

A Duality-Based Unified Approach to Bayesian Mechanism Design

Yang Cai*
McGill University, Canada
cai@cs.mcgill.ca

Nikhil R. Devanur†
Microsoft Research, USA
nikdev@microsoft.com

S. Matthew Weinberg‡
Princeton University, USA
smweinberg@princeton.edu.

April 15, 2018

Abstract

We provide a unified view of many recent developments in Bayesian mechanism design, including the black-box reductions of Cai et al. [CDW13b], simple auctions for additive buyers [HN12], and posted-price mechanisms for unit-demand buyers [CHK07]. Additionally, we show that viewing these three previously disjoint lines of work through the same lens leads to new developments as well. First, we provide a duality framework for Bayesian mechanism design, which naturally accommodates multiple agents and arbitrary objectives/feasibility constraints. Using this, we prove that either a posted-price mechanism or the Vickrey-Clarke-Groves auction with per-bidder entry fees achieves a constant-factor of the optimal revenue achievable by a Bayesian Incentive Compatible mechanism whenever buyers are unit-demand or additive, unifying previous breakthroughs of Chawla et al. [CHMS10] and Yao [Yao15], and improving both approximation ratios (from 30 to 24 and 69 to 8, respectively). Finally, we show that this view also leads to improved structural characterizations in the Cai et al. framework.

*Supported by NSERC Discovery RGPIN-2015-06127. Work done in part while the author was a Research Fellow at the Simons Institute for the Theory of Computing.

†Work done in part while the author was visiting the Simons Institute for the Theory of Computing.

‡Work done in part while the author was a Research Fellow at the Simons Institute for the Theory of Computing.

¹An earlier version of this work appeared under the same title in the proceedings of STOC 2016.

1 Introduction

In the past several years, we have seen a tremendous advance in the field of Bayesian Mechanism Design, based on ideas and concepts rooted in Theoretical Computer Science (TCS). For instance, due to a line of work initiated by Chawla, Hartline, and Kleinberg [CHK07], we now know that posted-price mechanisms are approximately optimal with respect to the optimal Bayesian Incentive Compatible² (BIC) mechanism whenever buyers are unit-demand,³ and values are independent⁴ [CHMS10, CMS15, KW12]. Due to a line of work initiated by Hart and Nisan [HN12], we now know that either running Myerson’s auction separately for each item or running the VCG mechanism with a per-bidder entry fee⁵ is approximately optimal with respect to the optimal BIC mechanism whenever buyers are additive, and values are independent [LY13, BILW14, Yao15]. Due to a line of work initiated by Cai et al. [CDW12a], we now know that optimal mechanisms are distributions over virtual welfare maximizers, and have computationally efficient algorithms to find them in quite general settings [CDW12b, CDW13a, CDW13b, BGM13, DW15, DDW15]. The main contribution of this work is a unified approach to these three previously disjoint research directions. At a high level, we show how a new interpretation of the Cai-Daskalakis-Weinberg (CDW) framework provides us a duality theory, which then allows us to strengthen the characterization results of Cai et al., as well as interpret the benchmarks used in [CHK07, CHMS10, CMS15, KW12, HN12, CH13, LY13, BILW14] as dual solutions. Surprisingly, we learn that *essentially the same dual solution* yields all the key benchmarks in these works. We show how to extend this dual solution to multi-bidder settings, and analyze the mechanisms developed in [CHMS10, Yao15] with respect to the resulting benchmarks. In both cases, our analysis yields improved approximation ratios.

1.1 Simple vs. Optimal Auction Design

It is well-known by now that optimal multi-item auctions suffer many properties that are undesirable in practice. For example, with just a single additive buyer and two items, the optimal auction could be randomized [Tha04, Pav11]. Moreover, there exist instances where the buyer’s two values are drawn from a correlated distribution where the optimal revenue achieves infinite revenue while the best deterministic mechanism achieves revenue ≤ 1 [BCKW10, HN13]. Even when the two item values are drawn independently, the optimal mechanism might offer uncountably many different randomized options for the buyer to choose from [DDT13]. Additionally, revenue-optimal multi-item auctions behave non-monotonically: there exist distributions F and F^+ , where F^+ stochastically dominates F , such that the revenue-optimal auction when a single additive buyer’s values for two items are drawn from $F \times F$ achieves strictly larger revenue than the revenue-optimal auction when a single additive buyer’s values are drawn from $F^+ \times F^+$ [HR12]. Finally, it is known that revenue-optimal auctions may not be DSIC [Yao17], and are also #P-hard to find [DDT14].

In light of the aforementioned properties, *simple* mechanisms are often used in lieu of *optimal* mechanisms in practice, and an active line of research coined “simple versus optimal” mechanism design [HR09] aims to rigorously understand when simple mechanisms are appropriate in practice. Still, prior work essentially shows that simple mechanisms are never *exactly* optimal, so the main goal of these works is to

²A mechanism is Bayesian Incentive Compatible (BIC) if it is in every bidder’s interest to tell the truth, assuming that all other bidders’ reported their true values. A mechanism is Dominant Strategy Incentive Compatible (DSIC) if it is in every bidder’s interest to tell the truth *no matter what reports the other bidders make*.

³A valuation is unit-demand if $v(S) = \max_{i \in S} \{v(\{i\})\}$. A valuation is additive if $v(S) = \sum_{i \in S} v(\{i\})$.

⁴That is, the random variables $\{v_{ij}\}_{i,j}$ are independent (where v_{ij} denotes bidder i ’s value for item j).

⁵By this, we mean that the mechanism offers each bidder the option to participate for b_i , which might depend on the other bidders’ bids but not bidder i ’s. If they choose to participate, then they play in the VCG auction (and pay any additional prices that VCG charges them).

understand when simple mechanisms are *approximately* optimal.⁶ Some of the most exciting contributions from TCS to Bayesian mechanism design have come from this direction, and include a line of work initiated by Chawla et al. [CHK07] for unit-demand buyers, and Hart and Nisan [HN12] for additive buyers.

In a setting with m heterogeneous items for sale and n *unit-demand* buyers whose values for the items are drawn independently, the state-of-the-art shows that a simple posted-price mechanism⁷ obtains a constant factor of the optimal BIC revenue (the revenue of the optimal BIC mechanism) [CHK07, CHMS10, CMS15, KW12]. The main idea behind these works is a multi- to single-dimensional reduction. They consider a related setting where each bidder is split into m separate copies, one for each item, with bidder i 's copy j interested only in item j . The value distributions are the same as the original multi-dimensional setting. One key ingredient driving these works is that the optimal revenue in the original setting is upper bounded by a small constant times the optimal revenue in the copies setting.

In a setting with m heterogeneous items for sale and n *additive* buyers whose values for the items are drawn independently, the state-of-the-art result shows that for all inputs, either running Myerson's optimal single-item auction for each item separately or running the VCG auction with a per-bidder entry fee obtains a constant factor of the optimal BIC revenue [HN12, LY13, BILW14, Yao15]. One main idea behind these works is a "core-tail decomposition", that breaks the revenue down into cases where the buyers have either low (the core) or high (the tail) values.

Although these two approaches appear different at first, we are able to show that they in fact arise from *basically the same dual* in our duality theory. Essentially, we show that a specific dual solution within our framework gives rise to an upper bound that decomposes into the sum of two terms, one that looks like the the copies benchmark, and one that looks like the core-tail benchmark. In terms of concrete results, this new understanding yields improved approximation ratios on both fronts. For additive buyers, we improve the ratio provided by Yao [Yao15] from 69 to 8. For unit-demand buyers, we improve the approximation ratio provided by Chawla et al. [CHMS10] from 30 to 24.

In addition to these concrete results, our work makes the following conceptual contributions as well. First, while the single-buyer core-tail decomposition techniques (first introduced by Li and Yao [LY13]) are now becoming standard [LY13, BILW14, RW15, BDHS15], they do not generalize naturally to multiple buyers. Yao [Yao15] introduced new techniques in his extension to multi-buyers termed " β -adjusted revenue" and " β -exclusive mechanisms," which are technically quite involved. Our duality-based proof can be viewed as a natural generalization of the core-tail decomposition to multi-buyer settings. Second, we use basically the same analysis for both additive and unit-demand valuations, meaning that our framework provides a unified approach to tackle both settings. Finally, we wish to point out that the key difference between our proofs and those of [CHMS10, BILW14, Yao15] are our duality-based benchmarks: we are able to immediately get more mileage out of these benchmarks while barely needing to develop new approximation techniques. Indeed, the bulk of the work is in properly decomposing our benchmarks into terms that can be approximated using ideas similar to prior work. All these suggest that our techniques are likely be useful in more general settings (and indeed, they have been: see Section 1.3).

⁶On this front, one should not interpret (say) an 8-approximation as suggesting that sellers should be happy with $1/8$ of the revenue they could potentially achieve. Rather, these guarantees are meant to be interpreted more qualitatively, and suggest claims like "If simple auction A guarantees a small constant-factor approximation in the worst-case, but simple auction B does not, maybe it's safer to use auction A in practice."

⁷A posted price mechanism visits each buyer one at a time and posts a price for each item. The buyer can then select any subset of items and pay the corresponding prices. Observe that such a mechanism is DSIC.

1.2 Optimal Multi-Dimensional Mechanism Design

Another recent contribution of the TCS community is the CDW framework for generic Bayesian mechanism design problems. Here, it is shown that Bayesian mechanism design problems for essentially any objective can be solved with black-box access just to an *algorithm* that optimizes a perturbed version of that same objective. That is, even though the original mechanism design problem involves incentives, the optimal BIC mechanism can be found via black-box queries to an algorithm (where the input is known/given and there are no incentives), but this algorithm optimizes a perturbed objective instead. One aspect of this line of work is computational: we now have computationally efficient algorithms to find the optimal (or approximately optimal) mechanism in numerous settings of interest (including the aforementioned cases of many additive/unit-demand buyers, but significantly more general as well). Another aspect is structural: we now know, for instance, that in all settings that fit into this framework, the revenue-optimal mechanism is a distribution over *virtual welfare optimizers*.⁸ A mechanism is a virtual welfare optimizer if it pointwise optimizes the *virtual welfare* (that is, on every input, it selects an outcome that maximizes the virtual welfare). The virtual welfare is given by a *virtual valuation/transformation*, which is a mapping from valuations to linear combinations of valuations.

The structural characterization from previous work roughly ends here: the guaranteed virtual transformations were randomized with no promise any additional properties beyond their existence (and that they could be found in poly-time). Our contribution to this line of work is to improve the existing structural characterization. Specifically, we show that every instance has a strong dual in the form of n disjoint flows, one for each agent. The nodes in agent i 's flow correspond to possible types of this agent,⁹ and non-zero flow from type $t_i(\cdot)$ to $t'_i(\cdot)$ captures that the incentive constraint between $t_i(\cdot)$ and $t'_i(\cdot)$ binds. We show how a flow induces a virtual transformation, and that the optimal dual gives a virtual valuation function such that:

1. This virtual valuation function is deterministic and can be found computationally efficiently.
2. The optimal mechanism has expected revenue = its expected virtual welfare, and every BIC mechanism has expected revenue \leq its expected virtual welfare.
3. The optimal mechanism optimizes virtual welfare pointwise (i.e. on every input, the virtual welfare maximizing outcome is selected).¹⁰

Here are a few examples of the benefits of such a characterization (which cannot be deduced from [CDW13b]). First, the promised virtual valuation function certifies the optimality of the optimal mechanism: every BIC mechanism has expected revenue \leq its expected virtual welfare, yet the optimal mechanism maximizes virtual welfare pointwise. Second, by looking at the promised virtual valuation function, we can immediately determine which incentive constraints “matter.” Specifically, if the flow corresponding to the promised virtual valuation function sends flow from $t_i(\cdot)$ to $t'_i(\cdot)$, then removing the constraint guaranteeing that $t_i(\cdot)$ prefers to tell the truth rather than report $t'_i(\cdot)$ (e.g. through some form of verification) would increase the optimal achievable revenue.¹¹ Such a characterization should prove a valuable analytical tool for multi-item auctions, akin to Myerson’s virtual values for single-dimensional settings [Mye81].

⁸Their reduction applies to objectives beyond revenue, such as makespan. The focus of the present paper is on revenue, so we only focus on the projection of their results onto this setting.

⁹Both the CDW framework and our duality theory only apply directly if there are finitely many possible types for each agent.

¹⁰This could be *randomized*; there is always a deterministic maximizer but in cases where the optimal mechanism is randomized, the virtual transformations are such that there are numerous maximizers, and the optimal mechanism selects one of them from a particular probability distribution.

¹¹This claim is only guaranteed to be true in non-degenerate instances with a unique optimal dual - and exactly results from the fact that relaxing tight constraints in non-degenerate LPs improves the optimal solution.

1.3 Related Work

1.3.1 Duality Frameworks

Recently, strong duality frameworks for a single additive buyer were developed in [DDT13, DDT15, DDT16, GK14, Gia14, GK15]. These frameworks show that the dual problem to revenue optimization for a single additive buyer can be interpreted as an optimal transport/bipartite matching problem. Work of Hartline and Haghpanah also provides an alternative “path-finding” duality framework for a single additive or unit-demand buyer, and has a more similar flavor to ours (as flows can be interpreted as distributions over paths) [HH15]. When they exist, these paths provide a witness that a certain Myerson-type mechanism is optimal, but the paths are not guaranteed to exist in all instances. Also similar is independent work of Carroll, which also makes use of a partial Lagrangian over incentive constraints, again for a single additive buyer [Car16]. In addition to their mathematical beauty, these duality frameworks also serve as tools to prove that mechanisms are optimal. These tools have been successfully applied to provide conditions when pricing only the grand bundle (give the buyer the choice only to buy everything or nothing) [DDT13], posting a uniform item pricing (post the same price on every item) [HH15], or even employing a randomized mechanism [GK15] is optimal when selling to a single additive or unit-demand buyer. However, none of these frameworks currently applies in multi-bidder settings, and to date have been unable to yield any approximate optimality results in the (single bidder) settings where they do apply.

We also wish to argue that our duality is perhaps more transparent than existing theories. For instance, it is easy to interpret dual solutions in our framework as virtual valuation functions, and dual solutions for multiple buyer instances just list a dual for each single buyer. In addition, we are able to re-derive and improve the breakthrough results of [CHK07, CHMS10, CMS15, HN12, LY13, BILW14, Yao15] *using essentially the same dual solution*. Still, it is not our goal to subsume previous duality theories, and our new theory certainly doesn’t. For instance, previous frameworks are capable of proving that a mechanism is *exactly* optimal when the input distributions are continuous. Our theory as-is can only handle distributions with finite support exactly.¹² However, we have demonstrated that there is at least one important domain (simple and approximately optimal mechanisms) where our theory seems to be more applicable.

1.3.2 Related Techniques

Techniques similar to ours have appeared in prior works as well. For instance, the idea to use Lagrangian multipliers/LP duality for mechanism design dates back at least to early work of Laffont and Robert for selling a single item to budget-constrained bidders [LR98], is discussed extensively for instance in [Mye97, Voh11], and also used for example in recent works as well [BGM13, Voh12]. It is also apparently informal knowledge among some economists that Myerson’s seminal result [Mye81] can be proved using some form of LP duality, and some versions of these proofs have been published as well (e.g. [MV04]). Still, we include in Section 4 a proof of [Mye81] in our framework to serve as a warm-up (and because some elements of the proof are simplified via our approach).

The idea to use “paths” of incentive compatibility constraints to upper bound revenue in single-bidder problems dates back at least to work of Rochet and Choné studying revenue optimization in general multi-item settings [RC98], and Armstrong [Arm96, Arm99], which studies the special case of a single bidder and two items. More recently, Cai et al. use this approach to prove hardness of approximation for a single bidder with submodular valuations for multiple items [CDW13b], Hartline and Haghpanah provide sufficient conditions for especially simple mechanisms to be optimal for a single unit-demand or additive bidder [HH15], and Carroll proves that selling separately is max-min optimal for a single additive buyer when only the

¹²Our theory can still handle continuous distributions arbitrarily well. See Section 2.

marginals are known but not the (possibly correlated) joint value distribution [Car16]. Indeed, many of these works also observe that the term “virtual welfare” is appropriate to describe the resulting upper bounds. Still, these works focus exclusively on providing conditions for certain mechanisms to be *exactly optimal*, and therefore impose some technical conditions on the settings where they apply. In comparison, our work pushes the boundaries by accommodating both approximation (in the sense that our framework can prove that simple mechanisms are approximately optimal and not just that optimal mechanisms are optimal) and unrestricted settings (in the sense that our framework isn’t restricted to a single buyer, or additive/unit-demand valuations).

Finally, we note that some of the benchmarks used in later sections can be derived without appealing to duality [CMS15]. Therefore, a duality theory is not “necessary” in order to obtain our benchmarks. Still, prior to our work it was unknown that these benchmarks were at all useful outside of the unit-demand settings for which they were developed. Additionally, both the primal and dual understanding of these benchmarks is valuable for extending the state-of-the-art, discussed in more detail below. Prior work has also obtained approximately optimal auctions via some sort of “benchmark decomposition” in the unrelated setting of digital goods [CGL15].

1.3.3 Approximation in Multi-Dimensional Mechanism Design

Finally, we provide a brief overview of recent work providing simple and approximately optimal mechanisms in multi-item settings. Seminal work of Chawla, Hartline, and Kleinberg proves that a posted-price mechanism gets a 3-approximation to the optimal deterministic mechanism for a single unit-demand buyer with independently drawn item values [CHK07]. Chawla et al. improve the ratio to 2, and prove a bound of 6.75 against the optimal deterministic, DSIC mechanism for multiple buyers [CHMS10]. Chawla, Malec, and Sivan show that the bound degrades by at most a factor of 5 when comparing to the optimal randomized, BIC mechanism [CMS15]. Kleinberg and Weinberg improve the bound of 6.75 to 6 [KW12]. Roughgarden, Talgam-Cohen and Yan provide a prior-independent “supply-limiting” mechanism in this setting, and also prove a Bulow-Klemperer [BK96] result: the VCG mechanism with additional bidders yields more expected revenue than the optimal deterministic, DSIC mechanism (with the original number of bidders), when bidder valuations are unit-demand, i.i.d., and values for items are regular and independent (possibly asymmetric) [RTCY12]. All of these results get mileage from the “OPT^{COPIES}” benchmark initiated in [CHK07].

More recent influential work of Hart and Nisan proves that selling each item separately at its Myerson reserve¹³ gets an $O(\log^2 m)$ approximation to the optimal mechanism for a single additive buyer and m independent (possibly asymmetric) items [HN12]. Li and Yao improve this to $O(\log m)$, which is tight [LY13]. Babaioff et al. prove that the better of selling separately and bundling together gets a 6-approximation [BILW14]. Bateni et al. extend this to a model of limited correlation [BDHS15]. Rubinfeld and Weinberg extend this to a single buyer with “subadditive valuations over independent items” [RW15]. Yao shows that the better of selling each item separately using Myerson’s auction and running VCG with a per-bidder entry fee gets a 69-approximation when there are many additive buyers and all values for all items are independent [Yao15]. Goldner and Karlin show how to use these results to obtain approximately optimal prior-independent mechanisms for many additive buyers [GK16]. More recently, Chawla and Miller show that a posted-price mechanism with per-bidder entry fee gets a constant-factor approximation for many bidders with “additive valuations subject to matroid constraints” [CM16]. All of these results get mileage from the “core-tail decomposition” initiated in [LY13].

¹³The Myerson reserve of a one-dimensional distribution refers to the revenue-optimal price to set if one seller is selling only this item to a single buyer.

The present paper unifies these two lines of work by showing that the $\text{OPT}^{\text{COPIES}}$ benchmark and the core-tail decomposition both arise from essentially the same dual in our duality theory. These benchmarks provide a necessary starting point for the above results, but proving guarantees against these benchmarks of course still requires significant work. We believe that our duality theory now provides the necessary starting point to extend these results to much more general settings, as evidenced by the follow-up works discussed below.

1.3.4 Subsequent Work

Since the presentation of an earlier version of this work at STOC 2016, numerous follow-up works have successfully made use of our framework to design (approximately) optimal auctions in much more general settings. For example, Cai and Zhao show that the better of a posted-price mechanism and an anonymous posted-price mechanism with per-bidder entry fee gets a constant-factor approximation for many bidders with “XOS valuations over independent items” [CZ17]. This extends the previous state-of-the-art [CM16] from Gross Substitutes to XOS valuations. Eden et al. show that the better of selling separately and bundling together gets an $O(d)$ -approximation for a single bidder with “complementarity- d valuations over independent items” [EFF⁺17b]. The same authors also prove a Bulow-Klemperer result: the VCG mechanism with additional bidders yields more expected revenue than the optimal randomized, BIC mechanism (with the original number of bidders), when bidder valuations are “additive subject to downward closed constraints,” i.i.d., and values for items are regular and independent (possibly asymmetric) [EFF⁺17a]. Brustle et al. design a simple mechanism that achieves $\frac{1}{2}$ of the optimal gains from trade in certain *two-sided markets*, such as bilateral trading and double auctions [BCWZ17]. Devanur and Weinberg provide an alternative proof of Fiat et al.’s solution to the “FedEx Problem,” and extend it to design the optimal auction for a single buyer with a private budget [FGKK16, DW17]. Finally, Liu and Psomas provide a Bulow-Klemperer result for *dynamic auctions* [LP16], and Fu et al. design approximately optimal BIC mechanisms for correlated bidders [FLLT17].

Organization. We provide preliminaries and notation below. In Section 3, we present our duality theory for revenue maximization in the special case of additive/unit-demand bidders. In Section 4, we present a duality proof of Myerson’s seminal result, and in Section 5 we present a canonical dual solution that proves useful in different settings. As a warm-up, we show in Section 6 how to analyze this dual solution when there is just a single buyer. In Section 7, we provide the multi-bidder analysis, which is more technical. In Section 8, we conclude with a formal statement of our duality theory in general settings.

2 Preliminaries

Optimal Auction Design. For the bulk of the paper, we will study the following setting (in Section 8, we will show that our duality theory holds much more generally). The buyers (we will use the terms buyer and bidder interchangeably) are either all unit-demand or all additive, with buyer i having value t_{ij} for item j . Recall that a valuation is unit-demand if $v(S) = \max_{i \in S} \{v(\{i\})\}$ and a valuation is additive if $v(S) = \sum_{i \in S} v(\{i\})$. We use $t_i = (t_{i1}, \dots, t_{im})$ to denote buyer i ’s values for all the goods and t_{-i} to denote every buyer except i ’s values for all the goods. T_{ij} is the set of all possible values of buyer i for item j , $T_i = \times_j T_{ij}$, $T_{-i} = \times_{i^* \neq i} T_{i^*}$ and $T = \times_i T_i$. All values for all items are drawn independently. We denote by D_{ij} the distribution of t_{ij} , $D_i = \times_j D_{ij}$, $D_{i,-j} = \times_{j^* \neq j} D_{ij^*}$, $D = \times_i D_i$, and $D_{-i} = \times_{i^* \neq i} D_{i^*}$, and f_{ij} (or $f_i, f_{i,-j}, f_{-i}$, etc.) the densities of these finite-support distributions (that is, $f_i(x) = \Pr_{t_i \sim D_i}[t_i = x]$).

We define \mathcal{F} to be a set system over $[n] \times [m]$ that describes all feasible allocations.¹⁴

A *mechanism* takes as input a reported type from each bidder and selects (possibly randomly) an outcome in \mathcal{F} , and payments to charge the bidders. A mechanism is Bayesian Incentive Compatible (BIC) if it is in each buyers interest to report their true type, assuming that the other buyers do so as well, and Bayesian Individually Rational (BIR) if each buyer gets non-negative utility for reporting their true type (assuming that the other bidders do so as well). The revenue of an auction is simply the expected sum of payments made when bidders drawn from D report their true values. The *optimal auction* optimizes expected revenue over all BIC and BIR mechanisms. For a given value distribution D , we denote by $\text{REV}(D)$ the expected revenue achieved by this auction, and it will be clear from context whether buyers are additive or unit-demand. For a specific BIC mechanism M , we will also use $\text{REV}^M(D)$ to denote the expected revenue achieved by M when bidders with valuations drawn from D report truthfully.

Reduced Forms. The reduced form of an auction stores for all bidders i , items j , and types t_i , the probability that bidder i will receive item j when reporting t_i to the mechanism (over the randomness in the mechanism and randomness in other bidders' reported types, assuming they come from D_{-i}) as $\pi_{ij}(t_i)$. It is easy to see that if a buyer is additive, or unit-demand and receives only one item at a time, that their expected value for reporting type t'_i to the mechanism is just $t_i \cdot \pi_i(t'_i)$ (where we treat t_i and $\pi_i(t'_i)$ as vectors, and \cdot denotes a vector dot-product). We say that a reduced form is *feasible* if there exists some feasible mechanism (that ex-post selects an outcome in \mathcal{F} with probability 1) that matches the probabilities promised by the reduced form. If $P(\mathcal{F}, D)$ is defined to be the set of all feasible reduced forms, it is easy to see (and shown in [CDW12a], for instance) that $P(\mathcal{F}, D)$ is closed and convex.

We will also use $p_i(t_i)$ to refer to the expected payment made by bidder i when reporting t_i to the mechanism (over the randomness in the mechanism and randomness in other bidders' reported types, assuming they come from D_{-i}).

Simple Mechanisms. Even though the benchmark we target is the optimal *randomized* BIC mechanism, the simple mechanisms we design will all be deterministic and satisfy DSIC. For a single buyer, the two mechanisms we consider are selling separately and selling together. Selling separately posts a price p_j on each item j and lets the buyer purchase whatever subset of items she pleases. We denote by $\text{SREV}(D)$ the revenue of the optimal such pricing. Selling together posts a single price p on the grand bundle, and lets the buyer purchase the entire bundle for p or nothing. We denote by $\text{BREV}(D)$ the revenue of the optimal such pricing. For multiple buyers the generalization of selling together is the VCG mechanism with an entry fee, which offers to each bidder i the opportunity to pay an entry fee $e_i(t_{-i})$ and participate in the VCG mechanism (paying any additional fees charged by the VCG mechanism). If they choose not to pay the entry fee, they pay nothing and receive nothing. We denote the revenue of the mechanism that charges the optimal entry fees to the buyers as $\text{BVCG}(D)$, and $\text{VCG}(D)$ the revenue of the VCG mechanism with no entry fees. The generalization of selling separately is a little different, and described immediately below.

Single-Dimensional Copies. A benchmark that shows up in our decompositions relates the multi-dimensional instances we care about to a single-dimensional setting, and originated in work of Chawla et. al. [CHK07]. For any multi-dimensional instance D we can imagine splitting bidder i into m different copies, with bidder i 's copy j interested only in receiving item j and nothing else. So in this new instance there are nm single-dimensional bidders, and copy (i, j) 's value for winning is t_{ij} (which is still drawn from D_{ij}). The set system \mathcal{F} from the original setting now specifies which copies can simultaneously win. We denote by $\text{OPT}^{\text{COPIES}}(D)$ the revenue of Myerson's optimal auction [Mye81] in the copies setting induced by D .¹⁵

¹⁴When bidders are additive, \mathcal{F} only allows allocating each item at most once. When bidders are unit-demand, \mathcal{F} contains all matchings between the bidders and the items.

¹⁵Note that when buyers are additive that $\text{OPT}^{\text{COPIES}}$ is exactly the revenue of selling items separately using Myerson's optimal

Continuous versus Finite-Support Distributions. Our approach explicitly assumes that the input distributions have finite support. This is a standard assumption when computation is involved. However, most existing works in the simple vs. optimal paradigm hold even for continuous distributions (including [CHK07, CHMS10, CMS15, HN12, LY13, BILW14, Yao15, RW15, BDHS15]). Fortunately, it is known that every D can be discretized into D^+ such that $\text{REV}(D) \in [(1 - \epsilon)\text{REV}(D^+), (1 + \epsilon)\text{REV}(D^+)]$, and D^+ has finite support. So all of our results can be made arbitrarily close to exact for continuous distributions. We conclude this section by proving this formally, making use of the following theorem proved in [RW15], which draws from prior works [HL10, HKM11, BH11, DW12]. Note that the theorem below holds for distributions over arbitrary valuation functions $t_i(\cdot)$, and not just additive/unit-demand.

Theorem 1. [RW15, DW12] *Let M be any BIC mechanism for values drawn from distribution D , and for all i , let D_i and D_i^+ be any two distributions, with coupled samples $t_i(\cdot)$ and $t_i^+(\cdot)$ such that $t_i^+(x) \geq t_i(x)$ for all $x \in \mathcal{F}$. If $\delta_i(\cdot) = t_i^+(\cdot) - t_i(\cdot)$, then for any $\epsilon > 0$, $\text{REV}^M(D^+) \geq (1 - \epsilon)(\text{REV}^M(D) - \frac{\text{VAL}(\delta)}{\epsilon})$, where $\text{VAL}(\delta)$ denotes the expected welfare of the VCG allocation when buyer i 's type is drawn according to the random variable $\delta_i(\cdot)$.*

To see how this implies that our duality is arbitrarily close to exact for continuous distributions, let D_i^ϵ be the distribution that first samples $t_i(\cdot)$ from D_i , then outputs $t_i^\epsilon(\cdot)$ such that $t_i^\epsilon(x) = t_i(x) \cdot \mathbb{1}(t_i([m]) \leq 1/\epsilon)$.¹⁶ It is easy to see that as $\epsilon \rightarrow 0$, $\text{REV}^M(D^\epsilon) \rightarrow \text{REV}^M(D)$: for every mechanism M and every $\eta > 0$, there exists an $\epsilon > 0$ such that a $(1 - \eta)$ fraction of M 's revenue when buyers' types are drawn from D comes from buyers with $t_i([m]) \leq 1/\epsilon$.¹⁷ For the chosen ϵ , M is still a BIC mechanism when buyers' types are drawn from D^ϵ , and its revenue under D^ϵ is at least $(1 - \eta)$ fraction of its revenue under D . So we can get arbitrarily close while only considering distributions that are bounded.

Now for any bounded distribution D_i , define $D_i^{+, \epsilon}$ to first sample $t_i(\cdot)$ from D_i , then output $t_i^{+, \epsilon}(\cdot)$ such that $t_i^{+, \epsilon}(x) = \epsilon^2 \cdot \lceil t_i(x)/\epsilon^2 \rceil$. Similarly define $D_i^{-, \epsilon}$ to first sample $t_i(\cdot)$ from D_i , then output $t_i^{-, \epsilon}(\cdot)$ such that $t_i^{-, \epsilon}(x) = \epsilon^2 \cdot (\lceil t_i(x)/\epsilon^2 \rceil - 1)$. Then it's clear that $D_i^{+, \epsilon}, D_i$, and $D_i^{-, \epsilon}$ can be coupled so that $t_i^{+, \epsilon}(x) \geq t_i(x) \geq t_i^{-, \epsilon}(x)$ for all x , and that taking either of the two consecutive differences results in a $\delta_i(\cdot)$ such that $\delta_i(x) \leq \epsilon^2$ for all x . So applying Theorem 1, we see that for any desired ϵ , we have $\text{REV}^M(D) \in [(1 - \epsilon)(\text{REV}^M(D^{-, \epsilon} - n\epsilon)), (1 + \epsilon)(\text{REV}^M(D^{+, \epsilon} + n\epsilon))]$ for all BIC mechanisms M , including the optimal mechanism. Finally, we just observe that $\text{REV}(D^{+, \epsilon}) = \text{REV}(D^{-, \epsilon}) + n\epsilon$ (the revenue of the optimal mechanism increases by exactly $n\epsilon$ going from $D^{-, \epsilon}$ to $D^{+, \epsilon}$), as every buyer values every outcome at exactly ϵ more in $D^{+, \epsilon}$ versus $D^{-, \epsilon}$. So as $\epsilon \rightarrow 0$, both approach $\text{REV}(D)$. Note that both $D^{+, \epsilon}$ and $D^{-, \epsilon}$ have finite support, so our theory will directly design M that achieve constant-factor approximations for $\text{REV}(D^{-, \epsilon})$. By the exact same reasoning above, we also have $\text{REV}^M(D^{-, \epsilon}) \rightarrow \text{REV}^M(D)$, and we achieve exactly the same guarantees for the original (possibly continuous) D .

3 Our Duality Theory

In this section we provide our duality framework, specialized to unit-demand/additive bidders. We begin by writing the linear program (LP) for revenue maximization (Figure 1). For ease of notation, assume that there is a special type \emptyset to represent the option of not participating in the auction. That means $\pi_i(\emptyset) = \mathbf{0}$ and $p_i(\emptyset) = 0$. Now a Bayesian IR (BIR) constraint is simply another BIC constraint: for any type t_i , bidder i

auction in the original setting.

¹⁶That is, if the value of t_i for the grand bundle satisfies $t_i([m]) \leq 1/\epsilon$, then $t_i^\epsilon(x) = t_i(x)$. Otherwise, $t_i^\epsilon(x) = 0$ for all x .

¹⁷Similarly, if M achieves infinite revenue, then for every $\eta > 0$, there exists an $\epsilon > 0$ such that the revenue of M when buyers' types are drawn from D coming from buyers with $t_i([m]) \leq 1/\epsilon$ is at least $1/\eta$. So our approach will still show that whenever $\text{REV}(D)$ is infinite, the revenue of the approximately optimal mechanisms we use is unbounded.

will not want to lie to type \emptyset . We let $T_i^+ = T_i \cup \{\emptyset\}$. To proceed, we will introduce a variable $\lambda_i(t, t')$ for each of the BIC constraints, and take the partial Lagrangian of LP 1 by Lagrangifying all BIC constraints. The theory of Lagrangian multipliers tells us that the solution to LP 1 is equivalent to the primal variables solving the partially Lagrangified dual (Figure 2).

Variables:

- $p_i(t_i)$, for all bidders i and types $t_i \in T_i$, denoting the expected price paid by bidder i when reporting type t_i over the randomness of the mechanism and the other bidders' types.
- $\pi_{ij}(t_i)$, for all bidders i , items j , and types $t_i \in T_i$, denoting the probability that bidder i receives item j when reporting type t_i over the randomness of the mechanism and the other bidders' types.

Constraints:

- $\pi_i(t_i) \cdot t_i - p_i(t_i) \geq \pi_i(t'_i) \cdot t_i - p_i(t'_i)$, for all bidders i , and types $t_i \in T_i, t'_i \in T_i^+$, guaranteeing that the reduced form mechanism (π, p) is BIC and BIR.
- $\pi \in P(\mathcal{F}, D)$, guaranteeing π is feasible.

Objective:

- Maximize: $\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i)$, the expected revenue.

Figure 1: A Linear Program (LP) for Revenue Optimization.

Variables:

- $\lambda_i(t_i, t'_i)$ for all $i, t_i \in T_i, t'_i \in T_i^+$, the Lagrangian multipliers for Bayesian IC constraints.

Constraints:

- $\lambda_i(t_i, t'_i) \geq 0$ for all $i, t_i \in T_i, t'_i \in T_i^+$, guaranteeing that the Lagrangian multipliers are non-negative.

Objective:

- Minimize: $\max_{\pi \in P(\mathcal{F}, D), p} \mathcal{L}(\lambda, \pi, p)$.

Figure 2: Partial Lagrangian of the Revenue Maximization LP.

Definition 2. Let $\mathcal{L}(\lambda, \pi, p)$ be the partial Lagrangian defined as follows:

$$\mathcal{L}(\lambda, \pi, p) = \sum_{i=1}^n \left(\sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) + \sum_{t_i \in T_i} \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \cdot \left(t_i \cdot (\pi(t_i) - \pi(t'_i)) - (p_i(t_i) - p_i(t'_i)) \right) \right) \quad (1)$$

$$\begin{aligned} &= \sum_{i=1}^n \sum_{t_i \in T_i} p_i(t_i) \left(f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \right) \\ &\quad + \sum_{i=1}^n \sum_{t_i \in T_i} \pi_i(t_i) \left(\sum_{t'_i \in T_i^+} t_i \cdot \lambda_i(t_i, t'_i) - \sum_{t'_i \in T_i} t'_i \cdot \lambda_i(t'_i, t_i) \right) \quad (2) \end{aligned}$$

3.1 Useful Properties of the Dual Problem

In this section, we make some observations about the dual problem to get some traction on what duals might induce useful upper bounds.

Definition 3 (Useful Dual). *A feasible dual solution λ is **useful** if $\max_{\pi \in P(\mathcal{F}, D), p} \mathcal{L}(\lambda, \pi, p) < \infty$.*

Lemma 4 (Useful Dual). *A dual solution λ is useful if and only if for each bidder i , λ_i forms a valid flow, i.e., iff the following satisfies flow conservation (flow in = flow out) at all nodes except the source and the sink:*

- *Nodes: A super source s and a super sink \emptyset , along with a node t_i for every type $t_i \in T_i$.*
- *Flow from s to t_i of weight $f_i(t_i)$, for all $t_i \in T_i$.*
- *Flow from t to t' of weight $\lambda_i(t, t')$ for all $t \in T$, and $t' \in T_i^+$ (including the sink \emptyset).*

Proof. Let us think of $\mathcal{L}(\lambda, \pi, p)$ using expression (2). Clearly, if there exists any i and $t_i \in T_i$ such that

$$f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \neq 0,$$

then since $p_i(t_i)$ is unconstrained (note that we do **not** include the constraints $p_i(t_i) \geq 0$ in Figure 1, so these variables are indeed unconstrained) and has a non-zero multiplier in the objective, $\max_{\pi \in P(\mathcal{F}, D), p} \mathcal{L}(\lambda, \pi, p) = +\infty$. Therefore, in order for λ to be useful, we must have

$$f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) = 0$$

for all i and $t_i \in T_i$. This is exactly saying what we described in the Lemma statement is a flow. The other direction is simple, whenever λ forms a flow, $\mathcal{L}(\lambda, \pi, p)$ only depends on π . Since π is bounded, the maximization problem has a finite value. \square

Definition 5 (Virtual Value Function). *For each λ , we define a corresponding virtual value function $\Phi^\lambda(\cdot)$, such that for every bidder i , every type $t_i \in T_i$, $\Phi_i^\lambda(t_i) = t_i - \frac{1}{f_i(t_i)} \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)(t'_i - t_i)$. Note that for all i , $\Phi_i^\lambda(\cdot)$ is a vector-valued function, so we use $\Phi_{ij}^\lambda(\cdot)$ to refer to the j^{th} component of $\Phi_i^\lambda(\cdot)$, and refer to this as bidder i 's virtual value for item j .*

Theorem 6 (Virtual Welfare \geq Revenue). *Let λ be any useful dual solution and $M = (\pi, p)$ be any BIC mechanism. The revenue of M is less than or equal to the virtual welfare of π w.r.t. the virtual value function $\Phi^\lambda(\cdot)$ corresponding to λ . That is:*

$$\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) \leq \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \Phi_i^\lambda(t_i).$$

Equality holds if and only if for all i, t, t' such that $\lambda_i(t, t') > 0$, the BIC constraint for bidder i between t and t' binds in M (that is, bidder i with type t is indifferent between reporting t and t'). Furthermore, let λ^ be the optimal dual variables and $M^* = (\pi^*, p^*)$ be the revenue-optimal BIC mechanism, then the expected virtual welfare with respect to Φ^* (induced by λ^*) under π^* equals the expected revenue of M^* , and*

$$\pi^* \in \operatorname{argmax}_{\pi \in P(\mathcal{F}, D)} \left\{ \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \Phi_i^*(t_i) \right\}.$$

Proof. When λ is useful, we can simplify $\mathcal{L}(\lambda, \pi, p)$ by removing all terms associated with p (because all such terms have a multiplier of zero, by Lemma 4), and replace the terms $\sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i)$ with $f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)$. After the simplification, we have $\mathcal{L}(\lambda, \pi, p) = \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \left(t_i - \frac{1}{f_i(t_i)} \cdot \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)(t'_i - t_i) \right)$, which equals $\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \Phi_i^\lambda(t_i)$, exactly the virtual welfare of π . Now, we only need to prove that $\mathcal{L}(\lambda, \pi, p)$ is greater than the revenue of M . Let us think of $\mathcal{L}(\lambda, \pi, p)$ using Expression (1). Since M is a BIC mechanism, $t_i \cdot (\pi(t_i) - p_i(t'_i)) - (p_i(t_i) - p_i(t'_i)) \geq 0$ for any i and $t_i \in T_i, t'_i \in T_i^+$. Also, all the dual variables λ are nonnegative. Therefore, it is clear that $\mathcal{L}(\lambda, \pi, p)$ is at least as large as the revenue of M . Moreover, if the BIC constraint for bidder i between t and t' binds in M for all i, t, t' such that $\lambda_i(t, t') > 0$, then we in fact have $\mathcal{L}(\lambda, \pi, p) = \sum_i \sum_{t_i} f_i(t_i) p_i(t_i)$, so the revenue of M is equal to its expected virtual welfare under $\Phi^\lambda(\cdot)$ (because the Lagrangian terms added to the revenue are all zero).

When λ^* is the optimal dual solution, by strong LP duality applied to the LP of Figure 1, we know $\max_{\pi \in P(\mathcal{F}, D), p} \mathcal{L}(\lambda^*, \pi, p)$ equals the revenue of M^* . But we also know that $\mathcal{L}(\lambda^*, \pi^*, p^*)$ is at least as large as the revenue of M^* , so π^* necessarily maximizes the virtual welfare over all $\pi \in P(\mathcal{F}, D)$, with respect to the virtual transformation Φ^* corresponding to λ^* . \square

To summarize: we have shown that every flow induces a finite upper bound on how much revenue a BIC mechanism can possibly achieve. We have also observed that this upper bound can be interpreted as the maximum virtual welfare obtainable with respect to a virtual valuation function that is decided by the flow. In the next two sections, we will instantiate this theory by designing specific flows, and obtain benchmarks that upper bound the optimal revenue.

4 Canonical Flow for a Single Item

In this section, we provide a canonical flow for single-item settings, and show that it implies the main result from Myerson's seminal work [Mye81]. Essentially, Myerson proposes a specific virtual valuation function and shows that for this virtual valuation function, the expected revenue of any BIC mechanism is always upper bounded by its expected virtual welfare. Moreover, he describes an ironing procedure to guarantee that this virtual valuation is monotone, and proves that the revenue-optimal mechanism simply awards the item to the bidder with the highest virtual value. Myerson's proof is quite elegant, and we are not claiming that our proof below is simpler.¹⁸ The purpose of the proof below is:

- Serve as a warm-up for the reader to get comfortable with flows and virtual valuations.
- Separate out parts of the proof that can be directly applied to more general settings (e.g. Theorem 6).
- Provide a specific flow that will be used in later sections to provide benchmarks in multi-item settings.

In this section, we will have $m = 1$ and drop the item subscript j . We begin with a definition of Myerson's (ironed) virtual valuation function, adapted to the discrete setting.

Definition 7 (Single-dimensional Virtual Value). *For a single-dimensional discrete distribution D_i , if $f_i(t_i) > 0$, let $\varphi_i^{D_i}(t_i) = t_i - \frac{(t'_i - t_i) \cdot \Pr_{t \sim D_i}[t > t_i]}{f_i(t_i)}$, where $t'_i = \min_{t > t_i, t \in T_i} \{t\}$. If $f_i(t_i) = 0$, let $\varphi_i^{D_i}(t_i) = 0$.¹⁹ If the distribution D_i is clear from context, we will just write $\varphi_i(t_i)$.*

¹⁸If one's goal is simply to understand Myerson's result and nothing more, the original proof and ours are comparable in simplicity. Many alternative comparably simple proofs exist as well, some of which are not much different than ours (e.g. [MV04]).

¹⁹Actually we could define $\varphi_i^{D_i}(t_i)$ arbitrarily and everything that follows will still hold.

The ironing procedure described below essentially finds any non-monotonicities in $\varphi_i(\cdot)$ and “irons” them out. Note that Steps 4 and 5 maintain that ironed virtual values are consistent within any ironed interval.

Definition 8 (Ironing). *Let \sim be an equivalence relation on the support of D_i , and $\tilde{\varphi}(\cdot)$ be the ironed virtual valuation function defined in the following way. We say that an interval $[t_i^*, t_i]$ is ironed if $t \sim t'$ for all $t, t' \in [t_i^*, t_i]$.*

1. Initialize $t_i = \max_{t \in T_i} \{t\}$, the highest un-ironed type.
2. For any $t \leq t_i$, define the average virtual value $a([t, t_i]) = \frac{\sum_{t' \in [t, t_i]} f_i(t') \cdot \varphi_i(t')}{\sum_{t' \in [t, t_i]} f_i(t')}$.
3. Let t_i^* maximize the average virtual value. That is, $t_i^* = \arg \max_{t \leq t_i} a([t, t_i])$ (break ties in favor of the maximum such t_i^*).
4. Update $\tilde{\varphi}_i(t) = a([t_i^*, t_i])$ for all $t \in [t_i^*, t_i]$.
5. Update $t \sim t'$ for all $t, t' \in [t_i^*, t_i]$.
6. Update $t_i = \max_{t < t_i^*, t \in T_i} \{t\}$, the highest un-ironed type.
7. Return to Step 2.

These definitions in the discrete case might be slightly different than what readers are used to in the continuous case. We provide some observations proving that this is “the right” definition for the discrete case briefly in Section 4.1. Our proof continues in Section 4.2.

4.1 Discrete Myersonian Virtual Values

In the continuous setting, Myerson’s virtual valuation is defined as $\varphi_i(v) = v - \frac{1 - F_i(v)}{f_i(v)}$, where F_i and f_i are the CDF and PDF of D_i . We first show that for any continuous distribution, discretizing it into multiples of ϵ and taking virtual valuations as in Definition 7, we recover Myerson’s virtual valuation in the limit as $\epsilon \rightarrow 0$.

Observation 9. *Let D_i be any continuous distribution, and D_i^ϵ be the discretization of D_i with point-masses at all multiples of ϵ . That is, $f_i^\epsilon(c\epsilon) = \int_{c\epsilon}^{(c+1)\epsilon} f_i(x) dx$ for all $c \in \mathbb{N}$. Then for all t_i ,*

$$\limsup_{\epsilon \rightarrow 0} \varphi_i^{D_i^\epsilon}(t_i) = t_i - \frac{1 - F_i(t_i)}{f_i(t_i)}.$$

Proof. For fixed t_i , consider the set of $\epsilon \in \{t_i/c \mid c \in \mathbb{N}\}$. Then clearly for all ϵ outside this set, $\varphi_i^{D_i^\epsilon}(t_i) = 0$. For any ϵ in this set, we have $\varphi_i^{D_i^\epsilon}(t_i) = t_i - \frac{\epsilon \cdot \Pr_{t \sim D_i}[t > t_i + \epsilon]}{\Pr_{t \sim D_i}[t \in [t_i, t_i + \epsilon]]}$. It’s also clear that as $\epsilon \rightarrow 0$, we have $\Pr_{t \sim D_i}[t > t_i + \epsilon] \rightarrow \Pr_{t \sim D_i}[t > t_i] = 1 - F_i(t_i)$, and $\frac{\epsilon}{\Pr_{t \sim D_i}[t \in [t_i, t_i + \epsilon]]} \rightarrow \frac{1}{f_i(t_i)}$ (the latter is simply the definition of probability density). \square

A second valuable property of Myersonian virtual values is that they capture the “marginal revenue.” That is, if the seller was selling to a single bidder at price just above (i.e. dv above) v , and decreased the price to just below (i.e. dv below) v , the revenue would go up by exactly $\varphi(v) \cdot f(v) dv$. We confirm that discrete virtual values as per Definition 7 satisfy this property as well.

Observation 10. *For any single-dimensional discrete distribution D_i , we have $t_i \cdot \Pr_{t \sim D_i}[t \geq t_i] - t'_i \cdot \Pr_{t \sim D_i}[t \geq t'_i] = f_i(t_i) \cdot \varphi_i(t_i)$, where $t'_i = \min_{t > t_i, t \in T_i} \{t\}$. In other words, $\varphi_i(t_i)$ captures the marginal change in revenue as we go from setting price t'_i to price t_i .*

Proof. This follows immediately from the definition of $\varphi_i(\cdot)$. But to be thorough:

$$\begin{aligned}
& t_i \cdot \Pr_{t \sim D_i}[t \geq t_i] \\
&= t_i \cdot (\Pr_{t \sim D_i}[t \geq t'_i] + f_i(t_i)) && \text{(using that } t'_i = \min_{t > t_i, t \in T_i} \{t\}\text{)} \\
&= t'_i \cdot (\Pr_{t \sim D_i}[t \geq t'_i] + f_i(t_i)) - (t'_i - t_i) \cdot (\Pr_{t \sim D_i}[t \geq t'_i] + f_i(t_i)) \\
&= t'_i \cdot \Pr_{t \sim D_i}[t \geq t'_i] + t_i \cdot f_i(t_i) - (t'_i - t_i) \cdot \Pr_{t \sim D_i}[t \geq t'_i] \\
&= t'_i \cdot \Pr_{t \sim D_i}[t \geq t'_i] + f_i(t_i) \cdot \varphi_i(t_i) && \text{(definition of } \varphi_i(\cdot)\text{)}
\end{aligned}$$

□

We need one more definition specific to discrete type spaces before we can get back to the proof. A little more specifically: Myerson's payment identity, which shows that allocation rules uniquely determine payments for any BIC mechanism over continuous type spaces, does not apply for *all* BIC mechanisms when types are discrete. Fortunately, the payment identity still holds for any mechanism that might possibly maximize revenue, but we need to be formal about this.

Definition 11. A BIC mechanism has **proper payments** if it is not possible to increase payments while keeping the allocation rule the same without violating BIC. Formally, a BIC mechanism $M = (\pi, p)$ has proper payments if for all i and all subsets $S \subseteq T_i$, and all $\epsilon > 0$, increasing $p_i(t_i)$ by ϵ for all $t_i \in S$ while keeping π the same does not result in a BIC mechanism. Note that all revenue-optimal mechanisms have proper payments.

Lemma 12. Let $\pi_i(\cdot)$ be monotone non-decreasing for all i . Then there exist $\{p_i(\cdot)\}_{i \in [n]}$ such that $M = (\pi, p)$ is BIC and has proper payments.

Proof. For ease of notation in the proof, label the types in T_i so that $0 = t_i^0 \leq t_i^1 < \dots < t_i^{|T_i|}$ ($t_i^j \in T_i$ for $j \in \{1, \dots, |T_i|\}$). Then for a given $\pi_i(\cdot)$, define:

$$p_i(t_i^j) = \sum_{k=1}^j t_i^k \cdot (\pi_i(t_i^k) - \pi_i(t_i^{k-1})).$$

We first claim that all t_i^j are indifferent between telling the truth and reporting t_i^{j-1} . This is clear, as by definition of $p_i(\cdot)$ we have $p_i(t_i^j) = p_i(t_i^{j-1}) + t_i^j \cdot (\pi_i(t_i^j) - \pi_i(t_i^{j-1}))$, which implies:

$$t_i^j \cdot \pi_i(t_i^j) - p_i(t_i^j) = t_i^j \cdot \pi_i(t_i^j) - p_i(t_i^{j-1}) - t_i^j \cdot \pi_i(t_i^j) + t_i^j \cdot \pi_i(t_i^{j-1}) = t_i^j \cdot \pi_i(t_i^{j-1}) - p_i(t_i^{j-1}).$$

Now, consider any set $S \subseteq T_i$ and any $\epsilon > 0$, and imagine raising the payments of all types $t \in S$ by ϵ . If we have $t_i^j \in S, t_i^{j-1} \notin S$ for any t_i^j , then increasing all payments in S by ϵ will cause t_i^j to prefer reporting t_i^{j-1} instead of telling the truth. So S must contain all of T_i , and in particular t_i^1 . But if we increase the payment of t_i^1 by $\epsilon > 0$, we violate individual rationality, as we defined $p_i(t_i^1) = t_i^1 \cdot \pi_i(t_i^1)$. So no such S, ϵ can exist, and $M = (\pi, p)$ has proper payments.

Finally, we just need to show that M is BIC. Notice that by definition of $p_i(\cdot)$, for any $j > k$, we have $p_i(t_i^j) - p_i(t_i^k) \in [t_i^{k+1} \cdot (\pi_i(t_i^j) - \pi_i(t_i^k)), t_i^j \cdot (\pi_i(t_i^j) - \pi_i(t_i^k))]$. The lower bound corresponds to the case that all the change from $\pi_i(t_i^k)$ to $\pi_i(t_i^j)$ occurs going from t_i^k to t_i^{k+1} (i.e. $\pi_i(t_i^{k+1}) = \pi_i(t_i^j)$), and the upper bound corresponds to where all the change occurs going from t_i^{j-1} to t_i^j (i.e. $\pi_i(t_i^k) = \pi_i(t_i^{j-1})$).

The lower bound directly implies that t_i^k prefers telling the truth to reporting t_i^j as:

$$t_i^k \cdot (\pi_i(t_i^j) - \pi_i(t_i^k)) \leq p_i(t_i^j) - p_i(t_i^k) \Rightarrow t_i^k \cdot \pi_i(t_i^k) - p_i(t_i^k) \geq t_i^k \cdot \pi_i(t_i^j) - p_i(t_i^j).$$

Similarly, the upper bound directly implies that t_i^j prefers telling the truth to reporting t_i^k as:

$$t_i^j \cdot (\pi_i(t_i^j) - \pi_i(t_i^k)) \geq p_i(t_i^j) - p_i(t_i^k) \Rightarrow t_i^j \cdot \pi_i(t_i^j) - p_i(t_i^j) \geq t_i^j \cdot \pi_i(t_i^k) - p_i(t_i^k).$$

As the above holds for any $j > k$, M is BIC. □

4.2 Proof of Myerson's Theorem via Duality

Now we return to our proof of Myerson's Theorem. Let us first quickly confirm that indeed the resulting $\tilde{\varphi}_i(\cdot)$ by our ironing procedure is monotone:

Observation 13 ([Mye81]). $t_i > t'_i \Rightarrow \tilde{\varphi}_i(t_i) \geq \tilde{\varphi}_i(t'_i)$.

Proof. First, all types in the same ironed interval share the same ironed virtual value. So if $\tilde{\varphi}_i(\cdot)$ is not monotone non-decreasing, there exists two adjacent ironed intervals $[x, y]$ and $[z, w]$ such that $x > w$ but $a([x, y]) < a([z, w])$. Note that $a([z, y]) = c \cdot a([x, y]) + (1 - c) \cdot a([z, w])$ for some $c \in (0, 1)$.²⁰ Since $a([x, y]) < a([z, w])$, we have $a([z, y]) > a([x, y])$. However, this contradicts with the choice of an ironed interval for y as specified in Step 3 of the ironing process. Hence, no such ironed intervals exist and $\tilde{\varphi}_i(\cdot)$ is monotone non-decreasing. □

And now, we can state Myerson's theorem applied to discrete type spaces. Afterwards, we will provide a proof using our new duality framework. One should map the theorem statement below to the statement of Theorem 6 and see that our proof will essentially follow by providing a flow λ that induces a virtual valuation function $\Phi_i^\lambda(\cdot) = \varphi_i(\cdot)$, another flow λ' inducing $\Phi_i^{\lambda'}(\cdot) = \tilde{\varphi}_i(\cdot)$, and understanding which edges have non-zero flow in each.

Theorem 14 ([Mye81]). *For any BIC mechanism $M = (\pi, p)$, the revenue of M is less than or equal to the virtual welfare of π w.r.t. the virtual valuation function $\varphi(\cdot)$, and is less than or equal to the virtual welfare of π w.r.t. the ironed virtual valuation function $\tilde{\varphi}(\cdot)$. That is:*

$$\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) \leq \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \varphi_i(t_i). \quad (3)$$

$$\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) \leq \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \tilde{\varphi}_i(t_i). \quad (4)$$

Equality holds in Equation (3) whenever M has proper payments. Equality holds in Equation (4) if and only if M has proper payments and $\pi_i(t_i) = \pi_i(t'_i)$ for all $t_i \sim t'_i$. Furthermore, the revenue-optimal BIC mechanism awards the item to the bidder with the highest non-negative ironed virtual value (if one exists), breaking ties arbitrarily but consistently across inputs.²¹ If no such bidder exists, the item remains unallocated.

²⁰In fact, $c = \frac{\sum_{t \in [x, y]} f_i(t)}{\sum_{t \in [z, y]} f_i(t)}$.

²¹Breaking ties consistently means that for any two different inputs, as long as they have the same ironed virtual value profile, the tie breaking should be the same.

We first provide a canonical flow inducing Myerson's virtual values as the virtual transformation. The first two lemmas below relate to proving Equation (3). The third relates to proving Equation (4).

Lemma 15. Define $\lambda_i(t'_i, t_i) = \Pr_{t \sim D_i}[t > t_i]$, where $t'_i = \min_{t > t_i, t \in T_i}\{t\}$, and $\lambda_i(t, t_i) = 0$ for all $t \neq t'_i$. Then λ is a useful dual, and $\Phi_i^\lambda(t_i) = \varphi_i(t_i)$.

Proof. That λ is a useful dual follows immediately by considering the total flow in and flow out of any given t_i . The total flow in is equal to $f_i(t_i) + \Pr_{t \sim D_i}[t > t_i]$. This is because t_i receives flow $\Pr_{t \sim D_i}[t > t_i]$ from t'_i , and $f_i(t_i)$ from the super source. The total flow out is equal to $\Pr_{t \sim D_i}[t \geq t_i]$, so the two are equal.

To compute $\Phi_i^\lambda(t_i)$, simply plug the choice of λ into Definition 5. \square

Lemma 16. In any BIC mechanism M with proper payments, bidder i with type t_i is indifferent between reporting t_i and $t'_i = \max_{t < t_i, t \in T_i}\{t\}$.

Proof. We first recall that if $M = (\pi, p)$ is BIC, then $\pi_i(\cdot)$ and $p_i(\cdot)$ are both monotone non-decreasing for all i [Mye81]. Because M is BIC, we know that $t'_i \cdot \pi_i(t'_i) - p_i(t'_i) \geq t'_i \cdot \pi_i(t) - p_i(t)$ for all $t \in T_i$. Therefore, we also have:

$$t'_i \cdot (\pi_i(t'_i) - \pi_i(t)) \geq p_i(t'_i) - p_i(t).$$

By monotonicity, the LHS above is non-negative whenever $t'_i > t$. Therefore, $t_i \cdot (\pi_i(t'_i) - \pi_i(t)) \geq p_i(t'_i) - p_i(t)$ as well. This directly says that t_i prefers reporting t'_i to reporting any $t < t'_i$. Assume now for contradiction that there is some t_i that is not indifferent between reporting t_i and t'_i . Then we have the following chain of inequalities:

$$\begin{aligned} & t_i \cdot \pi_i(t_i) - p_i(t_i) > t_i \cdot \pi_i(t'_i) - p_i(t'_i) \geq t_i \cdot \pi_i(t') - p_i(t'), \forall t' < t_i \\ \Rightarrow & t_i \cdot (\pi_i(t_i) - \pi_i(t')) > p_i(t_i) - p_i(t'), \forall t' \leq t'_i \\ \Rightarrow & t \cdot (\pi_i(t_i) - \pi_i(t')) > p_i(t_i) - p_i(t'), \forall t' < t_i, t \geq t_i \\ \Rightarrow & t \cdot \pi_i(t) - p_i(t) \geq t \cdot \pi_i(t_i) - p_i(t_i) > t \cdot \pi_i(t') - p_i(t'), \forall t' < t_i, t \geq t_i. \end{aligned}$$

The last line explicitly states that all $t \geq t_i$ strictly prefer telling the truth to reporting any $t' \leq t'_i$. Therefore, there exists a sufficiently small $\epsilon > 0$ such that if we increase all $p_i(t)$ by ϵ for all $t \geq t_i$ then M remains BIC, contradicting that M has proper payments. \square

At this point, we have proved Equation (3) and the related statements for unironed virtual values (but we will wrap up concretely at the end of the section). We now turn to Equation (4), and first show how to ‘‘fix’’ non-monotonicities in virtual values by adding cycles.

Lemma 17. Starting from any flow λ and induced virtual values Φ^λ , adding a cycle of flow x between bidder i 's type t_i and type $t'_i < t_i$:

- Increases $\Phi_i^\lambda(t_i)$.
- Decreases $\Phi_i^\lambda(t'_i)$.
- Preserves $f_i(t_i) \cdot \Phi_i^\lambda(t_i) + f_i(t'_i) \cdot \Phi_i^\lambda(t'_i)$.

Proof. Recall the definition of $\Phi_i^\lambda(t_i) = t_i + \frac{\sum_{t < t_i} (t_i - t) \lambda_i(t, t_i)}{f_i(t_i)}$. Adding a cycle of flow x between t_i and t'_i increases $\lambda_i(t'_i, t_i)$ and $\lambda_i(t_i, t'_i)$, but otherwise doesn't change λ . So $\Phi_i^\lambda(t_i)$ increases by $(t_i - t'_i)x / f_i(t_i)$. Similarly, $\Phi_i^\lambda(t'_i)$ decreases by exactly $(t_i - t'_i)x / f_i(t'_i)$. It is therefore clear that $f_i(t_i) \cdot \Phi_i^\lambda(t_i) + f_i(t'_i) \cdot \Phi_i^\lambda(t'_i)$ is preserved, as the changes are inversely proportional to the types' densities. \square

And we may now conclude that a flow exists inducing Myerson's ironed virtual values as well.

Corollary 18. *There exists a flow λ such that:*

- $\Phi_i^\lambda(t_i) = \tilde{\varphi}_i(t_i)$.
- $\lambda_i(t_i, t'_i) > 0$ whenever $t'_i = \max_{t < t_i, t \in T} \{t\}$.
- $\lambda_i(t'_i, t_i) > 0$ when $t'_i = \max_{t < t_i, t \in T} \{t\}$ **if and only if** t_i and t'_i are in the same ironed interval.
- $\lambda_i(t, t') = 0$ for all other t, t' .

Proof. Start with the λ inducing $\Phi_i^\lambda(t_i) = \varphi_i(t_i)$, and consider any ironed interval $[x, y]$. First, observe that we must have $\varphi_i(x) > a(x, y)$, as otherwise $[x, y]$ would not be an ironed interval (as x doesn't maximize $a(x, y)$, $a(z, y)$ would be at least as large where $z = \min_{t > x, t \in T_i} \{t\}$, and recall that we would break a tie in favor of z).

Now, add a cycle between x and $z = \min_{t > x, t \in T_i} \{t\}$. Per Lemma 17, this increases $\Phi_i^\lambda(z)$ and decreases $\Phi_i^\lambda(x)$. So increase the weight along this cycle until $\Phi_i^\lambda(x)$ decreases to $\tilde{\varphi}_i(x)$. At this point, either $[x, z]$ is the entire ironed interval, in which case this interval is “finished.” Or, maybe $z < y$. In this case, we claim that we can iterate the process with z . To see this, observe that by Lemma 17, we must have preserved $a(x, z)$. Observe again that we must have $a(x, z) > a(x, y)$ in order for $[x, y]$ to be an ironed interval, and we have just set $\Phi_i^\lambda(x) = \tilde{\varphi}_i(x) = a(x, y) < a(x, z)$, so we must have $\Phi_i^\lambda(z) > a(x, z) > a(x, y) = \tilde{\varphi}_i(z)$. So we can again add a cycle between z and $w = \min_{t > z, t \in T_i} \{t\}$ to decrease $\Phi_i^\lambda(z)$ to $\tilde{\varphi}_i(z)$ while preserving $a(x, w)$. Iterating this process all the way until y necessarily adds a cycle between all adjacent types and results in $\Phi_i^\lambda(t) = \tilde{\varphi}_i(t)$ for all $t \in [x, y]$, again by Lemma 17. Repeating this argument for all ironed intervals proves the corollary. \square

Now we may complete the proof of Theorem 14. The bulk of the proof is captured by Lemma 15 and Corollary 18, the only remaining work is to confirm the structure of the optimal mechanism.

Proof of Theorem 14: Inequality (3) now immediately follows from Theorem 6 and Lemma 15, and the condition for it to be an equality is implied by Lemma 16. Inequality (4) follows immediately from Theorem 6 and Corollary 18. Next, we argue why Inequality (4) is an equality when the stated condition holds. When a mechanism is BIC and $\pi_i(t_i) = \pi_i(t'_i)$, then $p_i(t_i) = p_i(t'_i)$, because $0 = t_i \cdot (\pi_i(t_i) - \pi_i(t'_i)) \geq p_i(t_i) - p_i(t'_i) \geq t'_i \cdot (\pi_i(t_i) - \pi_i(t'_i)) = 0$ (implied by the BIC constraints). Therefore, when the condition holds, any bidder i with type t_i is indifferent between reporting t_i and any type $t'_i \sim t_i$. Combining this observation with Lemma 16, we know that for the flow specified in Corollary 18, all BIC constraints bind between any two types t_i and t'_i with $\lambda(t_i, t'_i) > 0$. Thus, Inequality (4) is an equality when the condition holds due to Theorem 6.

Finally, to see that the optimal mechanism has the prescribed format, observe that the allocation rule that awards the item to the highest non-negative ironed virtual value clearly maximizes ironed virtual welfare. So we get that the revenue of the optimal BIC mechanism is upper bounded by the ironed virtual welfare of this allocation rule. Moreover, ironed virtual values are always monotone non-decreasing, so by Lemma 12, this allocation rule has corresponding proper payments that combine to a BIC mechanism. Finally, because the allocation rule by definition satisfies $\pi_i(t_i) = \pi_i(t'_i)$ whenever $t_i \sim t'_i$, we have that the revenue of this mechanism is equal to its expected ironed virtual welfare (again by Theorem 6), and is therefore optimal (as its expected ironed virtual welfare is an upper bound on the expected revenue of any BIC mechanism). \square

In summary, we have provided a duality-based proof of Myerson's Theorem [Mye81]. This gives some intuition for the flows we will develop in the following section. Also, it provides a different insight into the difference between ironed and non-ironed virtual values. Expected revenue is equal to expected virtual welfare for all BIC mechanisms (with proper payments) because the flow necessary to derive virtual values only sends non-zero flow along edges that correspond to BIC constraints that are always tight. On the other

hand, revenue is only upper bounded by ironed virtual welfare for all BIC mechanisms because the flow necessary to derive ironed virtual values sends non-zero flow along all edges between adjacent types in an equivalence class. So revenue is only equal to ironed virtual welfare if all of the corresponding BIC constraints are tight (and there exist truthful mechanisms for which this doesn't hold).

5 Canonical Flow and Virtual Valuation Function for Multiple Items

In this section, we present a canonical way to set the Lagrangian multipliers/flow that induces our benchmarks for multi-item settings. This flow will use similar ideas to Section 4. Informally, our approach for a single bidder first divides the entire type space of the bidder into regions based on their *favorite item* (that is, $\arg\max_j \{v_j\}$). We'll then use a different "Myerson-like" flow within each region, described in more detail shortly. For multiple bidders, we'll still divide the type space of each bidder into regions based on their favorite item, but define the "favorite" item slightly differently.

Specifically, let $P_{ij}(t_{-i})$ denote the price that bidder i could pay to receive exactly item j in the VCG mechanism against bidders with types t_{-i} .²² We will partition the type space T_i into $m+1$ regions: **(i)** $R_0^{(t_{-i})}$ contains all types t_i such that $t_{ij} < P_{ij}(t_{-i}), \forall j$; **(ii)** $R_j^{(t_{-i})}$ contains all types t_i such that $t_{ij} - P_{ij}(t_{-i}) \geq 0$ and j is the smallest index in $\arg\max_k \{t_{ik} - P_{ik}(t_{-i})\}$. This partitions the types into subsets based on which item provides the largest surplus (value minus price), and we break ties lexicographically. We'll refer to the largest surplus item as the bidder's favorite item. We'll refer to all other items as *non-favorite* items. For any bidder i and any type profile t_{-i} of everyone else, we define $\lambda_i^{(t_{-i})}$ to be the following flow. The flow as defined below will look similar to Myerson's (non-ironed) virtual values, and we will need to similarly iron it to accommodate irregular distributions.

Definition 19 (Initial Canonical Flow). *Define our initial canonical flow to be the following:*

1. $\forall j > 0$, any flow entering $R_j^{(t_{-i})}$ is from s (the super source) and any flow leaving $R_j^{(t_{-i})}$ is to \emptyset (super sink).
2. For every type t_i in region $R_0^{(t_{-i})}$, the flow goes directly to \emptyset . That is, $\lambda_i(t_i, \emptyset) = f_i(t_i)$ for all $t_i \in R_0^{(t_{-i})}$.
3. For every type t_i in region $R_j^{(t_{-i})}$, define type t'_i such that $t'_i = \max_{t < t_{ij}, t \in T_{ij}} \{t\}$, and $t'_{ik} = t_{ik}$ for all $k \neq j$.
 - If $t'_i \in R_j$ as well, then set $\lambda_i(t_i, t'_i) = \sum_{t_i^*: t_{ik}^* = t_{ik} \forall k \neq j \wedge t_{ij}^* \geq t_{ij}} f_i(t_i^*)$.
 - If $t'_i \notin R_j$, set $\lambda_i(t_i, \emptyset) = \sum_{t_i^*: t_{ik}^* = t_{ik} \forall k \neq j \wedge t_{ij}^* \geq t_{ij}} f_i(t_i^*)$.

Observe that this indeed defines a flow. All nodes in $R_0^{(t_{-i})}$ have no flow in (except from the super source), and send exactly this flow to the super sink. All other nodes get flow in from exactly one type (in addition to the super source) and send flow out to exactly one type, and the flow is balanced, just as in Lemma 15. At a high level, what we are doing is restricting attention to a single item and attempting to use the canonical single item flow for just this item. We restrict attention to different items for different types, depending on which item gives them highest utility at the prices $P_{ij}(t_{-i})$. Let's first study the induced virtual values $\Phi_{ik}^{(t_{-i})}(t_i)$ for when $t_i \notin R_k^{(t_{-i})}$.

Claim 20. *For any type $t_i \in R_j^{(t_{-i})}$, its corresponding virtual value $\Phi_{ik}^{(t_{-i})}(t_i)$ for item k is exactly its value t_{ik} for all $k \neq j$.*

²²Note that when buyers are additive, this is exactly the highest bid for item j from buyers besides i . When buyers are unit-demand, buyer i only ever buys one item, and this is the price she would pay for receiving j .

Proof. By the definition of $\Phi_i^{(t-i)}(\cdot)$, $\Phi_{ik}^{(t-i)}(t_i) = t_{ik} - \frac{1}{f_i(t_i)} \sum_{t'_i} \lambda_i^{(t-i)}(t'_i, t_i)(t'_{ik} - t_{ik})$. Since $t_i \in R_j$, by the definition of the flow $\lambda_i^{(t-i)}$, for any t'_i such that $\lambda_i^{(t-i)}(t'_i, t_i) > 0$, $t'_{ik} - t_{ik} = 0$ for all $k \neq j$, therefore $\Phi_{ik}^{(t-i)}(t_i) = t_{ik}$. \square

Let's now study the corresponding $\Phi_{ij}^{(t-i)}(t_i)$ for this flow when $t_i \in R_j$. This turns out to be closely related to the Myerson's virtual value function for single-dimensional distributions discussed in Section 4. For each i, j , we use $\varphi_{ij}(\cdot)$ and $\tilde{\varphi}_{ij}(\cdot)$ to denote the Myerson virtual value and ironed virtual value function for distribution D_{ij} respectively, as defined in Section 4.

Claim 21. *For any type $t_i \in R_j^{(t-i)}$, then the initial canonical flow induces virtual values satisfying:*
 $\Phi_{ij}^{(t-i)}(t_i) = \varphi_{ij}(t_{ij}) = t_{ij} - \frac{(t'_{ij} - t_{ij}) \cdot \Pr_{t \sim D_{ij}}[t > t_{ij}]}{f_{ij}(t_{ij})}$, where $t'_{ij} = \min_{t > t_{ij}, t \in T_{ij}} \{t\}$.

Proof. Let us fix $t_{i,-j}$, and prove this is true for all choices of $t_{i,-j}$. If t_{ij} is the largest value in T_{ij} , then there is no flow coming into it except the one from the source, so $\Phi_{ij}^{(t-i)}(t_i) = t_{ij}$. For every other value of t_{ij} , the flow coming from its predecessor $(t'_{ij}, t_{i,-j})$ is exactly (note below that several steps make use of the fact that $f_i(t_i) = \prod_j f_{ij}(t_{ij})$)

$$\begin{aligned} \sum_{t_i^*: t_{ik}^* = t_{ik} \forall k \neq j \wedge t_{ij}^* \geq t'_{ij}} f_i(t_i^*) &= \sum_{t_i^*: t_{ik}^* = t_{ik} \forall k \neq j \wedge t_{ij}^* \geq t'_{ij}} f_{ij}(t_{ij}^*) \cdot \prod_{k \neq j} f_{ik}(t_{ik}) = \prod_{k \neq j} f_{ik}(t_{ik}) \cdot \sum_{t_{ij}^* > t_{ij}} f_{ij}(t_{ij}^*) \\ &= \prod_{k \neq j} f_{ik}(t_{ik}) \cdot \Pr_{t \sim D_{ij}}[t > t_{ij}]. \end{aligned}$$

Now, we can compute according to Definition 5:

$$\Phi_{ij}^{(t-i)}(t_i) = t_{ij} - \frac{(t'_{ij} - t_{ij}) \cdot \prod_{k \neq j} f_{ik}(t_{ik}) \cdot \Pr_{t \sim D_{ij}}[t > t_{ij}]}{f_i(t_i)} = t_{ij} - \frac{(t'_{ij} - t_{ij}) \cdot \Pr_{t \sim D_{ij}}[t > t_{ij}]}{f_{ij}(t_{ij})} = \varphi_{ij}(t_{ij})$$

\square

Claims 20 and 21 show that our initial canonical flow induces virtual values such that the virtual value of each bidder for all of their *non-favorite* items is exactly their value, while their virtual value for their *favorite* item is exactly their Myersonian virtual value as per Definition 7. When D_{ij} is regular, this is the canonical flow we use. When the distribution is not regular, we also need to “iron” the virtual values as in Section 4. Essentially all we are doing is applying the same procedure as Definition 8, but we repeat it below to be clear exactly how the substitutions occur. Below, we use $\tilde{\varphi}_{ij}^*$ to denote the ironed virtual values instead of $\tilde{\varphi}_{ij}$ because we reserve $\tilde{\varphi}_{ij}$ to refer exactly to the ironed virtual values that result in the single item case, and we haven't yet proved that they are (essentially) the same.

Definition 22 (Ironed Canonical Virtual Values). *For a given bidder i , valuation vector $t_i \in R_j^{(t-i)}$, obtain the ironed canonical values, $\tilde{\varphi}_{ij}^*(\cdot)$ in the following manner: let X denote the minimum t_{ij} such that $(t_{ij}; t_{i,-j}) \in R_j^{(t-i)}$. Let \sim^* be an equivalence relation on T_{ij} . We say that an interval $[t_{ij}^*, t_{ij}]$ is ironed if $t \sim^* t'$ for all $t, t' \in [t_{ij}^*, t_{ij}]$.*

1. Initialize $t_{ij} = \max_{t \in T_{ij}} \{t\}$, the highest un-ironed type.

2. For any $t \in [X, t_{ij}]$, define the average virtual value $a([t, t_{ij}]) = \frac{\sum_{t' \in [t, t_{ij}]} f_{ij}(t') \cdot \varphi_{ij}(t')}{\sum_{t' \in [t, t_{ij}]} f_{ij}(t')}$.

3. Let t_i^* maximize the average virtual value. That is, $t_i^* = \arg \max_{t \in [X, t_{ij}]} \{a([t, t_{ij}])\}$ (break ties in favor of the maximum such t_i^*).
4. Update $\tilde{\varphi}_{ij}^*(t) = a([t_i^*, t_{ij}])$ for all $t \in [t_i^*, t_{ij}]$.
5. Update $t \sim^* t'$ for all $t, t' \in [t_i^*, t_{ij}]$.
6. Update $t_{ij} = \max_{t \in [X, t_i^*], t \in T_i} \{t\}$, the highest un-ironed type.
7. Return to Step 2.

Lemma 23. *The initial canonical flow can be ironed into a λ so that for all j , and all $t_i \in R_j$, $\Phi_{ij}^\lambda(t_i) = \tilde{\varphi}_{ij}^*(t_{ij})$ by only adding cycles between types $t_i, t'_i \in R_j^{(t-i)}$ satisfying $t_{i,-j} = t'_{i,-j}$.*

Proof. Exactly the same as Corollary 18, plus the observation that all types for which values for item j have identical values for items $\neq j$. \square

Now that we have a flow “ironing” one of the virtual values, we want to wrap up by observing that $\tilde{\varphi}_{ij}^*(v) \leq \tilde{\varphi}_{ij}(v)$ for all v (where $\tilde{\varphi}_{ij}(v)$ is Myerson’s ironed virtual value for the distribution D_{ij}).

Lemma 24. *For any i, j , $t_{-i}, t_{i,-j}, t_{ij}$, $\tilde{\varphi}_{ij}^*(t_{ij}) \leq \tilde{\varphi}_{ij}(t_{ij})$.*

Proof. Observe that the ironing procedure in Definition 22 is nearly identical to that of Definition 8. In fact, for all $x, y \geq X$ (X is defined in Definition 22), if $x \sim^* y$ as in Definition 22, then $x \sim y$ as in Definition 8. This immediately yields that $\tilde{\varphi}_{ij}(t_{ij}) = \tilde{\varphi}_{ij}^*(t_{ij})$ for all t_{ij} in ironed intervals $[x, y]$ (as in Definition 8) such that $x \geq X$. But the ironed virtual values might differ if t_{ij} lies inside an ironed interval that is “cut” by X in Definition 8. But observe that by the definition of ironing, we necessarily have $a([x, y]) \geq a([X, y])$ in order for $[x, y]$ to possibly be an ironed interval containing X . Therefore, we may immediately conclude that $\tilde{\varphi}_{ij}(t_{ij}) \geq \tilde{\varphi}_{ij}^*(t_{ij})$ for all t_{ij} , even those in ironed intervals cut by X . \square

Lemma 25. *There exists a flow $\lambda_i^{(t-i)}$ such that $\Phi_{ij}^{(t-i)}(t_i)$ satisfies the following properties:*

- For any $j > 0$, $t_i \in R_j^{(t-i)}$, $\Phi_{ij}^{(t-i)}(t_i) \leq \tilde{\varphi}_{ij}(t_{ij})$, where $\tilde{\varphi}_{ij}(\cdot)$ is Myerson’s ironed virtual value for D_{ij} .
- For any j , $t_i \in R_j^{(t-i)}$, $\Phi_{ik}^{(t-i)}(t_i) = t_{ik}$ for all $k \neq j$. In particular, $\Phi_i^{(t-i)}(t_i) = t_i$, $\forall t_i \in R_0^{(t-i)}$.

Proof. To see the first bullet, combine Lemmas 23 and 24. To see the second bullet, observe Claim 20, combined with the fact that the cycles added via Lemma 23 doesn’t affect virtual values for the non-favorite items. \square

At this point, for each $t_{-i} \in T_{-i}$, we have defined a different flow for bidder i . We have shown that this flow induces a virtual valuation function such that bidder i ’s virtual value for all non-favorite items is equal to their value for those items, and their virtual value for their favorite item is at most their Myersonian ironed virtual value for that item. Figure 3 contains a diagram illustrating our flow for a single bidder, and Figure 4 illustrates what the flow might look like for non-zero (t_{-i}) .

Finally, note that we’ve defined many possible flows for bidder i : each t_{-i} defines different $P_{ij}(t_{-i})$ s, which in turn define different $R_j^{(t-i)}$ s, which define different flows. But we only get to pick one flow for bidder i , and it cannot change depending on t_{-i} . The flow that we will finally use essentially averages these flows according to D_{-i} .

Definition 26 (Canonical Flow for Multiple Items). *Our flow for bidder i is $\lambda_i = \sum_{t_{-i} \in T_{-i}} f_{-i}(t_{-i}) \lambda_i^{(t-i)}$. Accordingly, the virtual value function Φ_i of λ_i is $\Phi_i(\cdot) = \sum_{t_{-i} \in T_{-i}} f_{-i}(t_{-i}) \Phi_i^{(t-i)}(\cdot)$.*




Figure 3: An example of λ (with ironing) for a single bidder.




Figure 4: An example of $\lambda_i^{(t-i)}$ for two items.

Intuition behind Our Flow: The social welfare is a trivial upper bound for revenue, which can be arbitrarily bad in the worst case. To design a good benchmark, we want to replace some of the terms that contribute the most to the social welfare with more manageable ones. The flow $\lambda_i^{(t_{-i})}$ aims to achieve exactly this. For each bidder i , we find the item j that contributes the most to the social welfare when awarded to i . Then we turn the **virtual value** of item j into its Myerson’s single-dimensional (ironed) virtual value, and keep the **virtual value** of all the other items equal to the value. This transformation is feasible only if we know exactly t_{-i} and could use a different dual solution for each t_{-i} . Since we can’t, a natural idea is to define a flow by taking an expectation over t_{-i} . This is indeed our flow.

We conclude this section with one final lemma and our main theorem regarding the canonical flow. Both proofs are immediate corollaries of the flow definition and Theorem 6. Note also that our flow only ever sends flow between types that are identical on all but one coordinate, and adjacent in the final coordinate (and that this coordinate is their “favorite” item - adjusted by t_{-i}). This means that our benchmark not only upper bounds the optimal revenue of any BIC mechanism, but it also upper bounds the optimal revenue of any (non-truthful) mechanism where bidder i with type t_i has no incentive to lie by misreporting their value for a single item to an adjacent value. A corollary of our work in the following sections is that this relaxation does not improve the optimal revenue by more than a constant factor.

Lemma 27. For all i, j, t_i , $\Phi_{ij}(t_i) \leq t_{ij} \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \notin R_j^{(v_{-i})}] + \tilde{\varphi}_{ij}(t_{ij}) \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \in R_j^{(v_{-i})}]$.

Theorem 28. Let M be any BIC mechanism with (π, p) as its reduced form. The expected revenue of M is upper bounded by the expected virtual welfare of the same allocation rule with respect to the canonical virtual value function $\Phi_i(\cdot)$. In particular,

$$\begin{aligned} \sum_i \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) &\leq \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \Phi_{ij}(t_i) \\ &\leq \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \left(t_{ij} \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \notin R_j^{(v_{-i})}] + \tilde{\varphi}_{ij}(t_{ij}) \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \in R_j^{(v_{-i})}] \right) \end{aligned} \quad (5)$$

6 Warm Up: Single Bidder

As a warm up, we start with the single bidder case. In this section, our goal is to show how to use the bounds obtained via Theorem 28 to prove that simple mechanisms are approximately optimal for a single additive or unit-demand bidder with independent item values, recovering results of [CMS15] and [BILW14]. Throughout this section, we keep the same notations but drop the subscript i and superscript (t_{-i}) whenever is appropriate.

Canonical Flow for a Single Bidder. Since the canonical flow and the corresponding virtual valuation functions are defined based on other bidders types t_{-i} , let us see how it is simplified when there is only a single bidder. First, the VCG prices are all 0, therefore λ is simply one flow instead of a distribution of different flows. Second, for the same reason, the region R_0 is empty and region R_j contains all types t with

$t_j \geq t_k$ for all k (see Figure 3 for an example). This simplifies Expression (5) to

$$\begin{aligned} & \sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot \left(t_j \cdot \mathbb{1}[t \notin R_j] + \tilde{\varphi}_j(t_j) \cdot \mathbb{1}[t \in R_j] \right) \\ = & \sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot t_j \cdot \mathbb{1}[t \notin R_j] \quad (\text{NON-FAVORITE}) \\ & + \sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot \tilde{\varphi}_j(t_j) \cdot \mathbb{1}[t \in R_j] \quad (\text{SINGLE}) \end{aligned}$$

Above, SINGLE refers to the bound coming from cases where $t \in R_j$. We name it ‘‘single’’ to reference the connection to single-dimensional settings. NON-FAVORITE refers to the bound coming from cases where $t \notin R_j$, and we name it ‘‘non-favorite’’ because this contribution only comes from non-favorite items. We bound SINGLE below, and NON-FAVORITE differently for unit-demand and additive valuations.

Lemma 29. *For any feasible $\pi(\cdot)$, $\text{SINGLE} \leq \text{OPT}^{\text{COPIES}}$.*

Proof. Assume M is the mechanism that induces $\pi(\cdot)$. Consider another mechanism M' for the Copies setting, such that for every type profile t , M' serves agent j iff M allocates item j in the original setting and $t \in R_j$. As M is feasible in the original setting, M' is clearly feasible in the Copies setting. When agent j 's type is t_j , its probability of being served in M' is $\sum_{t_{-j}} f_{-j}(t_{-j}) \cdot \pi_j(t_j, t_{-j}) \cdot \mathbb{1}[t \in R_j]$ for all j and t_j . Therefore, SINGLE is the ironed virtual welfare achieved by M' with respect to $\tilde{\varphi}(\cdot)$. Since the copies setting is a single dimensional setting, the optimal revenue $\text{OPT}^{\text{COPIES}}$ equals the maximum ironed virtual welfare, thus no smaller than SINGLE. Note that this proof makes use of the assumption that item values are independent, as otherwise Myerson's theory doesn't apply. \square

Upper Bound for a Unit-demand Bidder. As mentioned previously, the bulk of our work is in obtaining a benchmark and properly decomposing it. Now that we have a decomposition, we can use techniques similar to those of Chawla et al. [CHK07, CHMS10, CMS15] to approximate each term.

Lemma 30. *When the types are unit-demand, for any feasible $\pi(\cdot)$, $\text{NON-FAVORITE} \leq \text{OPT}^{\text{COPIES}}$.*

Proof. Indeed, we will prove that NON-FAVORITE is upper bounded by the revenue of the VCG mechanism in the Copies setting. Define $S(t)$ to be the second largest number in $\{t_1, \dots, t_m\}$. When the types are unit-demand, the Copies setting is a single item auction with m bidders. Therefore, if we run the Vickrey auction in the Copies setting, the revenue is $\sum_{t \in T} f(t) \cdot S(t)$. If $t \notin R_j$, then there exists some $k \neq j$ such that $t_k \geq t_j$, so $t_j \cdot \mathbb{1}[t \in R_j] \leq S(t)$ for all j . Therefore, $\sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot t_j \cdot \mathbb{1}[t \notin R_j] \leq \sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot S(t) \leq \sum_{t \in T} f(t) \cdot S(t)$. The last inequality is because the bidder is unit demand, so $\sum_j \pi_j(t) \leq 1$. \square

Combining Lemma 29 and Lemma 30, we recover the result of Chawla et al. [CMS15].²³

Theorem 31. *For a single unit-demand bidder, the optimal revenue is upper bounded by $2\text{OPT}^{\text{COPIES}}$.*

Upper Bound for an Additive Bidder. When the bidder is additive, we need to further decompose NON-FAVORITE into two terms we call CORE and TAIL. For simplicity of notation in the proofs that follow, define

²³This bound combined with [CHMS10] recovers the state-of-the-art 4-approximation via item-pricing.

$r = \text{SREV}$. Again, we remind the reader that most of our work is already done in obtaining our decomposition. The remaining portion of the proof is indeed inspired by prior work of Babaioff et al. [BILW14].²⁴

$$\begin{aligned}
\sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot t_j \cdot \mathbb{1}[t \notin R_j] &\leq \sum_{t \in T} \sum_j f(t) \cdot t_j \cdot \mathbb{1}[t \notin R_j] \\
&= \sum_j \sum_{t_j > r} f_j(t_j) \cdot t_j \cdot \sum_{t_{-j}} f_{-j}(t_{-j}) \cdot \mathbb{1}[t \notin R_j] + \sum_j \sum_{t_j \leq r} f_j(t_j) \cdot t_j \cdot \sum_{t_{-j}} f_{-j}(t_{-j}) \cdot \mathbb{1}[t \notin R_j] \\
&\leq \sum_j \sum_{t_j > r} f_j(t_j) \cdot t_j \cdot \Pr_{t_{-j} \sim D_{-j}}[t \notin R_j] \quad (\text{TAIL}) + \sum_j \sum_{t_j \leq r} f_j(t_j) \cdot t_j \quad (\text{CORE})
\end{aligned}$$

Before proceeding, let's parse term TAIL above (we'll parse CORE shortly after). TAIL captures contributions to the bound coming from *non-favorite* items whose value is at least SREV. In the term TAIL, the main idea is that we should expect $t_j \cdot \Pr[t \notin R_j]$ to be small when $t_j > \text{SREV}$. This is because t_j is already quite large, so we should expect the probability that we see another item with even larger value (a necessary condition for $t \notin R_j$) to be quite small. Lemma 32 captures this formally, and makes use of the fact that item values are independent.

Lemma 32. TAIL \leq SREV.

Proof. (Recall that we define $r = \text{SREV}$ for ease of notation in some places). By the definition of R_j , for any given t_j ,

$$\Pr_{t_{-j} \sim D_{-j}}[t \notin R_j] \leq \Pr_{t_{-j} \sim D_{-j}}[\exists k \neq j, t_k \geq t_j].²⁵$$

It is clear that by setting price t_j on each item separately, we can make revenue at least $t_j \cdot \Pr_{t_{-j} \sim D_{-j}}[\exists k \neq j, t_k \geq t_j]$, as the buyer will certainly choose to purchase something at price t_j whenever there is an item she values above t_j . So we see that therefore $\text{SREV} \geq t_j \cdot \Pr_{t_{-j} \sim D_{-j}}[t \notin R_j]$, for all t_j . Thus, TAIL $\leq \text{SREV} \cdot \sum_j \sum_{t_j > r} f_j(t_j) = \sum_j \text{SREV} \cdot \Pr_{t_j \sim D_j}[t_j > \text{SREV}] =$ the revenue of selling each item separately at price SREV, which by the same exact reasoning is also $\leq \text{SREV}$. \square

Now, let's parse CORE. CORE captures contributions to the bound coming from *non-favorite* items whose value is at most SREV. The main idea is that CORE is the expected sum of independent random variables, each supported on $[0, \text{SREV}]$. So maybe CORE = $O(\text{SREV})$, which is great. Or, maybe CORE $\gg \text{SREV}$, in which case it should concentrate (due to being the sum of "small" independent random variables). In the latter case, we should expect to have $\text{BREV} = \Omega(\text{CORE})$, which is also great. Lemma 33 states this formally, and also makes use of the fact that item values are independent.

Lemma 33. *If we sell the grand bundle at price $\text{CORE} - 2r$, the bidder will purchase it with probability at least $1/2$. In other words, $\text{BREV} \geq \frac{\text{CORE}}{2} - r$, or $\text{CORE} \leq 2\text{BREV} + 2\text{SREV}$.*

Proof. We will first need a technical lemma (also used in [BILW14], but proved here for completeness).

Lemma 34. *Let x be a positive single dimensional random variable drawn from F of finite support,²⁶ such that for any number a , $a \cdot \Pr_{x \sim F}[x \geq a] \leq \mathcal{B}$ where \mathcal{B} is an absolute constant. Then for any positive number s , the second moment of the random variable $x_s = x \cdot \mathbb{1}[x \leq s]$ is upper bounded by $2\mathcal{B} \cdot s$.*

²⁴The resulting 6-approximation is roughly the state of the art - subsequent works have improved the analysis of the same mechanism to guarantee a 5.2-approximation [MSL15].

²⁵If not for issues of tie-breaking, this would be equality.

²⁶The same statement holds for continuous distribution as well, and can be proved using integration by parts.

Proof. Let $\{a_1, \dots, a_\ell\}$ be the intersection of the support of F and $[0, s]$, and $a_0 = 0$.

$$\begin{aligned}
\mathbb{E}[x_s^2] &= \sum_{k=0}^{\ell} \Pr_{x \sim F}(x = a_k) \cdot a_k^2 \\
&= \sum_{k=1}^{\ell} (a_k^2 - a_{k-1}^2) \cdot \sum_{d=k}^{\ell} \Pr_{x \sim F}(x = a_d) \\
&\leq \sum_{k=1}^{\ell} (a_k^2 - a_{k-1}^2) \cdot \Pr_{x \sim F}[x \geq a_k] \\
&\leq \sum_{k=1}^{\ell} 2(a_k - a_{k-1}) \cdot a_k \cdot \Pr_{x \sim F}[x \geq a_k] \\
&\leq 2\mathcal{B} \cdot \sum_{k=1}^{\ell} (a_k - a_{k-1}) \\
&\leq 2\mathcal{B} \cdot s
\end{aligned}$$

The penultimate inequality is because $a_k \cdot \Pr_{x \sim F}[x \geq a_k] \leq \mathcal{B}$. \square

Now with Lemma 34, for each j define a new random variable c_j based on the following procedure: draw a sample v_j from D_j , if v_j lies in $[0, r]$, then $c_j = v_j$, otherwise $c_j = 0$. Let $c = \sum_j c_j$. It is not hard to see that we have $\mathbb{E}[c] = \sum_j \sum_{t_j \leq r} f_j(t_j) \cdot t_j$. Now we are going to show that c concentrates because it has small variance. Since the c_j 's are independent, $\text{Var}[c] = \sum_j \text{Var}[c_j] \leq \sum_j \mathbb{E}[c_j^2]$. We will bound each $\mathbb{E}[c_j^2]$ separately. Let $r_j = \max_x \{x \cdot \Pr_{t_j \sim D_j}[t_j \geq x]\}$. By Lemma 34, we can upper bound $\mathbb{E}[c_j^2]$ by $2r_j \cdot r$. On the other hand, it is easy to see that $r = \sum_j r_j$ (as this is exactly the definition of SREV), so $\text{Var}[c] \leq 2r^2$. By the Chebyshev inequality,

$$\Pr[c < \mathbb{E}[c] - 2r] \leq \frac{\text{Var}[c]}{4r^2} \leq \frac{1}{2}.$$

Therefore,

$$\Pr_{t \sim D} \left[\sum_j t_j \geq \mathbb{E}[c] - 2r \right] \geq \Pr[c \geq \mathbb{E}[c] - 2r] \geq \frac{1}{2}.$$

So $\text{BREV} \geq \frac{\mathbb{E}[c] - 2r}{2}$, as we can sell the grand bundle at price $\mathbb{E}[c] - 2r$, and it will be purchased with probability at least $1/2$. \square

Theorem 35. *For a single additive bidder, the optimal revenue is $\leq 2\text{BREV} + 4\text{SREV}$.*

Proof. Combining Lemma 29, 32 and 33, the optimal revenue is upper bounded by $\text{OPT}^{\text{COPIES}} + \text{SREV} + 2\text{BREV} + 2\text{SREV}$. It is not hard to see that $\text{OPT}^{\text{COPIES}} = \text{SREV}$, because the optimal auction in the copies setting just sells everything separately. So the optimal revenue is upper bounded by $2\text{BREV} + 4\text{SREV}$. \square

7 Multiple Bidders

In this section, we show how to use the upper bound in Theorem 28 to show that deterministic DSIC mechanisms can achieve a constant fraction of the (randomized) optimal BIC revenue in multi-bidder settings

when the bidders valuations are all unit-demand or additive. Before beginning, we remind the reader of some notation from Section 2: $\text{VCG}(D)$ refers to the revenue of the VCG mechanism when buyers have values drawn from D , and BVCG refers to the revenue of the optimal ‘‘VCG with entry fees’’ mechanism. $\text{OPT}^{\text{COPIES}}(D)$ refers to the optimal achievable revenue in the related single-dimensional ‘‘copies’’ setting, where each buyer has been split into m different buyers (one for each item).

Similar to the single bidder case, we first decompose the upper bound (Expression 5) into three components and bound them separately. In the last expression in what follows, we call the first term NON-FAVORITE , the second term UNDER and the third term SINGLE . We further break NON-FAVORITE into two parts, OVER and SURPLUS and bound them separately. The following are the approximation factors we achieve:

Theorem 36. *For multiple unit-demand bidders, the optimal revenue is upper bounded by $4\text{OPT}^{\text{COPIES}}$.*

Theorem 37. *For multiple additive bidders, the optimal revenue is upper bounded by $6\text{OPT}^{\text{COPIES}} + 2\text{BVCG}$.*

Note that a simple posted-price mechanism achieves revenue $\text{OPT}^{\text{COPIES}}/6$ when all buyers are unit-demand [CHMS10, KW12], and selling each item separately using Myerson’s auction achieves revenue $\text{OPT}^{\text{COPIES}}$ when buyers are additive. Therefore, the CHMS/KW [CHMS10, KW12] posted-price mechanism achieves a 24-approximation to the optimal BIC mechanism (previously, it was known to be a 30-approximation), and Yao’s approximation ratios [Yao15] are improved from 69 to 8. Some parts of the following analysis draw inspiration from prior works of Chawla et al. [CHMS10] and Yao [Yao15], however, much of the analysis also represents new techniques. In particular, it is worth pointing out that our proof of Theorem 37 looks similar to our single-bidder case, whereas Yao’s original proof required the entirely new machinery of ‘‘ β -adjusted revenue’’ and ‘‘ β -exclusive mechanisms.’’ Below is our decomposition, first into NON-FAVORITE , UNDER , and SINGLE , then further decomposing NON-FAVORITE into OVER and SURPLUS . Recall that \wedge refers to ‘‘AND’’ and \vee refers to ‘‘OR.’’

$$\begin{aligned}
& \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \left(t_{ij} \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \notin R_j^{(v_{-i})}] + \tilde{\varphi}_{ij}(t_{ij}) \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \in R_j^{(v_{-i})}] \right) \\
\leq & \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i} \in T_{-i}} t_{ij} f_{-i}(v_{-i}) \cdot \mathbb{1} \left[(\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})) \vee (t_{ij} < P_{ij}(v_{-i})) \right] \\
& + \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \pi_{ij}(t_i) \tilde{\varphi}_{ij}(t_{ij}) \Pr_{v_{-i} \sim D_{-i}} [t_i \in R_j^{(v_{-i})}] \\
\leq & \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \\
& \quad \sum_{v_{-i} \in T_{-i}} t_{ij} f_{-i}(v_{-i}) \cdot \mathbb{1} \left[(\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})) \wedge (t_{ij} \geq P_{ij}(v_{-i})) \right] \quad (\text{NON-FAVORITE}) \\
& + \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i} \in T_{-i}} t_{ij} \cdot f_{-i}(v_{-i}) \cdot \mathbb{1}[t_{ij} < P_{ij}(v_{-i})] \quad (\text{UNDER}) \\
& + \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \tilde{\varphi}_{ij}(t_{ij}) \cdot \Pr_{v_{-i} \sim D_{-i}} [t_i \in R_j^{(v_{-i})}] \quad (\text{SINGLE})
\end{aligned}$$

$$\begin{aligned}
\text{NON-FAVORITE} &\leq \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i} \in T_{-i}} P_{ij}(v_{-i}) f_{-i}(v_{-i}) \mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})] \quad (\text{OVER}) \\
&+ \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \\
&\quad \sum_{v_{-i} \in T_{-i}} (t_{ij} - P_{ij}(v_{-i})) \cdot f_{-i}(v_{-i}) \cdot \mathbb{1} \left[(\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})) \wedge (t_{ij} \geq P_{ij}(v_{-i})) \right] \quad (\text{SURPLUS})
\end{aligned}$$

Before continuing, let's try to parse these five terms:

- All terms sum over all bidders, all types, and all items, and take the density of that type times the interim probability that bidder receives that item when reporting that type, times some portion of the virtual valuation for that item.
- NON-FAVORITE takes the **value** for the item, times the probability that it is *not the bidder's favorite item*, as defined in Section 5, when v_{-i} is drawn from D_{-i} (roughly corresponds to items k such that $t_i \in R_j^{(v_{-i})}$ for $j \neq k$, but not perfectly).
- UNDER takes the **value** for the item, times the probability that the bidder is *not even willing to purchase the item* at the VCG prices defined by v_{-i} drawn from D_{-i} (roughly corresponds to when $t_i \in R_0^{(v_{-i})}$, but not perfectly).
- SINGLE takes the **Myerson Ironed Virtual Value** for the item, times the probability that it is *the bidder's favorite item* (corresponds to items j such that $t_i \in R_j^{(v_{-i})}$).
- OVER and SURPLUS split NON-FAVORITE in the following way:
 - OVER replaces the value in NON-FAVORITE with the VCG price induced by v_{-i} (and also upper bounds some probabilities by 1). This roughly corresponds to the revenue obtained by VCG (but not perfectly).
 - SURPLUS replaces the value in NON-FAVORITE with (value - VCG price induced by v_{-i}), and roughly corresponds to the bidder's utility for participating in the VCG auction (but not perfectly).
 - Observe that value = VCG price + (value - VCG price), so this is indeed a decomposition of Non-Favorite.

The plan of attack is as follows: SINGLE will be handled the same way as in Section 6. SURPLUS will be handled similarly to NON-FAVORITE from Section 6, and both parts yield the same approximation guarantees as their single-bidder counterparts. That leaves UNDER and OVER, which we will show each contribute at most an additional $\text{OPT}^{\text{COPIES}}$, and account for the “plus two” in transitioning from single-bidder to multi-bidder bounds. We now proceed to address these terms formally, beginning with SURPLUS.

Analyzing SURPLUS for Unit-demand Bidders: The proof of this lemma is similar in spirit to Lemma 30.

Lemma 38. *When the types are unit-demand, for any feasible $\pi(\cdot)$, $\text{SURPLUS} \leq \text{OPT}^{\text{COPIES}}$.*

Proof. Indeed, we will prove that SURPLUS is bounded above by the revenue of the VCG mechanism in the Copies setting. For any i define $S_i(t_i, v_{-i})$ to be the second largest number in $\{t_{i1} - P_{i1}(v_{-i}), \dots, t_{im} - P_{im}(v_{-i})\}$. Now consider running the VCG mechanism on type profile (t_i, v_{-i}) . An agent (i, j) is served in the VCG mechanism in the Copies setting, iff item j is allocated to i in the VCG mechanism in the original setting, which is equivalent to saying $t_{ij} - P_{ij}(v_{-i}) \geq 0$ and $t_{ij} - P_{ij}(v_{-i}) \geq t_{ik} - P_{ik}(v_{-i})$ for all k . The

Copies setting is single-dimensional, therefore any agent's payment is her threshold bid. For agent (i, j) , her threshold bid is $P_{ij}(v_{-i}) + \max\{0, \max_{k \neq j} t_{ik} - P_{ik}(v_{-i})\}$ which is at least $S_i(t_i, v_{-i})$. On the other hand, for any i , whenever $\exists j', t_{ij'} - P_{ij'}(v_{-i}) \geq 0$, there exists some j_i such that (i, j_i) is served in the VCG mechanism. Combining the two conclusions above, we show that on any profile (t_i, v_{-i}) , the payment in the VCG mechanism collected from agents in $\{(i, j)\}_{j \in [m]}$ is at least $S_i(t_i, v_{-i}) \cdot \mathbb{1}[\exists j', t_{ij'} - P_{ij'}(v_{-i}) \geq 0]$. So the total revenue of the VCG Copies mechanism is at least:

$$\sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \cdot S_i(t_i, v_{-i}) \cdot \mathbb{1}[\exists j', t_{ij'} - P_{ij'}(v_{-i}) \geq 0].$$

Next we argue for any j and (t_i, v_{-i}) , the following inequality holds.

$$\begin{aligned} (t_{ij} - P_{ij}(v_{-i})) \cdot \mathbb{1} \left[(\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i}) \geq 0) \right] \\ \leq S_i(t_i, v_{-i}) \cdot \mathbb{1}[\exists j', t_{ij'} - P_{ij'}(v_{-i}) \geq 0] \end{aligned} \quad (6)$$

We only need to consider the case when the LHS is non-zero. In that case, the RHS has value $S_i(t_i, v_{-i})$, and also there exists some k such that $t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})$, so $t_{ij} - P_{ij}(v_{-i}) \leq S_i(t_i, v_{-i})$.

So now we can rewrite SURPLUS and upper bound it with the revenue of the VCG mechanism in the Copies setting.

$$\begin{aligned} & \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \sum_{v_{-i} \in T_{-i}} (t_{ij} - P_{ij}(v_{-i})) \cdot f_{-i}(v_{-i}) \cdot \mathbb{1} \left[(\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i}) \geq 0) \right] \\ &= \sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \sum_j \pi_{ij}(t_i) \cdot (t_{ij} - P_{ij}(v_{-i})) \cdot \mathbb{1} \left[(\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i}) \geq 0) \right] \\ &\leq \sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \sum_j \pi_{ij}(t_i) \cdot S_i(t_i, v_{-i}) \cdot \mathbb{1}[\exists j', t_{ij'} - P_{ij'}(v_{-i}) \geq 0] \quad (\text{Inequality (6)}) \\ &\leq \sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \cdot S_i(t_i, v_{-i}) \cdot \mathbb{1}[\exists j', t_{ij'} - P_{ij'}(v_{-i}) \geq 0] \quad \left(\sum_j \pi_{ij}(t_i) \leq 1 \forall i, t_i \right) \end{aligned}$$

The last line is upper bounded by the revenue of the VCG mechanism in the Copies setting by our work above, which is clearly upper bounded by $\text{OPT}^{\text{COPIES}}$. \square

Analyzing SURPLUS for Additive Bidders: Similar to the single bidder case, we will again break the term SURPLUS into the CORE and the TAIL, and analyze them separately. Before we proceed, we first define the cutoffs. Let $r_{ij}(v_{-i}) = \max_{x \geq P_{ij}(v_{-i})} \{x \cdot \Pr_{t_{ij} \sim D_{ij}}[t_{ij} \geq x]\}$. The observant reader will notice that this is bidder i 's ex-ante payment for item j in Ronen's single-item mechanism [Ron01] conditioned on other bidders types being v_{-i} , but this connection is not necessary to understand the proof. Further let $r_i(v_{-i}) = \sum_j r_{ij}(v_{-i})$, $r_i = \mathbb{E}_{v_{-i} \sim D_{-i}}[r_i(v_{-i})]$ and $r = \sum_i r_i$, the expected revenue of running Ronen's mechanism separately for each item (again, the connection to Ronen's mechanism is not necessary to understand the proof). We first bound TAIL and CORE, using arguments similar to the single item case (Lemmas 32 and 33),

$$\begin{aligned}
\text{SURPLUS} &\leq \sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) \sum_j \sum_{t_{ij} \geq P_{ij}(v_{-i})} f_{ij}(t_{ij}) \cdot (t_{ij} - P_{ij}(v_{-i})) \cdot \sum_{t_{i,-j} \in T_{i,-j}} f_{i,-j}(t_{i,-j}) \cdot \\
&\quad \mathbb{1}[\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})] \\
&= \sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) \sum_j \sum_{t_{ij} \geq P_{ij}(v_{-i})} f_{ij}(t_{ij}) \cdot \\
&\quad (t_{ij} - P_{ij}(v_{-i})) \cdot \Pr_{t_{i,-j} \sim D_{i,-j}} [\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})] \\
&\leq \sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) \sum_j \sum_{t_{ij} > P_{ij}(v_{-i}) + r_i(v_{-i})} f_{ij}(t_{ij}) \cdot \\
&\quad (t_{ij} - P_{ij}(v_{-i})) \cdot \Pr_{t_{i,-j} \sim D_{i,-j}} [\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})] \quad (\text{TAIL}) \\
&+ \sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) \sum_j \sum_{t_{ij} \in [P_{ij}(v_{-i}), P_{ij}(v_{-i}) + r_i(v_{-i})]} f_{ij}(t_{ij}) \cdot (t_{ij} - P_{ij}(v_{-i})) \quad (\text{CORE})
\end{aligned}$$

Lemma 39. TAIL $\leq r$.

Proof. First, by union bound

$$\Pr_{t_{i,-j} \sim D_{i,-j}} [\exists k \neq j, t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})] \leq \sum_{k \neq j} \Pr_{t_{ik} \sim D_{ik}} [t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})].$$

By the definition of $r_{ik}(v_{-i})$, we certainly have $r_{ik}(v_{-i}) \geq (P_{ik}(v_{-i}) + t_{ij} - P_{ij}(v_{-i})) \cdot \Pr_{t_{ik} \sim D_{ik}} [t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})]$, so we can also derive:

$$\Pr_{t_{ik} \sim D_{ik}} [t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i})] \leq \frac{r_{ik}(v_{-i})}{P_{ik}(v_{-i}) + t_{ij} - P_{ij}(v_{-i})} \leq \frac{r_{ik}(v_{-i})}{t_{ij} - P_{ij}(v_{-i})}.$$

Using these two inequalities, we can upper bound TAIL:

$$\begin{aligned}
&\sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) \sum_j \sum_{t_{ij} > P_{ij}(v_{-i}) + r_i(v_{-i})} f_{ij}(t_{ij}) \cdot \sum_{k \neq j} r_{ik}(v_{-i}) \\
&\leq \sum_i \sum_{v_{-i}} f_{-i}(v_{-i}) \cdot \sum_j r_i(v_{-i}) \cdot \sum_{t_{ij} > P_{ij}(v_{-i}) + r_i(v_{-i})} f_{ij}(t_{ij}) \\
&\leq \sum_i \sum_{v_{-i}} f_{-i}(v_{-i}) \sum_j r_{ij}(v_{-i}) \quad (\text{Definition of } r_{ij}(v_{-i})) \\
&= r
\end{aligned}$$

□

Lemma 40. BVCG $\geq \frac{\text{CORE}}{2} - r$. In other words, $2r + 2\text{BVCG} \geq \text{CORE}$.

Proof. Fix any $v_{-i} \in T_{-i}$, let $t_{ij} \sim D_{ij}$, define two new random variables

$$b_{ij}(v_{-i}) = (t_{ij} - P_{ij}(v_{-i})) \mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})]$$

and

$$c_{ij}(v_{-i}) = b_{ij}(v_{-i}) \mathbb{1}[b_{ij}(v_{-i}) \leq r_i(v_i)].$$

Clearly, $c_{ij}(v_{-i})$ is supported on $[0, r_i(v_{-i})]$. Also, we have

$$\mathbb{E}_{t_{ij} \sim D_{ij}}[c_{ij}(v_{-i})] = \sum_{t_{ij} \in [P_{ij}(v_{-i}), P_{ij}(v_{-i}) + r_i(v_{-i})]} f_{ij}(t_{ij}) \cdot (t_{ij} - P_{ij}(v_{-i})).$$

So we can rewrite CORE as

$$\sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) \sum_j \mathbb{E}[c_{ij}(v_{-i})].$$

Now we will describe a VCG mechanism with per bidder entry fee. Define an entry fee function for bidder i depending on v_{-i} as $e_i(v_{-i}) = \sum_j \mathbb{E}[c_{ij}(v_{-i})] - 2r_i(v_{-i})$. We will show that for any i and other bidders types $v_{-i} \in T_{-i}$, bidder i accepts the entry fee $e_i(v_{-i})$ with probability at least $1/2$. Since bidders are additive, the VCG mechanism is exactly m separate Vickrey auctions, one for each item. So $P_{ij}(v_{-i}) = \max_{\ell \neq i} \{v_{\ell j}\}$, and for any set of S , its Clarke Pivot price for i to receive set S is $\sum_{j \in S} P_{ij}(v_{-i})$.

That also means $\sum_j b_{ij}(v_{-i})$ is the random variable that represents bidder i 's utility in the VCG mechanism when other bidders bids are v_{-i} . If we can prove $\Pr[\sum_j b_{ij}(v_{-i}) \geq e_i(v_{-i})] \geq 1/2$ for all v_{-i} , then we know bidder i accepts the entry fee with probability at least $1/2$.

It is not hard to see for any nonnegative number a ,

$$a \cdot \Pr[b_{ij}(v_{-i}) \geq a] \leq (a + P_{ij}(v_{-i})) \cdot \Pr[t_{ij} \geq a + P_{ij}(v_{-i})] \leq r_{ij}(v_{-i}).$$

Therefore, because each $c_{ij}(v_{-i}) \in [0, r_i(v_{-i})]$, by Lemma 34 we can again bound the second moment as: $\mathbb{E}[c_{ij}(v_{-i})^2] \leq 2r_i(v_{-i})r_{ij}(v_{-i})$. Since c_{ij} 's are independent,

$$\text{Var}[\sum_j c_{ij}(v_{-i})] = \sum_j \text{Var}[c_{ij}(v_{-i})] \leq \sum_j \mathbb{E}[c_{ij}(v_{-i})^2] \leq 2r_i(v_{-i})^2.$$

By Chebyshev inequality, we know

$$\Pr[\sum_j c_{ij}(v_{-i}) \leq \sum_j \mathbb{E}[c_{ij}(v_{-i})] - 2r_i(v_{-i})] \leq \frac{\text{Var}[\sum_j c_{ij}(v_{-i})]}{4r_i(v_{-i})^2} \leq 1/2.$$

Therefore, as $b_{ij}(v_{-i}) \geq c_{ij}(v_{-i})$, we can conclude:

$$\Pr[\sum_j b_{ij}(v_{-i}) \geq e_i(v_{-i})] \geq 1/2$$

So the entry fee is accepted with probability at least $1/2$ for all i and v_{-i} . So:

$$\text{BVCG} \geq \frac{1}{2} \sum_i \sum_{v_{-i} \in T_{-i}} f_{-i}(v_{-i}) (\mathbb{E}[c_{ij}(v_{-i})] - 2r_i(v_{-i})) = \frac{\text{CORE}}{2} - r.$$

□

Analyzing SINGLE, OVER and UNDER: First we consider SINGLE, which is similar to Lemma 29.

Lemma 41. For any feasible $\pi(\cdot)$, $\text{SINGLE} \leq \text{OPT}^{\text{COPIES}}$.

Proof. Assume M is the ex-post allocation rule that induces $\pi(\cdot)$. Consider another ex-post allocation rule M' for the copies setting, such that for every type profile t , if M allocates item j to bidder i in the original setting then M' serves agent (i, j) with probability $\Pr_{v_{-i} \sim D_{-i}}[t_i \in R_j^{(v_{-i})}]$. As M is feasible in the original setting, M' is clearly feasible in the Copies setting. When agent (i, j) has type t_{ij} , her probability of being served in M' is

$$\sum_{t_{i,-j}} f_{i,-j}(t_{i,-j}) \cdot \pi_{ij}(t_{ij}, t_{i,-j}) \cdot \Pr_{v_{-i} \sim D_{-i}} [(t_{ij}, t_{i,-j}) \in R_j^{(v_{-i})}]$$

for all j and t_{ij} . Therefore, SINGLE is the ironed virtual welfare achieved by M' with respect to $\tilde{\varphi}(\cdot)$. Since the copies setting is a single dimensional setting, the optimal revenue $\text{OPT}^{\text{COPIES}}$ equals the maximum ironed virtual welfare, thus no smaller than SINGLE. \square

Next, we move onto OVER. Recall that the terms in OVER are VCG prices times an indicator that the bidder's value exceeds the VCG prices. So we should hope to be able to cover OVER with some VCG-like mechanism. We begin with the following technical propositions:

Proposition 42. *Let $\pi(\cdot)$ be any reduced form of a BIC mechanism in the original setting. Define*

$$\Pi_{ij}(t_{ij}) = \mathbb{E}_{t_{i,-j} \sim D_{i,-j}} [\pi_{ij}(t_{ij}, t_{i,-j})].$$

Then $\Pi_{ij}(t_{ij})$ is monotone in t_{ij} .

Proof. In fact, for all $t_{i,-j}$, we must have $\pi_{ij}(\cdot, t_{i,-j})$ monotone increasing in t_{ij} . Assume for contradiction that this were not the case, and let $t_{ij} < t'_{ij}$ with $\pi_{ij}(t_{ij}, t_{i,-j}) > \pi_{ij}(t'_{ij}, t_{i,-j})$. Then $(t_{ij}, t_{i,-j}), (t'_{ij}, t_{i,-j})$ form a 2-cycle that violates cyclic monotonicity. This is because both types value all items except for j exactly the same.

To expand a bit for readers not familiar with cyclic monotonicity: observe that $t_{ij} < t'_{ij}$ but $\pi_{ij}(t_{ij}, t_{i,-j}) > \pi_{ij}(t'_{ij}, t_{i,-j})$ implies that

$$\begin{aligned} (t_{ij}, t_{i,-j}) \cdot \pi_i(t_{ij}, t_{i,-j}) + (t'_{ij}, t_{i,-j}) \cdot \pi_i(t'_{ij}, t_{i,-j}) &< (t'_{ij}, t_{i,-j}) \cdot \pi_i(t_{ij}, t_{i,-j}) + (t_{ij}, t_{i,-j}) \cdot \pi_i(t'_{ij}, t_{i,-j}) \\ \Rightarrow (t_{ij}, t_{i,-j}) \cdot \pi_i(t_{ij}, t_{i,-j}) - (t_{ij}, t_{i,-j}) \cdot \pi_i(t'_{ij}, t_{i,-j}) &< (t'_{ij}, t_{i,-j}) \cdot \pi_i(t_{ij}, t_{i,-j}) - (t'_{ij}, t_{i,-j}) \cdot \pi_i(t'_{ij}, t_{i,-j}). \end{aligned}$$

This directly implies that no matter what prices are set for $p_i(t_{ij}, t_{i,-j})$ and $p_i(t'_{ij}, t_{i,-j})$, if bidder i with type $(t_{ij}, t_{i,-j})$ is happy to tell the truth, then type $(t'_{ij}, t_{i,-j})$ strictly prefers to lie and report $(t_{ij}, t_{i,-j})$ than tell the truth. \square

Proposition 43. *For any $v \in T$, any $\pi(\cdot)$ that is a reduced form of some BIC mechanism,*

$$\text{OPT}^{\text{COPIES}} \geq \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot P_{ij}(v_{-i}) \cdot \mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})].$$

Proof. Recall from Proposition 42 that every BIC interim form $\pi(\cdot)$ in the original setting corresponds to a monotone interim form in the copies setting, $\Pi(\cdot)$. Let M be any (possibly randomized) allocation rule that induces $\Pi(\cdot)$, and $p(\cdot)$ a corresponding price rule (wlog we can let (M, p) be ex-post IR). Consider the following mechanism instead: on input t , first run (M, p) to (possibly randomly) determine a set of potential winners. Then, if (i, j) is a potential winner, offer (i, j) service at price $\max\{p_{ij}(t), P_{ij}(v_{-i})\}$. Whenever (i, j) is a potential winner, $t_{ij} \geq p_{ij}(t)$. It is clear that in the event that (i, j) is a potential winner, and $t_{ij} \geq$

$P_{ij}(t_{-i})$, (i, j) will accept the price and pay at least $P_{ij}(v_{-i})$. Therefore, for any t as long as (i, j) is served in M , then the payment from (i, j) in the new proposed mechanism is at least $P_{ij}(v_{-i})\mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})]$. That means the total revenue of the new mechanism is at least $\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot P_{ij}(v_{-i}) \cdot \mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})]$, which is upper bounded by $\text{OPT}^{\text{COPIES}}$. \square

Lemma 44. $\text{OVER} \leq \text{OPT}^{\text{COPIES}}$.

Proof. This can be proved by rewriting OVER and then applying Proposition 43.

$$\begin{aligned} \text{OVER} &= \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v \in T} P_{ij}(v_{-i}) f(v) \mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})] \\ &= \sum_{v \in T} f(v) \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot P_{ij}(v_{-i}) \cdot \mathbb{1}[t_{ij} \geq P_{ij}(v_{-i})] \\ &\leq \sum_{v \in T} f(v) \cdot \text{OPT}^{\text{COPIES}} = \text{OPT}^{\text{COPIES}} \end{aligned}$$

\square

Finally, we move on to UNDER . When there is only one bidder, UNDER is always 0. Here, $\text{UNDER} \leq \text{OPT}^{\text{COPIES}}$, and turns out to be the trickiest part to bound. Recall that UNDER contains terms that are (non-favorite) values times indicators that these values *do not* exceed the VCG prices. So the high-level hope is that the reason the VCG price for bidder i to receive item j exceeds t_{ij} is because someone else is paying at least t_{ij} for something, and we might hope to be able to come up with a clever charging argument. At a high level, this is indeed the plan, but the proof approach doesn't clearly map onto this intuition. We apply Proposition 45 (below) once for each type profile t , using the allocation of this mechanism on type profile t to specify (i_j, j) and let $x_j = t_{i_j j}$. Then taking the convex combination of the RHS of Proposition 45 for all profiles t with multipliers $f(t)$ gives $\text{UNDER} \leq \text{OPT}^{\text{COPIES}}$.

Proposition 45. Let $\{(i_j, j)\}_{j \in S \subseteq [m]}$ be a feasible allocation in the copies setting. For all choices $x_1, \dots, x_m \geq 0$, $\text{OPT}^{\text{COPIES}} \geq \sum_{v \in T} f(v) \cdot \sum_{j \in S} x_j \cdot \mathbb{1}[P_{i_j j}(v_{-i_j}) > x_j]$.

Proof. Before beginning the proof of Proposition 45, we will need the following definition and theorem due to Gul and Stacchetti [GS99].

Definition 46. Let $W_T(S)$ be the maximum attainable welfare using only bidders in T and items in S .

Theorem 47. ([GS99]) If all bidders in T have gross substitute valuations, then $W_T(\cdot)$ is a submodular function.

Now with Theorem 47, consider in the Copies setting the VCG mechanism with lazy reserve x_j for each copy (i, j) . Specifically, we will first solicit bids, then find the max-welfare allocation and call all (i, j) who get allocated temporary winners. Then, if (i, j) is a temporary winner, (i, j) is given the option to receive service for the maximum of their Clarke pivot price in the Copies setting and x_j . It is clear that in this mechanism, whenever any agent (i, j) receives service, the price she pays is at least x_j . Also, it is not hard to see that this is a truthful mechanism (for the Copies): for all other fixed bids, copy (i, j) can report a bid exceeding the maximum of x_j and their Clarke pivot price, or not. If they report a higher bid, they will receive service and pay the maximum of their Clarke pivot price and x_j . If they report a smaller number, they remain unserved. It's clear that bidding the Copy's true value is always optimal. Next, we argue for any $v \in T$ and $j \in S$, whenever $P_{i_j j}(v_{-i_j}) > x_j$, there exists some i such that (i, j) is served in the mechanism above.

By the definition of Clarke pivot price, we know

$$P_{ij}(v_{-i_j}) = W_{[n]-\{i_j\}}([m]) - W_{[n]-\{i_j\}}([m] - \{j\}).$$

First, we show that if item j is allocated to some bidder i in the max-welfare allocation in the original setting then $v_{ij} \geq P_{ij}(v_{-i})$. Assume S' to be the set of items allocated to bidder i . Since the VCG mechanism is truthful, the utility for winning set S' is better than winning set $S' - \{j\}$:

$$\begin{aligned} & \sum_{k \in S'} v_{ik} - (W_{[n]-\{i\}}([m]) - W_{[n]-\{i\}}([m] - S')) \\ & \geq \sum_{k \in S' - \{j\}} v_{ik} - (W_{[n]-\{i\}}([m]) - W_{[n]-\{i\}}([m] - S' + \{j\})). \end{aligned}$$

Rearranging the terms, we get

$$\begin{aligned} v_{ij} & \geq W_{[n]-\{i\}}([m] - S' + \{j\}) - W_{[n]-\{i\}}([m] - S') \\ & \geq W_{[n]-\{i\}}([m]) - W_{[n]-\{i\}}([m] - \{j\}) \quad (\text{Theorem 47}) \\ & = P_{ij}(v_{-i}). \end{aligned}$$

Now we still need to argue that whenever $P_{ij}(v_{-i_j}) > x_j$, item j is always allocated in the max-welfare allocation to some bidder i with $v_{ij} \geq x_j$.

1. If agent (i_j, j) is a temporary winner,

$$v_{ij} \geq P_{ij}(v_{-i_j}) > x_j.$$

Therefore, agent (i_j, j) will accept the price.

2. If agent (i_j, j) is not a temporary winner, let S' be the set of items that are allocated to bidder i_j in the welfare maximizing allocation in the original setting. Since $W_{[n]-\{i_j\}}([m] - S') - W_{[n]-\{i_j\}}([m] - S' - \{j\}) \geq W_{[n]-\{i_j\}}([m]) - W_{[n]-\{i_j\}}([m] - \{j\}) = P_{ij}(v_{-i_j})$ (by Theorem 47), and $P_{ij}(v_{-i_j}) > x_j$, the following are true: (i) item j is awarded to some bidder $i \neq i_j$ in the welfare maximizing allocation, because otherwise $W_{[n]-\{i_j\}}([m] - S')$ will have the same value as $W_{[n]-\{i_j\}}([m] - S' - \{j\})$; (ii) $v_{ij} > x_j$ because

$$v_{ij} \geq W_{[n]-\{i_j\}}([m] - S') - W_{[n]-\{i_j\}}([m] - S' - \{j\}) = P_{ij}(v_{-i_j}) > x_j.$$

So now we can conclude that for any $j \in S$ there is certainly some i such that (i, j) is served whenever $P_{ij} > x_j$, and therefore the revenue of this mechanism in the Copies setting is at least $\sum_{v \in T} f(v) \cdot \sum_{j \in S} x_j \cdot \mathbb{1}[P_{ij}(v_{-i_j}) > x_j]$, which is exactly the same as the sum in the proposition statement. \square

Lemma 48. UNDER \leq OPT^{COPIES}.

Proof. The idea is to interpret UNDER as the revenue of the following mechanism: let M be the mechanism that induces $\pi(\cdot)$. Sample t from D , let S be the set of agents that will be served in M for type profile t in the copies setting. Use t_{ij} to be the reserve price for j if $(i, j) \in S$, and use the mechanism in Proposition 45.

First, the inner sum

$$\sum_{v_{-i} \in T_{-i}} t_{ij} \cdot f_{-i}(v_{-i}) \cdot \mathbb{1}[t_{ij} < P_{ij}(v_{-i})]$$

only depends on t_i , so the maximum of UNDER is achieved by a $\pi(\cdot)$ induced by some deterministic mechanism. Wlog, we consider $\pi(\cdot)$ is induced by a deterministic mechanism whose ex-post allocation rule is $x(\cdot)$. Let us rewrite UNDER using $x(\cdot)$:

$$\begin{aligned}
& \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i} \in T_{-i}} t_{ij} \cdot f_{-i}(v_{-i}) \cdot \mathbb{1}[t_{ij} < P_{ij}(v_{-i})] \\
&= \sum_{t \in T} f(t) \sum_i \sum_j x_{ij}(t) \cdot t_{ij} \cdot \sum_{v \in T} f(v) \cdot \mathbb{1}[t_{ij} < P_{ij}(v_{-i})] \\
&= \sum_{t \in T} f(t) \cdot \sum_{v \in T} f(v) \sum_i \sum_j x_{ij}(t) \cdot t_{ij} \cdot \mathbb{1}[t_{ij} < P_{ij}(v_{-i})] \\
&\leq \sum_{t \in T} f(t) \cdot \text{OPT}^{\text{COPIES}} = \text{OPT}^{\text{COPIES}}
\end{aligned}$$

The penultimate inequality is because if we let $\{(i, j)\}_{j \in S}$ be the set of agents such that by $x_{ij}(t) = 1$, then

$$\sum_{v \in T} f(v) \sum_i \sum_j x_{ij}(t) \cdot t_{ij} \cdot \mathbb{1}[t_{ij} < P_{ij}(v_{-i})] = \sum_{v \in T} f(v) \cdot \sum_{j \in S} x_j \cdot \mathbb{1}[P_{ij}(v_{-i}) > x_j],$$

and by Proposition 45, this is upper bounded by $\text{OPT}^{\text{COPIES}}$. \square

Combining the above lemmas now yields our theorems:

Proof of Theorem 36: Combine Lemmas 38, 41, 44 and 48. \square

Proof of Theorem 37: Combining Lemmas 39, 40, 41, 44 and 48, we get the optimal revenue is upper bounded by

$$3\text{OPT}^{\text{COPIES}} + 3r + 2\text{BVCG}.$$

Since $\text{OPT}^{\text{COPIES}}$ is exactly the revenue of selling each item separately optimally using Myerson's auction, and r is the revenue of some mechanism (Ronen's) that sells the items separately, we have $\text{OPT}^{\text{COPIES}} \geq r$, proving the statement. \square

8 Duality Theory Beyond Additive Bidders

In this section we provide a statement of our duality theory that holds much more generally than when bidders are unit-demand or additive. The technical ideas are exactly the same as in Section 3 and just require updated notation.

Buyer Valuations. In this section, we will consider buyers with arbitrary valuation functions for subsets of items. That is, buyer i has some valuation function $t_i(\cdot)$ that takes as input a set of items and outputs a value. Buyer i 's type is drawn from some distribution D_i , and $D = \times_i D_i$ is the joint distribution over profiles of buyer types. We define \mathcal{F} to be a set system over $[n] \times [m]$ that describes all feasible allocations.

Implicit Forms. The implicit form of an auction stores for all bidders i , and pairs of types t_i, t'_i , what is the expected value that that agent i will receive when her real valuation function is t_i , but she reports t'_i to the mechanism instead (over the randomness in the mechanism and randomness in other agents' reported types, assuming they come from D_{-i}) as $\pi_i(t_i, t'_i)$. We say that an implicit form is *feasible* if there exists some feasible mechanism (that selects an outcome in \mathcal{F} with probability 1) that matches the expectations

promised by the implicit form. If $P(\mathcal{F}, D)$ is defined to be the set of all feasible implicit forms, it is easy to see (and shown in [CDW13b], for instance) that $P(\mathcal{F}, D)$ is closed and convex. Note that implicit forms are computed over the same randomness as reduced forms, but store directly the value that bidder i receives for having type t_i and reporting t'_i , instead of indirectly via interim probabilities.

We begin by writing the LP for revenue maximization in this more general setting (Figure 5). To proceed, we'll again introduce a variable $\lambda_i(t, t')$ for each of the BIC constraints, and take the partial Lagrangian of LP 5 by Lagrangifying all BIC constraints. The theory of Lagrangian multipliers tells us that the solution to LP 5 is equivalent to the primal variables solving the partially Lagrangified dual (Figure 6).

Definition 49. Let $\mathcal{L}(\lambda, \pi, p)$ be a the partial Lagrangian defined as follows:

$$\mathcal{L}(\lambda, \pi, p) = \sum_{i=1}^n \left(\sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) + \sum_{t_i \in T_i} \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \cdot \left(\pi_i(t_i, t_i) - \pi_i(t_i, t'_i) - (p_i(t_i) - p_i(t'_i)) \right) \right) \quad (7)$$

$$\begin{aligned} &= \sum_{i=1}^n \sum_{t_i \in T_i} p_i(t_i) \left(f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \right) \\ &\quad + \sum_{i=1}^n \sum_{t_i \in T_i} \left(\sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \cdot t_i(X_i(t_i)) - \sum_{t'_i \in T_i} t'_i \cdot \lambda_i(t'_i, t_i) \cdot t'_i(X_i(t_i)) \right) \quad (8) \end{aligned}$$

In Equation (8), we use $X_i(t_i)$ to denote the the random set awarded to bidder i when reporting type t_i to the mechanism. That is, $t_i(X_i(t'_i)) := \pi_i(t_i, t'_i)$.

Variables:

- $p_i(t_i)$, for all bidders i and types $t_i \in T_i$, denoting the expected price paid by bidder i when reporting type t_i over the randomness of the mechanism and the other bidders' types.
- $\pi_i(t_i, t'_i)$, for all bidders i , and types $t_i, t'_i \in T_i$, denoting the expected value that bidder i receives when her real type is type t_i but reports t'_i , over the randomness of the mechanism and the other bidders' types.

Constraints:

- $\pi_i(t_i, t_i) - p_i(t_i) \geq \pi_i(t_i, t'_i) - p_i(t'_i)$, for all bidders i , and types $t_i \in T_i, t'_i \in T_i^+$, guaranteeing that the implicit form mechanism (π, p) is BIC and BIR.
- $\pi \in P(\mathcal{F}, D)$, guaranteeing π is feasible.

Objective:

- Maximize: $\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i)$, the expected revenue.

Figure 5: A Linear Program (LP) for Revenue Optimization.

Lemma 4 immediately holds in this setting as well, and the proof is identical. That is, a feasible dual solution is still useful if and only if it induces a flow in the same graph. We will define virtual valuation functions in essentially the same way, just updating notation.

Variables:

- $\lambda_i(t_i, t'_i)$ for all $i, t_i \in T_i, t'_i \in T_i^+$, the Lagrangian multipliers for Bayesian IC constraints.

Constraints:

- $\lambda_i(t_i, t'_i) \geq 0$ for all $i, t_i \in T_i, t'_i \in T_i^+$, guaranteeing that the Lagrangian multipliers are non-negative.

Objective:

- Minimize: $\max_{\pi \in P(\mathcal{F}, D), p} \mathcal{L}(\lambda, \pi, p)$.

Figure 6: Partial Lagrangian of the Revenue Maximization LP.

Definition 50 (Virtual Value Function). *For each λ , we define a corresponding virtual value function $\Phi(\cdot)$, such that for every bidder i , every type $t_i \in T_i$, $\Phi_i(t_i)(\cdot) = t_i(\cdot) - \frac{1}{f_i(t_i)} \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)(t'_i(\cdot) - t_i(\cdot))$. That is, $\Phi_i(t_i)$ is a function that takes as input sets of items and outputs a value. For any set of items S , $\Phi_i(t_i)(S) = t_i(S) - \frac{1}{f_i(t_i)} \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)(t'_i(S) - t_i(S))$.*

We can now state the proper generalization of Theorem 6 in this general setting. The proof is identical to that of Theorem 6 and we omit it. In the theorem statement, $X_i(t_i)$ again denotes the random set allocated to bidder i when reporting type t_i , so that $\pi_i(t_i, t'_i) = t_i(X_i(t'_i))$.

Theorem 51 (Virtual Welfare \geq Revenue). *Let λ be any useful dual solution and $M = (\pi, p)$ any BIC mechanism. Then the revenue of M is \leq the virtual welfare of π w.r.t. the virtual value function $\Phi(\cdot)$ corresponding to λ . That is:*

$$\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) \leq \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \Phi_i(t_i)(X_i(t_i)).$$

Equality holds if and only if for all i, t, t' such that $\lambda_i(t, t') > 0$, the BIC constraint for bidder i between t and t' binds in M (that is, bidder i with type t is indifferent between reporting t and t'). Furthermore, let λ^ be the optimal dual variables and $M^* = (\pi^*, p^*)$ be the revenue optimal BIC mechanism, then the expected virtual welfare with respect to Φ^* (induced by λ^*) under π^* equals the expected revenue of M^* , and*

$$\pi^* \in \operatorname{argmax}_{\pi \in P(\mathcal{F}, D)} \left\{ \sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \Phi_i^*(t_i)(X_i(t_i)) \right\}.$$

9 Conclusion

We present a new duality framework for Bayesian mechanism design, and show how to recover and improve the state-of-the-art mechanisms for additive or unit-demand bidders with independent item values. Additionally, our proofs for the single-item, unit-demand, and additive settings are “unified” in the sense that we’ve separated out part of the proof that is (nearly) identical for all three settings (the duality-based upper bound, Section 5), so that the additional work necessary for each result is minimized (Sections 6 and 7).²⁷

²⁷One way to think of this is that if one wishes to exclusively understand the main result from a single one of [Mye81, CMS15, BILW14, Yao15], the quickest way to do so is probably still just to read the original papers. However, if one wishes to understand all of these results, a substantial fraction of the proofs overlap via our approach.

Additionally, our framework provides a principled starting point for future work (as evidenced by the numerous recent follow-ups of Section 1.3.4). Many of these works themselves explore new areas and present further open questions (such as competition complexity [EFF⁺17b, LP16], limited complementarity [EFF⁺17a], two-sided markets [BCWZ17], and “one-and-a-half” dimensional settings [FGKK16, DW17, DHP17]).

Our approach also yields insight into any incentives problem that can be captured by an LP with “incentive constraints” and “feasibility constraints.” Domains such as signaling (e.g. [Dug14]) and contract theory (e.g. [Car15]) are amenable to such LP formulations. Bayesian Persuasion seems an especially enticing domain, as the algorithmic CDW framework has already found application there [DX15].

10 Acknowledgements

We would like to thank Costis Daskalakis and Christos Papadimitriou for numerous helpful discussions during the preliminary stage of this work, and Jason Hartline and Rakesh Vohra for helpful discussion and pointers to related works.

References

- [Arm96] Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75, January 1996. [1.3.2](#)
- [Arm99] Mark Armstrong. Price discrimination by a many-product firm. *Review of Economic Studies*, 66(1):151–68, January 1999. [1.3.2](#)
- [BCKW10] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing Randomized Allocations. In *the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010. [1.1](#)
- [BCWZ17] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *The 18th ACM Conference on Economics and Computation (EC)*, 2017. [1.3.4](#), [9](#)
- [BDHS15] MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. Revenue maximization for selling multiple correlated items. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 95–105, 2015. [1.1](#), [1.3.3](#), [2](#)
- [BGM13] Anand Bhalgat, Sreenivas Gollapudi, and Kamesh Munagala. Optimal auctions via the multiplicative weight method. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, pages 73–90, 2013. [1](#), [1.3.2](#)
- [BH11] Xiaohui Bei and Zhiyi Huang. Bayesian incentive compatibility via fractional assignments. In *the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011. [2](#)
- [BILW14] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014. [1](#), [1.1](#), [1.3.1](#), [1.3.3](#), [2](#), [6](#), [6](#), [6](#), [27](#)

- [BK96] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, pages 180–194, 1996. [1.3.3](#)
- [Car15] Gabriel Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015. [9](#)
- [Car16] Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 2016. [1.3.1](#), [1.3.2](#)
- [CDW12a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012. [1](#), [2](#)
- [CDW12b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012. [1](#)
- [CDW13a] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing Revenue to Welfare Maximization : Approximation Algorithms and other Generalizations. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013. [1](#)
- [CDW13b] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013. [\(document\)](#), [1](#), [1.2](#), [1.3.2](#), [8](#)
- [CGL15] Ning Chen, Nikolai Gravin, and Pinyan Lu. Competitive analysis via benchmark decomposition. In *The 16th ACM Conference on Economics and Computation (EC)*, 2015. [1.3.2](#)
- [CH13] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013. [1](#)
- [CHK07] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007. [\(document\)](#), [1](#), [1.1](#), [1.3.1](#), [1.3.3](#), [2](#), [6](#)
- [CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010. [\(document\)](#), [1](#), [1.1](#), [1.3.1](#), [1.3.3](#), [2](#), [6](#), [23](#), [7](#)
- [CM16] Shuchi Chawla and J. Benjamin Miller. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 579–596, 2016. [1.3.3](#), [1.3.4](#)
- [CMS15] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015. [1](#), [1.1](#), [1.3.1](#), [1.3.2](#), [1.3.3](#), [2](#), [6](#), [6](#), [6](#), [27](#)
- [CZ17] Yang Cai and Mingfei Zhao. Simple Mechanisms for Subadditive Buyers via Duality. In *the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017. [1.3.4](#)

- [DDT13] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the Fourteenth ACM Conference on Economics and Computation (EC)*, 2013. 1.1, 1.3.1
- [DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014. 1.1
- [DDT15] 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, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 449–450, 2015. 1.3.1
- [DDT16] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 2016. 1.3.1
- [DDW15] Constantinos Daskalakis, Nikhil R. Devanur, and S. Matthew Weinberg. Revenue maximization and ex-post budget constraints. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 433–447, 2015. 1
- [DHP17] Nikhil Devanur, Nima Haghpanah, and Christos-Alexandros Psomas. Optimal multi-unit mechanisms with private demands. In *The 18th ACM Conference on Economics and Computation (EC)*, 2017. 9
- [Dug14] Shaddin Dughmi. On the hardness of signaling. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014. 9
- [DW12] Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012. 2, 1
- [DW15] Constantinos Daskalakis and S. Matthew Weinberg. Bayesian truthful mechanisms for job scheduling from bi-criterion approximation algorithms. In *the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015. 1
- [DW17] Nikhil Devanur and S. Matthew Weinberg. The optimal mechanism for a budget-constrained buyer: the general case. In *The 18th ACM Conference on Economics and Computation (EC)*, 2017. 1.3.4, 9
- [DX15] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. In *the 47th Annual ACM Symposium on Theory of Computing (STOC)*, 2015. 9
- [EFF⁺17a] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. In *The 18th ACM Conference on Economics and Computation (EC)*, 2017. 1.3.4, 9
- [EFF⁺17b] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. A simple and approximately optimal mechanism for a buyer with complements. In *The 18th ACM Conference on Economics and Computation (EC)*, 2017. 1.3.4, 9
- [FGKK16] 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. 1.3.4, 9

- [FLLT17] Hu Fu, Chris Liaw, Pinyan Lu, and Zhihao Gavin Tang. The Value of Information Concealment. *Manuscript*, 2017. [1.3.4](#)
- [Gia14] Yiannis Giannakopoulos. A note on optimal auctions for two uniformly distributed items. *CoRR*, abs/1409.6925, 2014. [1.3.1](#)
- [GK14] Yiannis Giannakopoulos and Elias Koutsoupias. Duality and optimality of auctions for uniform distributions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 259–276, 2014. [1.3.1](#)
- [GK15] Yiannis Giannakopoulos and Elias Koutsoupias. Selling two goods optimally. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, pages 650–662, 2015. [1.3.1](#)
- [GK16] Kira Goldner and Anna Karlin. A prior-independent revenue-maximizing auction for multiple additive bidders. In *The 12th Conference on Web and Internet Economics (WINE)*, 2016. [1.3.3](#)
- [GS99] Faruk Gul and Ennio Stacchetti. Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory*, 87(1):95–124, July 1999. [7](#), [47](#)
- [HH15] Nima Haghpanah and Jason Hartline. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, 2015. [1.3.1](#), [1.3.2](#)
- [HKM11] Jason Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility via matchings. In *the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011. [2](#)
- [HL10] Jason Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 2010. [2](#)
- [HN12] Sergiu Hart and Noam Nisan. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012. [\(document\)](#), [1](#), [1.1](#), [1.3.1](#), [1.3.3](#), [2](#)
- [HN13] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*, 2013. [1.1](#)
- [HR09] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234, 2009. [1.1](#)
- [HR12] Sergiu Hart and Philip J. Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. *Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem*, 2012. [1.1](#)
- [KW12] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012. [1](#), [1.1](#), [1.3.3](#), [7](#)
- [LP16] Siqi Liu and Christos-Alexandros Psomas. On bulow-klemperer type theorems for dynamic auctions. *Manuscript*, 2016. [1.3.4](#), [9](#)
- [LR98] Jean-Jacques Laffont and Jacques Robert. Optimal auction with financially constrained buyers, 1998. [1.3.2](#)

- [LY13] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013. [1](#), [1.1](#), [1.3.1](#), [1.3.3](#), [2](#)
- [MSL15] Will Ma and David Simchi-Levi. Reaping the benefits of bundling under high production costs. *Manuscript*, 2015. [24](#)
- [MV04] Alexey Malakhov and Rakesh V Vohra. Single and multidimensional optimal auctions - a network approach. *Discussion Paper Series 1397, The Center for Mathematical Studies in Economics and Management Science*, 2004. [1.3.2](#), [18](#)
- [Mye81] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981. [1.2](#), [1.3.2](#), [2](#), [4](#), [13](#), [14](#), [4.2](#), [4.2](#), [27](#)
- [Mye97] Roger Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997. [1.3.2](#)
- [Pav11] Gregory Pavlov. Optimal mechanism for selling two goods. *The B.E. Journal of Theoretical Economics*, 11(3), 2011. [1.1](#)
- [RC98] Jean-Charles Rochet and Philippe Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, July 1998. [1.3.2](#)
- [Ron01] Amir Ronen. On approximating optimal auctions. In *the Third ACM Conference on Electronic Commerce (EC)*, 2001. [7](#)
- [RTCY12] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Supply-limiting mechanisms. In *13th ACM Conference on Electronic Commerce (EC)*, 2012. [1.3.3](#)
- [RW15] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 377–394, 2015. [1.1](#), [1.3.3](#), [2](#), [1](#)
- [Tha04] John Thanassoulis. Hagglng over substitutes. *Journal of Economic Theory*, 117(2):217–245, August 2004. [1.1](#)
- [Voh11] Rakesh V Vohra. *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press, 2011. [1.3.2](#)
- [Voh12] Rakesh Vohra. Optimization and mechanism design. *Mathematical Programming*, pages 283–303, 2012. [1.3.2](#)
- [Yao15] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *SODA*, 2015. [\(document\)](#), [1](#), [1.1](#), [1.3.1](#), [1.3.3](#), [2](#), [7](#), [27](#)
- [Yao17] Andrew Chi-Chih Yao. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *The 18th ACM Conference on Economics and Computation (EC)*, 2017. [1.1](#)