1971), 17–33.
- [20] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2014. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [21] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2017. Strong Duality for a Multiple-Good Monopolist. *Econometrica* 85, 3 (2017), 735–767.
- [22] Constantinos Daskalakis, Christos H. Papadimitriou, and Christos Tzamos. 2016. Does Information Revelation Improve Revenue?. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, Vincent Conitzer, Dirk Bergemann, and Yiling Chen (Eds.). ACM, 233–250. <https://doi.org/10.1145/2940716.2940789>
- [23] Constantinos Daskalakis and S. Matthew Weinberg. 2012. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, June 4-8, 2012*. 370–387. <https://doi.org/10.1145/2229012.2229042>
- [24] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings. *Proc. Priv. Enhancing Technol.* 2015, 1 (2015), 92–112. <https://doi.org/10.1515/popets-2015-0007>
- [25] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundarajan. 2009. Revenue submodularity. In *Proceedings 10th ACM Conference on Electronic Commerce (EC-2009), Stanford, California, USA, July 6–10, 2009*, John Chuang, Lance Fortnow, and Pearl Pu (Eds.). ACM, 243–252. <https://doi.org/10.1145/1566374.1566409>
- [26] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, Shafi Goldwasser (Ed.). ACM, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [27] Cynthia Dwork and Christina Ilvento. 2019. Fairness Under Composition. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA (LIPIcs)*, Avrim Blum (Ed.), Vol. 124. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 33:1–33:20. <https://doi.org/10.4230/LIPIcs.ITCS.2019.33>
- [28] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 343. <https://doi.org/10.1145/3033274.3085115>
- [29] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. A Simple and Approximately Optimal Mechanism for a Buyer with Complements: Abstract. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 323. <https://doi.org/10.1145/3033274.3085116>
- [30] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631.
- [31] Sergiu Hart and Noam Nisan. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.
- [32] Sergiu Hart and Noam Nisan. 2017. Approximate revenue maximization with multiple items. *J. Economic Theory* 172 (2017), 313–347. <https://doi.org/10.1016/j.jet.2017.09.001>
- [33] Sergiu Hart and Philip J. Reny. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics* 10, 3 (2015), 893–922.
- [34] Jason D. Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*. 225–234.
- [35] Christina Ilvento, Meena Jagadeesan, and Shuchi Chawla. 2020. Multi-category fairness in sponsored search auctions. In *FAT* '20: Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, January 27-30, 2020*, Mireille Hildebrandt, Carlos Castillo, Elisa Celis, Salvatore Ruggieri, Linnet Taylor, and Gabriela Zanfir-Fortuna (Eds.). ACM, 348–358. <https://doi.org/10.1145/3351095.3372848>

- [36] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 3819–3828.
- [37] Pravesh Kothari, Divyarthi Mohan, Ariel Schwartzman, Sahil Singla, and S. Matthew Weinberg. 2019. Approximation Schemes for a Buyer with Independent Items via Symmetries. In *the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- [38] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Manag. Sci.* 65, 7 (2019), 2966–2981. <https://doi.org/10.1287/mnsc.2018.3093>
- [39] Sam Levin. 2016. A beauty contest was judged by AI and the robots didn't like dark skin.
- [40] Xinye Li and Andrew Chi-Chih Yao. 2013. On Revenue Maximization for Selling Multiple Independently Distributed Items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.
- [41] A. M. Manelli and D. R. Vincent. 2007. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory* 137, 1 (2007), 153–185.
- [42] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [43] Gregory Pavlov. 2011. Optimal Mechanism for Selling Two Goods. *The B.E. Journal of Theoretical Economics* 11, 3 (2011).
- [44] Lizzie Plaugic. 2017. FaceApp's creator apologizes for the app's lightening 'hot' filter.
- [45] Amir Ronen. 2001. On approximating optimal auctions. In *Proceedings 3rd ACM Conference on Electronic Commerce (EC-2001), Tampa, Florida, USA, October 14-17, 2001*, Michael P. Wellman and Yoav Shoham (Eds.). ACM, 11–17. <https://doi.org/10.1145/501158.501160>
- [46] Aviad Rubinstein and S. Matthew Weinberg. 2015. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 377–394. <https://doi.org/10.1145/2764468.2764510>
- [47] H. Ryu, H. Adam, and M. Mitchell. 2018. InclusiveFaceNet: Improving Face Attribute Detection with Race and Gender Diversity. In *Proceedings of FAT/ML*.
- [48] John Thanassoulis. 2004. Haggling over substitutes. *Journal of Economic Theory* 117 (2004), 217–245.
- [49] Antonio Torralba and Alexei A. Efros. 2011. Unbiased Look at Dataset Bias. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [50] William Vickrey. 1961. Counterspeculations, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 1 (1961), 8–37.
- [51] Angelina Wang, Arvind Narayanan, and Olga Russakovsky. 2020. ViBE: A Tool for Measuring and Mitigating Bias in Image Datasets. *CoRR* abs/2004.07999 (2020). arXiv:2004.07999 <https://arxiv.org/abs/2004.07999>
- [52] Andrew Chi-Chih Yao. 2017. Dominant-Strategy versus Bayesian Multi-item Auctions: Maximum Revenue Determination and Comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin (Eds.). ACM, 3–20. <https://doi.org/10.1145/3033274.3085120>
- [53] Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [54] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*. 325–333.

A ADDITIONAL DISCUSSION

We revisit the discussion of Section 1 with a brief technical highlight and concrete examples.

Brief Technical Highlight. Almost all prior work on simple vs. optimal multi-dimensional mechanism design consider selling separately as a simple mechanism, and are perfectly content to argue that the maximum between selling separately and some other simple mechanism achieves a constant-factor approximation [5, 6, 9, 11, 16, 28, 29, 46, 53]. These works proceed by proving elegant upper bounds on the optimal achievable revenue, and breaking these bounds into terms which can be approximated by simple mechanisms. In particular, there is usually a term that corresponds to “revenue achieved when a bidder has an unusually high value for some item” (e.g., “the tail” in [4, 5, 16, 40, 46], SINGLE in [6, 9, 11, 28, 29]), and this term is easily approximated by the revenue of selling separately (typically, this term is also the most straight-forward to approximate).

In our setting however, selling separately is not a symmetric mechanism, and in fact could be up to a factor of $\Omega(\#items)$ better than the optimal symmetric mechanism (See Example C.9). Therefore we need to target an upper bound that in some cases *is even tighter than the revenue achieved by selling separately*.

At a very high level, prior bounds leverage the fact that “the auctioneer cannot both extract revenue $\approx v$ when the buyer has a high value v for item j and revenue $\approx 2v$ when the buyer has an even higher value $2v$ for item j , because the buyer with value $2v$ can always lie and pretend that their value is v instead.” Our bound instead must leverage the fact that “the auctioneer cannot both extract revenue $\approx v$ when the buyer has a high value v for *some item* j and also revenue $\approx 2v$ when the buyer has an even higher value $2v$ for *some other item* ℓ .” This is because if the auctioneer extracts revenue $\approx v$ when the buyer has value v for item j , they must also extract revenue $\approx v$ when the buyer has value v for item ℓ (by item-symmetry), and then the buyer with value $2v$ for item ℓ can always pretend that their value is instead v . The main technical challenge is figuring out a way to leverage this intuition into a concrete bound. Once we find the right approach, the complete proof is fairly clean, and is able to leverage existing machinery in the simple vs. optimal literature for multi-dimensional mechanism design.

A.1 Interim vs. Ex-post fairness

First, we provide some examples illustrating that interim fairness better captures systemic fairness at a population level than ex-post fairness. We also demonstrate that unfairness can arise even in the single-bidder setting.

Example A.1 (Ex-post unfair, but interim fair). Consider the following (highly stylized) example with two items and three bidders. Item one represents the demographic of female users, and item two represents the demographic of male users. Buyer One represents the population of advertisers for STEM jobs, Buyer Two represents the population of advertisers for Maternity Clothes, and Buyer Three represents the population of advertisers for Men’s Shoes.

Concretely, D_{11} and D_{12} are both point masses at 1. This means that every STEM jobs advertiser has value 1 for displaying their ad, and they comply with anti-discrimination regulations by submitting an identical value for both demographics.

D_{21} is equal to 2 with probability $1/2$, and 0 with probability $1/2$. D_{22} is a point-mass at 0. This means that half of Maternity Clothes advertisers have a high value for displaying their ad to women, and half have no value for displaying their ads.

D_{31} is a point mass at 0. D_{32} is equal to 2 with probability $1/2$, and 0 with probability $1/2$. This means that half have a high value for displaying their ads to men, and half of Men’s Shoes advertisers have no value.

The revenue-optimal (and welfare-optimal) auction awards each item to the highest bidder, and charges them their full value (it is not hard to see that this is truthful. Conditioned on being truthful, it is clear that it is optimal because it achieves expected revenue equal to the full expected welfare).

This auction is interim symmetric with respect to bidder 1, and therefore interim fair as well. To see this, observe that $\pi_{11}(1, 1) = \pi_{12}(1, 1) = 1/2$. That is, the STEM jobs ad is displayed to each demographic with probability 1/2. However, this auction is *not* ex-post fair (and therefore not ex-post symmetric) with respect to bidder 1. To see this, observe that $x_{11}((1, 1), (2, 0), (0, 0)) = 0$, but $x_{12}((1, 1), (2, 0), (0, 0)) = 1$. That is, when the Maternity Clothes advertiser has high value for displaying their ad to women, but Men’s Shoes advertiser has no value for displaying their ad, men see the STEM jobs ad with probability 1, but women see it with probability 0.

Example A.1 is reminiscent of worst-case examples from [27, 35]: there exist instances of bids from other bidders that result in unfairness ex-post. However, Example A.1 is still fair at the population level, and this property is best captured by interim fairness.

Of course, it is still possible for optimal auctions to be interim unfair (again with examples reminiscent of [27, 35]). For example, if Example A.1 contained only bidders one and two (STEM jobs and Maternity Clothes), but not bidder three, the optimal auction is still to award each item to the highest bidder and charge their full value. But now the auction is interim unfair, as $\pi_{11}(1, 1) = 1$ while $\pi_{12}(1, 1) = 1/2$.

While the mathematics behind both examples is nearly identical, the precise definitions are important. Interim fairness better captures when unfairness manifests at the population level, rather than worst-case instances. We now note that when the auctioneer uses a revenue-optimal auction, unfairness can arise even when there is just a single bidder.

Example A.2 (Interim unfair with a single bidder). Consider the following example with two items and one bidder. Item one still represents the demographic of female users, and item two still represents the demographic of male users. There is a single population of all advertisers, and there are two types of advertisers in this population (each equally likely): STEM jobs, which have value 1 for both demographics, and Maternity clothes, which have value 3 for female users and 1 for male users.

This results in an instance where D_{11} takes value 1 with probability 1/2 and value 3 with probability 1/2. D_{12} takes value 1 with probability 1.

The revenue-optimal auction is to “sell separately” item one at price 1 and item two at price 3. That is, $x_{11}(1, 3) = x_{12}(1, 3) = 1$, $p_1(1, 3) = 3$, $x_{11}(1, 1) = 1$, $x_{12}(1, 1) = 0$, $p(1, 1) = 1$.⁹ Observe that this is interim unfair. Because there is just a single bidder, $\pi_{1j}(\vec{v}_1) = x_{1j}(\vec{v}_1)$, and therefore $\pi_{11}(1, 1) \neq \pi_{12}(1, 1)$.

Intuitively, what drives unfairness in Example A.2 is again the fact that the STEM jobs advertiser, which submits an identical bid for all demographics, is competing with the Maternity Clothes advertiser. In Example A.1 (or the interim unfair variant of Example A.1), this competition directly causes unfair viewing of STEM jobs ads, because supply constraints result in the “female view” item being sold to the Maternity Clothes advertiser instead. In Example A.2, this competition *indirectly* causes unfair viewing of STEM jobs ads, because the ad platform generates increased revenue by setting a higher price on the “female view” item.

Both types of competition are important. Prior work [12, 27, 35] focuses on competition caused by limited supply. We instead focus on fairness implications caused indirectly by the impact of

⁹To see why this is optimal, observe that this is essentially a single-item instance, because D_{11} is a point-mass. The optimal price for D_{12} is 3.

competition on pricing. We focus on the single-bidder setting to isolate this source of competition (as there are no supply constraints).

A.2 Fairness vs. Symmetries

Our definition of fairness (modulo the distinction between ex-post and interim) is motivated by individual fairness, and is (a special case of) the same one used in the works that introduced this direction [27]. We use symmetries as a technical lens to guarantee this definition of fairness, although it is worth briefly noting that symmetry may have standalone interest as a fairness concept. Specifically, fairness alone guarantees that whenever an advertiser places the same bid for two different demographics, those two demographics see the ad with the same probability. Symmetry further guarantees that swapping the bids of an advertiser for any two demographics swaps the probabilities that those demographics see the ad.

Beyond philosophical motivation, fairness properties are typically evaluated by their ability to guarantee *fair outcomes*. That is, systems are evaluated by their ability to mitigate bias. In our setting, the desirable outcome is that equally-qualified individuals from different demographics see protected ads with the same probability. A fair auction guarantees this outcome as long as the protected advertiser places the same bid on all relevant demographics. A symmetric auction also guarantees this outcome under the same circumstances, but also under the weaker condition that the protected advertiser’s bid vector is invariant under all permutations of the relevant demographics.

Therefore, while symmetry may be interesting as a standalone definition, it is not immediately clear what symmetric auctions might guarantee in terms of fair outcomes beyond what fair auctions already guarantee. As such, we view symmetries mainly as a technical lens to study fair auctions.

A.3 Ex-post vs. Interim Symmetries

Interim symmetric auctions have been previously studied within Bayesian mechanism design from a pure optimization perspective (e.g. [23, 37]). Indeed, this is because revenue-optimal auctions for symmetric distributions (including distributions for which the items are drawn i.i.d.) are known to be interim symmetric [23]. That is, if D is invariant under all item permutations, then $\text{SYMREV}(D) = \text{REV}(D)$.

Ex-post symmetries, on the other hand, have not been previously studied (to the best of our knowledge). This makes sense, as even when the entire input is i.i.d., the revenue-optimal auction need not be ex-post symmetric.¹⁰ This also suggests that from the perspective of Bayesian mechanism design, interim concepts may be better suited than ex-post concepts. Section A.1 argues this intrinsically from the definitions themselves, as these definitions best capture population-level phenomena. The above discussion (briefly) argues that from a technical perspective, interim definitions are likely to be more tractable/relevant than ex-post definitions.

B OMITTED PROOFS FROM SECTION 3

In this section, we provide the proof of Theorem 3.1. In Appendix B.1 we show how to upper bound $\text{SYMREV}(D)$ with $\text{REV}(D')$ for a modified distribution D' . In Appendix B.2 we upper bound $\text{REV}(D')$ with $24\text{BREV}(D) + 20\text{SSREV}(D)$.

B.1 Upper Bounding $\text{SYMREV}(D)$ with a Modified Distribution

To prove Theorem 3.2, we will first provide a modified distribution D' , show that its revenue is close to that of D , and then design a flow for D' . We first define our modified distribution.

¹⁰One such example is provided in [52], where there are two bidders and two items and the value of each bidder for each item is i.i.d. from a distribution supported on $\{1, 2\}$.

Definition B.1 (Modified Distribution). Let D be an additive single-buyer distribution over m independent items. Define the modified distribution D' to be the following distribution over $m + 1$ items (we will use index 0 to refer to the extra item):

- Draw v_j independently from D_j for all j .
- Let $j^* := \arg \max_j \{v_j\}$ (breaking ties lexicographically).
- Let $v'_0 := v_{j^*}$.
- Let $v'_j := v_j$, for all $j \neq j^*$.
- Draw v'_{j^*} from the distribution D_{j^*} , conditioned on $v'_{j^*} < v'_0$ (if $v'_0 = 0$, just set $v'_{j^*} = 0$).

Intuitively, D' ensures that the maximum-value item is always in the same coordinate (0), and this allows us to leverage a clean application of prior tools. We first need to argue that $\text{SYMREV}(D)$ is upper bounded by (an appropriate function of) $\text{REV}(D')$.

PROPOSITION B.2. For any $\varepsilon > 0$, $\text{SYMREV}(D) \leq \frac{1}{1-\varepsilon} \cdot \text{REV}(D') + \frac{2}{\varepsilon(1-\varepsilon)} \text{BREV}(D)$.

The proof of Proposition B.2 proceeds in several small steps, with each step moving closer from D to D' . Most steps in isolation should seem straight-forward to readers familiar with “nudge-and-round” arguments, although Lemma B.8 and Corollary B.9 are specific to our setting. We first consider the following:

Definition B.3. For a given distribution D over m items, define $D^{(1)}$ to be the distribution which draws $\vec{v} \leftarrow D$, and then concatenates $v_0 := 0$.

LEMMA B.4. $\text{SYMREV}(D) \leq \text{SYMREV}(D^{(1)})$.¹¹

PROOF. Consider any symmetric mechanism M for D , and view it by its menu (that is, the list of (\vec{x}, p) it allows the buyer to purchase). Recall that because M is symmetric, that for all (\vec{x}, p) on the menu, $(\sigma(\vec{x}), p)$ is also on the menu for all item permutations σ .

Consider now the menu M' which offers the option $((0, \vec{x}), p)$ for all (\vec{x}, p) on the menu for M , and also contains $(\sigma((0, \vec{x})), p)$ for all item permutations σ (for all (\vec{x}, p) on the menu for M). This menu is clearly symmetric by definition, and therefore M' is a symmetric mechanism.

We just need to figure out the revenue of M' compared to M . Consider any \vec{v} , and let (\vec{x}, p) denote their favorite option from the menu for M . We claim that $((0, \vec{x}), p)$ is the favorite option for $(0, \vec{v})$ from the menu for M' . Once we prove this, this will establish that $\text{REV}_M(D) = \text{REV}_{M'}(D^{(1)})$.

To see this, first observe that every option on the menu for M' awards (at least) one item with probability 0. Because $(0, \vec{v})$ has value 0 with item 0, their (weakly) favorite option from the menu for M' awards item 0 with probability 0. Conditioned on this, their utility for any option $((0, \vec{y}), q)$ is exactly $\vec{v} \cdot \vec{y} - q$, which is exactly their utility for the option (\vec{y}, q) on the menu for M . Because (\vec{x}, p) is their favorite such option in M , $((0, \vec{x}), p)$ remains their favorite option from M' . \square

We now take one more intermediate step. This step will require a “nudge-and-round” lemma from prior work:

LEMMA B.5 ([10, 46]). For a single bidder, let D and D' be coupled so that it is possible to jointly draw (\vec{v}, \vec{v}') from (D, D') so that $\mathbb{E}[\sum_j |v_j - v'_j|] \leq C$. Then for all $\varepsilon > 0$, and any mechanism M , there exists a mechanism M' such that:

$$\text{REV}_{M'}(D') \geq (1 - \varepsilon) \cdot \text{REV}_M(D) - \frac{1}{\varepsilon} \cdot C.$$

Moreover, if M is symmetric, then so is M' .¹²

¹¹Observe that $\text{REV}(D) \leq \text{REV}(D^{(1)})$ is trivial: simply ignore item 0. Lemma B.4 requires confirming that indeed symmetric mechanisms can also ignore item 0.

¹²In fact, M' is identical to M , after discounting all prices by a factor of $(1 - \varepsilon)$.

Definition B.6. For a given distribution D over m items, define $D^{(2)}$ to be the distribution which first draws $\vec{v} \leftarrow D$. Then, if $j^* := \arg \max_j \{v_j\}$, draw v_0 from the distribution D_{j^*} , conditioned on $v_0 < v_{j^*}$. We will let V' denote the random variable distributed according to the marginal of v_0 .

LEMMA B.7. For any $\varepsilon > 0$, $\text{SYMREV}(D^{(2)}) \geq (1 - \varepsilon)\text{SYMREV}(D^{(1)}) - \frac{1}{\varepsilon} \cdot \mathbb{E}[V']$.

PROOF. To see this, simply couple $D^{(1)}$ and $D^{(2)}$ so that the values for all items except for item 0 are equal. Then, the expected ℓ_1 distance between the two draws is exactly v_0 , which is distributed according to V' for $D^{(2)}$, and is a point mass at 0 for $D^{(1)}$. Therefore, we may take C in the hypothesis of Lemma B.5 to be $\mathbb{E}[V']$, and let M be the optimal symmetric mechanism for $D^{(1)}$. The lemma statement now follows, as Lemma B.5 proves the existence of a symmetric mechanism M' for $D^{(2)}$ with the desired revenue. \square

There are two remaining steps. First, we simply observe that $\text{SYMREV}(D^{(2)}) = \text{SYMREV}(D')$.

OBSERVATION 5. $\text{SYMREV}(D^{(2)}) = \text{SYMREV}(D')$.

PROOF. Observe that the distribution D' can be obtained by first drawing a vector from $D^{(2)}$, and then swapping coordinates 0 and j^* . This means that D' and $D^{(2)}$ can be coupled so that with probability one, the vector drawn from D' is a permutation of the one drawn from $D^{(2)}$. This immediately implies that all symmetric mechanisms achieve the same expected revenue from both D' and $D^{(2)}$. The observation now follows. \square

The final step is to reason about $\mathbb{E}[V']$. After this, we wrap up the proof of Proposition B.2. We begin with a slightly stronger lemma than necessary.

LEMMA B.8. $2\text{BREV}(D) \geq \mathbb{E}_{(\vec{v}, \vec{w}) \leftarrow D \times D} [\min\{\max_j \{v_j\}, \max_j \{w_j\}\}]$.

PROOF. Observe that $\mathbb{E}_{(\vec{v}, \vec{w}) \leftarrow D \times D} [\min\{\max_j \{v_j\}, \max_j \{w_j\}\}]$ is exactly the revenue of a second-price auction for a single item and two bidders whose values are drawn independently distributed according to $\max_j \{v_j\}$ (when $\vec{v} \leftarrow D$). Refer to this distribution as D^* . Therefore, the optimal revenue for two bidders from D^* is at least $\mathbb{E}_{(\vec{v}, \vec{w}) \leftarrow D \times D} [\min\{\max_j \{v_j\}, \max_j \{w_j\}\}]$, and this means that the optimal revenue for a single bidder from D^* is at least $\mathbb{E}_{(\vec{v}, \vec{w}) \leftarrow D \times D} [\min\{\max_j \{v_j\}, \max_j \{w_j\}\}] / 2$.¹³

Finally, observe that the distribution for $\sum_j v_j$ stochastically dominates D^* . Therefore, for any price p (including the revenue-optimal price for D^*), $\Pr[\sum_j v_j \geq p] \geq \Pr[\max_j \{v_j\} \geq p]$. Therefore, setting the revenue-optimal price for one bidder from D^* as a price on the grand bundle of all items witnesses that $\text{BREV}(D) \geq \text{REV}(D^*)$, and the lemma follows. \square

Interestingly, Lemma B.8 is tight. Consider the case of $m = 1$ item drawn from the equal revenue distribution. Then $\text{BREV}(D) = 1$, and the expected minimum of two draws from the equal revenue distribution is 2. We can also use it to upper bound $\mathbb{E}[V']$.

COROLLARY B.9. $\mathbb{E}[V'] \leq 2\text{BREV}(D)$.

PROOF. Observe that V' is stochastically dominated by $\mathbb{E}_{(\vec{v}, \vec{w}) \leftarrow D \times D} [\min\{\max_j \{v_j\}, \max_j \{w_j\}\}]$. Once we establish this, the corollary follows.

Observe that both $\min\{\max_j \{v_j\}, \max_j \{w_j\}\}$ and V' can be described as random variables of the form “first draw X from distribution F , then draw Y from distribution G , conditioned on $Y < X$.” Indeed, $\min\{\max_j \{v_j\}, \max_j \{w_j\}\}$ has F as the distribution of $\max\{\max_j \{v_j\}, \max_j \{w_j\}\}$, and G

¹³This follows by “Revenue Submodularity” [25]. It is also easy to see in this case that any truthful mechanism for two bidders from D^* can be projected to a mechanism for a single bidder from D^* , and that this mechanism retains half the original revenue.

as the distribution of $\max_j \{z_j\}$. V' has F as the distribution of $\max_j \{v_j\}$, and G as the distribution D_{j^*} .

It is easy to see that $\max\{\max_j \{v_j\}, \max_j \{w_j\}\}$ stochastically dominates $\max_j \{v_j\}$ (the former is the maximum of two draws from the latter). It is also easy to see that for any c , the distribution of $\max_j \{z_j\}$ conditioned on being $< c$ stochastically dominates D_{j^*} conditioned on being $< c$ (the former gets to draw values for *all* items conditioned on being $< c$ and take the largest, rather than just from item j^*). Taking these two claims together, we see that the distribution of $\min\{\max_j \{v_j\}, \max_j \{w_j\}\}$ stochastically dominates V' . \square

PROOF OF PROPOSITION B.2. Simply chain the following inequalities together. The first line follows from Lemma B.4. The second line follows from Lemma B.7. The third follows from Observation 5. The final follows from Corollary B.9 (and the fact that $\text{SYMREV} \leq \text{REV}$ always).

$$\begin{aligned} \text{SYMREV}(D) &\leq \text{SYMREV}(D^{(1)}) \\ &\leq \frac{1}{1-\varepsilon} \cdot \text{SYMREV}(D^{(2)}) + \frac{1}{\varepsilon(1-\varepsilon)} \mathbb{E}[V'] \\ &\leq \frac{1}{1-\varepsilon} \cdot \text{SYMREV}(D') + \frac{1}{\varepsilon(1-\varepsilon)} \mathbb{E}[V'] \\ &\leq \frac{1}{1-\varepsilon} \cdot \text{REV}(D') + \frac{2}{\varepsilon(1-\varepsilon)} \text{BREV}(D). \end{aligned}$$

\square

B.2 Upper Bounding $\text{REV}(D')$.

The main task of this section is to upper bound $\text{REV}(D')$. We provide below a proof based on tools used in [32]. For readers interested in a proof based on the duality framework of [9] (in the case where D'_0 is regular), see Appendix C. The ‘‘Marginal Mechanism’’ lemma is the only one we’ll use from prior work. Below, D_S denotes the marginals of distribution D onto items S , and the distribution $D_S|\vec{v}_S$ denotes the distribution of \vec{v}_S , assuming that \vec{v} is drawn from D , and conditioned on \vec{v}_S .

LEMMA B.10. [‘‘Marginal Mechanism’’ [10, 32]] *Let S, \bar{S} partition the items in $[m]$. Then for all (possibly correlated) distributions D :*

$$\text{REV}(D) \leq \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \in S} v_j \right] + \mathbb{E}_{\vec{v}_S \leftarrow D_S} [\text{REV}(D_{-\bar{S}}|\vec{v}_S)].$$

As an immediate corollary, we obtain the following bound on our modified distribution D' :

COROLLARY B.11. *For any modified distribution D' :*

$$\text{REV}(D') \leq \mathbb{E}_{\vec{v} \leftarrow D'} \left[\sum_{j \neq 0} v_j \right] + \mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|\vec{v}_{-0})].$$

Our remaining task is to upper bound the two terms on the right-hand side of Corollary B.11. We begin with the term $\mathbb{E}_{\vec{v} \leftarrow D'} [\sum_{j \neq 0} v_j]$. Our outline is similar to the bound of the ‘‘NON-FAVORITE’’ term in [9], with technical modifications to avoid upper bounding any terms with SREV (which suffices for their target bound, but not for us, as SREV is not symmetric). Upper bounding the term $\mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|\vec{v}_{-0})]$ is novel to our setting, and we complete this afterwards.

Our first step is to break the term into three parts:

$$\mathbb{E}_{\vec{v} \leftarrow D'} \left[\sum_{j \neq 0} v_j \right] = \mathbb{E}_{\vec{v} \leftarrow D'} \left[\sum_{j \notin \{0, j^*\}} v_j \cdot I(v_j \geq \text{SSREV}(D)) \right] \quad (\text{TAIL})$$

$$+ \mathbb{E}_{\vec{v} \leftarrow D'} \left[v_{j^*} \cdot I(v_{j^*} \geq \text{SSREV}(D)) \right] \quad (\text{SPECIAL})$$

$$+ \mathbb{E}_{\vec{v} \leftarrow D'} \left[\sum_{i \neq 0} v_i \cdot I(v_i < \text{SSREV}(D)) \right] \quad (\text{CORE}).$$

Indeed, the terms TAIL and CORE are quite similar to those in [9] (our cutoff of $\text{SSREV}(D)$ differs from their choice of $\text{SREV}(D)$), and the term SPECIAL is new for our approach.

LEMMA B.12. $\text{TAIL} \leq \text{SSREV}(D)$.

PROOF. Observe that the random variable $v_j \cdot I(v_j \geq \text{SSREV}(D)) \cdot I(j \notin \{0, j^*\})$ can alternatively be drawn by first drawing $v_j \leftarrow D_j$, and then independently drawing v_{-j} from D_{-j} . If v_j is the largest among these values, then in fact j is j^* . Otherwise, $j \notin \{0, j^*\}$, and we can just check whether $v_j \geq \text{SSREV}(D)$. Importantly, observe that each of these checks is independent. Therefore, we can write:

$$\begin{aligned} \text{TAIL} &= \mathbb{E}_{\vec{v} \leftarrow D'} \left[\sum_{j \notin \{0, j^*\}} v_j \cdot I(v_j \geq \text{BREV}(D)) \right] \\ &= \sum_{j > 0} \mathbb{E}_{v_j \leftarrow D_j} \left[v_j \cdot \Pr[\exists j' \neq j, v_{j'} > v_j] \cdot I(v_j \geq \text{BREV}(D)) \right] \\ &\leq \sum_{j > 0} \mathbb{E}_{v_j \leftarrow D_j} \left[\text{SSREV}(D) \cdot I(v_j \geq \text{SSREV}(D)) \right] \\ &\leq \text{SSREV}(D) \cdot \sum_{j > 0} \Pr_{v_j \leftarrow D_j} [v_j \geq \text{SSREV}(D)] \\ &\leq \text{SSREV}(D). \end{aligned}$$

Above, the first line is just the definition of TAIL. The second line observes that $j \notin \{0, j^*\}$ exactly when there is some other j' with larger value. The third line observes that for any p , we can set price p on all items. If any $j' \neq j$ has $v_{j'} > p$, then that item sells. v_j is one such possible price p , and all $v_{j'}$ are drawn independently of it, so $\text{SSREV}(D)$ is at least as good as the revenue of selling *only* items $\neq j$ at the same price v_j . The fourth line rewrites the expected value of an indicator variable as a probability. The final line again just observes that this is exactly the revenue of setting price $\text{SSREV}(D)$ on each item. \square

Next, we upper bound SPECIAL.

LEMMA B.13. $\text{SPECIAL} \leq 2\text{BREV}(D)$.

PROOF. Observe that $v_{j^*} \geq v_{j^*} \cdot I(v_{j^*} \geq \text{SSREV}(D))$. Observe further that Corollary B.9 *precisely* upper bounds $\mathbb{E}[v_{j^*}]$ (as V' is exactly this random variable). The lemma statement now follows immediately from these two facts. \square

Finally, we upper bound CORE. Recall that [5, 9] bound similar terms by connecting the variance of $\sum_j v_j$ to SREV . However, there is no (apparent) connection between the variance of $\sum_j v_j$ and SSREV , so instead we must take a different approach and appeal only to the fact that each v_j is an independent random variable guaranteed to be $\leq \text{SSREV}(D)$.

PROPOSITION B.14. $\max\{9\text{SSREV}(D), 4\text{BREV}(D)\} \geq \text{CORE}$.

PROOF. First, observe that $\sum_{j \neq 0} v_j \cdot I(v_j < \text{SSREV}(D))$ (when $\vec{v} \leftarrow D'$) is stochastically dominated by $\sum_j \min\{w_j, \text{SSREV}(D)\}$ (when $\vec{w} \leftarrow D$), which is stochastically dominated by $\sum_j w_j$ (when $\vec{w} \leftarrow D$). This means both that whenever $\sum_j \min\{w_j, \text{SSREV}(D)\} \geq p$, the buyer would purchase the grand bundle at price p , and also that $\text{CORE} \leq \mathbb{E}_{\vec{w} \leftarrow D}[\sum_j \min\{w_j, \text{SSREV}(D)\}]$. We analyze now this term, and refer to it as $C := \mathbb{E}_{\vec{w} \leftarrow D}[\sum_j \min\{w_j, \text{SSREV}(D)\}]$ for ease of notation.

Consider first the case that $9\text{SSREV}(D) \geq C \geq \text{CORE}$. Then clearly the proposition holds.

Now consider the case that $9\text{SSREV}(D) < C$. In this case, we have a sum of independent random variables, all bounded in $[0, \text{SSREV}(D)]$, whose expectation exceeds $9\text{SSREV}(D)$. We may then use a Chernoff bound to conclude that:¹⁴

$$\Pr_{\vec{w} \leftarrow D} \left[\sum_j w_j \geq C/3 \right] \geq 1 - e^{-(2/3)^2 \cdot (C/\text{SSREV}(D))/2} \geq 1 - e^{-2}.$$

In particular, this means that we can set price $C/3$ on the grand bundle, and it will sell with probability at least $1 - e^{-2}$, witnessing that $\text{BREV}(D) \geq (1 - e^{-2}) \cdot \text{CORE}/3 \geq \text{CORE}/4$. \square

This completes our analysis of the first term, and we now turn our attention to analyzing $\mathbb{E}_{\vec{v}_0 \leftarrow D'_0}[\text{REV}(D'_0 | \vec{v}_0)]$, and this part of the analysis is novel to our setting. To reason about this, we first need to understand the distribution $D'_0 | \vec{v}_0$.

LEMMA B.15. *For all fixed \vec{v}_0 , the distribution $D'_0 | \vec{v}_0$ is stochastically dominated by the distribution $D'_0 | (v_0 > \max_{j>0} \{v_j\})$.*

PROOF. We provide a proof in the case that each D_j is discrete. A proof when each D_j is continuous follows an identical outline, but with more tedious notation (essentially just replace probabilities with densities everywhere).

For a given v , let us first compute the probability of drawing v (this is where it is convenient to have discrete distributions, so that this is well-defined. Identical calculations would work with a PDF instead, but require more tedious notation). We get:

$$\begin{aligned} \Pr_{v_0 \leftarrow D'_0 | \vec{v}_0} [v_0 = v] &= \frac{\Pr_{\vec{w} \leftarrow D'} [w_0 = v \wedge \vec{w}_{-0} = \vec{v}_{-0}]}{\Pr_{\vec{w} \leftarrow D'} [\vec{w}_{-0} = \vec{v}_{-0}]} \\ &= \frac{\Pr_{w_0 \leftarrow D'_0} [w_0 = v] \cdot \prod_{j \geq 1} \Pr_{w_j \leftarrow D_j | w_j < v} [w_j = v_j]}{\Pr_{\vec{w} \leftarrow D'} [\vec{w}_{-0} = \vec{v}_{-0}]} \end{aligned}$$

The first line is just the definition of conditional probability. The second line observes that one way to draw from D' is to first draw w_0 , and then draw w_{-0} (which draws each w_j independently, conditioned on being $< w_0$).

Importantly, now observe that $\Pr_{v_0 \leftarrow D'_0 | \vec{v}_0} [v_0 = v]$ can be written as $\Pr_{w_0 \leftarrow D'_0} [w_0 = v] \cdot g(v, \vec{v}_{-0})$, where $g(v, \vec{v}_{-0})$ is monotone decreasing in v , for all \vec{v}_{-0} . That is, as v increases, the scaling factor for the probability of seeing v drawn from $D'_0 | \vec{v}_0$ decreases. This intuitively makes sense, as for any particular $v > v'$, there are more \vec{v}_{-0} that are consistent with $v_0 = v$ than $v_0 = v'$, so we should be (at least weakly) more likely to have $v_0 = v'$ than $v_0 = v$, conditioned on v_{-0} .

Now, let's do the same calculations for $\Pr_{v_0 \leftarrow D'_0 | (v_0 > \max_j \{v_j\})} [v_0 = v]$. In this case, it's easy to see that for all $v_0 > \max_j \{v_j\}$, that $\Pr_{v_0 \leftarrow D'_0 | (v_0 > \max_j \{v_j\})} [v_0 = v] = \Pr_{w_0 \leftarrow D'_0} [w_0 = v] / \Pr_{w_0 \leftarrow D'_0} [w_0 > \max_j \{v_j\}]$. Importantly, this is equal to $\Pr_{w_0 \leftarrow D'_0} [w_0 = v]$ times a constant scaling factor.

Taken together, this means that the two distributions have the same support (they are both supported on all $v > \max_{j \geq 1} \{v_j\}$), and also that the ratio $\Pr_{v_0 \leftarrow D'_0 | \vec{v}_0} [v_0 = v] / \Pr_{v_0 \leftarrow D'_0 | (v_0 > \max_j \{v_j\})} [v_0 = v]$ is monotonically decreasing. In particular, this implies that $1 - F_{D'_0 | \vec{v}_0}(v) \leq 1 - F_{D'_0 | (v_0 > \max_j \{v_j\})}(v)$

¹⁴After normalizing so that all variables fall in $[0, 1]$, this is the standard multiplicative Chernoff bound.

(because the two distributions have the same support, and the total probability mass in each is one), and therefore $D'_0|\vec{v}_{-0}$ is stochastically dominated by $D'_0|(v_0 > \max_j \{v_j\})$. \square

Now, we wish to upper bound the term $\mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|\vec{v}_{-0})]$ using Lemma B.15.

PROPOSITION B.16. $\mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|\vec{v}_{-0})] \leq 2\text{BREV}(D)$.

PROOF. First, we show that $\mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|\vec{v}_{-0})] \leq \mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|v_0 > \max_{j \geq 1} \{v_j\})]$ using Lemma B.15. This follows because $D'_0|(v_0 > \max_{j \geq 1} \{v_j\})$ stochastically dominates $D'_0|\vec{v}_{-0}$ for all \vec{v}_{-0} (and this implies that $\text{REV}(D'_0|v_0 > \max_{j \geq 1} \{v_j\}) \geq \text{REV}(D'_0|\vec{v}_{-0})$ pointwise).¹⁵

Now, consider the following procedure to draw \vec{v}_{-0} from D'_{-0} . First, draw $w_0 \leftarrow D'_0$, and then draw \vec{v}_{-0} from D'_{-0} conditioned on w_0 . Recall first that w_0 can be drawn by first drawing $w_j \leftarrow D_j$ independently, and then setting $w_0 := \max_j \{w_j\}$. Recall also that, conditioned on w_0 , \vec{v}_{-0} is distributed according to a product distribution, where each v_j is drawn from D_j conditioned on $v_j < w_0$.

But now, observe that $\max_{j \geq 1} \{v_j\}$, where each v_j is drawn independently from D_j conditioned on $v_j < w_0$, is *exactly* the distribution v_0 drawn from D'_0 , but conditioned on $v_0 < w_0$ (because v_0 is exactly set to $\max_j \{v_j\}$). Therefore, we have that in fact: $\mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|v_0 > \max_{j \geq 1} \{v_j\})] = \mathbb{E}_{w_0 \leftarrow D'_0, w'_0 \leftarrow D'_0 | w'_0 < w_0} [\text{REV}(D'_0|v_0 > w'_0)]$.

Finally, consider instead letting w_0 be the maximum of two independent draws from D'_0 , instead of a single draw. This only makes w_0 bigger (in a stochastically dominating manner), which therefore increases w'_0 (again, in a stochastically dominating manner), which therefore increases $\text{REV}(D'_0|v_0 > w'_0)$. Therefore, we can conclude the following chain of equalities:

$$\begin{aligned} \mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|\vec{v}_{-0})] &\leq \mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0|(v_0 > \max_{j>0} \{v_j\})] \\ &= \mathbb{E}_{w_0 \leftarrow D'_0, w'_0 \leftarrow D'_0 | w'_0 < w_0} [\text{REV}(D'_0|v_0 > w'_0)] \\ &\leq \mathbb{E}_{w_0 \leftarrow D'_0, w'_0 \leftarrow D'_0} [\text{REV}(D'_0|v_0 > \min\{w_0, w'_0\})]. \end{aligned}$$

The first line follows from Lemma B.15. The second follows from the previous paragraphs discussing another valid procedure to draw $\max_{j>0} \{v_j\}$. The third follows by the immediately preceding paragraph.

Now, we just need to bound $\mathbb{E}_{w_0 \leftarrow D'_0, w'_0 \leftarrow D'_0} [\text{REV}(D'_0|v_0 > \min\{w_0, w'_0\})]$. We claim, however, that this is simply the expected revenue of *some* truthful auction (in fact, Ronen's auction [45]) for two bidders drawn from D'_0 independently. Indeed, consider the auction that solicits a bid b_1, b_2 from the two bidders, and then sets the optimal reserve for $D'_0|v_0 > b_i$ to bidder $3 - i$. This auction is clearly truthful, and sells the item only to the highest bidder (if at all). In addition, the revenue achieved from the highest bidder is *exactly* $\text{REV}(D'_0|v_0 > \min\{b_1, b_2\})$. Therefore, $\mathbb{E}_{w_0 \leftarrow D'_0, w'_0 \leftarrow D'_0} [\text{REV}(D'_0|v_0 > \min\{w_0, w'_0\})]$ is the expected revenue of a truthful auction for two bidders drawn from D'_0 , and is therefore at most twice $\text{REV}(D'_0)$. Finally, observe that $\text{REV}(D'_0) \leq \text{BREV}(D)$ (this follows as $\sum_j v_j$ stochastically dominates $\max_j \{v_j\}$ when $\vec{v} \leftarrow D$. The former is the value of a bidder drawn from D for the grand bundle, the latter is the value of a bidder drawn from D'_0). So we may further conclude that:

$$\mathbb{E}_{w_0 \leftarrow D'_0, w'_0 \leftarrow D'_0} [\text{REV}(D'_0|v_0 > \min\{w_0, w'_0\})] \leq 2\text{BREV}(D),$$

which completes the proof. \square

And now, we can wrap up the proof of Theorem 3.2.

¹⁵This follows from revenue monotonicity. To quickly see this, consider setting the revenue-optimal price for $D'_0|\vec{v}_{-0}$ on distribution $D'_0|(v_0 > \max_{j \geq 1} \{v_j\})$. The item sells with (weakly) greater probability, so the revenue is larger.

PROOF OF THEOREM 3.2. We get the following chain of inequalities, by Proposition B.2 (line one), Corollary B.11 (line two), and Lemma B.12 and Lemma B.13 and Proposition B.14 and Proposition B.16 (line three)

$$\begin{aligned}
\text{SYMREV}(D) &\leq \frac{1}{1-\varepsilon} \cdot \text{REV}(D') + \frac{2}{\varepsilon(1-\varepsilon)} \cdot \text{BREv}(D) \\
&\leq \frac{1}{1-\varepsilon} \cdot \left(\text{TAIL} + \text{SPECIAL} + \text{CORE} + \mathbb{E}_{\vec{v}_{-0} \leftarrow D'_{-0}} [\text{REV}(D'_0 | \vec{v}_{-0})] \right) + \frac{2}{\varepsilon(1-\varepsilon)} \cdot \text{BREv}(D) \\
&\leq \frac{1}{1-\varepsilon} \cdot (\text{SSREV}(D) + 2\text{BREv}(D) + \max\{9\text{SSREV}(D), 4\text{BREv}(D)\}) \\
&\quad + \frac{2}{1-\varepsilon} \cdot \text{BREv}(D) + \frac{2}{\varepsilon(1-\varepsilon)} \cdot \text{BREv}(D)
\end{aligned}$$

Setting $\varepsilon = 1/2$ we get:

$$\text{SYMREV}(D) \leq 20\text{SSREV}(D) + 24\text{BREv}(D).$$

□

C ALTERNATIVE PROOF VIA DUALITY

In Section 3, we proved our main result, namely that BREv is a constant factor approximation to SYMREV . In this section, we provide an alternative proof based on the [9] duality framework (for the case when the distribution of the bidder's maximum value for the items is regular). We then also overview a naive attempt at applying their framework (using their ‘‘canonical flow’’), which helps provide intuition for our analysis.

THEOREM C.1. *Let D be any additive single-bidder distribution. Let D' be their modified distribution that ensures that the maximum-value item is always in the same coordinate (0) as per Definition B.1. If D'_0 is regular,*

$$\text{SYMREV}(D) \leq \max\{18\text{SSREV}(D), 4\text{BREv}(D)\} + 2\text{SSREV}(D) + 10\text{BREv}(D)$$

Finite-Support vs. Continuous Distributions. Our proof of Theorem C.1 in this section makes use of the [9] framework, which is designed for finite-support distributions. However, results in their framework also apply to continuous distributions (see [9, Section 2]). In our analysis, we will often compare two values drawn independently, and it will be convenient to break ties arbitrarily, but consistently. To be extra formal, whenever a single value v is drawn from a distribution labeled with a parameter j , we will formally draw a tuple (v, j) . Whenever we compare two values, we say that $(v, j) < (v', j')$ if $v < v'$, or $v = v'$ and $j < j'$. Importantly, note that this parameter is attached at the time v is drawn from D_j . This could be achieved alternatively by picking a sufficiently small $\varepsilon \rightarrow 0$ and adding $j\varepsilon$ to any value drawn from a distribution with parameter j . Importantly, observe that this means that any two values drawn from different distributions are never equal.

C.1 Brief Overview of [9]

We first provide the minimum preliminaries necessary to apply the [9] framework. The main concept in their framework is that of a *flow*, and a corresponding *virtual valuation function*. We specialize all definitions/theorems from their work to our setting (and refer the reader to [9] for the general statements).

Definition C.2 (Flow). For a single-bidder distribution D , supported on some set $T \subseteq \mathbb{R}_{\geq 0}^m$ of possible types, a *flow* $\lambda(\cdot, \cdot)$ defines a variable $\lambda(t, \vec{v}') \geq 0$ for all $t, t' \in T$.

A flow is *useful* if for all $\vec{v} \in T$: $f(\vec{v}) + \sum_{\vec{v}'} \lambda(\vec{v}', \vec{v}) \geq \sum_{\vec{v}'} \lambda(\vec{v}, \vec{v}')$.

To help parse this definition, one can interpret $\lambda(\vec{v}, \vec{v}')$ as “the flow going from t to t' .” A flow is then useful if the flow into \vec{v} (including additional flow $f(\vec{v})$ from some “super source”) is at least the flow out of \vec{v} , for all \vec{v} .

Definition C.3 (Virtual Valuation Function). For a useful flow λ , define the *virtual valuation function* $\Phi^\lambda(\cdot)$ associated with λ to satisfy the following for all j :

$$\Phi_j^\lambda(\vec{v}) := v_j - \frac{\sum_{\vec{v}'} (v'_j - v_j) \lambda(\vec{v}', \vec{v})}{f(\vec{v})}.$$

The main result of [9] is a framework for upper bounding terms like $\text{REV}(D)$, $\text{SYMREV}(D)$, etc. via useful flows.

THEOREM C.4 ([9]). *Let λ be a useful flow. Then for all truthful mechanisms M :*

$$\text{REV}_M(D) \leq \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \in [m]} \pi_j^M(\vec{v}) \cdot \Phi_j^\lambda(\vec{v}) \right].$$

Typically, applications of this theorem proceed by finding a useful flow, and then upper bounding $\mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \in [m]} \pi_j^M(\vec{v}) \cdot \Phi_j^\lambda(\vec{v}) \right]$ for any *not necessarily truthful* mechanism M . That is, a good choice of λ absolves the analysis from directly reasoning about incentives. In our setting, we will hope to leverage that M is symmetric.

Finally, a key concept that we leverage is the Myerson virtual value theory for single-item auctions. Below, for a discrete single-variable distribution D , $s^D(v)$ denotes the *successor* of v in the support of D : the minimum value v' in the support of D such that $v' > v$. If no such v' exists, define $s^D(v) := v + 1$.

Definition C.5 (Myersonian virtual value [42]). For a single-dimensional distribution D , we denote by $\varphi^D(v) := v - \frac{(1-F^D(v)) \cdot (s^D(v)-v)}{f^D(v)}$.

THEOREM C.6 ([42]). *For any truthful n -bidder single-item mechanism M and any single-item distribution D , its expected revenue equals its expected Myersonian virtual welfare. That is:*

$$\text{REV}_M(D) = \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{i \in [n]} x_i(\vec{v}) \cdot \varphi^{D_i}(v_i) \right].$$

C.2 A Failed Attempt: the CDW Canonical Flow

In order to get a sense of the technical challenges, we begin by overviewing a natural first attempt. [9] gives a canonical flow for a single bidder with independent items, that induces a particular virtual valuation function. For simplicity of notation below, we let R_j denote the set of m -dimensional vectors for which $v_j > v_{j'}$ for all $j' \neq j$ (tie-breaking as defined in Section 2).

THEOREM C.7 ([9]). *When D is an additive single-bidder distribution over any number of independent items, there exists a useful flow λ such that $\Phi_j^\lambda(\vec{v}) = v_j \cdot \mathbb{I}(\vec{v} \notin R_j) + \varphi^{D_j}(v_j) \cdot \mathbb{I}(\vec{v} \in R_j)$.*

In particular, this implies the following upper bound on $\text{SYMREV}(D)$:

COROLLARY C.8. *When D is an additive single-bidder distribution over any number of independent items:*

$$\begin{aligned} \text{SYMREV}(D) &\leq \max_{M, M \text{ is symmetric}} \left\{ \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \in [m]} \mathbb{I}(\vec{v} \in R_j) \cdot \varphi^{D_j}(v_j) \cdot x_j(\vec{v}) \right] \right\} && (\text{SYM}SINGLE) \\ &+ \max_{M, M \text{ is symmetric}} \left\{ \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \in [m]} \mathbb{I}(\vec{v} \notin R_j) \cdot v_j \cdot x_j(\vec{v}) \right] \right\} && (\text{SYM}NON-DOMINANT) \end{aligned}$$

Importantly, note above that we are not restricting M to be truthful, just that it has a symmetric allocation rule. The [9] proof that $6 \max\{\text{SREV}(D), \text{BREV}(D)\} \geq \text{REV}(D)$ indeed proceeds through this flow. They instead observe that $\text{REV}(D) \leq \text{SINGLE} + \text{NON-DOMINANT}$ (SINGLE and NON-DOMINANT are defined analogously to $\text{SYM}SINGLE$ and $\text{SYM}NON-DOMINANT$ above, but without restricting M to be symmetric). They establish that $\text{SINGLE} \leq \text{SREV}(D)$ by appealing to Myersonian virtual value theory, and also that $\text{NON-DOMINANT} \leq 5 \max\{\text{SREV}(D), \text{BREV}(D)\}$.

In our setting, it is not too hard to establish that $\text{SYM}NON-DOMINANT = O(\max\{\text{SSREV}(D), \text{BREV}(D)\})$, using ideas similar to those in [9]. The barrier to a successful analysis turns out to be $\text{SYM}SINGLE$. Consider the following example:

Example C.9. Let D_i sample 2^i with probability 2^{-i} , and 0 with probability $1 - 2^{-i}$. [31] establishes that for this distribution $\text{BREV}(D) \leq 4$. On the other hand, a symmetric allocation rule achieves $\text{SYM}SINGLE = \Omega(m)$.

To see this, consider the rule that awards item j if and only if $v_j > 0$. This rule is clearly symmetric. Also, for all j , with probability at least 2^{-j-1} : $\vec{v} \in R_j$ and $v_j = 2^j$. $\varphi^{D_j}(2^j) = 2^j$, and therefore we get that $\text{SYM}SINGLE \geq \sum_{j=1}^m 2^{-j-1} \cdot 2^j = m/2$.

Example C.9 does not necessarily mean that the CDW canonical flow is doomed, but it does mean that in order to possibly leverage it for analysis, we would need to directly invoke truthfulness of M (which in some sense defeats the purpose of the CDW framework).

Intuitively, the problem is that $\text{SYM}SINGLE$ is too big. The CDW canonical flow can't possibly shrink $\Phi_j^\lambda(\vec{v})$ to be smaller than $\varphi^{D_j}(v_j)$, but this is already too big. Our new flow directly leverages the fact that M must be symmetric in order to yield smaller virtual valuation functions.

C.3 A New Flow, and Proof of Theorem C.1

To prove Theorem C.1, we'll consider the same modified distribution D' from Definition B.1. The main task of this section is to upper bound $\text{REV}(D')$, and this is where we design a new flow. It is still inspired by the CDW canonical flow, modified to accommodate that 0 is always the "dominant item," and that D' is not a product distribution. Our flow will apply to any distribution "like D' ," which we formally define below.

Definition C.10 (Dominant Item). Item j is a *dominant item* for a distribution D over \mathbb{R}^k if:

- With probability 1, $v_j > v_\ell$ for all $\ell \neq j$ (breaking ties as defined in Section 2).
- For all $v_j > v'_j$, let \vec{v}_{-j} denote a draw from the distribution D_{-j} conditioned on v_j , and let \vec{v}'_{-j} denote a draw from the distribution D_{-j} conditioned on v'_j . Then \vec{v}_{-j} and \vec{v}'_{-j} can be coupled so that with probability one: $v_\ell \geq v'_\ell$ for all $\ell \neq j$.

Intuitively, item j is dominant when it is always the dominant item, and other item values are positively correlated (in a particular way) with v_j . We first observe that item 0 is always dominant for D' .

LEMMA C.11. *Item 0 is dominant for any modified distribution.*

PROOF OF LEMMA C.11. The proof of Lemma C.11 follows from two steps. First, we give another class of distributions that all have dominant items, and then show that every modified distribution is in this class.

Definition C.12 (Conditionally Independent). A distribution D over \mathbb{R}^k is *conditionally independent with respect to item j* if:

- With probability 1, $v_j > v_\ell$ for all $\ell \neq j$ (breaking ties as discussed in Section 2).
- For all v_j , the distribution of \vec{v}_{-j} conditioned on v_j is a product distribution.
- For all $v_j > v'_j$, and all $\ell \neq j$, the distribution of v_ℓ conditioned on v_j stochastically dominates the distribution of v_ℓ conditioned on v'_j .

LEMMA C.13. *If D is conditionally independent with respect to item j , then item j is dominant for D .*

PROOF. Because the distribution of v_ℓ conditioned on v_j stochastically dominates the distribution of v_ℓ conditioned on v'_j , they can be coupled so that $v_\ell \geq v'_\ell$ with probability one. Because all items are independent, conditioned on v_j , simply take the product of these couplings to form a coupling of \vec{v}_{-j} and \vec{v}'_{-j} . \square

Now, we simply observe that all modified distributions are conditionally independent with respect to item 0. To see this, observe first that indeed item 0 is always the dominant item. Next, observe that conditioned on v_0 , all items v_j are drawn from D_j conditioned on $v_j < v_0$ independently. This proves bullet two. Finally, observe that if $v_0 > v'_0$, then the distribution D_j conditioned on $v_j < v_0$ stochastically dominates the distribution D_j conditioned on $v_j < v'_0$, satisfying bullet three. \square

Now, we design a flow that applies to any distribution with a dominant item. We present the flow in two steps. We'll use the following notation:

- V_j denotes the possible values the bidder might have for item j (observe that this is the support of D_j — we phrase it this way to emphasize that D is not necessarily a product distribution), and use the notation $\text{PRED}^V(v)$ to denote the maximum element in V that is $< v$. If no such element exists, let $\text{PRED}^V(v) := \perp$.
- F_j still denotes the CDF for D_j (which is still the marginal distribution for v_j , even though D is not necessarily a product distribution), and $f_j(v)$ still denotes $\Pr_{v_j \leftarrow D_j}[v_j = v]$.
- When D has a dominant item j , let $g^{v_j}(\vec{v}_{-j})$ be the coupling promised between v_j and $\text{PRED}^{V_j}(v_j)$.¹⁶ That is:
 - When \vec{v}_{-j} is drawn from D_{-j} , conditioned on v_j , $g^{v_j}(\vec{v}_{-j})$ is a proper sample from D_{-j} , conditioned on $\text{PRED}^{V_j}(v_j)$.
 - For all $\ell \neq j$, $\vec{v}_{-j}, v_\ell \geq (g^{v_j}(\vec{v}_{-j}))_\ell$.

Definition C.14 (Symmetric Canonical Flow). Let j be a dominant item for distribution D . Let λ^* denote the flow such that for all \vec{v} :

- If $\text{PRED}^{V_j}(v_j) = \perp$, $\lambda^*(\vec{v}, \vec{v}') = 0$ for all \vec{v}' .
- Otherwise, $\lambda^*(\vec{v}, ((\text{PRED}^{V_j}(v_j); g^{v_j}(\vec{v}_{-j}))) = f(\vec{v}) \cdot \frac{(1 - F_j(\text{PRED}^{V_j}(v_j)))}{f_j(v_j)}$,¹⁷ and $\lambda^*(\vec{v}, \vec{v}') = 0$ for all other \vec{v}' .

¹⁶Because distributions have finite support, this coupling may be randomized/fractional.

¹⁷If $g^{v_j}(\vec{v}_{-j})$ is randomized/fractional, the flow is sent proportionally.

Intuitively, flow goes from types with strictly higher v_j to types with strictly lower v_j , and from weakly higher v_ℓ to weakly lower v_ℓ for all $\ell \neq j$. These are the two key properties used in the analysis. First, let's confirm that λ^* is indeed a useful flow.

PROPOSITION C.15. *When j is a dominant item for D , λ^* is a useful flow.*

PROOF. We've already explicitly defined how much flow leaves \vec{v} , for all \vec{v} , so we just need to compute how much flow enters.

To compute this, refer to *level* v as the set of all \vec{v} such that $v_j = v$. For simplicity of notation for the rest of this proof, for a fixed v_j let v'_j be such that $\text{PRED}^{V_j}(v'_j) = v_j$. Observe that for all \vec{v} the same level, that the term $\frac{1-F_j(\text{PRED}^{V_j}(v_j))}{f_j(v_j)}$ is the same. This means that the total flow *into* \vec{v} is just the total mass of types that send flow into \vec{v} , scaled by $\frac{1-F_j(v_j)}{f_j(v'_j)}$.

In order for \vec{v}' to send flow into \vec{v} , it must be that $v_j = \text{PRED}^{V_j}(v'_j)$, and also that $g^{v'_j}(\vec{v}'_{-j}) = \vec{v}_{-j}$. Recall that $g^{v'_j}(\cdot)$ is a coupling between D_{-j} conditioned on v'_j and D_{-j} conditioned on v_j . Observe also that the probability that \vec{v}_{-j} is drawn from D_{-j} conditioned on v_j is exactly $\frac{f(\vec{v})}{f_j(v_j)}$. Therefore, it must be the case that the total mass of types that send flow into \vec{v} is also a $\frac{f(\vec{v})}{f_j(v_j)}$ fraction of the $f_j(v'_j)$ mass of nodes at level v'_j . Therefore, we conclude that the total mass of types sending flow into \vec{v} is $\frac{f(\vec{v}) \cdot f_j(v'_j)}{f_j(v_j)}$, and therefore the total flow into \vec{v} is $\frac{f(\vec{v}) \cdot f_j(v'_j)}{f_j(v_j)} \cdot \frac{1-F_j(v_j)}{f_j(v'_j)} = f(\vec{v}) \cdot \frac{1-F_j(\vec{v})}{f_j(\vec{v})}$.

Finally, we observe simply that the total flow in, plus $f(\vec{v})$, is equal to the total flow out:

$$f(\vec{v}) + f(\vec{v}) \cdot \frac{1-F_j(v_j)}{f_j(v_j)} = f(\vec{v}) \cdot \frac{1-F_j(v_j) + f_j(v_j)}{f_j(v_j)} = f(\vec{v}) \cdot \frac{1-F_j(\text{PRED}^{V_j}(v_j))}{f_j(v_j)}.$$

□

Next, we compute the associated virtual valuations for λ^* .

PROPOSITION C.16. *Let j be a dominant item for D . Then for all \vec{v} , and all $\ell \neq j$:*

- $\Phi_j^{\lambda^*}(\vec{v}) = \varphi^{D_j}(v_j)$.
- $\Phi_\ell^{\lambda^*}(\vec{v}) \leq v_\ell$.

PROOF. To see the first bullet, observe that for all \vec{v}' , if $\lambda^*(\vec{v}', \vec{v}) > 0$, then $v'_j = s^{D_j}(v_j)$. Because the total flow into \vec{v} is $f(\vec{v}) \cdot \frac{1-F_j(v_j)}{f_j(v_j)}$, this means that we have:

$$\Phi_j^{\lambda^*}(\vec{v}) = v_j - \frac{f(\vec{v}) \cdot (1-F_j(v_j)) \cdot (s^{D_j}(v_j) - v_j)}{f_j(v_j) \cdot f(\vec{v})} = \frac{(1-F_j(v_j)) \cdot (s^{D_j}(v_j) - v_j)}{f_j(v_j)} = \varphi^{D_j}(v_j).$$

To see the second bullet, simply observe that for all \vec{v}' and all $\ell \neq j$: if $\lambda^*(\vec{v}', \vec{v}) > 0$, then $v'_\ell \geq v_\ell$ (by definition of dominant item). This means that $v'_\ell - v_\ell \geq 0$, and therefore $\Phi_\ell^{\lambda^*}(\vec{v}) \leq v_\ell$. □

This completes the analysis of our symmetric canonical flow. We observe that this canonical flow induces virtual values such that: (1) the virtual value of each bidder for all of their non-dominant items is at most their value, while (2) the virtual value for their dominant item zero is exactly the Myersonian virtual value of the value of the dominant item. We can now rewrite the following as:

$$\begin{aligned}
\mathbb{E}_{\vec{v} \leftarrow D}[p(\vec{v})] &\leq \mathbb{E}_{\vec{v} \leftarrow D}[\pi(\vec{v}) \cdot \Phi^\lambda(\vec{v})] \\
&\leq \mathbb{E}_{\vec{v} \leftarrow D}\left[\sum_{j \geq 0} \pi_j(\vec{v}) \cdot (v_j \cdot \mathbb{I}[j \neq 0] + \varphi_{D'_0}(v_0) \cdot \mathbb{I}[j = 0])\right] \\
&\leq \mathbb{E}_{\vec{v} \leftarrow D}\left[\sum_{j \geq 1} \pi_j(\vec{v}) \cdot v_j\right] + \mathbb{E}_{\vec{v} \leftarrow D}[\pi_0(\vec{v}) \cdot \varphi_{D'_0}(v_0)]
\end{aligned}$$

C.4 An improved bound for Regular distributions

For each type of the bidder, we have defined a flow. We have shown that this flow induces a virtual valuation function such that the bidder's virtual value for all non-dominant items is at most their value for those items, and their virtual value for their dominant item is the Myersonian virtual value of the highest value item. For the non-dominant items, the social welfare is a trivial upper bound for revenue. To get a better bound with our flow, we take the dominant item with value v_0 that contributes the most to the welfare, and turn its virtual value into its Myersonian virtual value. Recall that for any additive single-bidder distribution D , their modified distribution D' ensures that the maximum-value item is always in the same coordinate (0) as per Definition B.1.

OBSERVATION 6. $\mathbb{E}_{\vec{v} \leftarrow D}[\pi_0(\vec{v}) \cdot \varphi_{D'_0}(v_0)] \leq \mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}]$

This holds since $\max\{\varphi_{D'_0}(v_0), 0\}$ stochastically dominates $\varphi_{D'_0}(v_0)$. Let the allocation probability be 1 when φ is positive, and 0 when it's negative. The observation follows.

OBSERVATION 7. *When D'_0 is regular, selling only the favorite item gets expected revenue of $\mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}]$.*

PROOF. Myerson's theorem (expected revenue equals expected virtual surplus) states that maximizing revenue in expectation is equivalent to maximizing virtual surplus in expectation. When D'_0 is regular, consider the auction that sets a price of $p := \varphi_{D'_0}^{-1}(0)$, and awards the item only to the bidder when their value for the item (drawn from D'_0) exceeds this. Then the expected virtual surplus of this auction is exactly $\mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}]$.

In our setting, consider the auction which sets a price of p , and allows the bidder to pick *any single item* to purchase at price p . The buyer with values \vec{v} will purchase an item if and only if $v_0 \geq p$, and v_0 is drawn from D'_0 . Therefore, selling only the favorite item gets expected revenue $\mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}]$. □

LEMMA C.17. *When D'_0 is regular, $\mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}] \leq \text{BREV}(D)$.*

PROOF. By Observation 7, we know that $\mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}]$ can be achieved by setting some price p and letting the buyer purchase a single item at price p . Consider instead setting the same price p but letting the buyer purchase the grand bundle of all items. Then the expected revenue of this mechanism can only be larger, as $\sum_j v_j$ stochastically dominates $\max_j\{v_j\}$. □

This completes our analysis of the dominant item. For the non-dominant items, we need to further decompose into two terms we call CORE and TAIL. Define $r = \text{SSREV}(D)$.

$$\begin{aligned}
\mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} \pi_j(\vec{v}) \cdot v_j \right] &\leq \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} v_j \right] \\
&= \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} v_j \cdot \mathbb{I}(v_j \geq \text{SSREV}(D)) \right] && \text{(TAIL)} \\
&+ \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} v_j \cdot \mathbb{I}(v_j < \text{SSREV}(D)) \right] && \text{(CORE)}.
\end{aligned}$$

The term TAIL above captures contributions to the bound coming from non-dominant items whose value is at least $\text{SSREV}(D)$.

LEMMA C.18. $\text{TAIL} \leq \text{SSREV}(D)$.

PROOF.

$$\begin{aligned}
\text{TAIL} &= \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} v_j \cdot \mathbb{I}(v_j \geq \text{BREV}(D)) \right] \\
&\leq \sum_{j > 0} \mathbb{E}_{v_j \leftarrow D_j} [\text{SSREV}(D) \cdot \mathbb{I}(v_j \geq \text{SSREV}(D))] \\
&\leq \text{SSREV}(D) \cdot \sum_{j > 0} \Pr_{v_j \leftarrow D_j} [v_j \geq \text{SSREV}(D)] \\
&\leq \max_q \left\{ q \cdot \sum_{j > 0} \Pr_{v_j \leftarrow D_j} [v_j \geq q] \right\} \\
&= \text{SSREV}(D).
\end{aligned}$$

□

Above, the first line is just the definition of TAIL. The second line observes that for any p , we can set price p on all items. Any item with value at least p sells. In particular, for $p = v_j$, we find that $\text{SSREV}(D)$ is at least as good as the revenue of selling only item $j \neq 0$ at the same price v_j . The third line rewrites the expected value of an indicator variable as a probability. The fourth line again just observes that this is exactly the revenue of setting price $\text{SSREV}(D)$ on each item. The last line follows from the definition of $\text{SSREV}(D)$. This completes our analysis of the term TAIL.

The term CORE captures contributions to the bound coming from non-dominant items whose value is at most $\text{SSREV}(D)$. The idea is that CORE is the expected sum of independent random variables, each supported on $[0, \text{SSREV}(D)]$. So maybe $\text{CORE} \leq 12\text{SSREV}(D)$, which makes this bound trivial. Or, maybe $\text{CORE} > 12\text{SSREV}(D)$, in which case this should concentrate. In the latter case, we should expect to have $\text{CORE} \leq 3\text{BREV}(D)$, which also works out.

Definition C.19 (Chernoff bound). Let X_1, X_2, \dots, X_n be random variables such that $0 \leq X_i \leq 1$ for all i . Let $X = \sum_i X_i$ and let $\mu = E(X)$. Then, for all $0 < \delta < 1$:

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2}}$$

OBSERVATION 8. $\Pr_{\vec{v} \leftarrow D} [\sum_j v_j > (1 - \delta)\text{CORE}] \geq e^{-\frac{\delta^2}{6}}$ when $\text{CORE} \geq 12\text{SSREV}(D)$

PROOF. We observe that CORE is the expected sum of independent random variables supported on $[0, \text{SSREV}(D)]$. We define $\mu = \frac{\text{CORE}}{\text{SSREV}(D)} = E[\sum_j \frac{Y_j}{\text{SSREV}(D)}] = E[\sum_j X_j]$. This gives $E[X] = \mu = \frac{\text{CORE}}{\text{SSREV}(D)}$ and $E[Y] = E[\sum_j v_j] = \text{CORE} = \mu \cdot \text{SSREV}(D)$. In the case where

$\text{CORE} \leq k\text{SSREV}(D)$, for some constant k , our bound is trivially true. For the other case where $\text{CORE} > k\text{SSREV}(D)$, we get $\mu \geq k$ and Chernoff bounds give: $\Pr[X \leq (1 - \delta)\mu] \leq e^{-\frac{\delta^2\mu}{2}}$

$$\begin{aligned} \Pr[X \leq (1 - \delta)\mu] &= \Pr[Y \leq \text{SSREV}(D) \cdot (1 - \delta)\mu] \\ &= \Pr[Y \leq (1 - \delta)\text{CORE}] \\ &= \Pr_{\vec{v} \leftarrow D} \left[\sum_j v_j \leq (1 - \delta)\text{CORE} \right] \end{aligned}$$

We now plug in the value $k = 12$ and we get $\Pr[\sum_j v_j \leq (1 - \delta)\text{CORE}] \leq e^{-\frac{\delta^2\mu}{2}} \leq e^{-\frac{\delta^2k}{2}} = e^{-\frac{\delta^2}{6}}$ which is the probability that the grand bundle sells at price $(1 - \delta)\text{CORE}$. Now we're ready to prove the lemma below. \square

LEMMA C.20. $\text{CORE} \leq \max\{12\text{SSREV}(D), 2.568\text{BREV}(D)\}$

PROOF. Consider the mechanism M that sets price $p = (1 - \delta)\text{CORE}$ on the grand bundle. Then:

$$\begin{aligned} \text{BREV} &= \max_q \{q \cdot \Pr_{\vec{v} \leftarrow D} [\sum_j v_j \geq q]\} \\ &\geq (1 - \delta)\text{CORE} \cdot \Pr_{\vec{v} \leftarrow D} [\sum_j v_j > (1 - \delta)\text{CORE}] \\ &\geq ((1 - \delta)\text{CORE}) \cdot (1 - e^{-\frac{\delta^2k}{2}}) \\ &\geq \frac{\text{CORE}}{a}, \text{ where } a \text{ is some constant that we solve for below.} \end{aligned}$$

We now maximize the factor $(1 - \delta) \cdot (1 - e^{-\frac{\delta^2k}{2}})$ over all possible values of δ and $k < 12$. This gives $\delta \approx 0.4805$. We plug these values ($k = 12$ and $\delta = 0.48$) back into the expression above to get $(1 - 0.48)(1 - e^{-\frac{0.48^2}{2}k}) \geq 0.3894$. We now solve for the constant $a \leq \frac{1}{0.3894}$. And so $a \approx 2.568$, which gives $\text{CORE} \leq 2.568\text{BREV}(D)$. In particular, this means that we can set price 0.62CORE on the grand bundle, and it will sell with probability at least 0.75, witnessing that $\text{BREV}(D) \geq \text{CORE}/2.568$. \square

LEMMA C.21. $\mathbb{E}_{\vec{v} \leftarrow D} [\sum_{j \geq 1} \pi_j(\vec{v}) \cdot v_j] \leq \max\{12\text{SSREV}(D), 2.568\text{BREV}(D)\} + \text{SSREV}(D)$

PROOF. We use the upper bounds above on (CORE) from Lemma C.20 and (TAIL) from Lemma C.18 to get:

$$\begin{aligned} \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \geq 1} \pi_j(\vec{v}) \cdot v_j \right] &= \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} \pi_j(\vec{v}) \cdot v_j \right] \\ &\leq \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \neq 0} v_j \right] \\ &\leq (\text{CORE}) + (\text{TAIL}) \\ &\leq \max\{12\text{SSREV}(D), 2.568\text{BREV}(D)\} + \text{SSREV}(D) \end{aligned}$$

\square

Recall that we showed in Corollary B.9 that $\mathbb{E}[V'] \leq 2\text{BREV}(D)$, where V' denotes the random variable distributed according to the marginal of v_0 . We are now ready to combine these inequalities to get a bound, and prove the following theorem.

LEMMA C.22. For a single additive bidder, the optimal symmetric revenue for the modified distribution D' is upper bounded by $\max\{12\text{SSREV}(D), 2.568\text{BREV}(D)\} + \text{BREV}(D) + \text{SSREV}(D)$.

PROOF. We combine Lemma C.21 and Lemma C.17, and we find that the optimal symmetric revenue $\text{REV}(D')$ of the modified distribution D' is upper bounded by $E_{\vec{v} \leftarrow D}[\sum_{j \geq 1} \pi_j(\vec{v}) \cdot v_j] + \mathbb{E}_{\vec{v} \leftarrow D}[\varphi_{D'_0}(v_0)]$ which is upper bounded by $(\text{BREV}(D)) + (\max\{12\text{SSREV}(D), 2.568\text{BREV}(D)\} + \text{SSREV}(D)) = \max\{12\text{SSREV}, 2.568\text{BREV}\} + \text{BREV}(D) + \text{SSREV}(D)$. \square

And now, we can wrap up the proof of Theorem C.1.

PROOF OF THEOREM C.1. Similar to the proof of Theorem 3.2 and making use of Proposition B.2, we combine lemma C.22 and Corollary B.9, and we upper bound the optimal symmetric revenue of the original distribution D below, using $\varepsilon = 0.319$. Simply chain the following inequalities together. The first line follows from Observation 5. The second line follows from Corollary B.9 (and the fact that $\text{SYMREV} \leq \text{REV}$ always). The third line follows from Lemma C.22. The fourth line follows from plugging in $\varepsilon = 0.319$.

$$\begin{aligned} \text{SYMREV}(D) &\leq \frac{1}{1-\varepsilon} \cdot \text{SYMREV}(D') + \frac{1}{\varepsilon(1-\varepsilon)} \mathbb{E}[V'] \\ &\leq \frac{1}{1-\varepsilon} \cdot \text{REV}(D') + \frac{1}{\varepsilon(1-\varepsilon)} \cdot 2\text{BREV}(D) \\ &\leq \frac{1}{1-\varepsilon} \cdot (\max\{12\text{SSREV}, 2.568\text{BREV}(D)\} + \text{BREV}(D) + \text{SSREV}(D)) + \frac{1}{\varepsilon(1-\varepsilon)} \cdot 2\text{BREV}(D) \\ &\leq \max\{18\text{SSREV}(D), 4\text{BREV}(D)\} + 2\text{SSREV} + 10\text{BREV}(D) \end{aligned}$$

\square

We quickly discuss the main barrier in extending this approach to a proof of Theorem 3.2. In order to analyze the term $\mathbb{E}_{\vec{v} \leftarrow D}[\max\{\varphi_{D'_0}(v_0), 0\}]$ when D'_0 is irregular, we would need to *iron* the distribution. This would adjust the flow, and be completely fine for this part of the analysis. However, this would also require us to add cycles of flow between types, *which could therefore increase the virtual values for non-favorite items beyond their values*. Therefore, we would no longer be able to upper bound $\mathbb{E}_{\vec{v} \leftarrow D}[\sum_{j \geq 1} \pi_j(\vec{v}) \cdot v_j]$.

D WHY WE NEED AN EXTRA ITEM

In this section, we explain why our approach adds an extra item by showing that merely permuting the pre-existing m item values is not enough. We define the following permuted distribution D' , show that its revenue is close to that of D , and then design a flow for D' . We first define our permuted distribution.

Definition D.1 (Permuted Distribution). Let D be an additive single-buyer distribution over m independent items. Define the permuted distribution D' to be the following distribution over m items.

- Draw v_j independently from D_j for all j .
- Let $j^* := \arg \max_j \{v_j\}$ (breaking ties lexicographically).
- Let $v'_1 := v_{j^*}$.
- Let $v'_{j^*} := v_1$.
- Let $v'_j := v_j$, for all $j \neq 1, j^*$.

Intuitively, D' ensures that the maximum-value item is always in the same coordinate (1), and this allows us to leverage a clean application of prior tools. We first need to argue that $\text{SYMREV}(D)$ is upper bounded by (an appropriate function of) $\text{REV}(D')$.

OBSERVATION 9. $\text{SYMREV}(D) = \text{SYMREV}(D') \leq \text{REV}(D')$

PROOF. Consider any symmetric mechanism M for D , and view it by its menu (that is, the list of (\vec{x}, p) it allows the buyer to purchase). Recall that because M is symmetric, that for all (\vec{x}, p) on the menu, $(\sigma(\vec{x}), p)$ is also on the menu for all item permutations σ .

In particular, the permutation σ' that permutes the 1 and j^* indices is also on the menu. And so for any \vec{v} , let (\vec{x}, p) denote their favorite option from the menu for M . This option achieves revenue $\vec{v} \cdot \vec{x} - p$. Then $(\sigma'(\vec{x}), p)$ is the favorite option for \vec{v}' and achieves the same revenue $\vec{v}' \cdot \sigma'(\vec{x}) - p = \sigma'^{-1}(\vec{v}') \cdot \vec{x} - p = \vec{v} \cdot \vec{x} - p$. The equality follows. The inequality follows from the fact that for any distribution D , $\text{SYMREV}(D) \leq \text{REV}(D)$ always. \square

We now try to upper bound $\text{REV}(D')$. We provide below a proof based on tools used in [32]. Below, D_S denotes the marginals of distribution D onto items S , and the distribution $D_{\bar{S}}|\vec{v}_S$ denotes the distribution of $\vec{v}_{\bar{S}}$, assuming that \vec{v} is drawn from D , and conditioned on \vec{v}_S . Recall Lemma B.10 which states that for any S, \bar{S} partition the items in $[m]$, and any (possibly correlated) distributions D , we have:

$$\text{REV}(D) \leq \mathbb{E}_{\vec{v} \leftarrow D} \left[\sum_{j \in S} v_j \right] + \mathbb{E}_{\vec{v}_S \leftarrow D_S} [\text{REV}(D_{-\bar{S}}|\vec{v}_S)].$$

Similarly to Corollary B.11, we obtain the bound below on our permuted distribution D' .

COROLLARY D.2.

$$\text{REV}(D') \leq \mathbb{E}_{\vec{v} \leftarrow D'} \left[\sum_{j \neq 1} v_j \right] + \mathbb{E}_{\vec{v}_{-1} \leftarrow D'_{-1}} [\text{REV}(D'_1|\vec{v}_{-1})].$$

Now we upper bound the two terms on the right-hand side of Corollary D.2. For the first term, the sum of values of the non-favorite items can be bounded in the same way as before, except now we don't also have to worry about the extra item. For the second term, we try to bound the conditional revenue from the favorite item in a similar way to that of Section B.2. We focus on analyzing $\mathbb{E}_{\vec{v}_{-1} \leftarrow D'_{-1}} [\text{REV}(D'_1|\vec{v}_{-1})]$, and we will show that this term blows up for this permuted distribution.

PROPOSITION D.3. *There exists an instance of D for which $\mathbb{E}_{\vec{v}_{-1} \leftarrow D'_{-1}} [\text{REV}(D'_1|\vec{v}_{-1})] \geq \Omega(m)B\text{REV}(D)$.*

Example D.4. Consider the example for which each item i is drawn from D_i which samples 2^i w.p. 2^{-i} , and 2^{-i} with probability $1 - 2^{-i}$. When we condition on the non-favorite items, we actually know exactly what the value of the favorite item is, and we can get the full welfare as revenue. So the conditional revenue from the favorite item is too large compared to our distribution.

In particular, consider any vector $\vec{v}_{-1} \leftarrow D'_{-1}$. Then unless all v_i s have values 2^{-i} (which happens with probability at most $1/2$), we know that $v_1^* = 2^{j^*}$. And so, we know that $v_1^* = 2^{j^*}$ with probability at least $1/2$. Let mechanism M charge price 2^{j^*} on this item. Then, for each possible value of j^* , when $v_{j^*} = 2^{j^*}$ is the highest value, M generates a revenue of at least $2^{j^*} \cdot 1/2$. We compute the expected revenue of this mechanism, and we find that:

$$\mathbb{E}_{M, \vec{v}_{-1} \leftarrow D'_1} [\text{REV}(D'_1 | \vec{v}_{-1})] \geq \frac{1}{2} \sum_{j^* \neq 1} 2^{j^*} \cdot 2^{-j^*} = \frac{m-1}{2}$$

Recall that [31] showed that $\text{BREV}(D)$ achieves constant revenue for this particular distribution. We conclude that for this example, $\mathbb{E}_{\vec{v}_{-1} \leftarrow D'_1} [\text{REV}(D'_1 | \vec{v}_{-1})]$ is $\Omega(m)\text{BREV}(D)$, which is too large a bound.

Intuitively, we know that adding the extra item doesn't actually lower the revenue by that much, but it lowers the bound we get from our approach significantly: adding an extra item drawn from the same distribution as the favorite item, allows us to hide what the favorite value is, and therefore lowers the revenue bound.

E SSREV IS A LOG-REVENUE APPROXIMATION OF BREV

In this section we continue our comparison of BREV to SSREV . We prove our aforementioned Theorem 3.6 that states that SSREV is a log factor approximation to BREV . Recall that in our setting, BREV corresponds to the optimal revenue achieved by a mechanism which ignores demographic data entirely, while SSREV corresponds to the optimal revenue achieved by a mechanism which sets the same price to display an ad to each demographic, and allows each advertiser to choose which demographic views to purchase.

Using a similar approach to Appendix F in [5], we prove Theorem 3.6, namely that for any distribution D for a single buyer and m items (possibly correlated), $\text{BREV}(D) \leq 5 \log(m)\text{SSREV}(D)$.

Definition E.1. We say that an m -dimensional distribution D is a point-mass in sum distribution if there exists a p such that when \vec{v} is sampled from D , $\sum_j v_j = p$ with probability 1.

Definition E.2. An m -dimensional distribution D is symmetric if all marginals D_j are the same.

LEMMA E.3. Any m -dimensional distribution D has point-mass in sum m -dimensional distribution D' s.t. $\frac{\text{BREV}(D')}{\text{SSREV}(D')} \geq \frac{\text{BREV}(D)}{\text{SSREV}(D)}$

PROOF. Let's pick any D with PDF F , and let the optimal BREV price be p_B . Let $\Pr[\sum_j v_j > p_B] = q$. Then $\text{BREV}(D) = p_B \cdot q$. We define D'' with PDF F'' to be the distribution that is D transformed into a point-mass in sum distribution without decreasing $\frac{\text{BREV}(D)}{\text{SSREV}(D)}$ as follows: if $\sum_j v_j > p_B$, make it so $\sum_j v_j = p_B$, else set all $v_j = 0$. In particular, consider any p , then any item that is less than p under D is also less than p under D'' . In particular, the consumer is still willing to pay p_B w.p. q which means that $\text{BREV}(D'')$ is at least $\text{BREV}(D)$. Additionally, $F''(x) \leq F(x)$ for all x , and so values are lowered in a stochastically dominating way which means that $\text{SSREV}(D'')$ is at most $\text{SSREV}(D)$. We now combine these two inequalities to get:

$$\frac{\text{BREV}(D'')}{\text{SSREV}(D'')} \geq \frac{\text{BREV}(D)}{\text{SSREV}(D)}$$

Next, we define D' to be the distribution that is D'' conditioned on $\sum_j v_j = p_B$. D'' samples from D' w.p. q , else sets all values to 0. D' is a point-mass in sum distribution. We note that $\text{BREV}(D'') = q \cdot p_B$ which is a q fraction of $\text{BREV}(D')$. On the other hand, whatever price is set for each item sells w.p. exactly q times the probability it sells when the consumer is drawn from D' , and so $\text{SSREV}(D'')$ is also a q fraction of $\text{SSREV}(D')$. We now combine these inequalities to get:

$$\frac{\text{BREV}(D')}{\text{SSREV}(D')} = \frac{\text{BREV}(D'')}{\text{SSREV}(D'')} \geq \frac{\text{BREV}(D)}{\text{SSREV}(D)}$$

□

LEMMA E.4. Any m -dimensional distribution D , has a symmetric m -dimensional distribution D' s.t. $\frac{BREV(D')}{SSREV(D')} \geq \frac{BREV(D)}{SSREV(D)}$.

PROOF. Define distribution D' to be the following transformation of distribution D : sample $\vec{v} \sim D$, then randomly permute its components to form \vec{v}' . In particular, we have $BREV(D) = BREV(D')$ and $SSREV(D') = SSREV(D)$. We let D_j denote the j^{th} marginal of D and D'_k denote the k^{th} marginal of D' . We let v_j denote a sample from D_j , and v'_k a sample from D'_k , and so D'_k samples from each D_j w.p. $\frac{1}{m}$. Observe that $SSREV(D') = SREV(D') = \sum_k \max_p \{p \cdot \Pr[v'_k \geq p]\}$. Additionally, each D'_k samples each D_j w.p. $\frac{1}{m}$, so we achieve $\Pr[v'_k \geq p] = \sum_j \frac{1}{m} \Pr[v_j > p]$. We then get $SSREV(D') = SREV(D') = \sum_k \max_p \{p \sum_j \frac{1}{m} \Pr[v_j > p]\} = \max_p \{p \sum_j \Pr[v_j > p]\} = SSREV(D)$. We now combine these inequalities to get:

$$\frac{BREV(D')}{SSREV(D')} \geq \frac{BREV(D)}{SSREV(D)}$$

□

LEMMA E.5. Let D be any symmetric point-mass in sum distribution. Then $BREV(D) \leq 5 \log(m) SSREV(D)$

PROOF. Without loss of generality, scale D down to $SSREV(D) = SREV(D) = m$, and let $p = \text{Val}(D)$. We now want to see how large $\text{Val}(D)$ can be subject to these constraints. By symmetric point-mass constraint, each D_j is supported on $[0, p]$. And since $SREV(D) = m$ and D is symmetric, each D_j has expected revenue 1. This gives:

$$\text{Val}(D_j) = \int_0^p \Pr[v_j > x] dx \leq \int_0^1 1 dx + \int_1^p \frac{1}{x} dx = 1 + \log(p)$$

We know that $\text{Val}(D) = p$, and we know that $\text{Val}(D) = \sum_j \text{Val}(D_j) \leq m(1 + \log(p))$. Plugging those values back in gives $p \leq m + m \log(p)$ which simplifies to $p \leq 5m \log(m)$. We can now observe that $BREV(D) \leq 5 \log(m) SSREV(D)$ □

In our last step, we recall that for any distribution D , $SSREV(D) \geq \frac{1}{m} SREV(D)$, so if $SREV(D) \geq \alpha(m) BREV(D)$, then $SSREV(D) \geq \frac{\alpha(m)}{m} BREV(D)$. We now have the lemma below.

LEMMA E.6 (SSREV IS A $f(\alpha)$ APPROXIMATION OF BREV). Let $SREV = \alpha(m) BREV$ where $\alpha(m)$ can take values between $1/\Theta(\log(m))$ and $\Theta(m)$. And let $SSREV \geq f(\alpha(m)) BREV$. Then $f(\alpha)$ is always the better of $1/\Theta(\log(m))$ and $\Theta(\alpha(m)/m)$.

We observe that unless $\alpha(m)$ is really large (i.e. $\gg m^{1-\epsilon} \gg m/\log(m)$), it's not possible to achieve better than a $\log(m)$ approximation. We will also show examples where the bound $SSREV \leq f(\alpha) BREV$ is tight even when $SREV$ achieves better revenue than $BREV$, which shows that you can't get a good approximation. In other words, even when $SREV$ is better than $BREV$, $SSREV$ can still do much worse than $BREV$.

F TIGHT BOUNDS BETWEEN BREV AND SSREV

In this section, we first provide two examples of D illustrating that both bounds between $BREV(D)$ and $SSREV(D)$ are tight. Then, we merge these two examples into a canonical example of D to offer revenue guarantees of $SSREV(D)$ with respect to $BREV(D)$ using revenue guarantees of $SSREV(D)$ with respect to $SREV(D)$.

Example F.1 (BREV optimal example). Consider the following example where we define D^k to be the single buyer distribution with k values, where each v_j is drawn i.i.d. from an Equal-Revenue curve for all j . The Equal-Revenue distribution has CDF $F(x) = 0$ for $x \leq 1$, and $F(x) = 1 - 1/x$

for $x \geq 1$. That is, $\Pr[v_j \geq x] = 1/x$ for all $x \geq 1$ and all j . Here, the expected revenue obtained on each item by posting any price $p \geq 1$ is always one. And so, selling each item j separately at a price of $p_j \geq 1$ gets revenue one per item, for a total revenue of k for all items. In other words $\text{SREV}(D^k) = k$. Since any price achieves the same revenue, we observe that setting the same price $p \geq 1$ on each item also achieves revenue k . In other words $\text{SSREV}(D^k) = k$. On the other hand, bundling together achieves revenue $\Theta(k \log(k))$ as shown in [32]. We can conclude that SSREV being a log approximation of BREV is tight.

Example F.2 (SSREV optimal example). Consider the following example, similar to Example C.9 in the Appendix, where we define D^l to be the single buyer distribution with l values, where each v_j is 2^j with probability 2^{-j} and 0 with probability $1 - 2^{-j}$ independently. Then the optimal way to sell each item separately is to set a price 2^j on item j , which gives an expected revenue of one for each item, for a total revenue of l for all items. In other words $\text{SREV}(D^l) = l$. If we sell each item at the same price $p \in (2^j, 2^{j+1}]$, then the buyer will only purchase non-zero valued items indexed above j . This happens with probability at most $\sum_i 2^i < 2^j$. And so the revenue is $p \cdot 2^j$ which is at most $2^{j+1} \cdot 2^j = 2$. Then any price p generates revenue at most 2. In other words $\text{SSREV}(D^l) \leq 2$. Similarly, if we bundle the items together, then the buyer will only purchase if they have a non-zero value for some item $> j$, and so bundling together generates revenue at most 2. In other words $\text{BREV}(D^l) \leq 2$. We can conclude that BREV being a constant approximation of SSREV is tight.

Example F.3 (SSREV canonical example). Finally, consider the following example where we define D^m to be the single buyer distribution with m values, l of which are such that each v_j is either 2^j with probability 2^{-j} or 0 with probability $1 - 2^{-j}$ independently (we call this set S_l), and k of which are drawn i.i.d. from an Equal-Revenue curve (we call this set S_k). Then the optimal way to sell each item separately is to set a price 2^j on item j in S_l , which achieves expected revenue l , and to set any price $p \geq 1$ on each item in S_k , which achieves expected revenue k . In other words, $\text{SREV}(D^m) = l + k = m$. On the other hand, bundling together achieves expected revenue $\Theta(k \log(k))$ from the items in S_k and constant expected revenue from the items in S_l . In other words, $\text{BREV}(D^m)$ achieves revenue $\Theta(k \log(k))$. Finally, setting any same price p on each item achieves revenue k from the items in S_k and at most revenue 2 from the items in S_l . In other words $\text{SSREV}(D^m) \leq k + 2$.

We observe that in this example, $\text{SSREV}(D^m)$ is a k/m approximation of $\text{SREV}(D^m)$ and a $\log(k)$ approximation of $\text{BREV}(D^m)$. We now find an expression for k in terms of m .

Let $\text{SREV}(D^m) = \alpha(m)\text{BREV}(D^m)$, where $\alpha(m)$ can take values between $1/\Theta(\log(m))$ and $\Theta(m)$. And let $\text{SSREV}(D^m) \geq f(\alpha)\text{BREV}(D^m)$, where $f(\alpha)$ is a function of $\alpha(m)$. We now find k in terms of m to find tight bounds for $f(\alpha)$ and show that it's always the better of $1/\Theta(\log(m))$ and $\Theta(\alpha(m)/m)$.

Consider distribution D^m defined above in Example F.3 where $\alpha(m) = \text{SREV}(D^m)/\text{BREV}(D^m) = m/k \log(k)$. We rearrange the equality to get $k \log(k) = m/\alpha(m)$, which gives $k \sim \frac{m/\alpha(m)}{\log(m/\alpha(m))}$. Similarly, we know that $f(\alpha) \leq \text{SSREV}(D^m)/\text{BREV}(D^m) = k/k \log(k) = 1/\log(k)$. Plugging in the k approximation from above, we get $f(\alpha) = 1/\log(\frac{m/\alpha(m)}{\log(m/\alpha(m))})$.

Next, recall that $\alpha(m)$ can take values between $1/\Theta(\log(m))$ and $\Theta(m)$. When $\alpha(m)$ is on the order of $o(m)$, we achieve $f(\alpha) \sim 1/\log(\frac{m/\alpha(m)}{\log(m/\alpha(m))})$ which is on the order of $1/\Theta(\log(m/\alpha(m)))$. On the other hand, for values of $\alpha(m)$ closer to $\Theta(m)$, we get $f(\alpha)$ on the order of $\Theta(\alpha/m)$.

Finally, we double check this with the previous examples above, namely when $k = m$, which gives $D^m = D^k$ as in Example F.1, and $\alpha(m) = 1/\log(m)$ and $f(\alpha) \sim 1/\Theta(\log(m))$. This shows that when $\text{BREV}(D^m)$ is a $\log(m)$ factor of $\text{SSREV}(D^m)$ then $\text{SSREV}(D^m)$ and $\text{SREV}(D^m)$ achieve similar revenues. Similarly, when $k = 0$, we have $D^m = D^l$ as in Example F.2, and $\alpha(m) = m$ and

$f(\alpha) \sim \Theta(1)$. This shows that when $\text{BREV}(D^m)$ is a $1/m$ factor of $\text{SREV}(D^m)$, then $\text{SSREV}(D^m)$ and $\text{BREV}(D^m)$ achieve similar revenues.

An interesting observation is that for any m -dimensional distribution D , $\text{SSREV}(D)$ can achieve at least $1/m$ revenue of $\text{SREV}(D)$ by just setting the price on each item to be the price of the item that brings the most revenue to $\text{SREV}(D)$. And since $\text{SREV}(D)$ achieves an $\alpha(m)$ fraction of $\text{BREV}(D)$, then $\text{SSREV}(D)$ achieves at least an $\alpha(m)/m$ fraction of $\text{BREV}(D)$. Additionally, we previously showed in Theorem 3.6 that $\text{SSREV}(D)$ is a log approximation of $\text{BREV}(D)$, and so we can conclude that $\text{SSREV}(D)$ always achieves the better of $\log(m)$ and $\alpha(m)/m$ factor of $\text{BREV}(D)$.

We refer to Appendix G below for other interesting examples, one of which highlights the differences between a fair non-symmetric mechanism and a fair symmetric mechanism.

G FAIR BUT NON-SYMMETRIC MECHANISM

In this section, we present an example of a fair but non-symmetric mechanism and compare its revenue to that of a symmetric mechanism.

Consider the following example with two items and one bidder. Define D_1 , the distribution of the value of item one to be equal to 4 with probability $1/2$, and 1 with probability ε , and 0 with probability $1/2 - \varepsilon$. Define D_2 , the distribution of the value of item two to be equal to 4 with probability ε , and 1 with probability $1 - 2\varepsilon$, and 0 with probability ε . Let $\vec{v} = (v_1, v_2)$ be the buyer values, and (π_1, π_2, p) be the tuple where π_1 and π_2 are the respective allocation probabilities of items one and two, and p the price paid by the bidder for those allocations. Consider the following mechanism M that takes in a vector of values (v_1, v_2) and outputs allocations and price (π_1, π_2, p) :

$$\begin{aligned} M(4, 4) &= (1, 1, 3.5) \\ M(1, 1) &= (0.5, 0.5, 1) \\ M(0, 0) &= (0, 0, 0) \end{aligned}$$

$$\begin{aligned} M(4, 1) &= (1, 1, 3.5) \\ M(1, 4) &= (0, 1, 3) \end{aligned}$$

$$\begin{aligned} M(4, 0) &= (1, 0, 3) \\ M(0, 4) &= (0, 1, 3) \end{aligned}$$

$$\begin{aligned} M(1, 0) &= (0, 0, 0) \\ M(0, 1) &= (0, 1, 1) \end{aligned}$$

We observe that mechanism M is fair but not symmetric. The idea here is that if the buyer has highly asymmetric values for two items, then the buyer can't set different prices where the low value item is sold at a low price and the high value items is sold at a higher price, since symmetry would not allow this. Fairness, however, does allow it.

Now we consider the following mechanism M' :

$$\begin{aligned} M'(4, 4) &= (1, 1, \underline{5}) \\ M'(1, 1) &= (\underline{0}, \underline{1}, 1) \end{aligned}$$

$$M'(0, 0) = (0, 0, 0)$$

$$M'(4, 1) = (1, 1, \underline{5})$$

$$M'(1, 4) = (0, 1, \underline{3})$$

$$M'(4, 0) = (1, 0, \underline{4})$$

$$M'(0, 4) = (0, 1, \underline{3})$$

$$M'(1, 0) = (0, 0, 0)$$

$$M'(0, 1) = (0, 1, 1)$$

We observe that mechanism M' is not symmetric nor fair. But it achieves more revenue. In particular, for the case of value vector $(1, 1)$, the auctioneer would ideally not give item one but gives item two and charge price $p = 1$. This gives allocation-price vector $(0, 1, 1)$ which is not a fair allocation. This unfairness can be fixed by changing it to uniformly at random with probability $1/2$ of getting each item. Additionally, symmetry adds even more constraints. In particular, for the case of values $(1, 4)$ and $(4, 1)$, the auctioneer can't charge different prices (e.g. 1 vs. $7/2$).