A Duality Based Unified Approach to Bayesian Mechanism Design

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

Nikhil R. Devanur[†]
Microsoft Resarch, USA
nikdev@microsoft.com

S. Matthew Weinberg[‡]
Princeton University, USA
sethmw@cs.princeton.edu.

ABSTRACT

We provide a unified view of many recent developments in Bayesian mechanism design, including the black-box reductions of Cai et. al. [6, 7, 8, 9, 19], simple auctions for additive buyers [25, 32, 1, 38], and posted-price mechanisms for unit-demand buyers [11, 12, 13]. 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 VCG auction with per-bidder entry fees achieves a constant-factor of the optimal Bayesian IC revenue whenever buyers are unit-demand or additive, unifying previous breakthroughs of Chawla et. al. and Yao, and improving both approximation ratios (from 33.75 to 24 and 69 to 8). Finally, we show that this view also leads to improved structural characterizations in the Cai et. al. framework.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: General

General Terms

Theory, Economics, Algorithms

Keywords

Revenue, Simple and Approximately Optimal Auctions

[*]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 Microsoft Research Fellow at the Simons Institute for the Theory of Computing.

STOC '16, June 18 - 21, 2016, Cambridge, MA, USA

ISBN 978-1-4503-4132-5/16/06…$15.00

DOI: http://dx.doi.org/10.1145/2897518.2897645

1. INTRODUCTION

The past several years have seen a tremendous advance in the field of Bayesian Mechanism Design, based on ideas and concepts rooted in Theoretical Computer Science. For instance, due to a line of work initiated by Chawla et. al., we now know that posted-price mechanisms are approximately optimal with respect to the optimal Bayesian Incentive Compatible[1] (BIC) mechanism whenever buyers are unit-demand,[2] and values are independent [11, 12, 13, 30]. Due to a line of work initiated by Hart and Nisan [25], we now know that either running Myerson's auction separately for each item or running the VCG mechanism with a per-bidder entry fee[3] is approximately optimal with respect to the optimal BIC mechanism whenever buyers are additive, and values are independent [32, 1, 38]. Due to a line of work initiated by Cai et. al., we now know that optimal mechanisms are distributions over virtual welfare maximizers, and have computationally efficient algorithms to find them in quite general settings [6, 7, 8, 9, 4, 19, 17]. The main contribution of this work is a unified approach to all three of these 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 [11, 12, 13, 30, 25, 10, 32, 1] as dual solutions. Surprisingly, we learn that *essentially the same dual solution* yields all the key benchmarks in these works. This inspires us to use it to design the first non-trivial benchmark with respect to the optimal BIC revenue in settings considered in [12, 38], which we then analyze to achieve better approximation factors in both cases.

1.1 Simple vs. Optimal Auction Design

[1]A 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 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.*

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

[3]By 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).

It is well-known by now that the optimal auction suffers many properties that are undesirable in practice, including randomization, non-monotonicity, and others [26, 27, 5, 14, 15]. To cope with this, much recent work in multi-dimensional mechanism design has turned to designing simple mechanisms that are *approximately* optimal. 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. for unit-demand buyers, and Hart and Nisan 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 mechanism (i.e. a mechanism that visits each buyer one at a time and posts a price for each item) obtains a constant factor of the optimal BIC revenue [11, 12, 13, 30]. 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 the better of running Myerson's optimal 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 [25, 32, 1, 38]. 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 [38] from 69 to 8. For unit-demand buyers, we improve the approximation ratio provided by Chawla et. al. [12] from 33.75 to 24.

In addition to these concrete results, we believe our work makes the following conceptual contributions as well. First, while the single-buyer core-tail decomposition techniques (first introduced by Li and Yao [32]) are now becoming standard [32, 1, 36, 2], they do not generalize naturally to multiple buyers. Yao [38] 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. Indeed, the core-tail decomposition which required substantial work previously is obtained for free: it is as simple as breaking a summation into two parts. 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 [12, 1, 38] 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.

We view our major contribution as providing a duality based unified framework for designing simple and approximately optimal auctions. As an application, we provide a simpler and tighter analysis for both additive and unit-demand bidders. In particular, the fact that we achieve both results from the same dual solution provides strong evidence that even this particular dual solution (or at least the intuition behind it) is worthy of deeper study.

1.2 General Bayesian 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. 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. Another aspect is structural: we now know that in all settings that fit into this framework, the optimal mechanism is a distribution over *virtual objective optimizers*. A mechanism is a virtual objective optimizer if it pointwise maximizes the sum of the original objective and the *virtual welfare*. The virtual welfare is given by a *virtual valuation/transformation*, which is a mapping from valuations to linear combinations of valuations.

Our contribution to this line of work is to improve the existing structural characterization. Previously, these virtual transformations were thought to be randomized and arbitrary, having no clear connection to the objective at hand. Our duality theory can say much more about what these virtual transformations might look like: 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 valuations of this agent,[4] 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 single, deterministic virtual valuation function such that:

1. This virtual valuation function can be found computationally efficiently.
2. In the special case of revenue, 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 the original objective + virtual welfare pointwise.[5]

[4] Both the CDW framework and our duality theory only apply directly if there are finitely many possible types for each agent.
[5] This could be *randomized*; there is always a deterministic

1.3 Other Related Work

Recently, strong duality frameworks for a single additive buyer were developed in [14, 16, 21, 20, 22]. 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. More recent work of Hartline and Haghpanah [24] can also be interpreted as providing an alternative "path-finding" duality framework. 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. 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 [14], posting a uniform item pricing [24], or even employing a randomized mechanism is optimal [22] when selling to a single additive or unit-demand buyer. However, none of these frameworks currently apply 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 are just tuples of duals for single buyers. In addition, we are able to re-derive and extend the breakthrough results of [11, 12, 13, 25, 32, 1, 38] *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.[6] However, we have demonstrated that there is at least one important domain (simple, approximately optimal mechanisms) where our theory seems to be more applicable.

Organization. We provide preliminaries and notation below. In Section 3, we present our duality theory for revenue maximization, and in Section 4 we present a canonical dual solution that proves useful in different settings. As a warm-up, we show in Section 5 how to analyze this dual solution when there is just a single buyer. In Section 6, we provide the multi-bidder analysis, which is more technical. Due to space limitations, our extension of the CDW framework in settings beyond revenue can be found in the full version.

2. PRELIMINARIES

Optimal Auction Design. For this version of the paper, we restrict our attention to *revenue maximization* in the following setting (the full version contains our extension of the CDW framework in more general settings): there is one copy of each of m heterogenous goods for sale to n buyers. The buyers are either all additive or all unit-demand, with buyer i having value t_{ij} for item j. We use $t_i = (t_{i1}, \ldots, t_{im})$ 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} ($f_i, f_{i,-j}, f_{-i}$, etc.) the densities of these finite-support distributions. The *optimal auction* optimizes expected revenue over all BIC 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. We define $\mathcal{F}$ to be a set system over $[n] \times [m]$ that describes all feasible allocations.[7]

Reduced Forms. The reduced form of an auction stores for all bidders i, items j, and types t_i, what is the probability that agent i will receive item j when reporting t_i to the mechanism (over the randomness in the mechanism and randomness in other agents' 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 $t_i \cdot \pi_i(t_i')$. We say that a reduced form is *feasible* if there exists some feasible mechanism (that 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 [6], for instance) that $P(\mathcal{F}, D)$ is closed and convex.

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 no items. 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 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. [11]. For any multi-dimensional instance D we can imagine splitting bidder i into m different

maximizer but in cases where the optimal mechanism is randomized, the objective plus virtual welfare are such that there are numerous maximizers, and the optimal mechanism randomly selects one.

[6]Our theory can still handle continuous distributions arbitrarily well. See Section 2.

[7]When 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.

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 [34] in the copies setting induced by D.[8]

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 [11, 12, 13, 25, 32, 1, 38, 36, 2]). 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. The following theorem is shown in [36], drawing from prior works [29, 28, 3, 18].

THEOREM 1. *[36, 18] 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 $\epsilon > 0$, $\text{REV}(D^+) \geq (1-\epsilon)(\text{REV}(D) - \text{VAL}(\delta))$, where $\text{VAL}(\delta)$ denotes the 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 $t_i^\epsilon(x) = \min\{t_i^\epsilon(x), 1/\epsilon\}$. It is easy to see that as $\epsilon \to 0$, $\text{REV}(D^\epsilon) \to \text{REV}(D)$. 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 $t_i^{+,\epsilon}(x) = \epsilon \cdot \lceil t_i(x)/\epsilon \rceil$. Similarly define $D_i^{-,\epsilon}$ with the ceiling -1 instead of the ceiling. 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 $\delta(\cdot)$ such that $\delta(x) \leq \epsilon$ for all x. So applying Theorem 1, we see that for any desired ϵ, we have $\text{REV}(D) \in [(1-\epsilon)(\text{REV}(D^{-,\epsilon} - n\epsilon)), (1+\epsilon)(\text{REV}(D^{+,\epsilon} + n\epsilon))]$. Finally, we just observe that $\text{REV}(D^{+,\epsilon}) = \text{REV}(D^{-,\epsilon}) + n\epsilon$, as every buyer values every outcome at exactly ϵ more in $D^{+,\epsilon}$ versus $D^{-,\epsilon}$. So as $\epsilon \to 0$, the revenues are the same, and both approach $\text{REV}(D)$. Note that both $D^{+,\epsilon}$ and $D^{-,\epsilon}$ have finite support.

3. OUR DUALITY THEORY

We begin by writing the LP for revenue maximization (Figure 1). For ease of notation, assume that there is a special type $\varnothing$ to represent the option of not participating in the auction. That means $\pi_i(\varnothing) = \mathbf{0}$ and $p_i(\varnothing) = 0$. Now a Bayesian IR (BIR) constraint is simply another BIC constraint: for any type t_i, bidder i will not want to lie to type $\varnothing$. We let $T_i^+ = T_i \cup \{\varnothing\}$. To proceed, we'll 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). Lagrangian multipliers have been used for mechanism design before [33, 31, 37, 4], however, our results are the first to obtain useful approximation benchmarks from this approach.

[8]Note that when buyers are additive that $\text{OPT}^{\text{COPIES}}$ is exactly the revenue of selling items separately using Myerson's optimal auction in the original setting.

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:

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

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

DEFINITION 1. *Let $\mathcal{L}(\lambda, \pi, p)$ be a the partial Lagrangian defined as follows:*

$$\begin{aligned}
&\mathcal{L}(\lambda, \pi, p)\\
&= \sum_{i=1}^n \Big(\sum_{t_i \in T_i} f_i(t_i) \cdot p_i(t_i) + \sum_{t_i \in T_i} \sum_{t_i' \in T_i^+} \lambda_i(t_i, t_i')\\
&\qquad \cdot \Big(t_i \cdot \big(\pi(t_i) - \pi(t_i')\big) - \big(p_i(t_i) - p_i(t_i')\big) \Big) \Big) \quad (1)\\
&= \sum_{i=1}^n \sum_{t_i \in T_i} p_i(t_i) \big(f_i(t_i) + \sum_{t_i' \in T_i} \lambda_i(t_i', t_i) - \sum_{t_i' \in T_i^+} \lambda_i(t_i, t_i') \big)\\
&+ \sum_{i=1}^n \sum_{t_i \in T_i} \pi_i(t_i) \big(\sum_{t_i' \in T_i^+} t_i \cdot \lambda_i(t_i, t_i') - \sum_{t_i' \in T_i} t_i' \cdot \lambda_i(t_i', t_i) \big)\\
&\qquad (\pi_i(\varnothing) = \mathbf{0},\ p_i(\varnothing) = 0) \quad (2)
\end{aligned}$$

3.1 Useful Properties of the Dual Problem

DEFINITION 2 (USEFUL DUAL VARIABLES). *A feasible dual solution λ is **useful** if $\max_{\pi \in P(\mathcal{F}, D), p} \mathcal{L}(\lambda, \pi, p) < \infty$.*

LEMMA 1 (USEFUL DUAL VARIABLES). *A dual solution λ is useful iff for each bidder i, λ_i forms a valid flow, i.e., iff the following satisfies flow conservation at all nodes except the source and the sink:*

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

Figure 2: Partial Lagrangian of the Revenue Maximization LP.

- *Nodes: A super source s and a super sink $\varnothing$, 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).*

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

$$f_i(t_i) + \sum_{t_i' \in T_i} \lambda(t_i', t_i) - \sum_{t_i' \in T_i^+} \lambda(t_i, t_i') \neq 0,$$

then since $p_i(t_i)$ is 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

$$f_i(t_i) + \sum_{t_i' \in T_i} \lambda(t_i', t_i) - \sum_{t_i' \in T_i^+} \lambda(t_i, t_i') = 0$$

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. □

DEFINITION 3 (VIRTUAL VALUE FUNCTION). *For each λ, we define a corresponding virtual value function $\Phi(\cdot)$, such that for every bidder i, every type $t_i \in T_i$, $\Phi_i(t_i) = t_i - \frac{1}{f_i(t_i)} \sum_{t_i' \in T_i} \lambda_i(t_i', t_i)(t_i' - t_i)$.*

THEOREM 2 (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:*

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

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

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

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 1), and replace the terms $\sum_{t_i' \in T_i^+} \lambda(t_i, t_i')$ with $f_i(t_i) + \sum_{t_i' \in T_i} \lambda(t_i', t_i)$. 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 is equal to $\sum_{i=1}^n \sum_{t_i \in T_i} f_i(t_i) \cdot \pi_i(t_i) \cdot \Phi_i(t_i)$, 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, $\left(\pi(t_i) - \pi(t_i')\right) - \left(p_i(t_i) - p_i(t_i')\right) \geq 0$ 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.

When λ^* is the optimal dual soluation, by strong duality 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^*. □

4. CANONICAL FLOW AND VIRTUAL VALUATION FUNCTION

In this section, we present a canonical way to set the Lagrangian multipliers/flow that induces our benchmarks. We use $P_{ij}(t_{-i})$ to denote the price that bidder i could pay to receive exactly item j in the VCG mechanism against bidders with types t_{-i}.[9] 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 $\operatorname{argmax}_k \{t_{ik} - P_{ik}(t_{-i})\}$. This partitions the types into subsets based on which item has the largest surplus (value minus price), and we break ties lexicographically.

For any bidder i and any type profile t_{-i} of everyone else, we define $\lambda_i^{(t_{-i})}$ to be the following flow.

1. For every type t_i in region $R_0^{(t_{-i})}$, the flow goes directly to $\varnothing$ (the super sink).
2. $\forall j > 0$, any flow entering $R_j^{(t_{-i})}$ is from s (the super source) and any flow leaving $R_j^{(t_{-i})}$ is to $\varnothing$.
3. $\forall t_i$ and t_i' in $R_j^{(t_{-i})}$ $(j > 0)$, $\lambda_i(t_i, t_i') > 0$ only if t_i and t_i' only differs on the j-th coordinate.

We will now spend the majority of this section building our canonical flow and proving that it achieves certain desirable properties. We begin by establishing some nice properties of $\Phi_i^{(t_{-i})}(\cdot)$ of any flow $\lambda_i^{(t_{-i})}$ constructed according to the above partial description.

CLAIM 1. *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 $k \neq j$.*

[9] Note 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 $\Phi_i^{(t-i)}(\cdot)$, $\Phi_{ik}^{(t-i)}(t_i) = t_i - \frac{1}{f_i(t_i)} \sum_{t_i'} \lambda_i^{(t-i)}(t_i', t_i)(t_{ik}' - t_{ik})$. 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 $\Phi_{ik}^{(t-i)}(t_i) = t_i$. $\square$

Next, we study $\Phi_{ij}^{(t-i)}(t_i)$ for coordinate j when t_i is in $R_j^{(t-i)}$. This turns out the to be closely related to the ("ironed") virtual value function defined by Myerson [34] for a single dimensional distributions. 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.

CLAIM 2. *For any type $t_i \in R_j^{(t-i)}$, if we only allow flow from type t_i' to t_i, where $t_{ik} = t_{ik}'$ for all $k \neq j$ and t_{ij} is the successor of t_{ij}' (the largest value smaller than t_{ij}' in the support of D_{ij}), then $\Phi_{ij}^{(t-i)}(t_i) = \varphi_{ij}(t_{ij}) = t_{ij} - \frac{(t_{ij}' - t_{ij}) \cdot \Pr_{t \sim D_{ij}}[t > t_{ij}]}{f_{ij}(t_{ij})}$.*

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

$$\prod_{k \neq j} f_{ik}(t_{ik}) \cdot \sum_{v \in T_{ij}: v > t_{ij}} f_{ij}(v)$$

$$= \Pi_{k \neq j} f_{ik}(t_{ik}) \cdot \Pr_{t \sim D_{ij}}[t > t_{ij}].$$

This is because each type of the form $(v, t_{i,-j})$ with $v > t_{ij}$ is also in $R_j^{(t-i)}$. So all flow that enters these types will be passed down to t_i (and possibly further, before going to the sink), and the total amount of flow entering all of these types from the source is exactly $\Pi_{k \neq j} f_{ik}(t_{ik}) \cdot \sum_{v \in T_{ij}: v > t_{ij}} f_{ij}(v)$. Therefore, $\Phi_{ij}^{(t-i)}(t_i) = \varphi_{ij}(t_{ij})$. $\square$

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 like in Myerson's work. Indeed, we use the same procedure: first convexify the revenue curve, then take the derivates of the convexified revenue curve as the "ironed" virtual values. To convexify the revenue curve, we only need to add loops to the flow we described in Claim 2. The next Lemma states that there exists a flow that performs this procedure and the resulting virtual value function $\Phi_{ij}^{(t-i)}(t_i)$ is upper bounded by the Myerson's ironed virtual value function $\tilde{\varphi}_{ij}(t_{ij})$ if $t_i \in R_j^{(t-i)}$.

LEMMA 2 (IRONING). *For any i, t_{-i}, there exists a flow $\lambda_i^{(t-i)}$ such that the corresponding $\Phi_{ij}^{(t-i)}(t_i)$ satisfies: for any $j > 0$ any $t_i \in R_j^{(t-i)}$, $\Phi_{ij}^{(t-i)}(t_i) \leq \tilde{\varphi}_{ij}(t_{ij})$.*

PROOF. First, we show how to modify a flow to fix non-monotonicities in $\Phi_{ij}^{(t-i)}(\cdot)$. Then we show how to use this procedure to iron.

If we have two types, t_i and t_i' such that $\Phi_{ij}^{(t-i)}(t_i) > \Phi_{ij}^{(t-i)}(t_i')$, but $t_{ij} < t_{ij}'$ (and $t_{i,-j} = t_{i,-j}'$), let's consider

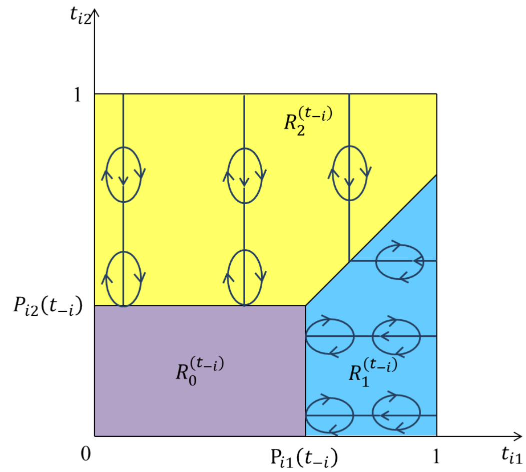

Figure 3: An example of $\lambda_i^{(t-i)}$ for two items.

adding a cycle between t_i and t_i' with weight w. Specifically, increase both $\lambda_i^{(t-i)}(t_i', t_i)$ and $\lambda_i^{(t-i)}(t_i, t_i')$ by w. What affect does this have on $\Phi_i^{(t-i)}(\cdot)$? First, it's clear that this is still a valid flow, as we've only added a cycle. Second, it's clear that we don't change $\Phi_i^{(t-i)}(t_i^*)$ at all, for any $t_i^* \notin \{t_i, t_i'\}$. Next, we see that we don't change $\Phi_{ik}^{(t-i)}(t_i)$ or $\Phi_{ik}^{(t-i)}(t_i')$ for any $k \neq j$. Finally, we see that $\Phi_{ij}^{(t-i)}(t_i)$ decreases by exactly $w(t_{ij}' - t_{ij})/f_i(t_i)$ and $\Phi_{ij}^{(t-i)}(t_i')$ increases by exactly $w(t_{ij}' - t_{ij})/f_i(t_i')$. So by setting w appropriately, we see that we can update $\lambda_i^{(t-i)}$ so that $\Phi_{ij}^{(t-i)}(t_i) = \Phi_{ij}^{(t-i)}(t_i')$, but without changing the average virtual value for item j among these two types, nor their virtual value for any other item, nor any other type's virtual values for any item.

Now, observe that Myerson ironing can always be implemented in the following way: pick a disjoint set of intervals $I_1, \ldots, I_k$ that we wish to iron. This is decided by the convex hull of the revenue curve for the corresponding distribution. In particular, inside each interval I_ℓ, the average virtual values of the highest N (for any N) types is no larger than the average virtual values in the whole interval. Iteratively, find two adjacent types $t_i, t_i' \in I_\ell$ (for any ℓ) such that $\Phi_{ij}^{(t-i)}(t_i) > \Phi_{ij}^{(t-i)}(t_i')$, but $t_{ij} < t_{ij}'$ (and $t_{i,-j} = t_{i,-j}'$). Then update each type's ironed virtual value to the average of their previous (ironed) virtual values. The end result will be that all types in I_j will have the same ironed virtual value, which is equal to the average virtual value on that interval. We have shown that we can certainly implement this procedure via the adjustments above.

The only catch between exact Myerson ironing and what we wish to do in our flow is that we are not ironing the entire support of D_{ij}, but only the portion above some cutoff, C. The only effect this has is that it possibly truncates some interval I_ℓ at C instead of its true (lower) lower bound. By the nature of ironing, we know that this necessarily implies that the average virtual value on $I_k \cap [0, C)$ is larger than the average virtual value on $I_k \cap [C, \infty)$ (recall: the ironing procedure is *only*

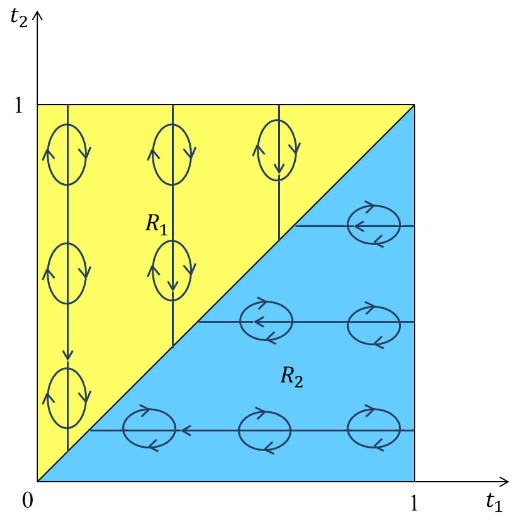

Figure 4: An example of λ (with ironing) for a single bidder.

to fix non-monotonicities. If the average virtual value on the lower interval were to be less than the average virtual value on the higher interval, we *wouldn't* iron them to the same ironed interval). So the virtual values we are left with after our procedure are certainly smaller than the true ironed virtual values, completing the proof. □

LEMMA 3. *There exists a flow $\lambda_i^{(t_{-i})}$ such that $\Phi_{ij}^{(t_{-i})}(t_i)$ satisfies the following properties:*

- *For any $j > 0$, $t_i \in R_j^{(t_{-i})}$, $\Phi_{ij}^{(t_{-i})}(t_i) \le \tilde{\varphi}_{ij}(t_{ij})$, 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. Combine Lemma 2 and Claim 1. □

Lemma 3 isn't exactly the flow we want to use: note that we've defined several flows that depend on t_{-i}, but we only get to select *one* flow for bidder i, and it doesn't get to change depending on t_{-i}. Below we define a single flow essentially by averaging across all t_{-i} according to the distributions.

DEFINITION 4 (FLOW). *The 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 $\Phi_i(\cdot) = \sum_{t_{-i}\in T_{-i}} f_{-i}(t_{-i})\Phi_i^{(t_{-i})}(\cdot)$.*

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 virtual value, and keep the **virtual value** of all the other items their values. 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 2.

LEMMA 4. *For all i, j, t_i, $\Phi_{ij}(t_i) \le t_{ij}\cdot \Pr_{v_{-i}\sim D_{-i}}[t_i \notin R_j^{(v_{-i})}] + \tilde{\varphi}_{ij}(t_{ij})\cdot \Pr_{v_{-i}\sim D_{-i}}[t_i \in R_j^{(v_{-i})}]$.*

THEOREM 3. *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,*

$$\begin{aligned}
&\sum_i \sum_{t_i\in T_i} f_i(t_i)\cdot p_i(t_i)\\
\le&\sum_i \sum_{t_i\in T_i} \sum_j f_i(t_i)\cdot \pi_{ij}(t_i)\cdot \Phi_{ij}(t_i)\\
\le&\sum_i \sum_{t_i\in T_i} \sum_j f_i(t_i)\cdot \pi_{ij}(t_i)\cdot\\
&\Big(t_{ij}\cdot \Pr_{v_{-i}\sim D_{-i}}[t_i\notin R_j^{(v_{-i})}]\\
&+\tilde{\varphi}_{ij}(t_{ij})\cdot \Pr_{v_{-i}\sim D_{-i}}[t_i\in R_j^{(v_{-i})}]\Big)
\end{aligned}\tag{3}$$

5. WARM UP: SINGLE BIDDER

As a warm up, we start with the single bidder case. 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 $t_j \ge t_k$ for all k (see Figure 4 for an example). This simplifies Expression (3) to

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

We bound SINGLE below, and NON-FAVORITE differently for unit-demand and additive valuations.

LEMMA 5. *For any feasible $\pi(\cdot)$, SINGLE $\le$ 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 $\sum_{t_{-j}} f_{-j}(t_{-j}) \cdot \pi_j(t_j, t_{-j}) \cdot \mathbb{I}[t \in R_j]$ 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. □

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. [11, 12, 13] to approximate each term.

LEMMA 6. *When the types are unit-demand, for any feasible $\pi(\cdot)$, 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 $\{t_1, \cdots, t_m\}$. 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 $t \in R_j$, then there exists some $k \neq j$ such that $t_k \geq t_j$, so $t_j \cdot \mathbb{I}[t \in R_j] \leq S(t)$ for all j. Therefore, $\sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot t_j \cdot \mathbb{I}[t \notin R_j] \leq \sum_{t \in T} \sum_j f(t) \cdot \pi_j(t) \cdot S(t) \leq \sum_{t \in T} f(t) \cdot S(t)$. The last inequality is because the bidder is unit demand, so $\sum_j \pi_j(t) \leq 1$. □

Combining Lemma 5 and Lemma 6, we recover the result of Chawla et al. [13]:

THEOREM 4. *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. Let $r = \text{SREV}$. 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. [1]. However, it is worth noting that the "core-tail decomposition" presented here is perhaps more transparent: we are simply splitting a sum into two parts depending on whether the buyer's value for item j is larger than some threshold.

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

LEMMA 7. TAIL $\leq r$.

PROOF. By the definition of R_j, for any given t_j,

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

It is clear that by setting price t_j on each item separately, we can make revenue at least $t_j \cdot \Pr_{t_{-j} \sim D_{-j}}[\exists k \neq j, t_k \geq t_j]$. 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 this term is upper bounded by r. Thus, TAIL $\leq r \cdot \sum_j \sum_{t_j > r} f_j(t_j) = \sum_j r \cdot \Pr_{t_j \sim D_j}[t_j > r] =$ the revenue of selling each item separately at price r, which is also $\leq r$. □

LEMMA 8. *If we sell the grand bundle at price 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 CORE $\leq 2\text{BREV} + 2\text{SREV}$.*

PROOF. We will first need a technical lemma (also used in [1], but proved here for completeness).

LEMMA 9. *Let x be a positive single dimensional random variable drawn from F of finite support,[10] 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 $x_s = x \cdot \mathbb{I}[x \leq s]$ is upper bounded by $2\mathcal{B} \cdot s$.*

PROOF. Let $\{a_1, \ldots, a_\ell\}$ 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 $a_k \cdot \Pr_{x \sim F}[x \geq a_k] \leq \mathcal{B}$. □

Now with Lemma 9, for each j define a new random variable c_j based on the following procedure: draw a sample r_j

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

from D_j, if r_j lies in $[0, r]$, then $c_j = r_j$, otherwise $c_j = 0$. Let $c = \sum_j c_j$. It is not hard to see that we have $\mathbb{E}[c] = \sum_j \sum_{t_j \le r} f_j(t_j) \cdot t_j$. Now we are going to show that c concentrates because it has small variance. Since the c_j's are independent, $\text{Var}[c] = \sum_j \text{Var}[c_j] \le \sum_j \mathbb{E}[c_j^2]$. We will bound each $\mathbb{E}[c_j^2]$ separately. Let $r_j = \max_x\{x \cdot \Pr_{t_j \sim D_j}[t_j \ge x]\}$. By Lemma 9, we can upper bound $\mathbb{E}[c_j^2]$ by $2r_j \cdot r$. On the other hand, it is easy to see that $r = \sum_j r_j$, so $\text{Var}[c] \le 2r^2$. By the Chebyshev inequality,

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

Therefore,

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

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

THEOREM 5. *For a single additive bidder, the optimal revenue is* $\le 2\text{BRev} + 4\text{SRev}$.

PROOF. Combining Lemma 5, 7 and 8, 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 $2\text{BRev} + 4\text{SRev}$. □

6. MULTIPLE BIDDERS

In this section, we show how to use the upper bound in Theorem 3 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. Similar to the single bidder case, we first decompose the upper bound (Expression 3) 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 6. *For multiple unit-demand bidders, the optimal revenue is upper bounded by* $4\text{OPT}^{\text{Copies}}$.

THEOREM 7. *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 buyers are unit-demand [12, 30], and selling each item separately using Myerson's auction achieves revenue $\text{OPT}^{\text{Copies}}$ when buyers are additive. Therefore, the CHMS/KW [12, 30] posted-price mechanism achieves a 24-approximation to the optimal BIC mechanism (previously, it was known to be a 33.75-approximation), and Yao's approximation ratios [38] are improved from 69 to 8. Some parts of the following analysis draw inspiration from prior works of Chawla et. al. [12] and Yao [38], however, much of the analysis also represents new techniques. In particular, it is worth pointing out that our proof of Theorem 7 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.

$$\begin{aligned}
&\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \Big(t_{ij} \cdot \Pr_{v_{-i}\sim D_{-i}}[t_i \notin R_j^{(v_{-i})}] \\
&\qquad + \tilde{\varphi}_{ij}(t_{ij}) \cdot \Pr_{v_{-i}\sim D_{-i}}[t_i \in R_j^{(v_{-i})}]\Big) \\
\le &\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i}\in T_{-i}} t_{ij} f_{-i}(v_{-i}) \\
&\quad \cdot \mathbb{I}\Big[\big(\exists k \ne j,\ t_{ik} - P_{ik}(v_{-i}) \ge t_{ij} - P_{ij}(v_{-i})\big) \\
&\qquad \vee \big(t_{ij} < P_{ij}(v_{-i})\big)\Big] \\
+ &\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i)\pi_{ij}(t_i)\tilde{\varphi}_{ij}(t_{ij}) \Pr_{v_{-i}\sim D_{-i}}[t_i \in R_j^{(v_{-i})}] \\
= &\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i}\in T_{-i}} t_{ij} f_{-i}(v_{-i}) \\
&\quad \cdot \mathbb{I}\Big[\big(\exists k \ne j,\ t_{ik} - P_{ik}(v_{-i}) \ge t_{ij} - P_{ij}(v_{-i})\big) \\
&\qquad \wedge \big(t_{ij} \ge P_{ij}(v_{-i})\big)\Big] \quad (\text{NON-FAVORITE}) \\
&+ \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i}\in T_{-i}} t_{ij} \\
&\qquad \cdot f_{-i}(v_{-i}) \cdot \mathbb{I}[t_{ij} < P_{ij}(v_{-i})] \quad (\text{UNDER}) \\
&+ \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \tilde{\varphi}_{ij}(t_{ij}) \\
&\qquad \cdot \Pr_{v_{-i}\sim D_{-i}}[t_i \in R_j^{(v_{-i})}] \quad (\text{SINGLE})
\end{aligned}$$

$$\begin{aligned}
&\text{NON-FAVORITE} \\
&\le \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \\
&\quad \cdot \sum_{v_{-i}\in T_{-i}} P_{ij}(v_{-i}) f_{-i}(v_{-i}) \mathbb{I}[t_{ij} \ge P_{ij}(v_{-i})] \quad (\text{OVER}) \\
&+ \sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot \sum_{v_{-i}\in T_{-i}} (t_{ij} - P_{ij}(v_{-i})) \cdot \\
&\quad f_{-i}(v_{-i}) \cdot \mathbb{I}\Big[\big(\exists k \ne j,\ t_{ik} - P_{ik}(v_{-i}) \ge t_{ij} - P_{ij}(v_{-i})\big) \\
&\qquad \wedge \big(t_{ij} \ge P_{ij}(v_{-i})\big)\Big] \quad (\text{SURPLUS})
\end{aligned}$$

Analyzing SURPLUS for Unit-demand Bidders: The proof of this lemma is similar in spirit to Lemma 6.

LEMMA 10. *When the types are unit-demand, for any feasible* $\pi(\cdot)$, $\text{SURPLUS} \le \text{OPT}^{\text{Copies}}$.

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 $\{t_{i1} - P_{i1}(v_{-i}), \cdots, t_{im} - P_{im}(v_{-i})\}$. 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 $S_i(t_i, v_{-i}) \cdot \mathbb{I}[\exists j',\ t_{ij'} - P_{ij'}(v_{-i}) \geq 0]$. So the total revenue of the VCG Copies mechanism is at least:

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

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

$$\begin{aligned}
&(t_{ij} - P_{ij}(v_{-i})) \\
&\quad \cdot \mathbb{I}\Big[\big(\exists k \neq j,\ t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i}) \geq 0\Big] \\
&\leq S_i(t_i, v_{-i}) \cdot \mathbb{I}[\exists j',\ t_{ij'} - P_{ij'}(v_{-i}) \geq 0] \qquad (4)
\end{aligned}$$

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.

$$\begin{aligned}
&\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \sum_{v_{-i} \in T_{-i}} (t_{ij} - P_{ij}(v_{-i})) \cdot \\
&f_{-i}(v_{-i}) \cdot \mathbb{I}\Big[\exists k \neq j,\ t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i}) \geq 0\Big] \\
&= \sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \sum_j \pi_{ij}(t_i) \cdot (t_{ij} - P_{ij}(v_{-i})) \\
&\quad \cdot \mathbb{I}\Big[\exists k \neq j,\ t_{ik} - P_{ik}(v_{-i}) \geq t_{ij} - P_{ij}(v_{-i}) \geq 0\Big] \\
&\leq \sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \sum_j \pi_{ij}(t_i) \cdot S_i(t_i, v_{-i}) \\
&\quad \cdot \mathbb{I}[\exists j',\ t_{ij'} - P_{ij'}(v_{-i}) \geq 0] \qquad \text{(Inequality (4))} \\
&\leq \sum_i \sum_{(t_i, v_{-i}) \in T_i} f(t_i, v_{-i}) \cdot S_i(t_i, v_{-i}) \\
&\quad \cdot \mathbb{I}[\exists j',\ t_{ij'} - P_{ij'}(v_{-i}) \geq 0] \quad \Big(\sum_j \pi_{ij}(t_i) \leq 1\ \forall i, t_i\Big)
\end{aligned}$$

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 $\text{OPT}^{\text{COPIES}}$. □

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 [35] 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 7 and 8),

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

LEMMA 11. TAIL $\leq r$.

PROOF. First, by union bound

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

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:

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

Using these two inequalities, we can upper bound TAIL:

$$\sum_i \sum_{v_{-i}\in T_{-i}} f_{-i}(v_{-i}) \sum_j \sum_{t_{ij}>P_{ij}(v_{-i})+r_i(v_{-i})} f_{ij}(t_{ij}) \cdot \sum_{k\neq j} r_{ik}(v_{-i})$$

$$\leq \sum_i \sum_{v_{-i}} f_{-i}(v_{-i}) \cdot \sum_j r_i(v_{-i}) \cdot \sum_{t_{ij}>P_{ij}(v_{-i})+r_i(v_{-i})} f_{ij}(t_{ij})$$

$$\leq \sum_i \sum_{v_{-i}} f_{-i}(v_{-i}) \sum_j r_{ij}(v_{-i}) \quad \text{(Definition of } r_{ij}(v_{-i}))$$

$$= r$$

□

LEMMA 12. 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

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

and

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

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

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

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 9 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,

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

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:

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

□

Analyzing SINGLE, OVER and UNDER: First we consider SINGLE, which is similar to Lemma 5.

LEMMA 13. *For any feasible* $\pi(\cdot)$, 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. □

Next, we move onto OVER. We begin with the following technical propositions:

PROPOSITION 1. *Let* $\pi(\cdot)$ *be any reduced form of a BIC mechanism in the original setting. Define*

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

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. □

PROPOSITION 2. *For any $v \in T$, any $\pi(\cdot)$ that is a reduced form of some BIC mechanism,*

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

PROOF. Recall from Proposition 1 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 $P_{ij}(v_{-i})\mathbb{I}[t_{ij} \geq P_{ij}(v_{-i})]$. That means the total revenue of the new mechanism is at least $\sum_i \sum_{t_i \in T_i} \sum_j f_i(t_i) \cdot \pi_{ij}(t_i) \cdot P_{ij}(v_{-i}) \cdot \mathbb{I}[t_{ij} \geq P_{ij}(v_{-i})]$, which is upper bounded by $\text{OPT}^{\text{COPIES}}$. □

LEMMA 14. OVER $\leq \text{OPT}^{\text{COPIES}}$.

PROOF. This can be proved by rewriting OVER and then applying Proposition 2.

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

□

When there is only one bidder, UNDER is always 0. Here, UNDER $\leq \text{OPT}^{\text{COPIES}}$. We apply Proposition 3 (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 3 for all profiles t with multipliers $f(t)$ gives UNDER$\leq \text{OPT}^{\text{COPIES}}$.

PROPOSITION 3. *Let $\{(i_j, j)\}_{j \in S \subseteq [m]}$ be a feasible allocation in the copies setting. For all choices $x_1, \ldots, x_m \geq 0$,* $\text{OPT}^{\text{COPIES}} \geq \sum_{v \in T} f(v) \cdot \sum_{j \in S} x_j \cdot \mathbb{I}[P_{i_j j}(v_{-i_j}) > x_j]$.

PROOF. Before beginning the proof of Proposition 3, we will need the following definition and theorem due to Gul and Stacchetti [23].

DEFINITION 5. *Let $W_T(S)$ be the maximum attainable welfare using only bidders in T and items in S.*

THEOREM 8. *([23]) If all bidders in T have gross substitute valuations, then $W_T(S)$ is submodular.*

Now with Theorem 8, 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 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 the allocation rule is monotone, thus this is a truthful mechanism. 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

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

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 $S' - \{j\}$:

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

Rearranging the terms, we get

$$\begin{aligned}
&v_{ij} \\
&\geq W_{[n]-\{i\}}([m] - S' + \{j\}) - W_{[n]-\{i\}}([m] - S') \\
&\geq W_{[n]-\{i\}}([m]) - W_{[n]-\{i\}}([m] - \{j\})) \quad (\textit{Theorem } 8) \\
&= P_{ij}(v_{-i}).
\end{aligned}$$

Now we still need to argue that whenever $P_{i_j j}(v_{-i}) \geq x_j$, 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,
$$v_{i_j j} \geq P_{i_j j}(v_{-i_j}) > x_j.$$
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 $W_{[n]-\{i_j\}}([m] - S') - W_{[n]-\{i_j\}}([m] - S' - \{j\}) \geq W_{[n]-\{i_j\}}([m]) - W_{[n]-\{i_j\}}([m] - \{j\}) = P_{i_j j}(v_{-i_j})$, and $P_{i_j j}(v_{-i_j}) > x_j$, that means (i) item j is awarded to some bidder $i \neq i_j$ in the welfare maximizing allocation, (ii) $v_{ij} > x_j$ because otherwise
$$W_{[n]-\{i_j\}}([m] - S') \leq W_{[n]-\{i_j\}}([m] - S' - \{j\}) + x_j,$$

contradiction.

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

LEMMA 15. UNDER $\leq$ OPT$^{\text{COPIES}}$.

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 3.

First, the inner sum

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

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 $x(\cdot)$:

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

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

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

and by Proposition 3, this is upper bounded by OPT$^{\text{COPIES}}$. □

Combining the above lemmas now yields our theorems:
Proof of Theorem 6: Combine Lemmas 10, 13, 14 and 15. □

Proof of Theorem 7: Combining Lemmas 11, 12, 13, 14 and 15, we get the optimal revenue is upper bounded by

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

Since OPT$^{\text{COPIES}}$= SMYERSON and SMYERSON $\geq r$ when bidders are additive; we proved our statement. □

7. ACKNOWLEDGEMENTS

We would like to thank Costis Daskalakis and Christos Papadimitriou for numerous helpful discussions during the preliminary stage of this work.

8. REFERENCES

[1] 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. (document), 1, 1.1, 1.3, 2, 5, 5

[2] 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, 2

[3] Xiaohui Bei and Zhiyi Huang. Bayesian incentive compatibility via fractional assignments. In *the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011. 2

[4] 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, 3

[5] 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

[6] 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. (document), 1, 2

[7] 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. (document), 1

[8] 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. (document), 1

[9] 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

[10] 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

[11] 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, 2, 5

[12] 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, 2, 5, 6

[13] 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. (document), 1, 1.1, 1.3, 2, 5, 5

[14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism Design via Optimal Transport. In the 14th Annual ACM Conference on Economics and Computation (EC), 2013. 1.1, 1.3

[15] 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

[16] 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

[17] 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

[18] 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

[19] 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. (document), 1

[20] Yiannis Giannakopoulos. A note on optimal auctions for two uniformly distributed items. *CoRR*, abs/1409.6925, 2014. 1.3

[21] 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

[22] 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

[23] Faruk Gul and Ennio Stacchetti. Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory*, 87(1):95–124, July 1999. 6, 8

[24] 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

[25] 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, 2

[26] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*, 2013. 1.1

[27] 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

[28] 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

[29] Jason Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 2010. 2

[30] Robert Kleinberg and S. Matthew Weinberg. Matroid Prophet Inequalities. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012. 1, 1.1, 6

[31] Jean-Jacques Laffont and Jacques Robert. Optimal auction with financially constrained buyers, 1998. 3

[32] 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. (document), 1, 1.1, 1.3, 2

[33] Roger Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997. 3

[34] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981. 2, 4

[35] Amir Ronen. On approximating optimal auctions. In *the Third ACM Conference on Electronic Commerce (EC)*, 2001. 6

[36] 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, 2, 1

[37] Rakesh Vohra. Optimization and mechanism design. *Mathematical Programming*, pages 283–303, 2012. 3

[38] 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, 2, 6