Selling Partially-Ordered Items: Exploring the Space between Single- and Multi-Dimensional Mechanism Design

Nikhil R. Devanur* Kira Goldner† Raghuvansh R. Saxena‡ Ariel Schvartzman§
S. Matthew Weinberg¶

June 15, 2018

Abstract

Most progress on auction design has been either in single-dimensional settings, such as selling a single item, or multi-dimensional, where the seller has multiple heterogenous items. With respect to revenue maximization, these settings sit at extreme ends of a spectrum: from simple and fully characterized (single-dimensional) to complex and nebulous (multi-dimensional).

In this paper, we identify a setting that sits in between these extremes. The items are *partially-ordered* in the following sense: a buyer has a particular item that we say is his "interest," and he is happy with any item that is higher in the partial order than his interest. The special case of this when the items are totally-ordered is equivalent to the the FedEx Problem [FGKK16]. This is equivalent to the class of single-minded valuations with known priors, where interests are bundles of items with the partial order defined by set inclusion.

We show formally that partially-ordered items lie in a space of their own, in between identical and heterogeneous items: there exist distributions over (value, interest) pairs for *three* partially-ordered items such that the menu complexity of the optimal mechanism is *unbounded*, yet for all distributions there exists an optimal mechanism of *finite* menu complexity. So this setting is vastly more complex than identical items (where the menu complexity is one), or the FedEx Problem [FGKK16] (where the menu complexity is at most 7, for three items), yet drastically more structured than heterogenous items (where the menu complexity can be uncountable [DDT15]).

Our approach is via Lagrangian duality on the continuous LP characterizing the optimal mechanism. We further apply our tools to conclude that the optimal mechanism is deterministic whenever the marginal value distribution for each interest satisfies declining marginal revenues. Finally, we apply our tools to conclude that optimal mechanisms for Multi-Unit Pricing (without a DMR assumption) can have unbounded menu complexity as well, thus resolving an open question from [DHP17].

*Microsoft Research (nikdev@microsoft.com)
†University of Washington (kgoldner@cs.washington.edu)
‡Princeton University (rrsaxena@princeton.edu)
§Princeton University (acohenca@princeton.edu)
¶Princeton University (smweinberg@princeton.edu). Supported by NSF CCF-1717899.

1 Introduction

Consider the problem of selling multiple items to a unit-demand buyer. The fundamental problem underlying much of mechanism design asks how the seller should maximize their revenue. If the items are all identical, then the setting is considered *single-dimensional.* In this case, seminal work of Myerson [1981] completely resolves this question with an exact characterization of the optimal mechanism. The optimal mechanism is a simple take-it-or-leave-it price, and the fact that there are multiple items versus just one is irrelevant. In contrast, if the items are heterogenous, then the setting is *multi-dimensional*, and unlike the single-dimensional setting, optimal mechanisms are no longer tractable in any sense, and numerous recent works identify various undesirable properties [Briest, Chawla, Kleinberg, and Weinberg, 2015; Hart and Nisan, 2013; Hart and Reny, 2012; Daskalakis, Deckelbaum, and Tzamos, 2013; Daskalakis et al., 2015].

Of particular relevance to this paper is the notion of *menu complexity*: the number of non-trivial options presented to the buyer. Through the lens of menu complexity, Myerson's seminal work shows that the optimal mechanism for single-dimensional settings has menu complexity 1, while Daskalakis et al. [2015] show that the optimal mechanism for the multi-dimensional setting might have *uncountable* menu complexity—this holds even for just two items, and even when the item values are drawn independently. This dichotomy serves as one fundamental difference between single-dimensional and multi-dimensional settings.

Very recently, work of Fiat et al. [2016] identifies a fascinating middle-ground. Imagine that the items are neither identical nor heterogeneous, but are instead varying qualities of the same item. To have an example in mind, imagine that you're shipping a package and the items are one-day, two-day, or three-day shipping. You obtain some value v for having your package shipped, but only if it arrives by your deadline (which is one, two, or three days from now). Viewed in the context of a unit-demand buyer, this means that the buyer will always have value v or 0 for every item, and the set of items which yield non-zero value is either $\{1\}$, $\{1, 2\}$, or $\{1, 2, 3\}$ (so we can think of the input as being a two-dimensional distribution over (value, deadline) pairs). In this "FedEx Problem," Fiat et al. show that the optimal mechanism has menu complexity at most $2^m - 1$ if there are m possible deadlines (and there exist example value distributions for which this is tight for every optimal mechanism [Saxena, Schvartzman, and Weinberg, 2017]).

In the context of this paper, we'll say that the FedEx problem has *totally-ordered* items: one-day shipping is at least as good as two-day shipping is at least as good as three-day shipping, and every buyer agrees. This paper explores settings with *partially-ordered* items, equivalent to single-minded valuations, and identifies a further fascinating realm between the totally-ordered and heterogeneous items regimes. To have an example in mind, imagine that a company offers internet, phone service, and cable TV. If the first item is internet, the second item is the bundle {internet, phone service}, and the third item is {internet, cable TV}, then it's clear that {internet, phone service} is at least as good as internet, and {internet, cable TV} is at least as good as internet, but {internet, phone service} and {internet, cable TV} are incomparable. You have a value, v, and are interested in either exclusively internet service, internet/phone service, or internet/cable. If you receive something at least as good as your interest, then you get value v, otherwise you get zero (so we can again think of the input distribution as being a two-dimensional distribution over (value, interest) pairs). More generally, one could imagine that a seller has m objects for sale and the "items" are subsets of objects. Then item S is at least as good as item T if $S \supseteq T$, but if neither S nor T contains the other, neither is better than the other. If the buyer is single-minded with a privately-known interest set S, then they get value v whenever they receive item S or a superset.

Observe that the following problem can also be interpreted as a partially-ordered setting. Suppose that each buyer has a publicly visible attribute which the seller can use to set a different price. E.g., the buyer could be a student, a senior, or general-admission, and the seller could set different prices for each category. Or, the buyer could be a "prime member" or a "non-prime member". However, buyers with certain attributes can disguise themselves as having other attributes, given by a partial order. For example, a prime member could disguise as a non-prime member, but not vice-versa. Then if item i is a movie ticket redeemable by anyone who can disguise themselves as having attribute i, the items are partially-ordered.

The main results of this paper are bounds (upper and lower) on the menu complexity for partially-ordered items. We show that even with just three items the menu complexity is *unbounded.* This implies that as soon as you have the simplest non-trivial (i.e., not totally-ordered) partial order, the optimal mechanism gets much more complicated. In particular, our lower bound implies that the optimal mechanism for single-minded valuations has unbounded menu-complexity.

On the other hand, we show that the menu complexity for the three item case is *always finite.* That is, for all M, there exists a distribution over (value, interest) pairs such that every optimal mechanism has menu complexity $\geq M$, but for any distribution, there exists an optimal menu of finite menu complexity. Contrast this with the single-dimensional setting (menu complexity 1), the totally-ordered setting (menu complexity 7), and the heterogeneous items setting (uncountable).

1.1 Techniques and Extensions

Like the FedEx problem, our problem can be formulated as a continuous[1] linear program with a strong dual[2], and both our upper and lower bound are achieved by analyzing dual solutions. As in Cai, Devanur, and Weinberg [2016], candidate dual solutions come with a nice interpretation as *virtual values*, which helps guide the analysis (but more on this in Section 3). To produce our lower bound (for all M, an example where every optimal mechanism has menu complexity $\geq M$), we need to accomplish a few tasks:

1. Pick a candidate distribution, primal $\vec{x}$ with menu complexity $\geq M$, and candidate dual $\vec{\lambda}$.
2. Prove that $\vec{x}$ and $\vec{\lambda}$ satisfy complementary slackness (CS). Therefore $\vec{x}$ (and $\vec{\lambda}$) is optimal.
3. Prove that any $\vec{y}$ that satisfies complementary slackness with $\vec{\lambda}$ has menu complexity $\geq M$.[3]

For our specific construction, steps two and three are "just linear algebra," while step one (as with any construction) requires some ingenuity to pick the right primal/dual pair.

Our primary technical contribution is novel methods for manipulating and optimizing dual variables beyond single-dimensional settings. Prior work touches on components of this: Fiat et al. [2016] reasoned about recovering the primal solution from the dual, but used a different way to reason about the dual. Devanur and Weinberg [2017] reinterpreted the FedEx solution via the lens of the Lagrangian relaxation, describing the structure of the solution via lack of ability to perform dual operations. In this work, we develop techniques for performing both components

[1] Note that in order for the claim "for any joint distribution, there exists an optimal menu of finite menu complexity" to possibly be interesting, we had better be studying continuous distributions and not distributions of finite support.

[2] Note that strong duality doesn't come for free in continuous duality, but holds for our setting [FGKK16].

[3] And therefore every optimal mechanism has menu complexity $\geq M$. This is because a primal and dual pair are optimal if and only if they satisfy complementary slackness. Because $\vec{\lambda}$ is optimal, every optimal primal must satisfy complementary slackness with it.

more in depth, reasoning about more complex dual operations, more complex primal recovery, and what the two combined together tell us about the structure of the optimal mechanism. It is clear from our results that these techniques are useful for gaining insight about the structure of optimal mechanisms, even when the setting is complex enough to require unbounded randomization.

We also wish to highlight one technical contribution of independent interest regarding step 1 (and postpone further details to Section 4). While searching for the right primal/dual pair is mathematically pretty interesting, the annoying part is reverse engineering from a candidate primal/dual pair to a candidate distribution for which they are a valid primal/dual pair, as the algebra that takes us from distributions to feasible duals is quite messy. On this front, we prove a "Master Theorem," stating (essentially) that every candidate dual is feasible for some input distribution (Theorem 2). This allows the user (of the theorem) to reason exclusively about primals and duals (the fun part), while letting the Master Theorem map the candidate pair back to an instance for which they are feasible. Often the feasible primal/dual pair is the only illuminating part of the input distribution anyway, so the Master Theorem formally separates the insightful analysis from the tedious part.

Of course, one should not expect this theorem to hold in general multi-dimensional settings (in particular, one key property that enables our Master Theorem is a "payment identity," which general multi-dimensional settings notoriously lack)—this is a further example of how our setting lies in-between single- and multi-dimensional, but the Master Theorem is quite generally applicable for problems in this intermediate range. In addition, because the Master Theorem takes care of guaranteeing that distributions corresponding to some dual will exist, this result also emphasizes the strength of reasoning about duals in similar settings.

For our upper bound (for every distribution, there exists an optimal primal with finite menu complexity), we need to at minimum find an optimal primal for every instance. For totally-ordered items, Fiat et al. [2016] found a closed form for the optimal primal/dual pair and prove complementary slackness directly. As will become clear to the reader, there's no hope of this (the fact that the menu complexity can be unbounded should be a good indicator). Instead, we do the following:

1. Define a class of "dual operations" which improve the quality of a dual solution.
2. Provide a "primal algorithm" that takes a dual and outputs a primal satisfying CS.
3. Prove that if the primal algorithm fails, then a dual operation can be performed on the dual. Conclude that every instance admits an optimal primal resulting from our algorithm.
4. Prove that every primal resulting from our algorithm has finite menu complexity.

This skeleton is similar to the approach used by Devanur and Weinberg [2017] for selling to a budget-constrained buyer, although the operations in step 1 and the algorithm in step 2 in our case are really themselves algorithms (and of course are tailored for our setting: there's no generic "dual operation" that works for all mechanism design problems). We'll again just highlight one interesting property of step 4 to give some intuition as to why the output menu can have unbounded but not infinite menu complexity. The algorithm produces a sequence of cutoffs $v_1 > v_2 > \ldots > v_i > \ldots$ such that all bidders interested in the worst item with value in $(v_{i+1}, v_i]$ will purchase the same outcome. It's actually possible for this sequence to be infinite, and therefore it initially seems like the algorithm might produce a menu of countably infinite size. But something magical happens: if the sequence of cutoffs is infinitely long *and* converges to some $\underline{v}$, then the algorithm can instead abort and just offer one non-trivial outcome to those interested in the worst item: receive the item for price $\underline{v}$. As all $v_i \geq 0$, the monotone convergence theorem implies that the sequence converges (because the sequence is decreasing and bounded). Therefore, either the sequence is finite (in which case our algorithm outputs a menu of finite size), or the algorithm can abort and offer an even simpler (finite) menu. The reader will have to wait for the technical sections to understand where these cutoffs come from, why aborting is possible, etc. But we hope this vignette gives some insight as to how menus can become unbounded but not infinite.

Beyond our main result (overviewed above), we prove two additional results using the same techniques. First, we apply the same techniques as in our lower bound to show that the menu complexity of the Multi-Unit Pricing problem Devanur et al. [2017] is also unbounded (Theorem 12). Multi-Unit Pricing is also a totally-ordered setting, where the items correspond to copies of a good (item one is one copy, item two is two copies, item three is three copies). The difference from FedEx is that if the buyer is interested in two copies but gets one, they get half their value (instead of zero). Second, we show that the optimal mechanism is deterministic for a number of simple partially-ordered settings when each marginal distribution (of values conditioned on interest) satisfies Declining Marginal Revenues (DMR).[4] This matches prior work for similar partially/totally-ordered settings [FGKK16; DW17; DHP17] and supports the notion that DMR is the "right" notion of regularity in these settings.

1.2 Summary and Roadmap

To recap: we analyze the revenue-optimal mechanisms for a three-item setting where the items are somewhere between identical and heterogeneous, which we call partially-ordered. We show that the menu complexity of optimal mechanisms in this setting is unbounded (Theorem 1), but always finite (Theorem 3). This implies that the menu complexity of optimal mechanisms in any partially-ordered setting (including single-minded valuations, etc.) is unbounded. Our "Master Theorem" (Theorem 2) is of independent interest for future work on mechanism design with totally- or partially-ordered items. We also use similar techniques to lower bound the menu complexity of optimal mechanisms for Multi-Unit Pricing (Theorem 12), and characterize optimal mechanisms for simple partially-ordered settings subject to DMR marginals (Theorem 13).

Immediately below, we overview the closest related work. In Section 3, we overview notation and preliminaries. Section 4 presents our lower bound, which also provides general insight to our upper bound. Section 5 presents our characterization of optimal mechanisms and a proof that they always have finite menu complexity. Due to space considerations, our auxiliary results appear in the appendices (Appendix E for Multi-Unit Pricing, and F for DMR).

2 Related Work

The most related line of works has already mostly been discussed. The FedEx Problem considers totally-ordered items (in our language), as does Multi-Unit Pricing and Budgets [FGKK16; DHP17; DW17]. The present paper is the first to consider partially-ordered items. In terms of techniques, we indeed draw on tools from prior work. All three prior works employ some form of duality. Our approach is most similar to that of Devanur and Weinberg [2017] in that we also perform "dual operations" rather than search for a closed form, but this is the extent of the similarities.

[4] A one-dimensional distribution F satisfies Declining Marginal Revenues if $v(1 - F(v))$ is concave.

Also related is a long line of work which aims to characterize optimal mechanisms beyond single-dimensional settings. Owing to the inherent complexity of mechanism design for heterogeneous items, results on this front necessarily consider restricted settings [Laffont, Maskin, and Rochet, 1987; Giannakopoulos and Koutsoupias, 2014; McAfee and McMillan, 1988; Daskalakis et al., 2013, 2015; Haghpanah and Hartline, 2015; Malakhov and Vohra, 2009; Devanur et al., 2017]. On this front, our work further pushes the limits where optimal mechanisms can be characterized without distributional assumptions.

Various forms of duality have also been used for multi-item mechanism design [Daskalakis et al., 2015; Giannakopoulos and Koutsoupias, 2015; Haghpanah and Hartline, 2015; Cai et al., 2016]. However, aside from the fact that these works all use duality, there is little technical similarity.

There is also a quickly growing body of work regarding the menu complexity of multi-item auctions. Much of this work focuses on settings with heterogeneous buyers [Briest et al., 2015; Hart and Nisan, 2013; Babaioff, Gonczarowski, and Nisan, 2017; Wang and Tang, 2014; Daskalakis et al., 2015]. Very recent work of [Saxena et al., 2017] considers the menu complexity of approximately optimal mechanisms for the FedEx Problem (for which [FGKK16] already characterized the menu complexity of exactly optimal mechanisms). On this front, our work places partially-ordered items (where the menu complexity is finite but unbounded) distinctly between totally-ordered items (where the menu complexity is bounded) [FGKK16], and heterogeneous items (uncountable) [Daskalakis et al., 2015]. Previously, no settings with this property were known.

3 Preliminaries

While this paper focuses on the three-item case, it's illustrative (and perhaps cleaner) to provide notation for general partially-ordered items. In general, there are m partially-ordered items. Item G can be better than, worse than, or incomparable to item G', and we'll use the relation $G \succ G'$ to denote that G is better than G'. We refer to the set of items as $\mathcal{G}$, and use $N^+(G)$ to denote the set of items $G' \in \mathcal{G}$ for which $G' \succ G$, but there is no G'' with $G' \succ G'' \succ G$ (i.e. the items "immediately better" than G, or the 1-out-neighborhood of G in a graphic representation). There is a single buyer with a (value, interest) pair (v, G), who receives value v if they are awarded an item $\succeq G$. An instance of the problem consists of a joint probability distribution over $[0, H] \times \mathcal{G}$, where H is the maximum possible value of any bidder for any item. We will use f to denote the density of this joint distribution, with $f_G(v)$ denoting the density at (v, G). We will also use $F_G(v)$ to denote $\int_0^v f_G(w)dw$, and q_G to denote the probability that the bidder's interest is G. Note that $F_G(H) = q_G < 1$, so $F_G(\cdot)$ is not the CDF of a distribution (although $F_G(\cdot)/q_G$ is the CDF of the marginal distribution of v conditioned on interest G).

We'll consider (w.l.o.g.) direct truthful mechanisms, where the bidder reports a (value, interest) pair and is awarded a (possibly randomized) item. For a direct mechanism, we'll define $a_G(v)$ to be the probability that item G is awarded to a bidder who reports (v, G), and $p_G(v)$ to be the expected payment charged. Then a buyer's utility for reporting any (v', G') where G' doesn't dominate G is $-p_{G'}(v')$, and the utility for reporting any (v', G') where G' dominates G is $v \cdot a_{G'}(v') - p_{G'}(v')$.

At this point, one can write a primal LP that maximizes expected revenue subject to incentive constraints, manipulate the LP, and consider a Lagrangian relaxation (and all of this is done in Fiat et al. [2016]; Devanur and Weinberg [2017]).

We note here some convenient properties among the highly-structured settings that fall in "the space between single- and multi-dimensional" that help to explain why duality techniques gain traction for producing general optimal characterizations. The following are commonalities of the partially-ordered setting, the totally-ordered setting of FedEx [FGKK16], the budget-constrained setting [DW17], and the multi-unit pricing setting [DHP17].

- The allocation rule "collapses": instead of an allocation probability for each item, the allocation can be written as a single probability of satisfaction.
- The number of incentive-compatibility (IC) constraints reduces.
- The payment identity applies.

Of course, these three properties are intertwined: the collapsible allocation rule is important to use the payment identity and to reduce the number of IC constraints.

In single-dimensional settings, all IC constraints are replaced by the payment identity, and there is one closed-form dual. In contrast, in "multi-dimensional" settings with heterogenous items, IC constraints exist between all pairs of m-parameter types; it is completely unclear which constraints might be tight and need attention when designing mechanisms. However, in these highly-structured settings, we do have the payment identity, which allows us to replace payment variables in the primal and work only with the allocation variables. In addition, we can significantly narrow down which IC constraints are important, but algorithmic reasoning is necessary to determine exactly which ones.

The following sections are for completeness for the reader not yet comfortable with Lagrangian duality; the experienced reader may wish to skip some subsections before Subsection 3.5. Additional preliminaries are located in Appendix A.

3.1 Formulating the Optimization Problem

As observed in Fiat et al. [2016], it is without loss of generality to only consider mechanisms that award bidders their declared item of interest with probability in $[0, 1]$, and all other items with probability 0.[5]

The "default" way to write the continuous LP characterizing the optimal mechanism would be to maximize $\sum_{G\in\mathcal{G}} \int_0^H f_G(v)p_G(v)dv$ (the expected revenue) such that everyone prefers to tell the truth than to report any other type. Also observed in Fiat et al. [2016] is that Myerson's payment identity holds in this setting as well, and any truthful mechanism must satisfy $p_G(v) = va_G(v) - \int_0^v a_G(w)dw$ (this also implies that the bidder's utility when truthfully reporting (v, G) is $u_G(v) = \int_0^v a_G(w)dw$). This allows us to drop the payment variables, and follow Myerson's analysis to recover:[6]

$$\mathbb{E}[\text{revenue}] = \sum_{G\in\mathcal{G}} \int_0^H f_G(v) \cdot p_G(v)dv = \sum_{G\in\mathcal{G}} \int_0^H f_G(v)a_G(v) \left(v - \frac{1 - F_G(v)}{f_G(v)}\right) dv$$

[5] To see this, observe that the bidder is just as happy to get nothing instead of an item that doesn't dominate their interest. See also that they are just as happy to get their interest item instead of any item that dominates it. It will also make this option no more attractive to any bidder considering misreporting. So starting from a truthful mechanism, modifying it to only award the item of declared interest or nothing cannot possibly violate truthfulness.

[6] For the familiar reader, this derivation is routine, so we omit it. The unfamiliar reader can refer to [Myerson, 1981; Hartline, 2013] for this derivation.

The experienced reader will notice that $v - \frac{1-F_G(v)}{f_G(v)}$ is exactly Myerson's virtual value for the conditional distribution $F_G(\cdot)/q_G$, which we'll denote by $\varphi_G(v)$. At this point, we still have a continuous LP with only allocation variables, but still lots of truthfulness constraints. Fiat et al. [2016] observe that many of these constraints are redundant, and in fact it suffices to only make sure that when the bidder has (value, interest) pair (v, G) they:

- Prefer to tell the truth rather than report any other (v', G). This is accomplished by constraining $a_G(\cdot)$ to be monotone non-decreasing (exactly as in the single-item setting).
- Prefer to tell the truth rather than report any other $(v, G' \in N^+(G))$. This is accomplished by constraining $\int_0^v a_G(w)dw \geq \int_0^v a_{G'}(w)dw$ (as the LHS denotes the utility of the buyer for reporting (v, G) and the RHS denote the utility of the buyer for reporting (v, G')).

All of these constraints together imply that (v, G) also does not prefer to report any other (v', G').[7] Below, we will now formulate the Primal LP and its Lagrangian relaxation. This derivation is not a new result, but important to understanding our approach. So we'll go through some of the steps to help provide some intuition for the reader, but omit any calculations and proofs.

3.2 The Primal

With the above discussion in mind, we can now formulate our primal continuous LP.

Variables: $a_G(v),\ \forall G \in \mathcal{G},\ v \in [0, H]$

Maximize $$\sum_{G \in \mathcal{G}} \int_0^H f_G(v) a_G(v) \varphi_G(v) dv$$

subject to

$$a'_G(v) \geq 0 \qquad \forall G \in \mathcal{G}\ \forall v \in [0, H] \text{ (dual variables } \lambda_G(v) \geq 0)$$

$$\int_0^v a_G(x)dx - \int_0^v a_{G'}(x)dx \geq 0 \qquad \forall G \in \mathcal{G},\ G' \in N^+(G)\ \forall v \in [0, H] \text{ (dual vars } \alpha_{G,G'}(v) \geq 0)$$

$$a_G(v) \in [0, 1] \qquad \forall G \in \mathcal{G},\ \forall v \in [0, H] \text{ (no dual variables)}$$

The first constraint requires that $a_G(\cdot)$ is monotone non-decreasing for all G. If an allocation rule is not monotone, it cannot possibly be part of a truthful mechanism. As discussed above, Myerson's payment identity combined with monotonicity guarantees that (v, G) will always prefer to report (v, G) instead of (v', G). The second constraint directly requires that the utility of (v, G) for reporting (v, G) is at least as high as for reporting (v, G') (also discussed above). The final constraint simply ensures that the allocation probabilities lie in $[0, 1]$.

3.3 The Lagrangian Dual and Virtual Values

The focus of this paper begins with the following Lagrangian relaxation. For a full derivation of how this came from the primal in subsection 3.2, see Appendix A.3. The primal variables are

[7] For example, if (v, G) prefers truthful reporting to reporting (v, G') where $G' \succ G$, and (v, G') prefers truthful reporting to reporting (v', G'), then since (v, G) gets the same utility for reporting (v, G') as type (v, G') does for truthfully reporting, (v, G) prefers truthful reporting to reporting (v', G').

$a_G(v)$ for all $G \in \mathcal{G}$, $v \in [0, H]$. The dual variables are $\lambda_G(v)$, $\alpha_{G,G'}(v)$.[8] The quality of a primal solution is measured by how well it solves the following Lagrangian relaxation induced by (λ, α). The quality of a dual solution is measured by the value of its induced Lagrangian relaxation. A dual is *better* if the value of its induced Lagrangian relaxation is *smaller*.

$$\begin{array}{ll} \text{Variables:} & a_G(v) \quad \forall G \in \mathcal{G},\ v \in [0, H] \\ \text{Maximize} & \sum_{G \in \mathcal{G}} \int_0^H f_G(v) \cdot a_G(v) \cdot \Phi_G^{\lambda,\alpha}(v) dv \\ \text{subject to} & a_G(v) \in [0, 1] \end{array}$$

where $\varphi_G(v) = v - \frac{1 - F_G(v)}{f_G(v)}$ and where $\Phi_G^{\lambda,\alpha}(v)$

$$:= \varphi_G(v) + \frac{1}{f_G(v)} \left[-\lambda'_G(v) + \sum_{G' \in N^+(G)} \int_v^H \alpha_{G,G'}(w) dw - \sum_{G' : G \in N^+(G')} \int_v^H \alpha_{G',G}(w) dw \right]. \quad (1)$$

Before continuing, let's parse the Lagrangian relaxation. The only remaining constraints are that $a_G(v) \in [0, 1]$, and the objective is a linear function of these variables. This immediately implies that the solution to this LP relaxation will set $a_G(v) = 1$ whenever $\Phi_G^{\lambda,\alpha}(v) > 0$, and $a_G(v) = 0$ whenever $\Phi_G^{\lambda,\alpha}(v) < 0$ (but we cannot draw immediate conclusions for $\Phi_G^{\lambda,\alpha}(v) = 0$). We'll refer to $\Phi_G^{\lambda,\alpha}(v)$ as the *virtual value* of (v, G) (which is different from Myerson's virtual value, $\varphi_G(v)$), and we'll also use $u_G(v) = \int_0^v a_v(w) dw$ to refer to the utility of (v, G).

After parsing the Lagrangian relaxation above, we can immediately conclude that the quality of a dual solution (λ, α) is exactly $\sum_{G \in \mathcal{G}} \int_0^H f_G(v) \cdot \max\{0, \Phi_G^{\lambda,\alpha}(v)\} dv$.

3.4 Quick Review of Complementary Slackness

For LPs with strong duals, a (primal, dual) pair is optimal if and only if the primal and dual satisfy complementary slackness. This implies, for instance, that if a dual (λ, α) satisfies complementary slackness with some primal, then another primal is optimal if and only if it satisfies complementary slackness with (λ, α). Let's review complementary slackness in our setting. A primal a and dual (λ, α) satisfy complementary slackness if and only if:[9]

1. (primal best response) $\Phi_G^{\lambda,\alpha}(v) > 0 \Rightarrow \alpha_G(v) = 1$. $\Phi_G^{\lambda,\alpha}(v) < 0 \Rightarrow \alpha_G(v) = 0$.
2. (dual best response) $\lambda_G(v) > 0 \Rightarrow a'_G(v) = 0$. $\alpha_{G,G'}(v) > 0 \Rightarrow \int_0^v (a_G(x) - a_{G'}(x)) dx = 0$.

3.5 Dual Terminology

Below we provide a list of terminology that we'll use throughout, depicted in Figures 1 and 2.

[8] For the experienced reader, the λ dual variables correspond to incentive constraints between types of the same interest but different value. The α dual variables correspond to incentive constraints between types of the same value but different interest.

[9] One can interpret these conditions as saying that the primal is an optimal solution to the Lagrangian relaxation, and the dual is the *worst* possible Lagrangian relaxation for the primal.

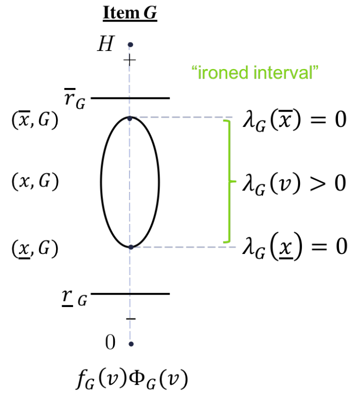

Figure 1: A pictorial interpretation of the dual variable $\lambda_G(v)$, in addition to the concepts of ironing, an ironed interval, the allocation in an ironed interval, and endpoints of the zero region. Complementary slackness dictates that if $\lambda_G(v) > 0$, then $a'_G(v) = 0$. That is, the allocation is constant in ironed intervals, and thus $a_G(x) = a_G(\underline{x})$.

- (Ironing) We say a type (x, G) is *ironed*, or that x is ironed in item G, if $\lambda_G(v) > 0$.
- (Ironed Intervals) For type (x, G), the ironed interval containing x in G is defined by the bottom end point $\underline{x}_G = \sup\{v \leq x \mid \lambda_G(v) = 0\}$ and the top end point $\bar{x}_G = \inf\{v \geq x \mid \lambda_G(v) = 0\}$.
- (Allocation in Ironed Intervals) Complementary slackness requires that if $\lambda_G(v) > 0$ then $a'_G(v) = 0$. Then for any x, an optimal allocation must satisfy that $a_G(x) = a_G(\underline{x}_G)$.
- (Endpoints of Zero Region) We define the bottom end point of the zero virtual value region in G by $\underline{r}_G = \inf\{v \mid \Phi_G^{\lambda,\alpha}(v) = 0\}$ and the top end point $\bar{r}_G = \sup\{v \mid \Phi_G^{\lambda,\alpha}(v) = 0\}$.

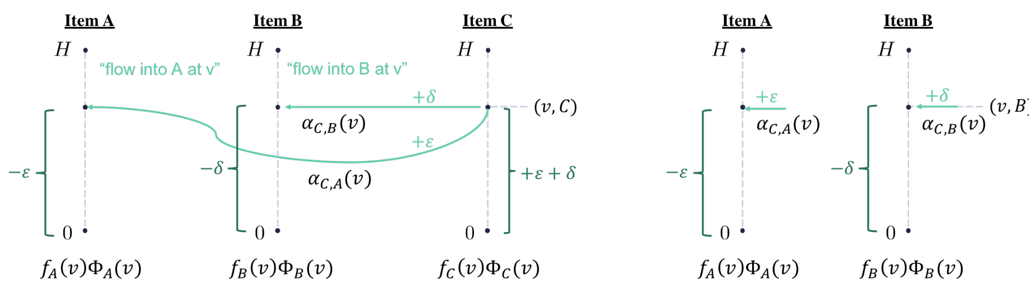

Figure 2: Left: A pictorial interpretation of the dual variable α, in addition to the concepts of flow, sending flow (for any δ and ϵ), preferable items, and equally preferable items. Right: A pictorial shorthand we will use, where flow is assumed to be coming from item C. Complementary slackness dictates that if $\alpha_{G,G'}(v) > 0$, then $u_G(v) = u_{G'}(v)$. That is, a bidder with type (v, G) gets the same utility as a bidder with type (v, G'). Because incentive compatibility requires $u_C(v) \geq u_A(v)$ and $u_C(v) \geq u_B(v)$ at all v, then if there is flow from C into A at v, that is, $\alpha_{C,A}(v) > 0$, then $u_C(v) = u_A(v)$, and thus item A must be preferable to item B. If there is flow from C into both A and B at v, then both items must be equally preferable, as $u_A(v) = u_B(v) = u_C(v)$.

- (Flow) We will call $\alpha_{G',G}(v)$ the "flow into (v, G)" or the "flow into G at v."
- (Sending Flow) If we increase $\alpha_{G,G'}(v)$ by ε, for all $v' \leq v$, $f_G(v')\Phi_G^{\lambda,\alpha}(v')$ decreases by ε and $f_{G'}(v')\Phi_{G'}^{\lambda,\alpha}(v')$ increases by ε. All other virtual values remain the same. This is immediate from the definition of $\Phi^{\lambda,\alpha}(\cdot)$ above in Equation 1.
- (Preferable Items) To satisfy complementary slackness, for any x such that $\alpha_{G,G'}(x) > 0$, we must have $u_{G'}(x) \leq u_{G''}(x) \quad \forall G'' \in N^+(G)$. This is because (a) $u_G(x) = u_{G'}(x)$ by complementary slackness and (b) $u_G(x) \leq u_{G''}(x) \quad \forall G'' \in N^+(G)$ by incentive compatibility.
- (Equally Preferable Items) By the above, to satisfy complementary slackness with any dual with $\alpha_{G,G'}(x) > 0$ and $\alpha_{G,G''}(x) > 0$, we must have $u_{G'}(x) = u_{G''}(x)$.

4 Menu Complexity is Unbounded

We begin by presenting our lower bound: for all M, there exists an instance with three items $\{A, B, C\}$ satisfying $A \succ C$, $B \succ C$, $A \not\succ B$, $B \not\succ A$, and a value distribution f such that every optimal mechanism has menu complexity at least M. Reasoning through candidate lower bounds will also help provide intuition for our upper bound in the following section. In order to prove the main result of this section, we require two things: a candidate dual instance and a proof that there exist distributions inducing this instance. Reasoning about the dual is the "fun part", and proving that matching distributions exist is handled by our Master Theorem.

Recall again that this menu-complexity lower bound applies to any partially-ordered setting that contains an item G with $|N^+(G)| \geq 2$, that is, with at least two incomparable items that dominate G (which includes every single-minded valuation setting with at least 3 items). On the flipside, for any instance where, for every item G, $|N^+(G)| \leq 1$, the FedEx closed-form solution applies[10], so the menu-complexity of the optimal mechanism is exponential [FGKK16; SSW17].

4.1 A Candidate Dual Instance

We define, in this section, for any M, a dual instance such that any feasible primal that satisfies complementary slackness with it has menu complexity at least M. Then, we prove that there exists a primal instance that satisfies complementary slackness with the defined dual. This proves that our dual is optimal and *any* optimal primal (*i.e.*, mechanism) must satisfy complementary slackness with it. Thus, any optimal mechanism for the instance defined must have menu complexity of at least M. We'll postpone whether any input distribution admits such a dual to the Master Theorem in the following section.

4.1.1 Intuition

We first provide a couple of illustrative examples that will give the intuition behind our construction. These examples deal with the 'base cases' in our construction (so the value of M is small).

In this light, consider the example depicted on the left in Figure 3, where $\overline{r}_B = \underline{r}_B = x$, say, and $\lambda_C(v) = 0$ for all v, *i.e.*, item C is unironed everywhere. However, the value x lies inside an

[10] This is not an immediate corollary of their theorem statements, but readers familiar with Fiat et al. [2016] will immediately observe that this follows.

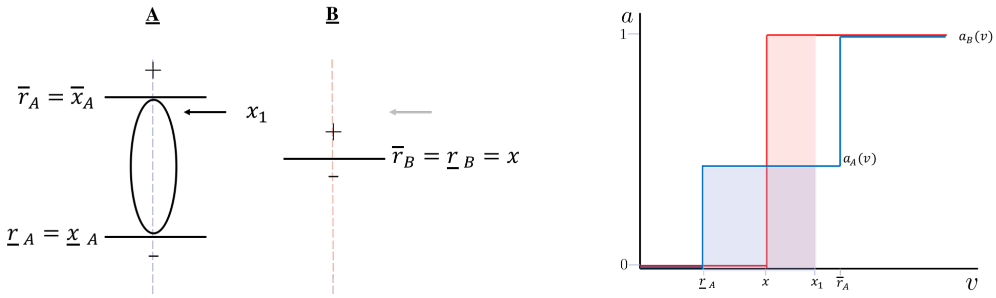

Figure 3: Left: Our first example that requires randomizing on A, containing an ironed interval $[\underline{r}_A, \overline{r}_A]$ (which forces that $a_A(x_1) = a_A(\underline{r}_A)$) and flow into both A and B at x_1 (which forces that $u_A(x_1) = u_B(x_1)$).
Right: The positive and negative virtual values require that item B is offered at a price of x, while A's allocation is 0 until $\underline{r}_A$ and 1 after $\overline{r}_A$. Complementary slackness requires that the allocation be constant in between. In order to have equal preferability (area under the allocation curve) with item B at x_1, A's allocation in between has a unique fractional solution.

ironed interval $[\underline{x}_A, \overline{x}_A]$ in the item A. Further, there is flow from C to *both* A and B at some point $x_1 \in [\underline{x}_A, \overline{x}_A]$.

To finish our example, we also require some minor conditions, *e.g.*, the intervals described are the only ironed intervals, the flow described is all the flow, and $\Phi_A^{\lambda,\alpha}(y) = 0$ everywhere except the interesting points mentioned above. We summarize all the above features of our example below:

- $\underline{r}_B = \overline{r}_B = x$ (thus, $\Phi_B^{\lambda,\alpha}(y) > 0$ for all $y > x$ and $\Phi_B^{\lambda,\alpha}(y) < 0$ for all $y < x$).
- $\underline{r}_A = \underline{x}_A \neq \overline{x}_A = \overline{r}_A$ (the value x is in an ironed interval in the item A and this ironed interval has all the zeros in item A).
- $\alpha_{C,A}(x_1) > 0$, $\alpha_{C,B}(x_1) > 0$ for some $x_1 \in [\underline{x}_A, \overline{x}_A]$ (forces $u_A(x_1) = u_B(x_1)$).
- $\alpha_{C,A}(v) = \alpha_{C,B}(v) = 0$ everywhere else.
- $\lambda_C(v) = 0$ for all v (item C is unironed everywhere).

We claim that the optimal mechanism for this dual instance must randomize when selling item A. The reasoning is the following: Since $\underline{r}_B = \overline{r}_B = x$, we have to set the price x for item B. Via complementary slackness, the ironed interval $[\underline{r}_A, \overline{r}_A]$ requires that $a_A(x_1) = a_A(\underline{r}_A)$, while the flow into both A and B at x_1 requires that $u_A(x_1) = u_B(x_1)$. Then to satisfy both of these constraints, as depicted on the right of Figure 3, the allocation to A at $\underline{r}_A$ cannot be 0 nor 1, thus the allocation must be randomized.

We now extend the above example in a way that forces us to randomize even more. This extension forms the 'inductive step' in our construction.

In the previous example, we had no option but to set the price x for the item B and randomize between $\underline{r}_A$ and $\overline{r}_A$ for the item A. This was due to the presence of the point x_1 that satisfied $u_A(x_1) = u_B(x_1)$. Our second example adds another point x_2 such that $u_A(x_2) = u_B(x_2)$ and we have to randomize on both items, A and B.

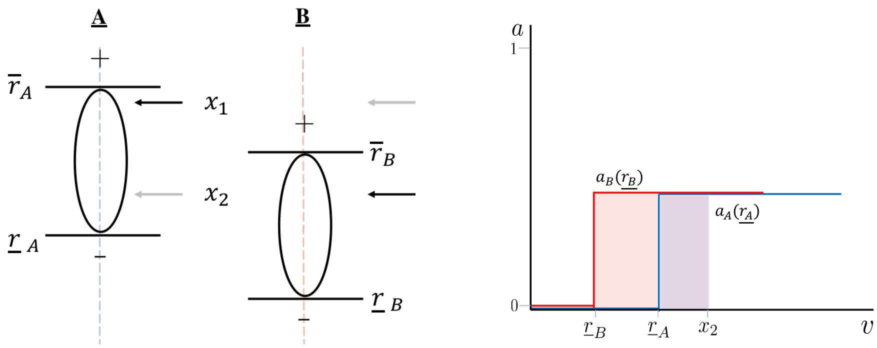

Figure 4: Left: Our second example, which requires randomization on both A and B. Right: If $a_A(\underline{r}_A) \leq a_B(\underline{r}_B)$, then $u_A(x_2) < u_B(x_2)$, which violates complementary slackness.

In this example, depicted in Figure 4, we again have an ironed interval $[\underline{r}_A, \overline{r}_A]$. We, however, replace the condition $\underline{r}_B = \overline{r}_B = x$ with an ironed interval $[\underline{r}_B, \overline{r}_B]$. These values satisfy $\underline{r}_B < \underline{r}_A < \overline{r}_B < \overline{r}_A$. The point x_1 lies in $[\overline{r}_B, \overline{r}_A]$ (as before) while the point x_2 lies in $[\underline{r}_A, \overline{r}_B]$. As before, we require $\alpha_{C,A}(x_1), \alpha_{C,B}(x_1) > 0$ and $\alpha_{C,A}(x_2), \alpha_{C,B}(x_2) > 0$, ensuring that $u_A(x_1) = u_B(x_1)$ and $u_A(x_2) = u_B(x_2)$.

We claim that this example requires us to randomize on both items. Intuitively, this is because we now have two constraints on utilities that must be satisfied, so two degrees of freedom seems necessary. As in the first example, the combination of the ironed interval $[\underline{r}_A, \overline{r}_A]$ and the flow into both A and B at x_1 forces the allocation $a_A(\underline{r}_A) > 0$. In this example, now the combination of the ironed interval $[\underline{r}_B, \overline{r}_B]$ and the flow into both A and B at x_2 similarly forces the allocation $a_B(\underline{r}_B) > 0$.

After reasoning downward that the allocation must be non-zero at the bottom of each ironed interval, we now reason upward that the allocation must be different at every point. On the right side of Figure 4, we show that for any nonzero choice of $a_B(\underline{r}_B)$, because $\underline{r}_A > \underline{r}_B$, in order to ensure that $u_A(x_2) = u_B(x_2)$ (as is required by complementary slackness), we must have $a_A(\underline{r}_A) > a_B(\underline{r}_B)$. This results in 3 distinct allocation probabilities ($a_B(\underline{r}_B)$, $a_A(\underline{r}_A)$, and 1), forcing a menu size of at least 3. A full proof is given in Appendix B.

4.1.2 Construction

It is possible to extend the examples above by continuing to interleave ironed intervals with flow coming in. This way, we can require any number of menu options. The combination of the equal preferability constraints in conjunction with the inability to increase the allocation in the middle of an ironed interval is what requires us to randomize differently within each interval.

We'll postpone a text definition of our construction (and proof of Theorem 1) to Appendix B, but Figure 5 contains a clear pictorial description, and we discuss the intuition. As in the two examples, we can reason from the top downward that the allocation at the bottom of every ironed interval must be positive, and reason from the bottom upward that the allocation must strictly increase for each new overlapping ironed interval we encounter, yielding all different menu options. We formally define this interleaving structure and call it a "chain," depicted in Figure 6. As another sanity check: each new point in the chain induces a new equality that has to be satisfied. So if the chain is of length M, intuition suggests that we should need M degrees of freedom to possibly

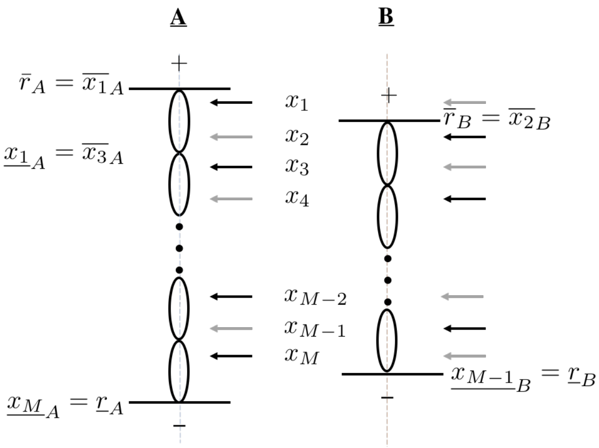

Figure 5: Our candidate dual instance: a top chain that spans the entire region of zero virtual values for both A and B with no gaps between the ironed intervals that comprise the chain. For each point in the chain with flow in (the black arrows), we also have flow into the other side (the gray arrows).

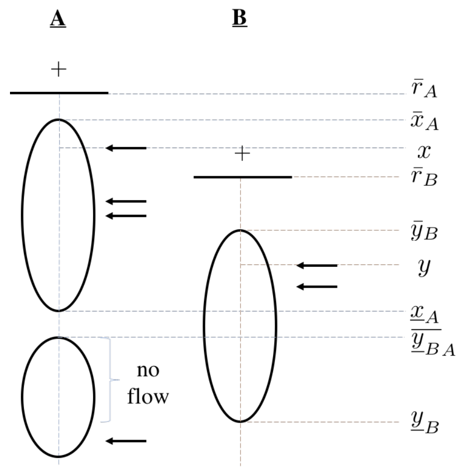

Figure 6: This is an example of a chain that consists of the points $\{(x, A), (y, B)\}$. It is a top chain as $x > \bar{r}_B$. Note that (y, B) is preceded by (x, A) as there is flow into B at y and $y > \underline{x}_A$, and there is no flow into B for any $v \in (y, \bar{y}_B]$. The chain terminates at (y, B) since there is no flow into A for any $v \in [\underline{y}_B, \overline{\underline{y}_B}_A]$.

satisfy complementary slackness (but this is just intuition).

Definition 1 (Top chain)**.** A sequence $(x_1, A), (x_2, B), (x_3, A), \cdots$ of points that switch between items A and B is called a *chain* if the following hold:

- $\Phi_A^{\lambda,\alpha}(x) = 0$ for all (x, A) in the chain and $\Phi_B^{\lambda,\alpha}(y) = 0$ for all (y, B) in the chain.
- $\alpha_{C,A}(x) > 0$ for all (x, A) in the chain and $\alpha_{C,B}(y) > 0$ for all (y, B) in the chain.
- $\underline{x_i}_A < x_{i+1} < x_i$ if (x_i, A) is in the chain and $\underline{x_i}_B < x_{i+1} < x_i$ if (x_i, B) is in the chain.

We call a chain the *top chain* if $x_1 > \bar{r}_B$.

Theorem 1. *Any mechanism that satisfies complementary slackness with the candidate dual instance, which contains a top chain of length M, has menu complexity at least M, and there is such a mechanism.*

4.2 A Master Theorem

All of the analysis in the previous section started from a candidate dual solution, and showed that such duals are optimal (as in, there is a feasible primal satisfying complementary slackness). The missing step is ensuring that there exists an input distribution for which these duals are feasible. To save ourselves (and future work) the tedium of hand-crafting an actual distribution for which these duals are feasible, we prove a general Master Theorem, essentially stating that for a wide class of duals (essentially, anything dictated by ironed intervals, positive/negative regions, and flow in), there exists a distribution for which this dual is feasible.

Theorem 2 (Master Theorem)**.** *Suppose we are given a partial order over $\mathcal{G}$, for each item $G \in \mathcal{G}$ candidate endpoints of zero region (bounded away from 0) $\bar{r}_G, \underline{r}_G$, a finite set of candidate ironed intervals (bounded away from zero) $[\underline{x}_{i,G}, \overline{x}_{i,G}]$ with $\underline{r}_G \leq \underline{x}_{i,G} \leq \overline{x}_{i,G} \leq \bar{r}_G$, and for each pair of items $G' \succ G$ a finite set of candidate flow-exchanging points (bounded away from zero) $y_{i,G,G'}$* not in $(\underline{x}_{i,G}, \overline{x}_{i,G}]$ *for any candidate ironed interval. Then there exists a joint distribution over (value, interest) pairs with a feasible dual (λ, α) such that:*

- *the endpoints of the zero region for $\Phi_G^{\lambda;\alpha}$ are $\underline{r}_G$ and $\bar{r}_G$.*
- *the ironed intervals of $\Phi_G^{\lambda;\alpha}$ are exactly to the intervals $[\underline{x}_{i,G}, \overline{x}_{i,G}]$ (no others).*
- *$\alpha_{G,G'}(y) > 0 \Leftrightarrow y = y_{i,G,G'}$ for some i.*

5 Characterizing the Optimal Mechanism via Duality

In this section, we'll characterize the optimal mechanism for three items $\{A, B, C\}$ with structure $A \succ C$, $B \succ C$, and $A \not\succ B, B \not\succ A$. While our approach will be algorithmic, our focus isn't to actually run this algorithm or analyze its runtime. We'll merely use the algorithms to deduce structure of the optimal mechanism. We prove essentially that the interleaving of ironed intervals used in the construction of the previous section is the worst case (in terms of menu complexity of the optimal mechanism). Still, in order to possibly prove this, we need to at minimum find an optimal mechanism for every possible instance.

Our approach is the following: we propose a *primal recovery algorithm* that, given a dual (λ, α), produces a primal solution that (1) satisfies complementary slackness with the dual and (2) has finite menu complexity. Obviously, the algorithm can't possibly succeed for every input dual (as some duals are simply not optimal for any instance). But we show that whenever the algorithm fails, the dual has some strange structure, in the form of a double swap or an upper swap. We then show that the best dual (which is optimal and always exists, definition below) never admits these strange structures, and therefore the algorithm always succeeds when given the best dual as input.

Definition 2 (Best Dual)**.** We define the best dual of an instance with three partially-ordered items to be the (λ, α) satisfying the following:

1. First, (λ, α) is optimal: $(\lambda, \alpha) \in \arg\min\{\sum_{G \in \{A,B,C\}} \int_0^H f_G(v) \cdot \max\{0, \Phi_G^{\lambda;\alpha}(v)\} dv\}$.
2. Among (λ, α) satisfying (1), (λ, α) has the *fewest ironed intervals* of virtual value zero. That is, (λ, α) minimizes $|\mathcal{I}(\lambda, \alpha)| = |\{\underline{x}_G \mid (x, G) \in [0, H] \times \{A, B, C\}, \Phi_G^{\lambda;\alpha}(v) = 0\}|$.
3. Among (λ, α) satisfying (2), (λ, α) has the *lowest positives* (lexicographically ordered). That is, (λ, α) minimizes $\bar{r}_A$, followed by $\bar{r}_B$, followed by $\bar{r}_C$.

Definition 3. A *double swap* exists when there are consecutive points (x, A) and (y, B) in a chain, and there is flow into A for $v \in [\underline{x}_A, y)$. See Figure 7.

Definition 4. An *upper swap* occurs when there is flow into (x, A) and (y, B) where $x > \bar{r}_A > y > \bar{r}_B$. See Figure 9.

Proposition 1. The best dual has no double swaps or upper swaps.

The full proof of Proposition 1 is in Appendix D. The high-level approach is that whenever a double swap or upper swap exists, we can modify the dual to find a better one. To illustrate our techniques, we prove the following lemma here.

Lemma 2. *The optimal dual that has the minimal number of ironed intervals does not contain any double swaps.*

First, we discuss why this structure would cause a problem for how we're used to satisfying complementary slackness conditions. Complementary slackness forces that in the ironed intervals $[\underline{z}_B, \overline{z}_B]$ and $[\underline{y}_A, \overline{y}_A]$, the allocation is constant, and thus utility in these regions is linear. However, no linear utility functions can satisfy the preferability constraints of having utility that is higher for item A, then B, then A, as illustrated on the left in Figure 7.

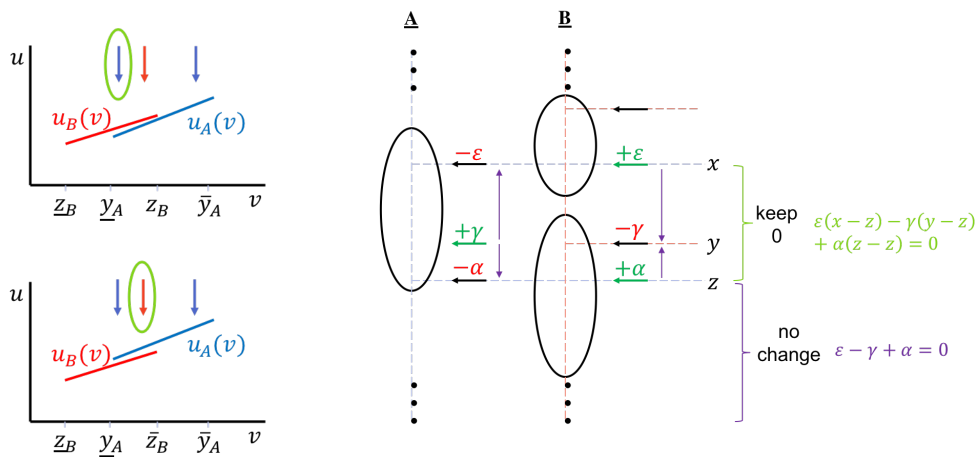

Figure 7: Left: Complementary slackness forces linear utility in ironed intervals. For any choice of linear utility functions, we cannot satisfy the preferability constraints imposed by the double swap for item A, then B, then A in this region. Right: The operation used in the proof of Lemma 2, using a double swap to maintain virtual welfare and create fewer ironed intervals.

Proof. Proof by contradiction. Suppose that somewhere in the top chain, some point in the chain (x, A) is succeeded by (y, B) and $\alpha_{C,A}(z) > 0$ for some $z \in (\underline{x}_A, y)$, creating a double swap. We consider the following operation (depicted on the right in Figure 7) that pushes flow down within the ironed interval $[\underline{x}_A, \overline{x}_A]$ and does the reverse on B, yet negates the change in flow at z to maintain the virtual values below here. Move ε flow from (x, A) to (x, B). Move γ flow from (y, B) to (y, A). Move α flow from (z, A) to (z, B). We will set

$$\alpha = \left(\frac{x-y}{y-z}\right)\varepsilon \qquad \text{and} \qquad \gamma = \left(1 + \frac{x-y}{y-z}\right)\varepsilon.$$

First, this ensures that $\varepsilon - \gamma + \alpha = 0$, and thus for $v \le z$, $\hat{\Phi}_A^{\lambda,\alpha}(v) = \Phi_A^{\lambda,\alpha}(v)$ as well as $\hat{\Phi}_B^{\lambda,\alpha}(v) = \Phi_B^{\lambda,\alpha}(v)$. Second, this ensures that $\varepsilon(x-z) - \gamma(y-z) = 0$, keeping the average virtual value from

z to x the same for both items.

$$\begin{aligned}\int_z^x f_A(v)\hat{\Phi}_A^{\lambda,\alpha}(v)dv &= \int_z^y f_A(v)(\Phi_A^{\lambda,\alpha}(v)+\varepsilon-\gamma)dv + \int_y^x f_A(v)(\Phi_A^{\lambda,\alpha}(v)-\gamma)dv \\ &= \int_z^x f_A(v)\Phi_A^{\lambda,\alpha}(v)dv + \varepsilon(x-z) - \gamma(y-z) \\ &= \int_z^x f_A(v)\Phi_A^{\lambda,\alpha}(v)dv\end{aligned}$$

However, the virtual values in $[y, x]$ are increasing for item A and decreasing for item B, and likewise those in $[z, x)$ are decreasing for item A and increasing for item B. If we choose ε small enough as to not uniron the interval $[\underline{x}_A, \bar{x}_A]$, the change gets spread around the interval and the interval remains all zeroes. However, for item B, the interval $[\underline{y}_B, \bar{y}_B]$ becomes positive while the region above becomes negative. Since the average of both regions is the same and there is now a non-monotonicity, the regions will be ironed together, creating a larger ironed interval with virtual value zero.

Since the virtual welfare of the dual hasn't changed, but we have reduced the number of ironed intervals, then we did not start with an optimal dual with the fewest possible ironed intervals, deriving a contradiction. □

Theorem 3. *For any dual solution that does not contain a double swap or an upper swap, we can find a primal with finite menu complexity that satisfies complementary slackness (and is therefore optimal).*

A full proof appears in Appendix D, but the high-level approach is explained in the following.

Proof Sketch of Theorem 3. (No bad structures exist in best duals.) First, we try to satisfy the necessary complementary slackness system of equations as in Section 4, and identify all possible barriers to solutions existing. These barriers are exactly double swaps or upper swaps.

(Inductive primal recovery algorithm.) Without these barriers, an inductive argument shows that we can indeed find an allocation rule that satisfies all of the complementary slackness conditions. Every dual has a (possibly empty) top chain, and each point in the chain has another set of preferability constraints for that item, along with the constraint that the allocation is constant. We use induction to handle one point in the chain at a time, as illustrated in Figure 8. We take the partially-constructed allocation that satisfies the constraints for the chain so far, scale it down (and thus continue to satisfy the constraints), and then solve for the allocation probability that will satisfy the new constraints given by this point in the chain. As shown in Section 4, this requires choosing a different allocation probability at the bottom of each ironed interval in the chain, but we show that this is sufficient, giving menu complexity at most the length of the chain.

(Finite menu complexity.) The other interesting part not addressed in Section 4 is what to do if there is a chain of countably infinite length (which can certainly exist). Naively following our algorithm would indeed result in a primal of countably infinite menu complexity. But, because the sequence of chain points is monotonically decreasing (and lower bounded by zero), they must converge to some value v. If they converge, and the chain is indeed infinitely long, then neither A nor B can possibly be ironed at v, and we can simply set price v for both items instead. □

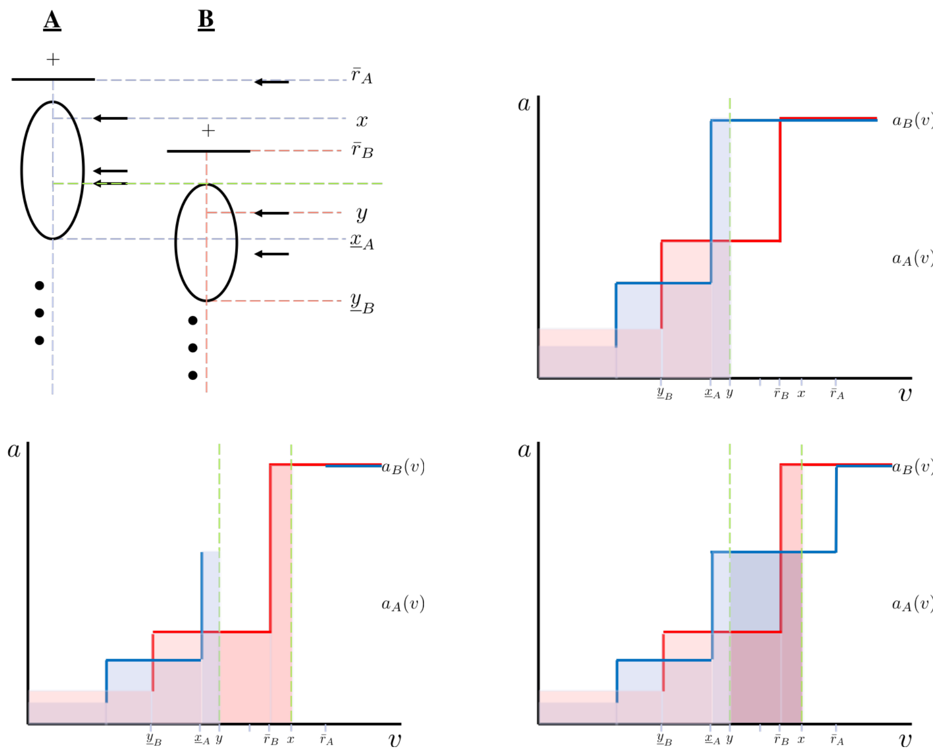

Figure 8: Top Left: A top chain from a candidate dual. We use the inductive hypothesis on the chain of one size smaller (below the green line). Top Right: The allocation rule from the inductive hypothesis that satisfies all CS constraints on the smaller chain (below the green line). Bottom Left: The scaled allocation rule, requiring preferability of A between the green lines. Bottom Right: The allocation rule that satisfies these preferability constraints.

6 Conclusions

We study optimal mechanisms for three partially-ordered items, and show that the menu complexity of optimal mechanisms is unbounded but finite. Recall that for identical items, the menu complexity is 1, for totally-ordered items the menu complexity is at most 7, and for heterogeneous items the menu complexity is uncountable. So our setting fits nicely "in between" totally-ordered and heterogeneous by this measure. By fuzzier measures of complexity, the same is true too: for identical items, the optimal mechanism has a clean closed-form description. For totally-ordered items, the optimal dual has a closed form, and the primal can be recovered by a simple algorithm as a function of this dual. For partially-ordered items, the optimal dual is unlikely to have a closed form, but can be characterized algorithmically (and the primal can still be recovered algorithmically as a function of this dual). For heterogeneous items, optimal mechanisms are pure chaos.

We also provide extensions (menu complexity of MUP, simple partial orderings subject to DMR) proving the usefulness of our techniques beyond our setting.

The most interesting direction for future work is to continue exploring this fascinating space between identical (where optimal mechanisms are extremely tractable, but not rich) and heterogeneous items (where optimal mechanisms are extremely rich, but not tractable), pushing the limits of settings where optimal mechanisms are rich yet tractable. It's perhaps difficult to try and pose a concrete open problem on this front, but it would be wonderful to discover another setting squeezed further in between partially-ordered and heterogeneous items, perhaps where optimal mechanisms can have countable but not uncountable menu complexity.

References

Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 869–877, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4528-6. doi: 10.1145/3055399.3055426. URL `http://doi.acm.org/10.1145/3055399.3055426`.

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

Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 926–939, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4132-5. doi: 10.1145/2897518.2897645. URL `http://doi.acm.org/10.1145/2897518.2897645`.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 269–286, 2013. URL `http://doi.acm.org/10.1145/2482540.2482593`.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 449–450. ACM, 2015.

Nikhil R. Devanur and S. Matthew Weinberg. The Optimal Mechanism for Selling to a Budget Constrained Buyer: The General Case. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 39–40, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4527-9. doi: 10.1145/3033274.3085132. URL `http://doi.acm.org/10.1145/3033274.3085132`.

Nikhil R. Devanur, Nima Haghpanah, and Christos-Alexandros Psomas. Optimal Multi-Unit Mechanisms with Private Demands. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 41–42, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4527-9. doi: 10.1145/3033274.3085122. URL `http://doi.acm.org/10.1145/3033274.3085122`.

Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. The FedEx Problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 21–22, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0. doi: 10.1145/2940716.2940752. URL `http://doi.acm.org/10.1145/2940716.2940752`.

Yiannis Giannakopoulos and Elias Koutsoupias. Duality and Optimality of Auctions for Uniform Distributions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, EC '14, pages 259–276, 2014. URL `http://doi.acm.org/10.1145/2600057.2602883`.

Yiannis Giannakopoulos and Elias Koutsoupias. Selling Two Goods Optimally. In *ICALP (2)*, volume 9135 of *Lecture Notes in Computer Science*, pages 650–662. Springer, 2015.

Nima Haghpanah and Jason Hartline. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 757–758. ACM, 2015. Updated version: http://arxiv.org/abs/1404.1341.

Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 565–566, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1962-1. doi: 10.1145/2482540.2482544. URL `http://doi.acm.org/10.1145/2482540.2482544`.

Sergiu Hart and Philip J Reny. Maximizing revenue with multiple goods: Nonmonotonicity and other observations. hebrew university of jerusalem. *Center for Rationality DP-630*, 2012.

Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

Jean-Jacques Laffont, Eric Maskin, and Jean-Charles Rochet. Optimal nonlinear pricing with two-dimensional characteristics. *Information, Incentives and Economic Mechanisms*, pages 256–266, 1987.

Alexey Malakhov and Rakesh V Vohra. An optimal auction for capacity constrained bidders: a network perspective. *Economic Theory*, 39(1):113–128, 2009.

R. Preston McAfee and John McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335 – 354, 1988. URL `http://dx.doi.org/10.1016/0022-0531(88)90135-4`.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. URL `http://dx.doi.org/10.1287/moor.6.1.58`.

Raghuvansh R. Saxena, Ariel Schvartzman, and S. Matthew Weinberg. The menu-complexity of one-and-a-half dimensional mechanism design. In *To appear in Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, 2017.

Zihe Wang and Pingzhong Tang. Optimal mechanisms with simple menus. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 227–240, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2565-3. doi: 10.1145/2600057.2602863. URL `http://doi.acm.org/10.1145/2600057.2602863`.

A Preliminaries Appendix

A.1 Review of Dual Properties

- (Rerouting Flow Among $N^+(G)$) If $G', G'' \in N^+(G)$ and we decrease $\alpha_{G,G'}(v)$ by ε and increase $\alpha_{G,G''}(v)$ by ε, then $v' \leq v$, $f_{G'}(v')\Phi^{\lambda,\alpha}_{G'}(v')$ decreases by ε and $f_{G''}(v')\Phi^{\lambda,\alpha}_{G''}(v')$ increases by ε. All other virtual values, including all of those within G, remain the same.

- (Utility based on the dual) We can often simplify how utility is written in terms of the dual and complementary slackness constraints. If $\underline{x}_G < x < y < \overline{x}_G$, then allocation in ironed intervals implies $u_G(y) = u_G(x) + a_G(y)(y - x)$.

- (Allocation to Nonzero Virtual Values) As shown above in Subection 3.3, the dual variables (1) determine the virtual welfare functions $\Phi^{\lambda,\alpha}(\cdot)$ and (2) are chosen to minimize the maximum virtual welfare under $\Phi^{\lambda,\alpha}(\cdot)$. For an optimal dual solution, the optimal mechanism will simply be the corresponding virtual welfare maximizer that satisfies complementary slackness. Parts of this mechanism are easy to predict if the virtual value functions are sign-monotone, which we will later ensure that they are. Assuming this, we can talk about the virtual values in terms of three regions: positives, negatives, and zeroes.

- (Ironing and Proper Monotonicity.) We say that a dual satisfies *proper monotonicity* if $f_G \cdot \Phi^{\lambda,\alpha}_G(\cdot)$ is monotone non-decreasing (note the multiplier of f_G). As shown in [FGKK16; DW17], for all α, there exists a λ such that (λ, α) is properly monotone.

- (Boosting can only improve the dual.) Given any dual with properly monotone virtual values, if there exists v such that $f_G(v)\Phi^{\lambda,\alpha}_G(v) < 0$, then for any $G' \in N^+(G)$, incrementing $\alpha_{G,G'}(v)$ by $f_G(v)\Phi^{\lambda,\alpha}_G(v)$ only improves the dual. By proper monotonicity, for all $v' \leq v$, $f_G(v')\Phi^{\lambda,\alpha}_G(v') < f_G(v)\Phi^{\lambda,\alpha}_G(v) < 0$, hence increasing $\alpha_{G,G'}(v)$ will not create any positives within G, not hurting the dual objective. Sending flow into an item G' can only help by making positives less so, and does not increase any virtual values (but it's possible that it doesn't strictly help). This operation is coined *boosting* in [DW17]. While it is clear that G should send the flow, the remaining question is *which* $G' \in N^+(G)$ should the flow be sent to. This is the bulk of our analysis.

- By sign monotonicity, $v > \bar{r}_G$ has a positive virtual value, and thus the allocation rule must set $a_G(v) = 1$, otherwise it is not maximizing virtual welfare.

- Similarly, for values with negative virtual values, that is, $v < \underline{r}_G$, it must be that $a_G(v) = 0$.

From these observations, we can conclude that the flow out of C is identical to the flow out of the root node (day n) in the FedEx solution. That is,

$$\alpha_{C,A}(v) + \alpha_{C,B}(v) = \begin{cases} 0 & v > \bar{r}_C \\ -\hat{R}''_C(v)/f_C(v) & v \leq \bar{r}_C. \end{cases}$$

where $R_C(\cdot)$ is defined as in Definition 6, $\hat{R}_C(\cdot)$ is the least concave upper bound on $R_C(\cdot)$, and $\hat{R}''_C(\cdot)$ is the second derivative of this function with respect to v.

We conclude with a fundamental result from [FGKK16].

Theorem 4 (Proper Ironing [FGKK16]). *Given all dual variables α, suppose $\lambda_G(v) = 0$ for all (v, G). Then $f_G(v)\Phi_G^{\lambda;\alpha}(v)$ is defined for all (v, G). We define $\Gamma_G(v) = -\int_0^v f_G(x)\Phi_G^{\lambda;\alpha}(x)dx$, and $\hat{\Gamma}_G(\cdot)$ is the least concave upper bound on this function. Then setting $\lambda_G(v) = \hat{\Gamma}_G(v) - \Gamma_G(v)$ defines a continuous and differentiable $\lambda_G(\cdot)$ that, with the update of $\Phi_G^{\lambda;\alpha}(\cdot)$ based on $\lambda_G(\cdot)$, results in the proper monotonicity of $f_G(\cdot)\Phi_G^{\lambda;\alpha}(\cdot)$.*

A.2 Other definitions

We conclude this section with some standard definitions and observations.

Definition 5 (DMR). We say that a marginal distribution of values F_G satisfies *declining marginal revenues* (DMR) if $v(1-F_G(v))$ is concave, or equivalently, if $f_G(v)\varphi_G(v)$ is increasing where $\varphi_G(v)$ is Myerson's virtual value.

Definition 6 (Revenue Curve). The *revenue curve* of a distribution $\mathcal{F}$ is a function $R_{\mathcal{F}}$ such that $R_{\mathcal{F}}(v) = v(1 - F(v))$, where F is the CDF of $\mathcal{F}$. We say a revenue curve is feasible if there exists a distribution that induces it.

Note that $R'_{\mathcal{F}}(v) = 1 - F(v) - vf(v) = -\varphi(v)f(v)$, where $\varphi(v)$ is Myerson's virtual value function. When clear from context we will omit the subindex.

A.3 Derivation of the Partial Lagrangian Dual

The primal stated in subsection 3.2 is:

Variables: $a_G(v), \ \forall G \in \mathcal{G}, \ v \in [0, H]$

Maximize $$\sum_{G\in\mathcal{G}} \int_0^H f_G(v)a_G(v)\varphi_G(v)dv$$

subject to

$$a'_G(v) \geq 0 \qquad \forall G \in \mathcal{G} \ \forall v \in [0, H] \text{ (dual variables } \lambda_G(v) \geq 0)$$

$$\int_0^v a_G(x)dx - \int_0^v a_{G'}(x)dx \geq 0 \qquad \forall G \in \mathcal{G}, \ G' \in N^+(G) \ \forall v \in [0, H] \text{ (dual vars } \alpha_{G,G'}(v) \geq 0)$$

$$a_G(v) \in [0, 1] \qquad \forall G \in \mathcal{G}, \ \forall v \in [0, H] \text{ (no dual variables).}$$

Moving the first two types of constraints to the objective function with multipliers $\lambda_G(v)$ and $\alpha_{G,G'}(v)$ respectively gives the partial Lagrangian primal:

$$\max_{a: a_G(v)\in[0,1]\,\forall G\in\mathcal{G},\forall v\in[0,H]} \ \min_{\lambda,\alpha\geq 0} \mathcal{L}(a; \lambda, \alpha)$$

where

$$\mathcal{L}(a; \lambda, \alpha) = \sum_{G\in\mathcal{G}} \int_0^H \left[f_G(v)a_G(v)\varphi_G(v) + \sum_{G'\in N^+(G)} \alpha_{G,G'}(v) \cdot \left[\int_0^v a_G(x)dx - \int_0^v a_{G'}(x)dx \right] + \lambda_G(v)a'_G(v) \right] dv.$$

This gives the corresponding partial Lagrangian dual of

$$\min_{\lambda,\alpha\geq 0} \max_{a:a_G(v)\in[0,1]\,\forall G\in\mathcal{G},\forall v\in[0,H]} \mathcal{L}(a;\lambda,\alpha).$$

Note however that we can rewrite $\mathcal{L}(a;\lambda,\alpha)$ by using integration by parts on the $a'_G(v)$ term to get $a_G(v)$ terms, using that $a_G(0)=0$ and $\lambda_G(H)=0$ without loss:

$$\int_0^H \lambda_G(v)a'_G(v) = \lambda_G(v)a_G(v)\mid_0^H - \int_0^H \lambda'_G(v)a_G(v)dv = -\int_0^H \lambda'_G(v)a_G(v)dv$$

As in [Fiat et al., 2016], this uses the facts that $\lambda_G(\cdot)$ is continuous and equal to 0 at any point that $a'_G(v)=\infty$, which occurs at only countably many points. Then, collecting the $a_G(v)$ terms gives:

$$\begin{aligned}
\mathcal{L}(a;\lambda,\alpha) &= \sum_{G\in\mathcal{G}} \int_0^H \Bigg[f_G(v)a_G(v)\varphi_G(v) \\
&\quad + \sum_{G'\in N^+(G)} \alpha_{G,G'}(v)\cdot\left[\int_0^v a_G(x)dx - \int_0^v a_{G'}(x)dx\right] - \lambda'_G(v)a_G(v)\Bigg]dv \\
&= \sum_{G\in\mathcal{G}} \int_0^H f_G(v)a_G(v)\Phi_G^{\lambda,\alpha}(v)dv.
\end{aligned}$$

where we define

$$\Phi_G^{\lambda,\alpha}(v) := \varphi_G(v) + \frac{1}{f_G(v)}\cdot\left[\sum_{G'\in N^+(G)}\int_v^H \alpha_{G,G'}(x)dx - \sum_{G':G\in N^+(G')}\int_v^H \alpha_{G',G}(v)dx\right] - \frac{1}{f_G(v)}\lambda'_G(v).$$

Then we can write that the Lagrangian dual problem is

$$\min_{\lambda,\alpha\geq 0} \quad \max_{a:a_G(v)\in[0,1]\,\forall G\in\mathcal{G},\forall v\in[0,H]} \quad \sum_{G\in\mathcal{G}}\int_0^H f_G(v)a_G(v)\Phi_G^{\lambda,\alpha}(v)dv$$

as in subsection 3.3.

B Missing Details from Section 4:

The purpose of this section is to provide a complete proof of Theorem 1. First, we provide a construction of our candidate dual, which is depicted in Figure 5. The instance uses definition 1 of a top chain.

Construction of candidate dual instance:

- Let there exist no point at which A and B both have virtual value zero and both are unironed, that is, there is no v such that $\Phi_A^{\lambda,\alpha}(v) = \Phi_B^{\lambda,\alpha}(v) = 0$ and $\lambda_A(v)=\lambda_B(v)=0$.
- Let $\overline{r}_A > x_1 > \overline{r}_B > x_2 > x_3 > \cdots > x_M > \underline{r}_B > \underline{r}_A$. The dual has a top chain of length M defined by $(x_1, A), (x_2, B), \ldots, (x_M, A)$.

- In addition, we have flow into the other item at each point in the chain: let $\alpha_{C,B}(x_i) > 0$ for all (x_i, A) in the chain as well as $\alpha_{C,A}(x_i) > 0$ for all (x_i, B) in the chain.
- Let $\lambda_C(v) = 0$ for all v, *i.e.*, item C is unironed everywhere.
- For all v where α has not already been defined, let $\alpha_{C,A}(v) = \alpha_{C,B}(v) = 0$.

We first make some remarks that follow directly from our construction. All the remarks below (only) talk about our dual and any feasible primal that satisfies complementary slackness with our dual.

Remark 5. For all $i \in \{1, 3, \cdots, M-2\}$, we have $x_i, x_{i+1} \in [\underline{x_i}_A, \overline{x_i}_A] = [\underline{x_{i+1}}_A, \overline{x_{i+1}}_A]$. Since this interval is ironed, we have $\lambda_A(v) > 0 \implies a'_A(v) = 0$ for v in this interval. Thus, $a_A(x_i) = a_A(x_{i+1})$.

Remark 6. For all $i \in \{2, 4, \cdots, M-1\}$, we have $x_i, x_{i+1} \in [\underline{x_i}_B, \overline{x_i}_B] = [\underline{x_{i+1}}_B, \overline{x_{i+1}}_B]$. Since this interval is ironed, we have $\lambda_B(v) > 0 \implies a'_B(v) = 0$ for v in this interval. Thus, $a_B(x_i) = a_B(x_{i+1})$.

Remark 7. For all $i \in \{1, 2, \cdots, M\}$, we have $u_A(x_i) = u_B(x_i)$.

We now prove a lemma that forms the backbone of our inductive argument:

Lemma 3. *For all $i \in \{1, 2, \cdots, M-1\}$, we have $a_A(x_i) > a_B(x_i) \iff a_A(x_{i+1}) < a_B(x_{i+1})$. Similarly, we have $a_A(x_i) < a_B(x_i) \iff a_A(x_{i+1}) > a_B(x_{i+1})$*

Proof. Note that either $a_A(x_i) = a_A(x_{i+1})$ or $a_B(x_i) = a_B(x_{i+1})$ by Remark 5 and Remark 6. We only prove $a_A(x_i) > a_B(x_i) \iff a_A(x_{i+1}) < a_B(x_{i+1})$ for the case $a_A(x_i) = a_A(x_{i+1})$ and omit the other (symmetric) cases. Since $a_A(x_i) = a_A(x_{i+1})$, we have

$$u_A(x_i) = u_A(x_{i+1}) + a_A(x_i) \cdot (x_i - x_{i+1}) = a_A(x_{i+1}) \cdot (\underline{x_i}_B - x_{i+1}) + a_A(x_i) \cdot (x_i - \underline{x_i}_B).$$

We also have, by the structure of the ironed intervals for the item B,

$$u_B(x_i) = u_B(x_{i+1}) + a_B(x_{i+1}) \cdot (\underline{x_i}_B - x_{i+1}) + a_B(x_i) \cdot (x_i - \underline{x_i}_B)$$

Now, since the utilities at all points x_i is the same for both items A and B (Remark 7), the fact that $a_A(x_i) > a_B(x_i)$ is equivalent to $a_A(x_i) \cdot (x_i - \underline{x_i}_B) > a_B(x_i) \cdot (x_i - \underline{x_i}_B)$ which is equivalent to $a_A(x_{i+1}) \cdot (\underline{x_i}_B - x_{i+1}) < a_B(x_{i+1}) \cdot (\underline{x_i}_B - x_{i+1})$ which, in turn, is equivalent to $a_A(x_{i+1}) < a_B(x_{i+1})$. □

Finally, we prove Theorem 1.

Proof of Theorem 1. At x_M, we have that

$$u_A(x_M) = a_A(x_M) \cdot (x_M - \underline{x_M}_A) \qquad \text{and} \qquad u_B(x_M) = a_B(x_M) \cdot (x_M - \underline{x_M}_B).$$

Since $\underline{x_M}_B > \underline{x_M}_A$ and $a_A(x_M) > 0$, then to ensure that $u_A(x_M) = u_B(x_M)$ (Remark 7), we must have $a_B(x_M) > a_A(x_M)$. However, with this fact, Lemma 3 says that $a_B(x_i) > a_A(x_i)$ and $a_B(x_{i+1}) > a_A(x_{i+1})$ in alternation.

Since $a_A(\cdot)$ and $a_B(\cdot)$ are non-decreasing sequences, they can only alternate if they have $\Omega(M)$ distinct elements.

By Theorem 3, there exists a feasible primal that satisfies complementary slackness. The primal algorithm constructs a mechanism with menu complexity at least M and satisfies complementary slackness, hence this dual is in fact optimal. □

Corollary 4. This idea gives a lower bound for Multi-Unit Pricing as well.

We expand on this on Appendix E.

C Missing Details from Section 4.2: Proof of Theorem 2

In this section we provide a complete proof of Theorem 2. On our way to prove this theorem, we generalize a result of Saxena et al. [2017], in which they show that for totally ordered preferences, one can always find a discrete distribution that produces a well-enough-behaved revenue curve. They use this result to show that there exist instances for which the menu complexity is the worst possible, exponential in the number of items. Here we extend their construction and show that for any well-enough-behaved set of continuous revenue curves for the partially ordered setting, there exist distributions that induce them.

The first step is to generalize the result of Saxena et al. [2017] from discrete distributions to continuous distributions.

Lemma 5 (Revenue Theorem for Continuous Curves)**.** *Given a continuous curve $R : [1, H]$ differentiable everywhere except at countably many points, such that $R(1) = 1$ and $|R'(x)_+|, |R'(x)_-| \leq \frac{1}{2H} \forall x \in [1, H]$, there exists a distribution $\mathcal{F}$ such that R is the revenue curve that arises from selling to a single bidder with a valuation drawn from $\mathcal{F}$.*

Proof. Consider the following distribution

$$F(x) = 1 - \frac{R(x)}{x}, x \in [1, H]$$

and $F(x) = 0$ for $x \leq 1$, $F(x) = 1$ for $x \geq H$. In order to show that this is a valid distribution, it suffices to show that it is monotonic non-decreasing. For that, we consider its derivative and show it is non-negative everywhere:

$$F'(x) = \frac{-xR'(x) + R(x)}{x^2}.$$

It suffices to show that the numerator, $R(x) - R'(x)x$, is always non-negative. Note that for $x \geq 1$, $R(x) \geq \frac{1}{2}$ (since $R(1) = 1$ and the derivative doesn't change fast enough) and $|R'(x)_+| \leq \frac{1}{2H}$. Since $x \leq H$, the claim follows.

It remains to show that indeed the revenue from this distribution matches the curve $R(x)$. Consider setting a price of x, then the revenue of selling at x is exactly $x(1 - F(x)) = R(x)$. □

Now we want to extend this to say we can find distributions for revenue curves with specific properties that will be useful.

Theorem 8 (Master Theorem for Single Item)**.** *Given candidate endpoints of zero region x_+, x_- and candidate ironed interval endpoints $[\underline{x}_i, \overline{x}_i]_{i=1}^k$ (where $x_- \leq \underline{x}_i \leq \overline{x}_i \leq x_+$) there is a distribution $\mathcal{F}$ such that the revenue curve induced by a bidder whose valuation is drawn from $\mathcal{F}$ satisfies*

- *$\Phi^{\lambda,\alpha}(x)f(x)$ is negative for $x < x_-$ (i.e. x_- is the lower endpoint of the zero region),*
- *$\Phi^{\lambda,\alpha}(x)f(x)$ is positive for $x > x_+$ (i.e. x_+ is the upper endpoint of the zero region) and,*
- *the ironed intervals correspond exactly to the intervals $[\underline{x}_i, \overline{x}_i]$ for $i = 1$ to k.*

Proof. We will reduce the problem of finding a valid distribution to that of constructing a revenue curve that will guarantee these properties and then apply Lemma 5. Consider the following revenue curve

$$R(x) = \begin{cases} x & 0 \leq x \leq 1, \\ 1 + \frac{x}{2H} & 1 \leq x \leq x_-, \\ 1 + \frac{x_-}{2H} & x_- \leq x \leq \underline{x}_1 \\ 1 + \frac{x_- + \underline{x}_1 - x}{2H} & \underline{x}_1 \leq x \leq \frac{\underline{x}_1 + \overline{x}_1}{2} \\ 1 + \frac{x_- + x - \overline{x}_1}{2H} & \frac{\underline{x}_1 + \overline{x}_1}{2} \leq x \leq \overline{x}_1 \\ \dots & \\ 1 + \frac{x_-}{2H} & \overline{x}_{i-1} \leq x \leq \underline{x}_i \\ 1 + \frac{x_- + \underline{x}_i - x}{2H} & \underline{x}_i \leq x \leq \frac{\underline{x}_i + \overline{x}_i}{2} \\ 1 + \frac{x_- + x - \overline{x}_i}{2H} & \frac{\underline{x}_i + \overline{x}_i}{2} \leq x \leq \overline{x}_i \\ \dots & \\ 1 + \frac{x_-}{2H} & \overline{x}_k \leq x \leq x_+ \\ 1 + \frac{x_-}{2H} - \frac{x - x_+}{2H(H - x_+)}(x_- + 1) & x_+ \leq x \leq H. \end{cases}$$

This revenue curve is such that $R(1) = 1$ and $|R'(x)| \leq \frac{1}{2H}$ for $x \in [1, H]$. This allows us to claim that there is a distribution that induces this revenue curve. Moreover, from the way we constructed this revenue curve, the derivative is positive from 0 to x_-, negative from x_+ to H, goes from negative to positive for the intervals $[\underline{x}_i, \overline{x}_i]$ and is 0 elsewhere. We will show that these conditions are sufficient to make the virtual values take the signs we intend them to.

It suffices to note that the sign of the derivative of the revenue at x is the opposite of the sign of the virtual value at x (noted in Definition 6). By construction, our revenue curve has negative slope on values higher than x_+ and positive slope on points below x_-. The intervals in between will be ironed and turn into 0 slope intervals. □

Remark 9. It is possible to relax the condition that all ironed intervals are between x_-, x_+. It is not hard to see how to adapt the proof to have ironed intervals either in $[1, x_-]$ or $[x_+, H]$. It is sufficient to add dimpled intervals, like the ones in our construction, as the revenue curve is increasing or decreasing. We don't need them for our main result, hence don't worry about this more general result. Likewise, the revenue curve R could be made differentiable everywhere if we used a smoother function to transition between the ironed and non-ironed intervals, as opposed to straight lines.

Proof of Theorem 2. If the constraint over flows wasn't there, the problem would be a direct application of Theorem 8. Unfortunately, the flow constraints may affect the virtual values of neighboring items. It is not hard to predict how outgoing and incoming flow will change the virtual values for the different items. From the study of duality in this context we know that if there is ε-flow leaving from (y_i, G) to (y_i, G') (where $G' \in N^+(G)$), then the virtual values of all points of item G with $y \leq y_i$ will increase by ε and all points of item G' with $y \leq y_i$ will decrease by ε. Thus, given that we know what we want the revenue curves to look like after all flow has been sent, we can reverse engineer what they must look like in order to make that happen. In particular, since the flows shift the virtual values by a constant it will suffice to subtract a function whose value is 0 before y_i and becomes a line with small, negative slope at x_i (say, slope $\varepsilon = \frac{1}{2H}$) from the "suggested" (by Theorem 8) revenue curve for item G (since these will increase by ε after the flow is sent) and add positive slope functions of the same value at x_i on item $G_{i,G}$ from its suggested revenue curve (since these will decrease by ε after the flow is sent). This is sufficient because of the connection between virtual values and revenue curves argued before: the derivative corresponds to changes in the virtual value. So for a constant change in virtual value, the matching change would be adding a linear term to the revenue curve of opposite sign. The order in which we do these changes is by processing items from leaves to the root (i.e. only process a node once all its children have been processed) and within an item G, address the flow-exchange values from smallest to largest. □

We abuse this opportunity to prove a similar result for the multi-unit pricing setting.

Theorem 10 (Master Theorem for MUP)**.** *Suppose we are given a MUP instance where the buyer can get up to n copies of an item. Let G_i for $1 \leq i \leq n$ be the item corresponding to i copies. For each item G_i we are given candidate endpoints of the zero region x_{-i}, x_{+i} and a set of candidate ironed interval endpoints $[\underline{x}_{j,i}, \overline{x}_{j,i}]_{j=1}^{k_i}$ with $x_{-i} \leq \underline{x}_{j,i} \leq \overline{x}_{j,i} \leq x_{+i}$. Moreover, for each tuple $(i, i+1)$ and $(i, i-1)$, we are given a set of candidate flow-exchanging points $y_{j,i,i+1}$ and $y_{j,i,i-1}$ not in $(\underline{x}_{j,i}, \overline{x}_{j,i}]$ for any candidate ironed interval. Then, there exists distributions $\mathcal{F}_G$ for all items G such that:*

- *the endpoints of the zero region for G_i correspond to x_{-i}, x_{+i},*
- *the ironed intervals correspond exactly to the intervals $[\underline{x}_{j,i}, \overline{x}_{j,i}]_{j=1}^{k_i}$ (and no other),*
- *the dual of the problem is such that there $\alpha_{G_i,G_{i+1}}(y_{j,i,i+1}) \geq 0$ (i.e. there is flow sent from G_i at y_i to G_{i+1} into $y_{j,i,i+1}$ and no other flow from i to $i+1$).*
- *the dual of the problem is such that there $\alpha_{G_i,G_{i-1}}(y_{j,i,i-1}) \geq 0$ (i.e. there is flow sent from G_i at y_i to G_{i-1} into $\frac{i-1}{i} y_{j,i,i-1}$ and no other flow from i to $i-1$).*

Proof. This proof is similar to that of 2 with the exception that on the former, increasing the flow from (v, G) to (v, G') (with $G' \in N^+(G)$) by a little bit increases and decreases the virtual values below v by the same amount. This is no longer true since we are moving from $(y_{j,i,i-1}, G_i)$ to $(\frac{i-1}{i} y_{j,i,i-1}, G_{i-1})$. In this case, sending ε flow from $(y_{j,i,i-1}, G_i)$ to $(\frac{i-1}{i} y_{j,i,i-1}, G_{i-1})$ increases the virtual values below $(y_{j,i,i-1}, G_i)$ by ε but decreases the ones on the other end by only $\frac{i-1}{i}\varepsilon$. So, in order to reverse engineer the change in virtual value induced by this setting we need to add the same functions as in the proof of Theorem 2 to the revenue curve suggested for G_i and add a $\frac{i}{i-1}$-scaled version of it for the receiving item at the point $(\frac{i-1}{i} y_{j,i,i-1}, G_{i-1})$ on the revenue curve for G_{i-1}. The order in which these we do these changes is by processing items from leaves to root (i.e. from G_n to G_1) and within a item G_i, address the flow-exchange points from smallest to largest. □

D Missing Details from Section 5

We begin below by reviewing properties of the dual previously observed in Fiat et al. [2016]; Devanur and Weinberg [2017]. Throughout this section we'll reference the "best" dual. While multiple optimal duals might exist, we'll be interested in a specific tie-breaking among them (and refer to the one that satisfies these conditions as "best").

Theorem 11 (Devanur and Weinberg [2017]). *The best dual* (λ, α) *satisfies the following:*

- *(Proper monotonicity)* $(f_G \cdot \Phi_G^{\lambda,\alpha})(\cdot)$ *is monotone non-decreasing, for all* v.
- *(No-boosting)* $\Phi_G^{\lambda,\alpha}(v) \geq 0$ *for all* G *such that there exists a* $G' \succ G$.
- *(No-rerouting)* $\Phi_G^{\lambda,\alpha}(v) > 0 \Rightarrow \alpha_{G,G'}(v) = 0$ *for all* G'.
- *(No-splitting)* $\lambda_G(v) > 0 \Rightarrow \alpha_{G,G'}(v) = 0$ *for all* G'.

Returning to our three-item example, prior work nicely characterizes the flow coming out of C in the optimal dual: No-boosting tells us that we must always send flow out of (v, C) into somewhere whenever $\Phi_C^{\lambda,\alpha}(v) < 0$ (in order to bring it up to 0). No-rerouting tells us that we can never send flow out of (v, C) if $\Phi_C^{\lambda,\alpha}(v) > 0$. No-splitting tells us that we never send flow out of the middle of an ironed interval. But, we still need to decide whether to send this flow into A or B. This is the novel part of our analysis.

Proof of Proposition 1. By Definition 2, we know that a best dual has the minimum number of ironed intervals amongst all optimal duals. Similarly, a best dual has the lowest positives amongst all optimal duals. We prove the proposition using two lemmas. The first lemma proves that a best dual can't have double swaps:

The second lemma proves that a best dual can't have upper swaps:

Lemma 6. *The optimal dual that has the lowest positives does not contain any upper swaps.*

Proof. Proof by contradiction. Suppose an upper swap exists. Then (as depicted on the right in Figure 9) we can push up α flow from (y, B) to (x, B), causing $f_B(v)\hat{\Phi}_B^{\lambda,\alpha}(v) = f_B(v)\Phi_B^{\lambda,\alpha}(v) - \alpha$ for $v \in [y, x]$ and improving virtual welfare by $\alpha(x - y)$. To leave the flow out of item C unchanged, we balance this out by pushing α flow down from (x, A) to (y, A), causing $f_A(v)\hat{\Phi}_A^{\lambda,\alpha}(v) = f_A(v)\Phi_A^{\lambda,\alpha}(v) + \alpha$ for $v \in [y, x]$.

If y is unironed at A, that is, $\bar{y}_A = y$, or if $\bar{y}_A < \bar{r}_A$, then by choosing $\alpha = -f_A(\bar{y}_A)\Phi_A^{\lambda,\alpha}(\bar{y}_A)$, this will cause $\hat{\bar{r}}_A = \bar{y}_A$, lowering the positives.

Alternatively, if y is ironed up to $\bar{r}_A$ such that $\bar{y}_A = \bar{r}_A$, then we can choose a very small α to keep the interval $[\underline{y}_A, \bar{r}_A]$ ironed, making the whole interval positive and causing $\hat{\bar{r}}_A = \underline{y}_A$, lowering the positives. The dual will only increase by $\alpha(x - y)$, even when the values are ironed around, as ironing preserves virtual welfare. This is canceled out by the improvement in virtual welfare from item B. Then we have maintained virtual welfare but lowered the positives, showing that this dual solution could not have had the lowest positives. □

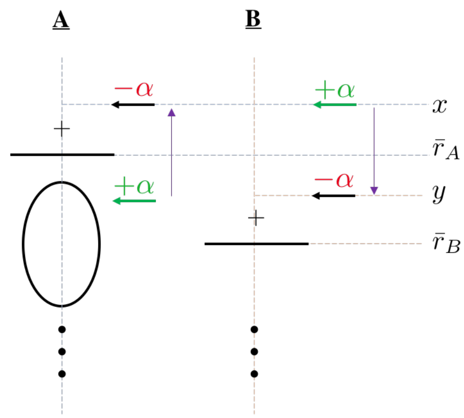

Figure 9: The operation used in the proof of Lemma 6, using a upper swap to maintain virtual welfare and create lower positives.

Lemma 2 and Lemma 6 comprise the proof of Proposition 1. □

Now we prove that our primal recovery algorithm always succeeds in finding an optimal primal (that satisfies complementary slackness) when given a best dual.

Proof of Theorem 3. First, consider the case where there exists some point v where $\Phi_A^{\lambda,\alpha}(v) = \Phi_B^{\lambda,\alpha}(v) = 0$, and v is unironed both in A and in B. Then we simply set v as a price for both A and B, automatically satisfying the complementary slackness conditions of flow into A or B, as both are equally preferable. Since both items A and B, have the same allocation rule, the instance degenerates into a FedEx instance. Thus, an optimal allocation rule for the item C can be determined.

Otherwise, the dual solution contains no point v as described in the first case, meaning that ironed intervals interleave throughout the region of zero virtual values. This means that, if without loss of generality $\bar{r}_A > \bar{r}_B$, that $\bar{r}_B = x$ must sit in an ironed interval $[\underline{x}_A, \bar{x}_A]$ on A.

If the top chain is empty, then we have $\bar{r}_A > \bar{r}_B > \underline{x}_A$ with no flow into A for any $v \in [\underline{x}_A, \bar{x}_A]$. Then, setting

$$a_A(v) = \begin{cases} 1 & v \geq \bar{x}_A \\ \frac{\bar{r}_A - \bar{r}_B}{\bar{r}_A - \underline{x}_A} & v \in [\underline{x}_A, \bar{x}_A) \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad a_B(v) = \begin{cases} 1 & v \geq \bar{r}_B \\ 0 & \text{otherwise} \end{cases}$$

makes both options equally preferable for all v except for $v \in [\underline{x}_A, \bar{x}_A]$, where reporting B is strictly preferable, but this does not violate complementary slackness by the assumption that the top chain is empty.

Otherwise, the top chain is non-empty. A dual gives a system of utility inequalities via complementary slackness which the allocation rule must satisfy. Instead, we can solve a system of utility equalities given by the chain via induction on the length of the top chain, and this will imply a solution that satisfies all of the inequalities. More specifically, the following will hold for top chains of all lengths:

1. The allocation rule will only increase at the bottom of ironed intervals in the chain. That is, if the allocation rule increases at z, so $a'_A(z) > 0$, then z must be the bottom of an ironed interval for a point (x, A) in the top chain, thus $z = \underline{x}_A$, and $a_A(x) = a_A(\underline{x}_A)$.
2. We will fully allocate to all positive virtual values. That is, $a_A(\bar{r}_A) = a_B(\bar{r}_B) = 1$.
3. If (x, A) is followed by (y, B) in the chain, then $a_A(x) = a_A(\underline{x}_A) > a_B(y) = a_B(\underline{y}_B)$.
4. At any point (x, A) in the top chain, we will have $u_A(x) = u_B(x)$.
5. An alternative solution can, for the first point in the chain (x, A), vary $a_A(\underline{x}_A)$ such that the utility constraint is a strict inequality $u_A(x) > u_B(x)$, and instead we have equality at $\bar{r}_A$: $u_A(\bar{r}_A) = u_B(\bar{r}_A)$. This gives an equal expected price for the two items, and equal utility for all values $v \geq \bar{r}_A$.

To satisfy complementary slackness, for any type (x, A) with flow in, it must be that $u_A(x) \geq u_B(x)$. We now show why (3-4) imply that complementary slackness will be satisfied everywhere.

Consider a subsequence of points in the chain: $(x, B), (y, A), (z, B)$, hence $y > \underline{x}_B$ and $z > \underline{y}_A$. Then $a_B(x) > a_A(y) > a_B(z)$ by (3). Since $u_A = u_B$ for every point in the chain and a larger allocation rule implies a larger change in utility, we can deduce that $u_A(v) \geq u_B(v)$ for all $v \in [z, y]$.

- For $v \in (\underline{y}_A, \underline{x}_B)$, we have that $a_A(v) > a_B(v)$, and since $u_A(z) = u_B(z)$, then $u_A(v) \geq u_B(v)$ in this region.
- For $v \in (\underline{x}_B, y)$, we have that $a_B(v) > a_A(v)$, and since $u_A(y) = u_B(y)$, then $u_A(v) \geq u_B(v)$ in this region.
- By definition of a double swap, there is no $v \in [\underline{y}_A, z)$ such that there is flow into (v, A). Likewise, there is no $v \in [\underline{x}_B, y)$ such that there is flow into (v, B).

Hence all possible complementary slackness conditions are satisfied.

We now show that these sufficient properties hold by induction. As a base case, consider when there is one point in the top chain, which without loss is (x, A). By definition of the top chain, $\bar{r}_A > x > \bar{r}_B > \underline{x}_A$ and there is flow into item A at x, which is in ironed interval $[\underline{x}_A, \bar{x}_A]$. We can set $a_A(\underline{x}_A) = \frac{x - \bar{r}_B}{x - \underline{x}_A}$ and set $a_A(\bar{r}_A) = a_B(\bar{r}_B) = 1$. Then

$$u_A(x) = a_A(\underline{x}_A) \cdot (x - \underline{x}_A) = \frac{x - \bar{r}_B}{x - \underline{x}_A} \cdot (x - \underline{x}_A) = 1(x - \bar{r}_B) = u_B(x).$$

Then conditions (1-4) are met. To satisfy (5), we can instead set $a_A(\underline{x}_A) = \frac{\bar{r}_A - \bar{r}_B}{\bar{r}_A - \underline{x}_A}$. Then

$$u_A(\bar{r}_A) = a_A(\underline{x}_A) \cdot (\bar{r}_A - \underline{x}_A) = \frac{\bar{r}_A - \bar{r}_B}{\bar{r}_A - \underline{x}_A} \cdot (\bar{r}_A - \underline{x}_A) = 1(\bar{r}_A - \bar{r}_B) = u_B(\bar{r}_A).$$

For the inductive hypothesis, suppose for any chain of length n, we have allocation rules such that (1-5) hold.

Now consider a chain of length $n + 1$. Without loss of generality, let (x, A) be the top point in the chain, where x sits in the ironed interval $[\underline{x}_A, \bar{x}_A]$, and this point is proceeded by (y, B) which sits in $[\underline{y}_B, \bar{y}_B]$, hence $\bar{r}_A > x > \bar{r}_B$ and $y > \underline{x}_A$ by definition of the chain.

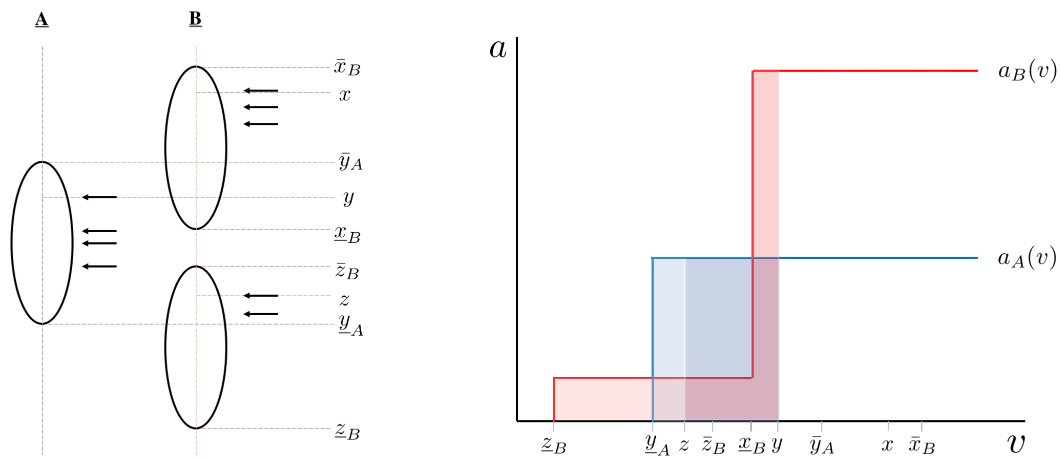

Figure 10: Left: A candidate dual (with no double or upper swaps); part of a chain. Right: An allocation that satisfies complementary slackness up to value y, satisfying equal preferability at z and y and preferability at all points with flow in.

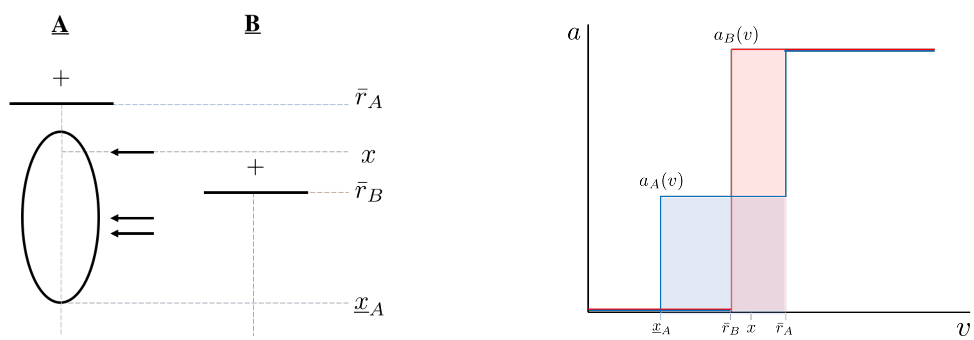

Figure 11: Left: The base case of a candidate dual with an empty chain. Right: An allocation that satisfies complementary slackness.

By the inductive hypothesis, we can come up with allocation rules $a_A(\cdot)$ and $a_B(\cdot)$ that satisfy complementary slackness to the same chain without the highest point (x, A), and will have $a_A(\underline{x}_A) = a_B(\bar{r}_B) = 1$. We construct an allocation rule $\hat{a}$ for the top chain of size $n+1$ as follows; this is depicted in Figure 8 in Section 4. Let $\lambda = \frac{x-\bar{r}_B}{x-y-a_B(\underline{y}_B)(\bar{r}_B-y)} < 1$. Then let

$$\hat{a}_A(v) = \begin{cases} 1 & v \geq \bar{r}_A \\ \lambda a_A(v) & \text{otherwise} \end{cases} \quad \text{and} \quad \hat{a}_B(v) = \begin{cases} 1 & v \geq \bar{r}_B \\ \lambda a_B(v) & \text{otherwise.} \end{cases}$$

This clearly satisfies (1-3). To show that (4) holds, we observe that at any previous point of concern $v < \bar{r}_B$, we had $u_A(v) = u_B(v)$. Now at those points, we have $\hat{u}_A(v) = \int_0^v \hat{a}_A(v)dv = \lambda \int_0^v a_A(v)dv = \lambda u_A(v)$. This holds for $\hat{u}_B(v) = \lambda u_B(v)$ as well. Thus, complementary slackness is still satisfied at all previous points $v \leq \bar{r}_B$; we only need to check equal utility at x.

$$\hat{u}_A(x) = \hat{u}_A(y) + \hat{a}_A(\underline{x}_A)(x-y) = \lambda u_A(y) + \lambda \cdot 1 \cdot (x-y)$$
$$\hat{u}_B(x) = \hat{u}_B(y) + \hat{a}_B(\underline{y}_B)(\bar{r}_B - y) + \hat{a}_B(\bar{r}_B)(x - \bar{r}_B) = \lambda u_B(y) + \lambda \cdot a_B(\underline{y}_B)(\bar{r}_B - y) + 1 \cdot (x - \bar{r}_B)$$

Then to have $\hat{u}_A(x) = \hat{u}_B(x)$, since $u_A(y) = u_B(y)$, we require that

$$\lambda(x-y) = \lambda \cdot a_B(\underline{y}_B)(\bar{r}_B - y) + 1 \cdot (x - \bar{r}_B).$$

The solution here is exactly the λ defined above.

Alternatively, by replacing x with $\bar{r}_A$, thus setting $\lambda = \frac{\bar{r}_A - \bar{r}_B}{\bar{r}_A - y - a_B(\underline{y}_B)(\bar{r}_B - y)}$, we get a solution that has $u_A(x) > u_B(x)$ and $u_A(\bar{r}_A) = u_B(\bar{r}_A)$ as required in (5).

Thus we have ensured that for top chains of all lengths, we can give an allocation rule that satisfies complementary slackness for all values from the bottom to the top of the chain. For v below the chain, $u_B(v) = u_A(v) = 0$, so we automatically satisfy complementary slackness. Above the chain, if we have used the alternate solution that (5) guarantees exists, we automatically satisfy complementary slackness for $v \geq \bar{r}_A$. This would only fail if there is flow into item B for $v \in [x, \bar{r}_A)$—that is, if the dual contains a upper swap, but by assumption it does not. Then for any dual solution with no double swaps or upper swaps, this algorithm gives an allocation rule that satisfies complementary slackness.

We prove that the menu complexity of the mechanism output by this algorithm is finite below:

Claim. The menu complexity is always finite.

Proof. Proof by contradiction. Suppose that there exists an instance such that the mechanism output by the algorithm has infinite menu complexity.

Note that this can only happen if the length of the top chain is infinity. Thus, there exists a sequence of points $(x_1, A), (x_2, B), (x_3, A), \cdots$ such that the point (x_i, A) is inside an ironed interval $[\underline{x_i}_A, \overline{x_i}_A]$ and $x_{i+1} \geq \underline{x_i}_A$. Analogous claims hold for an element (x_{i+1}, B) in the chain.

Thus, we have

$$x_1 \geq \bar{r}_B \geq \overline{x_2}_B \geq x_2 \geq \underline{x_1}_A \geq \overline{x_3}_A \geq x_3 \geq \cdots$$

Since the infinite sequence $x_1, x_2, \cdots$ is monotone and bounded, it converges to a limit, say x^*. Observe that x^* satisfies $\Phi_A^{\lambda,\alpha}(x^*) = \Phi_B^{\lambda,\alpha}(x^*) = 0$ and is unironed. This is because points arbitrarily close to it are unironed and are zeroes of $\Phi_A^{\lambda,\alpha}(\cdot)$ and $\Phi_B^{\lambda,\alpha}(\cdot)$. However, in this case, our algorithm just sets the price x^* and thus has constant menu complexity, a contradiction.

□

□

E A Candidate Dual for a Lower Bound on Menu-Complexity for the Multi-Unit Pricing Problem

Consider an MUP instance where the buyer can get one, two, or three copies of a given item. The relevant complementary slackness constraints in this setting go from

- **Rightwards.** For all v, from $(v,1) \to (v,2)$ and $(v,2) \to (v,3)$. This is because a buyer can always misreport and get more items.
- **Leftwards.** For all v, from $(v,2) \to (v/2,1)$ and $(v,3) \to (2v/3,2)$. This is because a buyer would prefer getting fewer items if they are available for much cheaper.

As shown in [DHP17], a buyer of type (v,C)'s utility for reporting $(v/2,A)$ is given by $u_A(v/2) = \int_0^{v/2} a_A(x)dx$. The same buyer's utility for reporting (v,B) is given by $u_B(v) = 2\int_0^v a_B(x)dx$.

To construct a lower bound for the MUP instance, we adapt our construction from the partially ordered case. We describe our construction formally below, but note here all the relevant differences. Observe that the incentive compatibility constraints for the MUP instance described above hide a partially ordered instance inside them. Indeed, the 'item' 2 is analogous to the item C, while the items A and B are the items 1 and 3 respectively. Just like the partially ordered instance, there are incentive compatibility constraints from $(v,2) \to (v,3)$ for all v. The only difference is that the constraints from $(v,2) \to (v,1)$ have been replaced by those from $(v,2) \to (v/2,1)$. Also, there are 'new' constraints from $(v,1) \to (v,2)$ and $(v,3) \to (2v/3,2)$.

We claim that, despite these changes, the essence of our argument there still holds. Roughly speaking, our argument there involved constructing a top-chain (see Definition 1) oscillating between items A and B. For any value x in this chain, we had flow coming from C to *both* A and B. Reasoning about complementary slackness constraints, then, gave us our lower bound.

For the MUP case, we can still do all the above things with the caveat that the value $(v/2,1)$ has to be treated as if it were $(v,1)$. An analogous master theorem can still be proved as the effect of the 'diagonal' flow on the virtual values is predictable. Using the master theorem, we can construct (essentially) any dual we want. Thus, we can have a feasible dual with a top-chain of an arbitrary length M oscillating between items 1 and 3. Also, we have flow from the item 2 to both 1 and 3 at all values in this top-chain. Chasing through the complementary slackness constraints in this dual again gives us a lower bound.

To highlight this analogy, in what follows, we use C instead of 2, A instead of 1, and B instead of 3.

Formally, we construct given an integer $M > 0$, a dual containing a top chain among A and B of length M. That is, a sequence of points $(x_1, A), (x_2, B), \ldots, (x_M, A)$ such that

$$\underline{x_M/2}_A < \underline{x_{M-1}}_B/2 < \cdots < \underline{x_2}_B/2 < \underline{x_1/2}_A.$$

In this dual, we have no extra space between the ironed intervals:

- $\underline{r}_A = \underline{x_M}_A/2$, $\overline{r}_A = \overline{x_1}_A/2$, and for i such that (x_i, A) and (x_{i+2}, A) are in the chain, $\underline{x_i/2}_A = \overline{x_{i+2}/2}_A$.
- $\underline{r}_B = \underline{x_M}_B$, $\overline{r}_B = \overline{x_2}_B$, and for i such that (x_i, B) and (x_{i+2}, B) are in the chain, $\underline{x_i}_B = \overline{x_{i+2}}_B$.

Recall that by definition of the $\bar{\cdot}$ and $\underline{\cdot}$ operators, $(\underline{x}_G, \overline{x}_G]$ is ironed in G. Also by our definitions, $f_G(v)\Phi_G^{\lambda,\alpha}(v) > 0$ for $v \geq \bar{r}_G$; $f_G(v)\Phi_G^{\lambda,\alpha}(v) = 0$ for $v \in [\underline{r}_G, \bar{r}_G]$; $f_G(v)\Phi_G^{\lambda,\alpha}(v) < 0$ for $v \leq \underline{r}_G$.

We will also define C to be DMR (and thus have no ironed intervals) with $\underline{r}_C = 2\underline{r}_A$ and $\bar{r}_C = 2\bar{r}_A$.

We adapt the flow from the partially ordered lower bound example: for any (x, G) in the chain, $\alpha_{C,A}(x \to x/2) > 0$ and $\alpha_{C,B}(x \to x) > 0$.

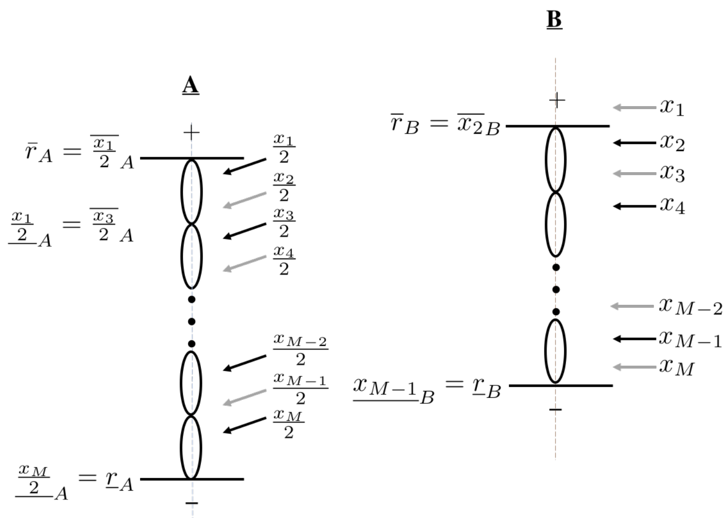

Figure 12: The analogue of the partially ordered candidate dual, adjusted for the Multi-Unit Pricing problem.

Theorem 12. *To satisfy complementary slackness with the candidate dual, the allocation requires M distinct allocation probabilities; the menu complexity is at least M.*

Proof. The proof is almost identical to that of Theorem 1. Using the constraint that the allocation can't increase in the middle of an ironed interval and that $u_A(x/2) = u_B(x)$ for all (x, G) in the chain, we show that the allocations must be non-zero throughout the chain.

Then, we show that for consecutive points in a chain $(x_i, A), (x_{i+1}, B)$ that $(1/2)a_A(x_i/2) > 2a_B(x_{i+1})$, and similarly, for $(x_i, B), (x_{i+1}, A)$, that $2a_B(x_i) > (1/2)a_A(x_{i+1}/2)$

This is enough to show that all of the menu options must be distinct, requiring meu complexity $\geq M$. □

F An Exact Characterization Under the Assumption of DMR

Recall from Subsection A.1 that when the distributions are DMR, $\lambda_G(v) = 0$ for all (v, G). Our main result is the following:

Theorem 13. *Consider a partially-ordered items setting that can be described by the following structure: There is one item G and an $n-1$ items H_i and $H_i \succ G$ for all i. If the marginal distribution for all the items is DMR, the optimal mechanism is deterministic.*

The key idea is to track a point in the lowest region of zeroes, r, and to always send flow to the items with lower zeroes. When the regions with virtual value zero overlap, we set r to a value where the smallest items are all zero, and continue to raise this point by splitting flow between all such items. This results in an optimal dual with which a deterministic mechanism satisfies complementary slackness.

We will define $S \subseteq N^+(G)$ to be the items such that for all $H \in S$, there exists v such that $\alpha_{G,H}(v) > 0$. We will define r such that, for all $H \in S$, r can always be set as a price for item H. Then $f_H(r)\Phi_H^{\lambda,\alpha}(r) = 0$.

Initially, $r = \min_{H \in N^+(G)} \bar{r}_H$, and $S = \arg\min_{H \in N^+(G)} \bar{r}_H$.

In every step, we decrease the amount of flow to send and increase r. We terminate when there is no flow left to send. The point r only increases and the set S only increases. We maintain the above properties.

First, we set the flow out of G:

$$\sum_{H \in N^+(G)} \alpha_{G,H}(v) = \begin{cases} 0 & v > \bar{r}_G \\ -\hat{R}''_G(v)/f_G(v) & v \leq \bar{r}_G. \end{cases}$$

Lemma 7. *We can always send α flow from G to $N^+(G)$ such that*

1. *If $\alpha_{G,H}(x) > 0$ for any x then $H \in S$.*
2. *If $H \in S$ then $f_H(r)\Phi_H^{\lambda,\alpha}(r) = 0$.*
3. *If $H \in N^+(G) \setminus S$ then $r < \underline{r}_H$ and thus $f_H(r)\Phi_H^{\lambda,\alpha}(r) < f_H(\underline{r}_H)\Phi_H^{\lambda,\alpha}(\underline{r}_H) = 0$.*

Proof. Suppose we have α flow to send at v.

Let $H' = \arg\min_{H' \in N^+(G) \setminus S} \underline{r}_{H'}$.

Choose ε such that by sending α flow to all items in S with correct proportions, we will maintain S and raise r by ε. That is,

$$\sum_{H \in S} f_H(r+\varepsilon)\Phi_H^{\lambda,\alpha}(r+\varepsilon) = \alpha.$$

If $r + \varepsilon < \underline{r}_{H'}$, we can send this flow without growing S.

By setting

$$\alpha_{G,H}(v) = f_H(r+\varepsilon)\Phi_H^{\lambda,\alpha}(r+\varepsilon) \quad \forall H \in N^+(G)$$

we ensure that after this update, $f_H(r+\varepsilon)\Phi_H^{\lambda,\alpha}(r+\varepsilon) = 0$ for all $H \in S$. Update $r \leftarrow r+\varepsilon$. Clearly (2) holds, and (3) holds since $r < \underline{r}_{H'} < \underline{r}_H$ for all $H \in N^+(G) \setminus S$.

Otherwise, $r + \varepsilon \geq \underline{r}_{H'}$ and $v \geq \underline{r}_{H'}$. Then we instead choose $\varepsilon = \underline{r}_{H'} - r$ and make the same update described above, then add H' to S. Note that we have sent positive flow but $< \alpha$. After the update, we will have $r \leftarrow r + \varepsilon = \underline{r}_{H'}$ and $f_H(r)\Phi_H^{\lambda,\alpha}(r) = 0$ for all $H \in S$, including H'. Then clearly (2) holds, and (3) holds since $r = \underline{r}_{H'} < \underline{r}_H$ for all $H \in N^+(G) \setminus S$.

Finally, (1) holds in both cases as we only send flow to elements of S and S is non-decreasing. □

Lemma 8. *For any v, even after setting α variables, $\lambda_G(v) = 0$.*

Proof. Since the flow out of G is chosen exactly to bring all virtual values to 0 below $\bar{r}_G$, no non-monotonicities are caused. □

Lemma 9. *For any $H \in \{H_i\}$, even after setting α variables, $\lambda_H(v) = 0$ for all v.*

Proof. Suppose we get flow α into H at x. Every value $v \leq x$ has $f_H(v)\Phi_H^{\lambda,\alpha}(v)$ decrease by α while this remains unchanged for $v > x$, causing no non-monotonicities.

□

Lemma 10. *The following deterministic allocation rule always satisfies complementary slackness with the dual:*

- *If* $H \in S$*, set a price of* r*.*
- *If* $H \in N^+(G) \setminus S$*, set a price of* $\bar{r}_H$*.*
- *Set a price of* $\min\{r, \bar{r}_G\}$ *for item* G*.*

this is because if $\alpha_{G,H} > 0$ *then* H *is preferable among* $N^+(G)$*.*

Proof. From DMR and our setting of λ, we will have $\lambda_G(v) = 0$ for all (v, G), automatically satisfying complementary slackness for these variables. Further, even after sending α flow, $f_G(\cdot)\Phi_G^{\lambda,\alpha}(\cdot)$ will be properly monotone for all G by Lemma 8 and Lemma 9.

First, we verify that the when we set a price, the virtual values are 0 at that price, so we have the freedom to do so.

By Lemma 7, $f_H(r)\Phi_H^{\lambda,\alpha}(r) = 0$ for all $H \in S$. Of course, by definition of $\bar{r}$, $f_H(\bar{r}_H)\Phi_H^{\lambda,\alpha}(\bar{r}_H) = 0$. In addition, by definition of the flow out of G, $f_G(v)\Phi_G^{\lambda,\alpha}(v) = 0$ for all $v \le \bar{r}_G$ so $f_G(r)\Phi_G^{\lambda,\alpha}(r) = 0$. Then all of the prices posted are viable.

It remains to choose a mechanism that satisfies complementary slackness with the α variables. If $\alpha_{G,H}(v) > 0$ for some v then we know that (1) $H \in S_G$ and (2) $v < \bar{r}_G$.

Case 1: $r \le \bar{r}_G$. If $\alpha_{G,H}(v) > 0$ for some v then $H \in S_G$, and thus $u_G(v) = u_H(v)$ for all v, so we automatically satisfy complementary slackness.

Case 2: $\bar{r}_G < r$. Then for any $v < \bar{r}_G$ where $\alpha_{G,H}(v) > 0$, $u_G(v) = u_H(v) = 0$, so there is tight (zero) utility respecting this complementary slackness constraint. □