

## Prior-free Auctions for Budgeted Agents

NIKHIL R. DEVANUR, Microsoft Research, Redmond, WA.

BACH Q. HA, Northwestern University, Evanston, IL.

JASON D. HARTLINE, Northwestern University, Evanston, IL.

We consider prior-free auctions for revenue and welfare maximization when agents have a common budget. The abstract environments we consider are ones where there is a downward-closed and symmetric feasibility constraint on the probabilities of service of the agents. These environments include *position auctions* where slots with decreasing click-through rates are auctioned to advertisers. We generalize and characterize the envy-free benchmark from Hartline and Yan [2011] to settings with budgets and characterize the optimal envy-free outcomes for both welfare and revenue. We give prior-free mechanisms that approximate these benchmarks. A building block in our mechanism is a clinching auction for position auction environments. This auction is a generalization of the multi-unit clinching auction of Dobzinski et al. [2008] and a special case of the polyhedral clinching auction of Goel et al. [2012]. For welfare maximization, we show that this clinching auction is a good approximation to the envy-free optimal welfare for position auction environments. For profit maximization, we generalize the random sampling profit extraction auction from Fiat et al. [2002] for digital goods to give a 10.0-approximation to the envy-free optimal revenue in symmetric, downward-closed environments. Even without budgets this revenue maximization question is of interest and we obtain an improved approximation bound of 7.5 (from 30.4 by Ha and Hartline [2012]).

Categories and Subject Descriptors: F.0 [Theory of Computation]: General

General Terms: Economics, Theory, Algorithms

Additional Key Words and Phrases: prior-free, envy-free, budget, mechanism design

### 1. INTRODUCTION

Economic mechanisms that are less dependent on the assumptions of the environment are more likely to be relevant [Wilson 1987]. The area of *prior-free mechanism design* attempts to give mechanisms that guarantee a good approximation to the designer's objective without dependence on distributional assumptions on the agents' preferences. The main open questions in prior-free mechanism design center around the departure from the ideal single-dimensional and linear model of preferences, i.e., where an agent's utility is given by her value for service less the payment she is required to make. A paradigmatic example of a non-linear agent utility is one that is linear up to a given budget that restricts the agent's maximum possible payment. In this paper we consider the designer's objectives of revenue and welfare (separately) when agents have budgets, and we give simple prior-free auctions that approximate a natural prior-free benchmark (for each objective).

Mechanism design studies optimization on inputs that are the private information of strategic agents who may misreport their information if it benefits them. Agents will not misreport only if they have no incentive to, i.e., if their utilities are maximized by truthful reporting. The key challenge in designing mechanisms for strategic agents, then, is that incentive constraints bind across possible agent misreports. A mechanism must therefore

---

This work is supported in part by NSF CAREER Award CCF 0846113 and NSF CCF 0830773, and in part while the authors visiting Microsoft Research (Redmond, WA).

Permission to make digital or hardcopies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

EC'13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

trade-off performance on one input versus another. For general objectives, e.g., welfare with budgeted agents, profit, or makespan (for unrelated machines), there is no single mechanism that is simultaneously optimal on all inputs.<sup>1</sup> There are two approaches for addressing the non-point-wise-optimality of mechanisms. The Bayesian approach, which is standard in economics, assumes that the agents' preferences (inputs) are drawn from a known distribution and the performance of the mechanism across different inputs can be traded off so as to optimize its expected performance with respect to this given distribution. The Bayesian optimal mechanism, therefore, depends on the distribution. The prior-free approach, which is currently being developed in computer science, instead looks for a single mechanism that approximates an economically motivated prior-free benchmark in worst-case over all inputs.

The first step in developing prior-free mechanisms is to identify an appropriate prior-free benchmark. Hartline and Yan [2011] recently observed that a simple and intuitive prior-free benchmark can be defined based on a relaxation of the no-misreporting incentive constraint to a no-envy constraint. The advantage of the no-envy constraint is that it binds point-wise on each input instead of across all inputs like incentive constraints; therefore, there is always a point-wise optimal envy-free outcome. Furthermore, as Hartline and Yan [2011] pointed out, often this benchmark is an upper bound on the optimal performance on the Bayesian optimal mechanism for any distribution; in these cases approximating it point-wise gives a very strong performance guarantee. Our first contribution is a generalization of the revenue-optimal envy-free benchmark without budgets to the objectives of revenue and welfare with budgets.

A mechanism must optimize its objective subject to incentive constraints (discussed above), feasibility constraints (i.e., constraints on how agents can be served together), and budget constraints. It is most instructive to classify feasibility constraints in terms of the sophistication required of constrained optimization of a weighted sum of the set of agents served (or, for randomized environments, probabilities of service). An environment, like that of digital good auctions, may be *unconstrained*. An *ordinal environment*, like those of position auctions (as popularized by advertising on Internet search engines), is one where the optimal algorithm is greedy on agents ordered by weight. In a general *cardinal environment*, like those of single-minded combinatorial auctions, the weights of the agents are necessary for optimization. An environment is *symmetric* if the feasibility constraint respects all permutations of the agent identities. While feasibility constraints limit which agents are served, budget constraints limit the prices that agents pay.

Recent results of Dobzinski et al. [2008] and Goel et al. [2012] have shown that a generalization of the Ausubel [2004] *clinchng auction* is the only *Pareto optimal* mechanism in ordinal environments. At a high-level, the clinching auction is given by an ascending price at which each agent is allowed to claim any of the supply that would be left over if that agent were given the last choice. Pareto optimality is the condition that there is no other feasible outcome where some participant (including the designer) can be made strictly better off without making some participant strictly worse off. Pareto optimality is a condition not an objective which means that it is not clear what an approximation to Pareto optimality would mean. It is also not true that all reasonable auctions must satisfy Pareto optimality. The Bayesian welfare-optimal auction for budgeted agents is not generally Pareto optimal (see Sections 3 and 4); and moreover, there does not generally exist Pareto-optimal auctions for budgeted agents in cardinal environments [Goel et al. 2012].

---

<sup>1</sup>For these objectives, the designer's objective and the agents' objectives are fundamentally at odds and the incentive constraints of the agents do not permit the designer to obtain the same performance possible when the inputs are public. These objectives contrast starkly to the objective of welfare maximization without budgets where there is no conflict in the designer's objective and the agents' objectives and the Vickrey-Clarke-Groves (VCG) mechanism is point-wise optimal [Vickrey 1961; Clarke 1971; Groves 1973].

Our first goal, given the limits of Pareto optimality, is to identify a prior-free auction that approximates the envy-free optimal welfare when agents have budgets. The outcome of the clinching auction (for ordinal environments) is envy free; however, it is not the welfare-optimal envy-free outcome. Moreover, given distribution over agent values, the clinching auction is not the Bayesian optimal auction for welfare either. We give a simple closed form expression for the clinching auction in symmetric ordinal environments with a common budget and we show that it is a 2-approximation to the envy-free benchmark.<sup>2</sup>

Our second goal is to identify a prior-free auction that approximates the envy-free optimal revenue when agents have budgets. For revenue maximization without budgets Hartline and Yan [2011] and Ha and Hartline [2012] recently gave general approaches for approximating the optimal envy-free revenue. The former extends a standard random sampling approach (for digital good auctions) from Goldberg et al. [2001]; the latter extends an approach based on “consensus estimates” and “profit extraction” from Goldberg and Hartline [2003]. Our approach is based on an extension of the *random sampling profit extraction* auction from Fiat et al. [2002].<sup>3</sup> Not only is our mechanism the only one that is readily compatible with budget constraints, but also the approximation factors we obtain, relative to the revenue-optimal envy-free benchmark, are the best known. We show a 10.0-approximation to the envy-free optimal revenue in symmetric cardinal environments. Moreover, without budgets, our techniques give a 7.5-approximation which improves on the 30.4 approximation of Ha and Hartline [2012].

*Summary of Results.* Our main conceptual contribution is the adaptation of the prior-free mechanism design framework initiated by the literature on digital goods (i.e., unconstrained environments), e.g., Goldberg et al. [2001] to the structurally rich environments of Hartline and Yan [2011] (including symmetric ordinal environment and cardinal environments)<sup>4</sup> when agents’ preferences are non-linear as given by a common budget constraint. Our technical results are as follows:

- We give a characterization of the envy-free benchmark for welfare and revenue when agents have a common budget. This characterization is via an extension of the characterization of Bayesian optimal auctions for agents with budgets of Laffont and Robert [1996] to general distributions.<sup>5</sup>
- We give a closed-form characterization of the clinching auction of Goel et al. [2012] in symmetric ordinal environments with a common budget.
- We prove that the clinching auction is a 2-approximation to the envy-free optimal welfare in symmetric ordinal environments with a common budget.
- We extend the random sampling profit extraction auction from Fiat et al. [2002] to symmetric cardinal environments with a common budget. This generalization gives a 10.0-approximation to the envy-free benchmark. This is the first prior-free approximation of an economically well motivated benchmark for agents with budgets.

---

<sup>2</sup>While both Dobzinski et al. [2008] and Goel et al. [2012] allow agents to have distinct budgets, the envy-free benchmark is only economically well motivated in symmetric environments therefore a common budget is required. Our restriction to symmetric environments and in particular a common budget is reasonable as designing prior-free auctions for asymmetric environments is a challenge in itself even without budgets; for asymmetric environments only a few positive results are known, see, e.g., Balcan et al. [2008] and Leonardi and Roughgarden [2012], both of which are for unconstrained environments (e.g., digital goods).

<sup>3</sup>Therefore all of the leading approaches for digital good auctions extend to more general environments.

<sup>4</sup>Hartline and Yan [2011] refer to symmetric ordinal environments equivalently as *position environments* and *matroid permutation environments* and to symmetric cardinal environments as *downward-closed permutation environments*.

<sup>5</sup>The Laffont and Robert [1996] characterization holds for monotone hazard rate distributions with increasing density [Pai and Vohra 2008]; under such assumptions, many of the novel properties of optimal auctions with budgets do not arise.

- The clinching auction is not well defined in cardinal environments; the above auction converts the symmetric cardinal environment to a symmetric ordinal environment where the clinching auction can be run and its objective is close to optimal (for the original cardinal environment).

These results are best contrasted with the literature on prior-free revenue maximization without budgets (which is equivalent to a common budget of infinity). Our main result for agents without budgets is:

- Our random sampling profit extraction auction for agents without budgets is a 7.5 to the envy-free revenue benchmark; this improves on the best known auction and bound of Ha and Hartline [2012] of 30.4.

*Related Work.* The theory of Bayesian optimal auctions for welfare or revenue when agents have budgets (a form of non-linear utility) is more complex than that of revenue when agents have linear utility. In the latter, e.g., the revenue-optimal mechanism is given by optimizing *virtual values* which are given by a simple distribution-dependent function of agents' values [Myerson 1981]. In the former, under some restrictive distributional assumptions, a similar Lagrangian virtual value approach gives the optimal mechanism (subject to careful choice of the Lagrangian variable, see Laffont and Robert [1996]).

Hartline and Yan [2011] defined the envy-free benchmark as a relaxation of the Bayesian optimal auction that can be optimized point-wise. Our characterization of the envy-free benchmark for welfare and revenue for agents with budgets combines and extends the results of Laffont and Robert [1996] and Hartline and Yan [2011].

There are three main techniques for designing revenue maximizing prior-free auctions for digital goods (i.e., where there is no feasibility constraint). The *random sampling optimal price* auction was defined by Goldberg et al. [2001]. The *consensus estimate profit extraction* auction was defined by Goldberg and Hartline [2003]. The *random sampling profit extraction* auction was defined by Fiat et al. [2002]. The first two approaches were generalized to symmetric cardinal environments by Hartline and Yan [2011] and Ha and Hartline [2012], respectively. We generalize the third approach to these environments. Our generalization gives the best known approximation factor (to the envy-free benchmark) without budgets (of 7.5) and the first approximation with budgets (of 10.0).

Our mechanisms are based on the clinching auction of Ausubel [2004] generalized to multi-unit environments (a special case of ordinal environments) with budgets by Dobzinski et al. [2008] and ordinal environments by Goel et al. [2012]. There are two dimensions on which we can compare our results to these prior studies of the clinching auction, (a) whether agents have a common budget or distinct budgets and (b) the feasibility constraint of the designer. With respect to (a), our results are weaker as we require a common budget, with respect to (b) our symmetric ordinal environment is between multi-unit environments and the general ordinal environments where the latter allows for asymmetry. The advantage of our restriction to environments that are symmetric in budget (a) and feasibility (b), is that we are able to derive a closed-form formula for the outcome of the clinching auction. Finally, Goel et al. [2012] show that the clinching auction does not generally extend to arbitrary cardinal environments; however, we show that any symmetric cardinal environment contain a symmetric ordinal environment for which the clinching auction performs well (with respect to the objective on the original cardinal environment). Moreover, we can effectively find this ordinal environment and run the clinching auction on it without compromising the agent incentives.

*Organization.* We give formal definitions of the model in Section 2. In Section 3 we characterize the envy-free benchmarks for agents with budgets. In Section 4 we characterize the clinching auction in position environments and show that it is a 2-approximation to the envy-free optimal welfare. In Section 5 we define the biased sampling profit extraction

auction and prove that it is a 10.0-approximation to the envy-free optimal revenue when the agents have a common budget. When the agents do not have a budget constraint, the auction can be improved to a 7.5-approximation and this improvement is given in the full version of the paper [Devanur et al. 2012].

## 2. PRELIMINARIES

*Incentives.* We study auction problems for  $n$  single-dimensional agents with a common budget. Each agent  $i$  has a value  $v_i$  for the service. A mechanism maps reported values  $\mathbf{v} = (v_1, \dots, v_n)$  to a probability that agent  $i$  wins,  $x_i(\mathbf{v})$ , and a payment  $p_i(\mathbf{v})$ .<sup>6</sup> The agents are financially constrained by a budget  $B$  but otherwise are risk neutral. Agent  $i$ 's utility from the mechanism on reports  $\mathbf{v}$  is  $v_i x_i(\mathbf{v}) - p_i(\mathbf{v})$  if  $p_i(\mathbf{v}) \leq B$  and negative infinity otherwise.

A mechanism is *budget respecting (BR)* if no agent pays more than the budget on any valuation profile, i.e., for all  $i$  and  $\mathbf{v}$ ,  $p_i(\mathbf{v}) \leq B$ . A mechanism is *individually rational (IR)* if a risk-neutral agent weakly prefers to participate in the mechanism than not.  $\forall i, \mathbf{v}, v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq 0$ . We say that a mechanism is *incentive compatible (IC)* if a risk-neutral agent maximizes her utility by bidding her true value. I.e.,  $\forall i, \mathbf{v}, z, v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i})$ . [Myerson 1981] characterized incentive compatible mechanisms for single dimensional agents (without budgets) as follows.

**THEOREM 2.1 (MYERSON 1981).** *A mechanism is incentive compatible if and only if the allocation is monotonically non-decreasing in the reported values, i.e., for all  $i$ ,  $x_i(z, \mathbf{v}_{-i})$  is monotone non-decreasing in  $z$  and the expected payments satisfy  $p_i(z, \mathbf{v}_{-i}) = v_i x_i(z, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) dz$ .*

*Feasibility.* As described above, an auction produces a randomized outcome for each agent with probabilities denoted by  $\mathbf{x} = (x_1, \dots, x_n)$ . We assume there is a feasibility constraint that governs the set of such allocations that can be produced. We denote the space of feasible allocations by  $\mathcal{X} \subset [0, 1]^n$ . Our only requirement on this space is that it is symmetric, convex, and downward-closed.<sup>7</sup> Moreover, all we need from our feasibility constraint is that there is an algorithm that (approximately) optimizes a linear sum of weights of the agents served subject to it (and that any agent served can be instead rejected); in these cases we instead view  $\mathcal{X}$  as the induced allocation of the algorithm.

Given this algorithmic view, we partition the classes of feasibility constraints by the kinds of algorithms that work. If “greedy by weight” is optimal then we refer to the feasibility constraint as *ordinal* as only the order of the weights matters and not the actual cardinal weights. We refer to the more general case as *cardinal*. Importantly the ordinal, symmetric case is identical to the position auction environment under common study. A *position environment* is given by a decreasing sequence of position weights  $w_1 \geq \dots \geq w_n$  and each agent can be matched to at most one position. The cardinal, symmetric case includes problems considered in the literature such as the (symmetric restriction) of the *polyhedral environments* of Goel et al. [2012] and the *downward-closed permutation environments* of Hartline and Yan [2011].

*Envy-free Benchmarks.* The goal of prior-free mechanism design is to give a mechanism with a performance guarantee that holds point-wise, i.e., in worst case, on valuation profiles. Such a prior-free guarantee requires comparison to a prior-free benchmark which is also defined point-wise on valuation profiles. Prior-free benchmarks that do not take into account

<sup>6</sup>For clarity we will equate randomized mechanisms with deterministic mechanisms outputting fractional assignments and deterministic payments (both equal to their expectations).

<sup>7</sup>Our envy-free benchmark only makes sense in symmetric environments, mechanism design spaces are always convex if randomization is allowed, and downward closure says that if  $\mathbf{x} \in \mathcal{X}$  then  $\mathbf{x}_{-i} \in \mathcal{X}$  where  $\mathbf{x}_{-i} = (x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n)$ .

the incentive constraints of the mechanism design problem are often inapproximable, but consideration of incentive constraints is non-trivial because incentive constraints bind on possible agent misreports and not point-wise on the valuation profile.

Hartline and Yan [2011] recently demonstrated that *envy-freedom* (EF) constraints are a reasonable point-wise relaxation of incentive constraints. Formally, an outcome  $(\mathbf{x}, \mathbf{p})$  is envy free for valuation profile  $\mathbf{v}$  if for all  $i$  and  $j$ , agent  $i$  does not prefer to swap allocation and payment with agent  $j$ , i.e.,  $v_i x_i - p_i \geq v_i x_j - p_j$ . The envy-free benchmark (with budgets) is defined by optimizing over all envy-free outcomes (that are budget respecting). The following lemma characterizes envy-free outcomes for valuation profiles  $\mathbf{v}$  that are, without loss of generality, indexed by value, i.e.,  $v_i$ 's are monotonically non-increasing in  $i$ .

LEMMA 2.2 (HARTLINE AND YAN 2011). *Allocation  $\mathbf{x}$  has prices for which it is envy free if and only if it is swap monotone, i.e.,  $x_i \geq x_{i+1}$ . The minimum and maximum payments for which such an  $\mathbf{x}$  is envy-free are  $p_i^{\min} = \sum_{j=i+1}^n (x_{j-1} - x_j)v_j$  and  $p_i^{\max} = \sum_{j=i}^n (x_j - x_{j+1})v_j$ , respectively.*

Notice that envy-free payments are monotone non-decreasing in agent values so an envy-free outcome is budget feasible if and only if the highest valued agent (i.e., agent 1) has payment  $p_1 \leq B$ .

The envy-free optimal benchmarks for welfare and revenue with budgets are defined by optimizing over all envy-free outcomes with respect to the respective objective. (These benchmarks are further characterized in Section 3.)

$$\begin{aligned} \text{EFO}^{\text{W}}(\mathbf{v}, B) &= \max \left\{ \sum_i v_i x_i : (\mathbf{x}, \mathbf{p}) \text{ is EF, IR and BR} \right\} \\ \text{EFO}^{\text{R}}(\mathbf{v}, B) &= \max \left\{ \sum_i p_i : (\mathbf{x}, \mathbf{p}) \text{ is EF, IR and BR} \right\} \end{aligned}$$

A prior-free guarantee about a mechanism's performance is defined as follows. A mechanism  $\mathcal{M}$  is a  $\beta$ -approximation to an envy-free benchmark if its expected performance  $\mathcal{M}(\mathbf{v}, B)$  is at least  $\text{EFO}(\mathbf{v}, B)/\beta$  for all  $\mathbf{v}$  and  $B$ . For technical reasons, we slightly modify the envy-free benchmark for revenue and instead approximate  $\text{EFO}^{\text{R}}(\mathbf{v}^{(2)}, B)$  where  $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \dots, v_n)$ . This is necessary because, e.g., when  $v_1 \gg nv_2$ , it is impossible to approximate  $\text{EFO}^{\text{R}}(v_1, B)$ . When the context is clear, we will remove the superscripts and the budget and write  $\text{EFO}(\mathbf{v})$  for readability.

### 3. THE ENVY-FREE BENCHMARK

In this section we characterize welfare-optimal envy-free outcomes for agents with a common budget in symmetric, cardinal environments; at the end of the section we adapt the characterization to the objective of revenue (Section 3.3). This characterization and construction has three main ingredients, *Lagrangian virtual values*, *ironed intervals*, and *partial ironing*. While Lagrangian virtual values and ironed intervals are standard in the literature on optimal mechanism design (e.g., Myerson [1981] and Myerson and Satterthwaite [1983]), partial ironing is new and necessary for agents with budgets.

Lagrangian virtual values for welfare and their associated ironed intervals can be calculated as follows; their derivation and relevance to welfare maximization for agents with budgets is described below in Section 3.1. (Recall: agents are indexed by non-increasing value.)

*Definition 3.1.* For Lagrangian parameter  $\lambda$ , *Lagrangian virtual values* (for welfare), *ironed intervals*, and *ironed Lagrangian virtual values* are calculated as follows:

- (1) The Lagrangian virtual values are  $\phi_1^\lambda = v_1 - \lambda v_2$  and  $\phi_i^\lambda = v_i + \lambda(v_i - v_{i+1})$  for  $i \geq 2$ .
- (2) The Lagrangian welfare curve is  $R^\lambda(j) = \sum_{i=1}^j \phi_i^\lambda = \sum_{i=1}^j v_i - \lambda v_{j+1}$  (with  $R^\lambda(0) = 0$ ).

- (3) The ironed Lagrangian welfare curve  $\bar{R}^\lambda$  is the smallest concave function that is point-wise larger than the Lagrangian welfare curve  $R^\lambda$ .
- (4) The ironed intervals are given by sequences  $\{i, \dots, j\}$  of consecutive agents where  $\bar{R}^\lambda(i-1) = R^\lambda(i-1)$ ,  $\bar{R}^\lambda(j) = R^\lambda(j)$ , and  $\bar{R}^\lambda(k) > R^\lambda(k)$  for  $k \in \{i, \dots, j-1\}$ .
- (5) The ironed Lagrangian virtual values are  $\bar{\phi}_i = \bar{R}^\lambda(i) - \bar{R}^\lambda(i-1)$ , i.e., the left-slope of the ironed Lagrangian welfare curve.

The welfare-optimal outcome for agents with a common budget is a Lagrangian virtual value optimizer; however, there may be several such outcomes. The following construction, with the appropriate parameters, picks from among these outcomes the one that meets the budget constraint with equality (which, as we will see below in Section 3.2, is optimal).

*Definition 3.2.* For Lagrangian parameter  $\lambda \in [0, \infty)$  and partial ironing parameter  $\rho \in [0, 1]$  the *Lagrangian partially ironed outcome* is as follows:

- (1) Construct Lagrangian virtual values (for welfare) from  $\mathbf{v}$  as  $(\phi_1^\lambda, \dots, \phi_n^\lambda)$ , ironed intervals, and ironed Lagrangian virtual values  $(\bar{\phi}_1, \dots, \bar{\phi}_n)$  (Definition 3.1).
- (2) With probability  $\rho$ , merge consecutive ironed intervals with equal ironed Lagrangian virtual value (otherwise, with probability  $1-\rho$ , these consecutive ironed intervals remain separate).
- (3) Iron values  $\mathbf{v}$  to obtain  $\bar{\mathbf{v}}$  by averaging over ironed intervals.
- (4) Find the feasible outcome  $\mathbf{x}$  that maximizes ironed Lagrangian virtual welfare  $\sum_i \bar{\phi}_i x_i$  with tie-breaking by ironed welfare  $\sum_i \bar{v}_i x_i$ .<sup>8</sup>

Notice that above we are taking the convex combination of a minimal ironing and maximal ironing that are consistent with the ironed Lagrangian virtual values and the ironed intervals. A pair of consecutive ironed intervals with the same Lagrangian ironed virtual value by this construction will be *partially ironed*. This partial ironing is absent in existing characterizations of optimal mechanisms. Partial ironing is never necessary for revenue maximization without budgets (for which the ironing technique was first developed) and prior work on welfare or revenue maximization with budgets has restricted attention to benevolent distributions where there is only a single ironed interval (if any) that includes the highest-valued agent (and therefore there is no partial ironing).

The main theorem of this section shows the correctness of this construction; intuition is given but the proof is deferred to the full version of the paper [Devanur et al. 2012].

**THEOREM 3.3.** *For all symmetric cardinal environments and any valuation profile  $\mathbf{v}$  and budget  $B$ , there exist parameters  $\lambda$  and  $\rho$  such that the Lagrangian partially ironed outcome (Definition 3.2) is the welfare-optimal envy-free outcome.*

### 3.1. Ironed Lagrangian Virtual Values

Recall that an allocation  $\mathbf{x} = (x_1, \dots, x_n)$  is envy free if and only if it is swap-monotone and its minimum payments are given by the formula  $p_i^{\min} = \sum_{j=i+1}^n v_j(x_{j-1} - x_j)$  for all agents  $i$  (Lemma 2.2). Note that in order to maximize welfare subject to a budget constraint, picking the minimum envy-free payments is clearly optimal.<sup>9</sup> Further, as envy-free payments are monotone, it is sufficient to impose the budget constraint only on the payment of the top agent, i.e.,  $p_1$ . Therefore, the welfare-optimal envy-free allocation can be captured by the

<sup>8</sup>For ordinal environments where the optimal solution is given by the greedy algorithm, this last step sorts the agents in order of ironed virtual value (which is the same order as values) and then randomly permutes the order of agents in ironed intervals.

<sup>9</sup>Consider any envy-free welfare maximizing allocation subject to the budget. Suppose  $p_i > p_i^{\min}$ . Changing all agents' payments  $p_i$  to  $p_i^{\min}$  is envy-free, has the same welfare, and, since the payments are only lower, satisfies the budget constraint.

following linear program (LP).

$$\begin{aligned}
 \max \quad & \sum_{i=1}^n v_i x_i \\
 \text{s.t.} \quad & x_i \geq x_{i+1} \quad \forall i \\
 & p_1 = \sum_{i=2}^n v_i (x_{i-1} - x_i) \leq B. \\
 & \mathbf{x} \text{ is feasible.}
 \end{aligned} \tag{1}$$

The relaxation of this LP obtained by Lagrangifying the budget constraint is as follows.

$$\begin{aligned}
 \max \quad & \sum_{i=1}^n v_i x_i - \lambda \left( \sum_{i=2}^n v_i (x_{i-1} - x_i) \right) + \lambda B \\
 \text{s.t.} \quad & x_i \geq x_{i+1} \quad \forall i \\
 & \mathbf{x} \text{ is feasible.}
 \end{aligned} \tag{2}$$

LEMMA 3.4. *An allocation is optimal for LP (1) if and only if, either*

- *for some choice of  $\lambda > 0$ , the allocation is optimal for the Lagrangian relaxation (2) with  $\lambda$ , and it satisfies the budget constraint with equality,  $\sum_{i=2}^n v_i (x_{i-1} - x_i) = B$ , or*
- *the allocation is optimal for the Lagrangian relaxation (2) with  $\lambda = 0$  and satisfies the budget constraint,  $\sum_{i=2}^n v_i (x_{i-1} - x_i) \leq B$ .*

PROOF. The statement of the theorem is equivalent to complementary slackness conditions characterizing optimal solutions of an LP.  $\square$

We now consider the optimization problem given by the Lagrangian relaxation (2) for a fixed choice of  $\lambda$ . The objective function of the Lagrangian relaxation (2) can be rewritten as  $\sum_i \phi_i^\lambda x_i$  where  $\phi_1^\lambda = v_1 - \lambda v_2$  and  $\phi_i^\lambda = v_i + \lambda(v_i - v_{i+1})$  for  $i \geq 2$ . The Lagrangian relaxation is now simply the problem of finding the Lagrangian virtual welfare optimal allocation, subject to feasibility, swap-monotonicity, and, if  $\lambda > 0$  then  $p_1 = B$ .

Optimizing the Lagrangian virtual welfare  $\sum_i \phi_i^\lambda x_i$  of non-monotone virtual values subject to swap monotonicity (of the allocation) can be simplified via the technique of *ironing* [Myerson 1981]. The resulting ironed virtual values are monotone and, therefore, ironed virtual welfare maximization without a swap-monotonicity constraint on the allocation will always give an allocation that is swap monotone. Explicit in the ironing construction of Definition 3.1 are ironed intervals (consecutive sequences of agents) on which ironing corresponds to averaging.

LEMMA 3.5. *An allocation  $\mathbf{x}$  is optimal for the Lagrangian relaxation (2) if and only if the allocation maximizes the ironed Lagrangian virtual welfare and the allocation is constant over agents in the same ironed interval.*

Lemma 3.5 has the same proof as the corresponding lemma of Hartline and Yan [2011] for (non-Lagrangian) ironed virtual welfare maximization.

### 3.2. Partial Ironing

Ironed Lagrangian virtual welfare maximization must inherently address ties. Notice ties may arise because agents within the same ironed interval have same ironed virtual value, because agents in consecutive ironed intervals may have the same ironed virtual value, and because several sets of agents may have the same cumulative ironed virtual value. The first kind of tie must be broken uniformly at random, the latter two kinds of ties must be broken so as to meet the budget constraint with equality. We now describe a tie-breaking procedure for Lagrangian ironed virtual welfare maximization that (a) serves agents in the same ironed



interval with the same probability (as per Lemma 3.5) and (b) meets the budget constraint with equality (as per Lemma 3.4).

The tie-breaking rule we will give is based on agents' values. Notice that the objective of (2) is the difference between the social welfare and  $\lambda p_1$  (the payment of the top agent scaled by  $\lambda$ ). Therefore, when there are ties in Lagrangian virtual welfare, it must be that the tied allocation with the maximum (resp. minimum) welfare minimizes (resp. maximizes) the payment of the top agent. This maximum payment must be over budget and the minimum payment must be under budget. Therefore, the appropriate convex combination of these two allocations has payment exactly equal to the budget. The outcome produced is welfare optimal for budgeted agents.

The following approach optimizes Lagrangian ironed virtual welfare with tie-breaking to maximize or minimize the welfare subject to (a) swap monotonicity and (b) agents within the same ironed interval receiving the same probability of service. To maximize welfare, average the values of agents within each ironed interval and tie-break to maximize this averaged welfare. This ensures that the agents in the same ironed interval are treated the same, but otherwise allows the mechanism to optimize welfare over sets of agents with tied Lagrangian ironed virtual welfare. To minimize welfare we would like to optimize the negative of the welfare over allocations with equal Lagrangian ironed virtual welfare. However, this could result in failure of swap monotonicity as agents in successive ironed intervals with the same Lagrangian ironed virtual value will be ranked in the opposite order as required for swap monotonicity. Therefore, to minimize welfare, average the values of agents with equal Lagrangian ironed virtual value (this includes the averaging of agents within the same ironed interval, but additionally averages agents in successive ironed intervals that have the same Lagrangian ironed virtual value), and tie-break to minimize this averaged welfare.

### 3.3. Revenue Maximization

The characterization of the revenue-optimal envy-free outcome for agents with a common budget is the same as above except for (a) the specific formula for Lagrangian virtual values and (b) the tie-breaking procedure. The tie-breaking procedure is a bit more complex than for the welfare objective.

Virtual values are derived starting from the maximum envy-free payments  $p_i^{\max} = \sum_{j=i}^n v_j(x_j - x_{j+1})$  from Lemma 2.2. The objective revenue (without budgets) is given by maximization of the virtual welfare for (non-Lagrangian) virtual values  $iv_i - (i-1)v_{i-1}$  [Hartline and Yan 2011]. Relaxing the budget constraint gives Lagrangian virtual values  $\phi_1^\lambda = v_1(1-\lambda)$  and  $\phi_i^\lambda = (i-\lambda)v_i - (i-1-\lambda)v_{i-1}$  for  $i \geq 2$ . The Lagrangian revenue curve is  $R^\lambda(i) = (i-\lambda)v_i$  with  $R^\lambda(0) = 0$ .

For tie breaking, notice that the analogous objective of the Lagrangian relaxation (2) is revenue minus  $\lambda p_1$ . Therefore, among allocations with the same Lagrangian ironed virtual welfare, the one with the highest revenue has the highest payment of the top agent and the one with the lowest revenue has the lowest payment of the top agent. Revenue is, of course, equal to the (non-Lagrangian) virtual welfare. Whereas for maximizing and minimizing welfare the monotonicity of values implies that we should either prefer to iron as little or as much as possible, for maximizing revenue, the virtual values may not be monotone. Therefore, ironing (i.e., averaging) can be good and bad for both maximizing and minimizing revenue. The following process averages the (non-Lagrangian) virtual values correctly. Consider a set of consecutive Lagrangian ironed intervals with the same Lagrangian ironed virtual value. Average the (non-Lagrangian) virtual values within each interval, calculate the induced revenue curve (by summing prefixes of these averaged virtual values), consider the two-dimensional convex hull of the point set that defined this revenue curve. For maximum revenue, iron as for the upper convex hull; for minimum revenue, iron as for the lower convex hull. Optimizing Lagrangian ironed virtual welfare with tie-breaking by averaged

virtual welfare (from the averaged virtual values calculated above) gives the outcomes with the minimum and maximum payment of the top agent. Mixing between these appropriately to meet the budget constraint with equality gives the revenue-optimal envy-free outcome.

#### 4. WELFARE APPROXIMATION FOR AGENTS WITH A COMMON BUDGET

In this section we study the (polyhedral) clinching auction of Goel et al. [2012] in position environments with a common budget. The outcome of the clinching auction is fundamentally simpler in structure than those of the optimal incentive-compatible auction and optimal envy-free outcome. A fundamental construct in incentive-compatible and envy-free optimization is *ironing*, that is, randomizing between agents whose values fall within a given interval. In Section 3 we characterized welfare-optimal envy-free outcomes as having multiple disjoint ironed intervals. Our first task of this section is to give a similar description of the outcome of the clinching auction. In these terms, the clinching auction has (essentially) one ironed interval and it always contains the top agent. This ironed interval is partially ironed with the singleton interval containing the next highest-valued agent. We give a simple closed-form expression for calculating exactly how this partial ironing is performed.

The clinching auction is not welfare-optimal in two respects. First, given a Bayesian prior distribution, the clinching auction’s expected welfare is not generally optimal among all incentive compatible mechanisms. Second, though the outcome of the clinching auction is envy free, it is not the welfare-optimal envy-free outcome. Nonetheless, we show that the clinching welfare is a two-approximation to the envy-free optimal welfare.

##### 4.1. The clinching auction for position environments

Goel et al. [2012] generalize the clinching auction for budgeted agents to ordinal environments. In this section, we characterize the outcome of this process for symmetric ordinal environments, a.k.a., position environments. At a high level, the clinching auction is described by an ascending price-clock with agents clinching some of the supply at the price as it increases. By the symmetry of the environment and the fact that values of agents above an offered price do not affect the allocation, the budget and partial allocations are identical for each agent that remains in the auction as the price increases. The clinching auction can thus be formulated as follows:

*Definition 4.1 (Clinching Auction).* The *clinching auction* maintains an allocation and price-clock that start from zero. The price-clock ascends continuously and the allocation and budget are adjusted as follows.

- (1) Agents whose values are less than the price-clock are removed and their allocation is frozen.
- (2) The *demand* of any remaining agent is the remaining budget divided by the price clock.
- (3) Each remaining agent *clinches* (and adds to their current allocation) an amount that corresponds to the largest fraction of their demand that can be satisfied when all other remaining agents are first given as much of their demand as possible (subject to the feasibility constraint).<sup>10</sup>
- (4) The budget and allocation are updated to reflect the amount clinched in the previous step.

The auction ends when everyone is removed or the remaining budget is zero.

The reason that the clinching auction is relatively simple to describe in position environments with a common budget is that the feasibility constraint imposed by clinching auctions is one where allocation probability of top positions can be shifted to bottom positions (e.g.,

---

<sup>10</sup>This step is vague in general environments; however, in ordinal environments, i.e., where the greedy algorithm is optimal, it is precise.

by randomizing), but not vice versa. Therefore, an allocation  $\mathbf{x}$  (in decreasing order) is feasible for position weights  $(w_1, \dots, w_n)$  (in decreasing order) if the cumulative allocation  $X_i = \sum_{j \leq i} x_j$  at each coordinate  $i$  is at most the cumulative position weight  $W_i = \sum_{j \leq i} w_j$ .

**PROPOSITION 4.2.** *The clinching auction is incentive compatible and Pareto optimal in position environments with a common budget.*

Pareto optimality means that there is no other reallocation of goods and money that makes an agent strictly better off and no agents are worse off. Proposition 4.2, which is a special case of a more general result of Goel et al. [2012] implies the following structure on the outcome. This structural theorem generalizes one from Dobzinski et al. [2008] (for single-item auctions). It shows that, essentially, the clinching auction is ironing only the top agents.

**THEOREM 4.3.** *Order the agents and positions in decreasing order of their values and let  $\kappa$  be the highest-valued agent who pays strictly less than the budget. Then,*

- (1) *the auction terminates the moment the price-clock exceeds  $v_\kappa$ ,*
- (2) *agents with higher values than  $\kappa$  each receive the same service probability (and pay the budget),*
- (3) *agent  $\kappa$  receives at least the service probability of her corresponding position,*
- (4) *agents with lower values than  $\kappa$  each receive exactly the service probability of their corresponding positions, and*
- (5) *the outcome is envy free.*

**PROOF.** An agent drops out of the clinching auction when her value is exceeded; otherwise, the auction terminates with the price clock below her value when the remaining budget is zero. Let  $\kappa$  be the last agent to drop out when her value is exceeded. By the definition of the clinching auction and symmetry, all higher-valued agents pay the budget and receive the same probability of service. Again by the symmetry of the process there is no envy.

Now consider the agents  $\kappa, \dots, n$  who are paying strictly less than the budget. Assume that initially all excess service probability from the top  $\kappa - 1$  is given to agent  $\kappa$ . Feasibility implies that service probability cannot be shifted up from low-valued agents to high-valued agents and Pareto optimality implies that service probability cannot be shifted down. Consider any probability shifted down from a higher valued agent to a lower valued agent, as these agents are not paying their budget, a Pareto improvement would be for the higher valued agent to buy this shifted service probability from the lower valued agent (at a per-unit price equal to her value). Consequently, agents  $\kappa + 1, \dots, n$  get their corresponding position weight and agent  $\kappa$  gets at least her corresponding position weight.

Finally, we show that the price clock stops immediately after it exceeds  $v_\kappa$ . Assume that the allocation probability of  $\kappa$  is strictly higher than her corresponding position weight (the case of equality is addressed by Theorem 4.4) and suppose that the price clock continues to rise. The actual values of the agents who have not retired are never taken into account in the behavior of the clinching auction. Therefore we can lower  $v_{\kappa-1}$  to just below price clock and at which point agent  $\kappa - 1$  would retire (and not pay her full budget). However, now we have both  $\kappa$  and  $\kappa - 1$  not paying their full budget and  $\kappa$  is getting strictly more service probability than her corresponding weight which contradicts the other results of this theorem.  $\square$

It is easy to see from Theorem 4.3 that the service probability of each of the top  $\kappa - 1$  agents is a little less than the average weight of the top  $\kappa - 1$  positions, and a little more than the average weight of the top  $\kappa$  positions. The service probability of  $\kappa$  is between the




Fig. 1. The outcome of the clinching auction is completely specified by this figure. For agents  $i < \kappa$ , the allocation rule  $x_i(z, \mathbf{v}_{-i})$  of the clinching auction is depicted; the payment of these agents is equal to the area of shaded region which is equal to the budget  $B$ . Agents  $i \geq \kappa$  share this allocation rule on  $z \leq v_i$  which is the relevant portion of the allocation rule for calculating payments. In the envy-free outcome that irons the top  $i$  agents (and ignores the budget constraint) the payment of the top agent, denoted  $B_i$ , is depicted for  $i \in \{\kappa - 1, \kappa\}$ . It is clear from the picture these are monotone in  $i$  and  $B_\kappa < B \leq B_{\kappa-1}$ .

weight of her corresponding position and the average weight of the top  $\kappa$  positions. In fact, the exact service probabilities (and corresponding payments) can be precisely calculated.

The execution of the clinching auction can be described by two phases. In the first phase, the position weights and values are binding; in the second phase, the budget is binding but the position weights are not. Consider ironing the top  $i$  agents, the associated minimum envy-free payments, and the minimum budget  $B_i$  for which the solution is budget respecting. In such a solution, agents  $j \leq i$  are served with probability  $\bar{w}_i = (w_1 + \dots + w_i)/i$  and the *minimum feasible budget* is equal to their payment  $B_i = v_{i+1}(\bar{w}_i - w_i) + \sum_{j=i+1}^n v_j(w_{j-1} - w_j)$ . The minimum feasible budget  $B_i$  is decreasing in  $i$  (Figure 1), zero for  $i = n$ , and (if the budget is binding) greater than the budget for  $i = 1$ . Let  $\kappa$  be an index such that  $B_\kappa \geq B > B_{\kappa+1}$ . In the first phase each of the bottom  $n - \kappa$  agents will clinch their corresponding positions. In the second phase, the clinching auction will behave exactly like the clinching auction for multi-unit environments: the budget starts to bind at a price clock at most  $v_\kappa$  and then the instant the price-clock exceeds  $v_\kappa$  the remaining supply is evenly clinched by the highest  $\kappa - 1$  agents. Figure 1 depicts the resulting outcome and Theorem 4.4 formalizes the observed structure; its proof is given in the full paper [Devanur et al. 2012].

**THEOREM 4.4.** *For any position environment given by position weights  $(w_1, \dots, w_n)$  and budget  $B$  satisfying  $B_\kappa < B \leq B_{\kappa-1}$  for some  $\kappa$ , the polyhedral clinching auction would allocate with:*


- (1)  $w_i$  to every  $i \geq \kappa + 1$ , and
- (2)  $\kappa \bar{w}_\kappa$  split among the top  $\kappa$  agents evenly except for agent  $\kappa$  obtaining  $\delta$  less,

where  $\delta$  is a simple function of  $v_{\kappa+1}; v_\kappa$ ; the remaining budget and the un-clinched supply after agent  $\kappa + 1$  drops out.

#### 4.2. Welfare approximation for ordinal environments

We now show that the clinching auction which (essentially) irons only the top positions, is a two-approximation to the envy-free optimal welfare which may come from ironing an arbitrary number of consecutive positions; moreover, this bound is tight.

**THEOREM 4.5.** *For any position environment with common budgets, the welfare obtained by the clinching auction is a 2-approximation to the envy-free optimal welfare. Furthermore, this ratio is tight even for the single-item environment.*



(a) The upper bound (4) is depicted pictorially. In the clinching auction, the payment of the highest valued agent (cross-hatched) is equal to the budget and at most the rectangle (striped) whose area is  $v_\kappa \bar{w}_{\kappa-1}$ .

(b) The lower bound (5) is depicted pictorially. The payment of the highest valued agent (striped) is equal to the budget and at least the rectangle (cross-hatched) whose area is  $v_\kappa(x_1 - x_\kappa)$ .

Fig. 2. Proofs by picture of the upper and lower bounds on the budget  $B$ .

PROOF. Let  $\kappa$  be the highest-valued agent who does not pay the budget in the clinching auction. Recall from Theorem 4.3 that, relative to the outcome of the clinching auction, if we iron the top  $\kappa$  agents (to get average service probability  $\bar{w}_\kappa = \sum_{i \leq \kappa} w_i / \kappa$ ) then agent  $\kappa$  gets slightly more service probability at the expense of lowering the service probability of the top  $\kappa - 1$  agents; overall there is a net decrease in welfare. Denote the social welfare obtained by the clinching auction on  $\mathbf{v}$  as  $\text{Clinching}(\mathbf{v})$ . We have,

$$\text{Clinching}(\mathbf{v}) \geq \sum_{i=1}^{\kappa} v_i \bar{w}_\kappa + \sum_{i=\kappa+1}^n v_i w_i. \quad (3)$$

Let  $\mathbf{x}$  be the optimal envy-free allocation. We know two things about  $\mathbf{x}$ . First, it is feasible, which means, in particular, that  $\sum_{i \leq \kappa} x_i \leq \kappa \bar{w}_\kappa$ , i.e., the cumulative allocation is at most the cumulative supply. Second, the payment of the highest-valued agent, i.e.,  $p_1$ , (which is given by the “area above the allocation rule” as specified by the minimum envy-free payment identity of Lemma 2.2) is at most the budget. We use these two bounds to show that  $x_1 \leq 2\bar{w}_\kappa$ .

The clinching auction ends when the price-clock just exceeds  $v_\kappa$ , consequently the per-unit cost of service is bounded by  $v_\kappa$ . The probability of service clinched by the top  $\kappa - 1$  agents is slightly lower than  $\bar{w}_{\kappa-1} = \frac{1}{\kappa-1} \sum_{i < \kappa} w_i$ . Therefore an upper bound on the maximum payment (and therefore the budget) is:

$$v_\kappa \bar{w}_{\kappa-1} \geq B. \quad (4)$$

In the envy-free optimal outcome the payment of the top agent (and therefore the budget) is at least:

$$B = \sum_{i=2}^n (x_{i-1} - x_i) v_i \geq v_\kappa (x_1 - x_\kappa). \quad (5)$$

The bounds (4) and (5) combine to give a bound on the probability of service  $x_1$  of the top agent (and thus any agent) in the envy-free outcome.

$$x_1 \leq x_\kappa + \bar{w}_{\kappa-1}. \quad (6)$$

The feasibility constraint of the position environment restricts the envy-free outcome so that

$$\kappa \bar{w}_\kappa \geq \sum_{i \leq \kappa} x_i \geq x_1 + (\kappa - 1)x_\kappa.$$

Solving for  $x_1$  we get a second upper bound.

$$x_1 \leq \kappa \bar{w}_\kappa - (\kappa - 1)x_\kappa. \tag{7}$$

Add  $(\kappa - 1)$  times (6) to (7) to get:

$$\kappa x_1 \leq \kappa \bar{w}_\kappa + (\kappa - 1)\bar{w}_{\kappa-1}. \tag{8}$$

We conclude that  $x_1 \leq 2\bar{w}_\kappa$  as desired.

For the optimal envy-free welfare problem, if the budget constraint is replaced by the weaker constraint of  $x_i \leq 2\bar{w}_\kappa$ , the welfare can only get better. Furthermore, the optimal allocation for this relaxed problem would shift as little service probability down from top slots to lower slots as possible so as to meet the allocation constraint that  $x_i \leq 2\bar{w}_\kappa$ . As the average weight of the top  $\kappa$  positions is  $\bar{w}_\kappa$  the probability of service for agent  $\kappa$  (which is the least of the top agents) can only be at most the average. Therefore, no additional weight is shifted down to lower agents  $j > \kappa$  so,

$$\text{EFO}(\mathbf{v}) \leq \sum_{i=1}^{\kappa} 2v_i \bar{w}_\kappa + \sum_{i=\kappa+1}^n v_i w_i \leq 2 \text{Clinching}(\mathbf{v}),$$

where the last inequality follows from (3).

To show that the 2-approximation is tight, consider the following single-item scenario with a common budget of  $B = 1$ . There are  $N + 1$  agents; the highest valuation is  $N^3$ , the middle  $N - 1$  valuations are  $N$ , and the last valuation is  $N - \epsilon$  where  $\epsilon$  is a small positive number.

The welfare-optimal envy-free allocation would serve the bottom  $N$  agents with equal probability  $x_L$  and the top agent with probability  $x_H > x_L$ . By optimizing the welfare  $N^3 x_H + N^2 x_L$  with the budget constraint  $N(x_H - x_L) \leq 1$ , and the supply constraint  $x_H + N x_L \leq 1$ , we have  $x_H = \frac{2}{N+1}$  while  $x_L = \frac{N-1}{N(N+1)}$ . Thus the optimal envy-free welfare for this case is  $2N^2 - N$ .

The clinching auction would not let anybody clinch as long as the price clock is below  $N$  since there are  $N + 1$  agents who demand at least  $\frac{1}{N}$  while we only have 1 item. However, as soon as the price-clock reaches  $N$ , the bottom agent drops out, and we have  $N$  agents left who demand  $\frac{1}{N}$  each. Thus, each of the top  $N$  agents would receive exactly  $\frac{1}{N}$  and pay the budget. Thus the welfare for clinching in this case is  $N^2 + N - 1$ .

In the limit as  $N$  approaches  $\infty$ , the ratio between the two welfares approaches 2.  $\square$

## 5. REVENUE APPROXIMATION FOR AGENTS WITH A COMMON BUDGET

The main approaches to prior-free auctions for digital goods generalize to symmetric cardinal environments (without budgets). Hartline and Yan [2011] generalized the random sampling auction, and Ha and Hartline [2012] generalized the consensus estimate profit extraction auction. In this section, we generalize the random sampling profit extraction auction from Fiat et al. [2002] for digital good environments to symmetric cardinal environments with a common budget. The random sampling profit extraction auction splits the agents into a market and a sample, estimates the optimal profit from the sample, and then attempts to extract that profit from the market.

A profit extractor is a mechanism that is given some extra information and, if that information is correct, is able to extract a corresponding profit. For symmetric cardinal

environments, Ha and Hartline [2012] gave a profit extractor that is parameterized by an estimated valuation profile and is able to extract profit of at least the envy-free optimal revenue for the estimated valuation profile when that estimate is a coordinate-wise lower bound on the true valuation profile. Our profit extractor below is a simplification of Ha and Hartline [2012] generalized to the case where agents have budgets.

*Definition 5.1.* The *clinchng profit extractor*,  $\text{PE}^{\tilde{\mathbf{v}}}$ , is parameterized by non-increasing valuation profile  $\tilde{\mathbf{v}}$ . It calculates the optimal envy-free outcome  $\tilde{\mathbf{x}}$  for  $\tilde{\mathbf{v}}$  and then runs the clinching auction for position weights  $\tilde{\mathbf{x}}$  on the true valuation profile  $\mathbf{v}$ .

Assume that  $\mathbf{v}$  and  $\tilde{\mathbf{v}}$  are in non-increasing order. Define  $\mathbf{v}$  as *one-ahead after index  $\eta$*  for  $\tilde{\mathbf{v}}$  if  $\eta$  is the lowest index for which all  $i > \eta$  satisfy  $v_{i+1} \geq \tilde{v}_i$ . When  $\eta = 0$  define  $\mathbf{v}$  as *one-ahead dominating*  $\tilde{\mathbf{v}}$ , denoted  $\mathbf{v}_{-1} \geq \tilde{\mathbf{v}}$ . The following lemma shows that the clinching profit extractor on  $\mathbf{v}$  is able to obtain the contribution to the optimal envy-free revenue for  $\tilde{\mathbf{v}}$  from agents  $\{\eta + 1, \dots, n\}$ .

**LEMMA 5.2.** *If  $\mathbf{v}$  one-ahead dominates  $\tilde{\mathbf{v}}$  then the clinching profit extractor revenue is at least the estimated envy-free optimal revenue, i.e.,  $\text{PE}^{\tilde{\mathbf{v}}}(\mathbf{v}) \geq \text{EFO}(\tilde{\mathbf{v}})$ ; moreover, if  $\mathbf{v}$  is one-ahead after index  $\eta$  for  $\tilde{\mathbf{v}}$  then the contribution to the profit extractor revenue from agent  $i > \eta$  is at least the contribution to the estimated envy-free optimal revenue from  $i$ , i.e.,  $\text{PE}_i^{\tilde{\mathbf{v}}}(\mathbf{v}) \geq \text{EFO}_i(\tilde{\mathbf{v}})$ .*

**PROOF.** We will prove the second part of the lemma which implies the first. Consider an  $i > \eta$ . The maximum  $i$  could pay in any outcome is the budget, so if  $i$  pays her budget in the clinching auction then the bound holds. Suppose instead that  $i$  pays strictly less than her budget in the clinching auction. Consider the following sequence of inequalities with explanation below (where  $\mathbf{x}$  is the allocation of the clinching auction and  $\tilde{\mathbf{x}}$  is the envy-free optimal outcome for  $\tilde{\mathbf{v}}$ ).

$$\begin{aligned} \text{PE}_i^{\tilde{\mathbf{v}}}(\mathbf{v}) &\geq \sum_{j=i+1}^n (x_{j-1} - x_j)v_j \geq \sum_{j=i+1}^n (\tilde{x}_{j-1} - \tilde{x}_j)v_j \\ &\geq \sum_{j=i+1}^n (\tilde{x}_{j-1} - \tilde{x}_j)\tilde{v}_{j-1} = \text{EFO}_i(\tilde{\mathbf{v}}). \end{aligned}$$

The first inequality follows from envy freedom of the clinching auction and the formula for minimum envy-free payments (Lemma 2.2). For  $j > i$ , Theorem 4.3 implies that  $x_j = \tilde{x}_j$  because all but the highest-valued agent who does not pay her budget are allocated with exactly their corresponding position weight; the theorem also implies that  $x_i \geq \tilde{x}_i$  as agent  $i$  also does not pay her budget (but she might be the highest such agent). The second equality then follows as the service probabilities are unaffected by the swap from  $x_j$  to  $\tilde{x}_j$  except for  $x_i$  which only appears positively and is at least  $\tilde{x}_i$ . The third inequality comes from the fact that  $i$  is greater than  $\eta$  so one-ahead dominance holds at  $i$  and higher indices. The final equality follows from the formula for maximum envy-free payments (Lemma 2.2).  $\square$

We now define a simple biased sampling procedure and show that it ensures one ahead dominance with significant probability.

*Definition 5.3 (Biased Sampling).* Parameterized by a probability  $p$ , the *biased sampling* process assigns each agent into the sample  $S$  independently with probability  $p$ , otherwise the market  $M$ .

Let  $\mathbf{v}^M$  and  $\mathbf{v}^S$  be the sorted valuation vectors of  $M$  and  $S$  respectively, and assume that  $\mathbf{v}^M$  and  $\mathbf{v}^S$  are padded with 0's to be equal in length for comparison convenience. The

biased sampling process has the following probabilistic properties (proof given at the end of the section).

LEMMA 5.4. *For the biased sampling process with  $p < 0.5$  and  $\eta$  being a random variable for the index after which  $\mathbf{v}^M$  is one-ahead for  $\mathbf{v}^S$ ,*

- (1)  $\Pr[\mathbf{v}^M \not\geq \mathbf{v}^S] \leq \frac{p}{1-p}$ ,
- (2)  $\Pr[\mathbf{v}^M \not\geq \mathbf{v}^S \mid 1 \in M] \leq \left(\frac{p}{1-p}\right)^2$ , and
- (3)  $\sum_{i=1}^n i \Pr[\eta = i \mid 1 \in M] \leq \frac{p}{(1-2p)^2}$ .

Furthermore, all inequalities are tight in the limit as  $n$  approaches  $\infty$ .

LEMMA 5.5. *The optimal envy-free revenue of a random sample  $S$  whose elements are selected i.i.d. with probability  $p$  satisfies  $\mathbf{E}[\text{EFO}(\mathbf{v}^S)] \geq p \text{EFO}(\mathbf{v})$ .*

PROOF. Consider the envy-free optimal outcome for  $\mathbf{v}$ . Clearly if we restrict attention only to the agents in  $S$  there is still no envy. Therefore,  $\text{EFO}(\mathbf{v}^S) \geq \text{EFO}_S(\mathbf{v})$  where  $\text{EFO}_S(\mathbf{v})$  is a short-hand notation for the contribution from the agents in  $S$  to the envy-free optimal revenue on  $\mathbf{v}$ . Of course,  $\mathbf{E}[\text{EFO}_S(\mathbf{v})] = p \text{EFO}(\mathbf{v})$ .  $\square$

*Definition 5.6 (BSPE<sub>p</sub>).* The biased sampling profit extraction auction parameterized by  $p < 0.5$  for a common budget  $B$  works as follows.

- (1) Partition the set of agents into  $\mathbf{v}^M$  and  $\mathbf{v}^S$  using biased sampling parameterized by  $p$ .
- (2) Run the clinching profit extractor parameterized by  $\mathbf{v}^S$  on  $\mathbf{v}^M$  and budget  $B$ .

Incentive compatibility of BSPE<sub>p</sub> comes straight from that of the profit extractor. We have the following revenue guarantee.

LEMMA 5.7. *For all  $p < 0.5$  the revenue of BSPE<sub>p</sub> satisfies,*

$$\text{BSPE}_p(\mathbf{v}) \geq (1-p)p \text{EFO}(\mathbf{v}_{-1}) - \frac{p(1-p)}{(1-2p)^2} \text{EFO}(v_2).$$

PROOF. Condition on the case that the highest-valued agent, i.e., 1, is in  $M$  and let  $\eta$  is the index after which  $\mathbf{v}^M$  is one-ahead for  $\mathbf{v}^S$ . From Lemma 5.2, the profit extractor's revenue conditioned on  $\eta = i$  is  $\mathbf{E}[\text{PE}^{\mathbf{v}^S}(\mathbf{v}^M) \mid \eta = i] \geq \mathbf{E}[\text{EFO}(\mathbf{v}^S) \mid \eta = i] - \sum_{j=1}^i \mathbf{E}[\text{EFO}_j(\mathbf{v}^S) \mid \eta = i]$ . This inequality holds since we would extract the payment from all agents that are lower than  $i$ ; or equivalently, we would extract the full payment (the first term of the right-hand side) minus those from the  $i$  highest agents (the second term). Using the observation that  $\text{EFO}_j(\mathbf{v}^S) \leq \text{EFO}(v_2)$  for all  $j$  (because agent 1 is in  $M$ ), we have:

$$\mathbf{E}[\text{PE}^{\mathbf{v}^S}(\mathbf{v}^M) \mid \eta = i] \geq \mathbf{E}[\text{EFO}(\mathbf{v}^S) \mid \eta = i] - i \text{EFO}(v_2).$$

Summing these revenue guarantees over all  $\eta$ , we have:

$$\begin{aligned} \text{BSPE}_p(\mathbf{v}) &= \sum_{i=1}^{\infty} \mathbf{E}[\text{PE}^{\mathbf{v}^S}(\mathbf{v}^M) \mid \eta = i] \Pr[\eta = i] \\ &\geq \sum_{i=1}^{\infty} \mathbf{E}[\text{EFO}(\mathbf{v}^S) \mid \eta = i] \Pr[\eta = i] - \text{EFO}(v_2) \sum_{i=1}^{\infty} i \Pr[\eta = i] \\ &= \mathbf{E}[\text{EFO}(\mathbf{v}^S)] - \text{EFO}(v_2) \frac{p}{(1-2p)^2} \geq p \text{EFO}(\mathbf{v}_{-1}) - \text{EFO}(v_2) \frac{p}{(1-2p)^2}. \end{aligned}$$

The last inequality comes from Lemma 5.5 on  $\mathbf{v}_{-1}$ . Finally, we remove the conditioning on  $1 \in M$  by multiplying the above quantity by the probability  $1-p$ .  $\square$

*Definition 5.8 (Pseudo-Vickery).* The pseudo-Vickrey auction finds the feasible outcome  $\mathbf{x}$  that optimizes  $x_1$  with  $x_j = 0$  for  $j \neq 1$  and runs the clinching auction with position weights  $\mathbf{x}$ .



[Hartline and Yan 2011] observe that EFO is sub-additive (without budgets); it continues to be sub-additive with budgets. Thus,  $\text{EFO}(\mathbf{v}_{-1}) + \text{EFO}(v_2) \geq \text{EFO}^{(2)}(\mathbf{v})$ .<sup>11</sup> Furthermore, since pseudo-Vickrey obtains at least  $\text{EFO}(v_2)$ , we have the following result.

**THEOREM 5.9.** *The convex combination of the pseudo-Vickrey auction (with probability  $\frac{q}{1+q}$ ) and BSPE<sub>p</sub> (with probability  $\frac{1}{1+q}$ ), where  $q = (1-p)p + \frac{p(1-p)}{(1-2p)^2}$ , approximates  $\text{EFO}^{(2)}(\mathbf{v})$  within a factor of  $1 + \frac{1}{(1-p)p} + \frac{1}{(1-2p)^2}$ . This ratio is minimized at 10.0 when  $p = 0.211$ .*

**PROOF OF LEMMA 5.4.** Consider the following infinite random walk on a straight line: starting from position 0, with probability  $p$ , move backward one step; otherwise, move forward one step. The position of this random walk describes precisely the difference between the number of agents in  $M$  and  $S$ , where positive value means  $M$  has more agents than  $S$ . We will show the results as equalities by a “probability of ruin” analysis of an infinite random walk; inequalities follow for random walks that terminate after a finite number  $n$  of steps.

- (1) The event  $\mathbf{v}^M \not\geq \mathbf{v}^S$  happens when there exists a time that  $M$  has fewer agents than  $S$ . Let  $r$  be the probability of ruin, i.e., the random walk eventually takes one step backward from the initial position, we have  $r = p + (1-p)r^2$ . The first component is the probability of taking one step backward in the first step, and the second component is the probability of the first step being a forward step, then eventually take two steps backward. Solving this equation for  $r \in (0, 1)$  gives  $r = p/(1-p)$ .
- (2) When we condition on  $1 \in M$ , our initial position is 1, not 0, and the probability of ruin is  $r^2$ .
- (3) We will first derive  $\Pr[\eta = i \mid 1 \in M]$  for  $i \geq 1$ . Since  $i$  is the lowest index after which  $\mathbf{v}^M$  is one-ahead for  $\mathbf{v}^S$ , we must have (a) an equal partition amongst the top  $2i$  agents, (b)  $v_{2i+1}$  is in  $M$ , and (c) from this point on, the number of agents assigned to  $M$  is never fewer than that from  $S$ . Thus, by conditioned on the highest value agent already in  $M$ , we have:

$$\begin{aligned} \Pr[\eta = i \mid 1 \in M] &= \binom{2i-1}{i} p^i (1-p)^{i-1} (1-p) \left(1 - \frac{p}{1-p}\right) \\ &= \binom{2i}{i} [p(1-p)]^i \frac{1-2p}{2(1-p)}. \end{aligned} \quad (9)$$

The Taylor’s series expansion of  $\frac{1}{\sqrt{1-4z}}$  for any  $0 < z < 1/4$  gives us  $\sum_{i=0}^{\infty} \binom{2i}{i} z^i = \frac{1}{\sqrt{1-4z}}$ . By differentiating both sides with respect to  $z$ , then multiplying them with  $z$ , we have  $\sum_{i=1}^{\infty} i \binom{2i}{i} z^i = \frac{2z}{(1-4z)\sqrt{1-4z}}$ . For  $z = p(1-p)$ , we have  $\sqrt{1-4z} = 1-2p$ . Hence, this equality translates to

$$\sum_{i=1}^{\infty} i \binom{2i}{i} [p(1-p)]^i = \frac{2p(1-p)}{(1-2p)^3}. \quad (10)$$

Putting these all together,

$$\begin{aligned} \sum_{i=1}^{\infty} i \Pr[\eta = i \mid 1 \in M] &= \sum_{i=1}^{\infty} i \binom{2i}{i} [p(1-p)]^i \frac{1-2p}{2(1-p)} && \text{from (9)} \\ &= \frac{2p(1-p)}{(1-2p)^3} \frac{1-2p}{2(1-p)} = \frac{p}{(1-2p)^2}. && \text{from (10)} \end{aligned}$$

□

<sup>11</sup> $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(\mathbf{v}^{(2)})$  where  $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \dots, v_n)$ .

REFERENCES

- AUSUBEL, L. M. 2004. An efficient ascending-bid auction for multiple objects. *American Economic Review* 94, 5, 1452–1475.
- BALCAN, M.-F., BLUM, A., HARTLINE, J. D., AND MANSOUR, Y. 2008. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.* 74, 8, 1245–1270.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11, 17–33. 10.1007/BF01726210.
- DEVANUR, N. R., HA, B. Q., AND HARTLINE, J. D. 2012. Prior-free auctions for budgeted agents. *arXiv preprint arXiv:1212.5766*.
- DOBZINSKI, S., LAVI, R., AND NISAN, N. 2008. Multi-unit auctions with budget limits. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*. FOCS '08. IEEE Computer Society, Washington, DC, USA, 260–269.
- FIAT, A., GOLDBERG, A., HARTLINE, J., AND KARLIN, A. 2002. Competitive generalized auctions. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*. ACM, 72–81.
- GOEL, G., MIRROKNI, V., AND PAES LEME, R. 2012. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the 44th symposium on Theory of Computing*. ACM, 107–122.
- GOLDBERG, A. AND HARTLINE, J. 2003. Competitiveness via consensus. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '03. Society for Industrial and Applied Mathematics, 215–222.
- GOLDBERG, A., HARTLINE, J., AND WRIGHT, A. 2001. Competitive auctions and digital goods. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '01. Society for Industrial and Applied Mathematics, 735–744.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 4, pp. 617–631.
- HA, B. AND HARTLINE, J. 2012. Mechanism design via consensus estimates, cross checking, and profit extraction. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 887–895.
- HARTLINE, J. AND YAN, Q. 2011. Envy, truth, and profit. In *Proceedings of the Twelfth ACM conference on Electronic Commerce*. EC '11. ACM, 243–252.
- LAFFONT, J.-J. AND ROBERT, J. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2, 181–186.
- LEONARDI, S. AND ROUGHGARDEN, T. 2012. Prior-free auctions with ordered bidders. In *Proceedings of the 44th Symposium on Theory of Computing*. ACM, 427–434.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- MYERSON, R. AND SATTERTHWAITE, M. 1983. Efficient mechanisms for bilateral trading. *Journal of economic theory* 29, 2, 265–281.
- PAI, M. AND VOHRA, R. 2008. Optimal auctions with financially constrained bidders.
- VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 1, 8–37.
- WILSON, R. 1987. Game theoretic analysis of trading processes. *Advances in Economic Theory*.