Envy, Truth, and Profit

Jason D. Hartline* Qiqi Yan†

November 4, 2010

Abstract

We consider (profit maximizing) mechanism design in general settings that include, e.g., position auctions (for selling advertisements on Internet search engines) and single-minded combinatorial auctions. We analyze optimal envy-free pricing in these settings and give economic justification for using optimal envy-free revenue as a benchmark for prior-free mechanism design and analysis. In addition to its economic justification, the envy-free revenue has a very simple structure and a strong connection to incentive compatibility constraints in mechanism design.

As a first example of the connection between envy-free pricing and incentive compatible mechanism design, because the structures of optimal pricings and optimal mechanisms are similar, we give a reduction from structurally rich environments including position auctions (and environments with a matroid structure) to multi-unit auction environments (i.e., auctioning k identical units to n unit-demand agents). For instance, via this reduction we are able to extend all prior-free digital good auctions to position auctions with a factor of two of loss in the approximation factor.

As a second example we extend a variant of the random sampling auction to downward closed settings. To prove that its revenue (as an incentive compatible mechanism) is a good approximation to the envy-free bechmark, we consider its envy-free revenue instead. The envy-free revenue of a mechanism is closely tied to its incentive compatible revenue, but is much easier to analyze because it is defined pointwise on valuation profiles. Our analysis shows that the envy-free revenue of the random sampling auction is a constant approximation to the optimal. Also, we show that its IC revenue is at least half of its envy-free revenue. The random sampling auction, therefore, is a prior-free constant approximation in downward-closed environments.

*Electrical Engineering and Computer Science, Northwestern University, Evanston, IL. Email: hartline@eecs.northwestern.edu.

†Department of Computer Science, Stanford University, Stanford, CA 94305. Supported by Stanford Graduate Fellowship. Email: qiqiyan@cs.stanford.edu.

1 Introduction

Envy freedom, the constraint that no agent would prefer the fate of another to his own, has seen extensive consideration in the recent algorithmic pricing literature where it generalizes the economic consideration that prevents the sale of the same good at different prices. *Incentive compatibility*, the solution concept that predicts that selfish agents will truthfully report their preferences if each agent's utility is optimized by truthful report, on the other hand, is the de facto standard one in the extensive classic and current literature on mechanism design. Envy-free pricings have been of interest because of their connection to prior-free optimal mechanism design [16], at least in unlimited supply settings such as *digital goods* [13]. This paper extends this connection to more general settings and supports the informal thesis: *the revenue of envy-free pricings can be approximated by incentive-compatible mechanisms.*

As an example setting, consider *position auctions environments*, a popular model for advertising auctions on Internet search engines. Such an environment is specified by n non-increasing click-through weights $\mathbf{w} = (w_1, \ldots, w_n) \in [0,1]^n$ with w_i corresponding to the probability that the searcher will click on an advertisement shown in position i. A mechanism makes a partial assignment of n advertisers to these n positions. In position auction environments, the literature and practice have predominantly considered reserve pricing as a means to boost the auctioneer's revenue (e.g., with the Vickrey-Clarke-Groves (VCG) [25, 5, 15] mechanism or generalized second price (GSP) auction [24, 9]). This approach is justified by the optimality of reserve pricing in benign settings where revenues satisfy a natural regularity assumption which implies that reserve pricing is optimal. For settings where there are several distinct kinds of agents, e.g., advertisers for zoos and cars both competing for the keyword "jaguar," reserve pricing, however, is not optimal in general.

The classical economic approach to revenue maximization in mechanism design assumes that the agents values, e.g., value for a click in the position auction example, are drawn from a known distribution. In such a Bayesian setting, the optimal mechanism is the one that maximizes expected revenue, among all mechanisms, for the known distribution. For many reasons, a recent branch of literature has been exploring prior-free optimal mechanism design, i.e., where the agent values are worst case. There is no single optimal prior-free mechanism, so prior-free mechanism design looks to find a mechanism that approximates the revenue of a reasonable benchmark for any profile of agent values.

The prior-free environment that has received the most attention is that of digital goods auctions. This is equivalent to the position auction setting described above with $w_i = 1$ for all i, i.e., all agents can be served. In this structurally benevolent setting, a natural and economically well justified benchmark, "the optimal revenue from a posted price (with at least two winners)," can be approximated by a prior-free mechanism, e.g., [14]. Notice first that such a benchmark is envy-free: any agent can choose to accept the price and receive the item or not, therefore no agent is envious of any other agent. Furthermore, in the Bayesian setting where the distribution of values is known, a posted price is optimal. Therefore, in approximating this benchmark, a prior-free auction simultaneously approximates the Bayesian optimal mechanism for any distribution.

A centerpiece of this paper is a generalization of the posted price benchmark (which is envy-free) to structurally rich environments. We obtain the general benchmark by writing down the envy-freedom constraints and solving for the envy-free optimal pricing. In position auction environments such an envy-free optimal pricing can be described as a grouping of positions into bins with a price on each bin and an assignment of agents to bins, such that each agent prefers a uniform random position from her assigned bin at its price to a uniform random position from another bin at the other bin's price.

In many relevant environments the optimal envy-free pricing, which is prior-free, gives an upper bound on the revenue of the Bayesian optimal mechanism for any distribution. Such an upper bound holds for position auctions and more generally for environments with a matroid feasibility constraint. (It does not hold, for example, in non-downward-closed environments such as public projects, i.e., we serve either all agents or none.)

The first example we give to showcase the importance of envy-freedom in incentive compatible mechanism design is in a reduction that reduces position auction environments (and environments given by matroid feasibility constraints) to simple multi-unit auction environments, i.e., selling k identical units to n unit-demand agents ($k \leq n$). E.g., from any multi-unit auction that is a prior-free β-approximation to the envy-free benchmark, we can derive a position auction that gives a β-approximation. Furthermore, from any β-approximation auction for digital good environments, a 2β-approximation auction for multi-unit environments can be derived. Therefore, from each of the many digital goods auctions discussed in the literature, we can derive a position auction with at most twice the approximation factor.

The second example we give considers general downward-closed environments such as single-minded combinatorial auctions. In single-minded combinatorial auctions, e.g., [20], there is a set of items available and each agent has a private value for a known bundle of items. Essentially, it is a weighted set packing problem. We consider an intermediary between the Bayesian setting, where the agents' values are drawn i.i.d. from a distribution, and the prior-free setting, where the agents' values are worst case, by assuming that the agents' values are randomly permuted with respect to the desired bundles of items. In such a prior-free setting we show that a random sampling auction approximates the optimal envy-free pricing benchmark. This result is made possible by a close connection between the revenue in an incentive compatible mechanism and the revenue from the pricing for which the outcome of the mechanism is envy-free. (Note: the envy-free payments and the incentive compatible payments may generally be distinct.) For implication of this result and background on revenue benchmarks, we refer the readers to Section C.

Overview We formally present the setting we consider and review the classical theory of Bayesian optimal mechanism design in Section 2. We develop a parallel theory of optimal envy-free pricing in Section 3. In section 4 we give a reduction from position auctions and matroid environments to multi-unit auctions (and digital good environments). In Section 5 we formally relate payments in an incentive compatible mechanism to payments in an envy-free pricing. Finally, in Section 6 we adapt the random sampling auction to general downward closed environments. Conclusions are given in Section 7.

Related work This paper follows from a line of work that studies prior-free revenue properties of the random sampling auction of [13]. The tightest analysis of the random sampling auction for digital good settings is given by Alaei et al. [1]. For the limited supply version of this problem (i.e., k-unit auctions), Hartline and Roughgarden [18] proposed a benchmark for prior-free analysis that is derived from Bayesian optimal auctions. With this benchmark, Devanur and Hartline [7] extended the analysis from [1] to limited supply settings. In this context, the present paper further extends the benchmark of [18] and the analysis of [7] to settings with general downward-closed feasibility constraints.

Our derivation of optimal envy-free pricings closely mirrors Myerson's theory of Bayesian optimal auctions [22]. Connections between envy-free pricings and prior-free mechanism design have been made before, e.g., in [16, 3].

The paper most related to ours is Dhangwatnotai et al. [8], which shows that in Bayesian settings when agents' values are drawn i.i.d. from an unknown distribution, a mechanism based on a "single-sample approach" gives an 8-approximation for downward-closed set systems and distributions that satisfy a standard *monotone hazard rate* condition, and a 2-approximation for structurally nicer matroid set systems and all distributions that satisfy a standard *regularity* condition. However, as we show, a single-sample cannot give better than a logarithmic approximation for the fully general distributional setting. In this context, we give a constant approximation for matroid settings and general i.i.d. distributions on agent values. This is important as many distributions, e.g., bimodal, do not satisfy the regularity condition required by [8].

2 Bayesian Optimal Mechanism Design (Briefly)

In this section we review Bayesian optimal mechanism design in a general downward-closed single-dimensional environments. The agents' preferences are drawn from a known distribution and the designer's goal is the mechanism with maximum expected profit (in equilibrium). The derivation here is based on Myerson [22] and refinements by Bulow and Roberts [4].

There are $n \geq 2$ agents. Each agent i has a valuation v_i for receiving an abstract service. The *valuation profile* is $\mathbf{v} = (v_1, \ldots, v_n)$. The values are drawn independently and identically from distribution F (assume the distribution is continuous with distribution function $F(z)$ and density function $f(z)$). An agent i who is served with probability x_i and charged price p_i obtains utility $u_i = v_i x_i - p_i$.

An *allocation* is the vector $\mathbf{x} = (x_1, \ldots, x_n) \in [0, 1]^n$ where x_i is the probability that agent i is served. The set of feasible allocations contains the origin, is convex, and is downward closed in the sense that if $\mathbf{x}$ is feasible and $\mathbf{x}' \leq \mathbf{x}$ (i.e., $x'_i \leq x_i$ for all i), then $\mathbf{x}'$ is feasible. Downward closure implies, for instance, that the service provided to any agent can always be degraded without affecting any other agents. These general environments include digital good auctions, multi-unit auctions, position auctions, matroid environments, and single-minded combinatorial auctions.

A *mechanism* is specified by an allocation rule and a payment rule. The *allocation rule* $\mathbf{x}(\mathbf{v}) = (x_1(\mathbf{v}), \ldots, x_n(\mathbf{v}))$ maps a valuation profile to a feasible allocation. A *payment rule* $\mathbf{p}(\mathbf{v}) = (p_1(\mathbf{v}), \ldots, p_n(\mathbf{v}))$ maps a valuation profile to a non-negative payment for each agent.

An allocation and payment rule pair is (ex post) *incentive compatible* if no agent prefers the outcome when misreporting her value to the outcome when reporting the truth. Formally,

$$\forall i, z, \mathbf{v}, \quad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i}),$$

where $(z, \mathbf{v}_{-i})$ represents the valuation profile $\mathbf{v}$ with v_i replaced with z. A payment rule is *individually rational* if each agent's utility is non-negative.

An allocation rule is *value monotone* if the probability that an agent is served is monotone non-decreasing in her value, i.e., for all agents i, $x_i(z, \mathbf{v}_{-i})$ is non-decreasing in z. The following well-known theorem characterizes ex post IC mechanisms.

Theorem 2.1 *[22] An allocation rule $\mathbf{x}$ admits a non-negative and individually rational payment rule $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is incentive compatible if and only if $\mathbf{x}$ is value monotone, and the uniquely determined payment rule is:*

$$p_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) dz = \int_0^{v_i} z x'_i(z, \mathbf{v}_{-i}),$$

where $x'_i(z, \mathbf{v}_{-i})$ is the derivatives of $x_i(z, \mathbf{v}_{-i})$ with respect to z.

Revenue curves and *virtual values* are important and related constructs in optimal mechanism design. Revenue curves correspond the revenue a seller can obtain from selling to a single agent as a function of the probability of sale. The revenue curve for a single agent with value distributed according to distribution function F is $\mathrm{R}^F(q) = q \cdot F^{-1}(1-q)$ which is explicitly a product between the offer price, $F^{-1}(1-q)$, and the probability of sale at that price, q. The *ironed revenue curve* $\bar{\mathrm{R}}^F(q)$ is defined as the minimum concave function that upper-bounds $\mathrm{R}^F(q)$. Intuitively, when $\bar{\mathrm{R}}(q) \neq \mathrm{R}(q)$, the best way of selling with probability q is to offer a lottery pricing instead of price $F^{-1}(1-q)$. *Virtual values* and *ironed virtual values* are the derivative of the respective revenue curves, but are specified in value space but not probability space, i.e., $\Phi^F(v) = \frac{d\,\mathrm{R}^F(q)}{dq}$ and $\bar{\Phi}^F(v) = \frac{d\bar{\mathrm{R}}^F(q)}{dq}$ with $q = 1 - F(v)$.[1] Importantly, the concavity of ironed revenue curves implies that ironed virtual valuation functions are monotone non-decreasing.

The importance of revenue curves and virtual values is summarized by the following theorem.

Theorem 2.2 *[22] The non-negative and individually rational payment rule* $\mathbf{p}$ *for monotone allocation rule* $\mathbf{x}$ *satisfies,*

$$\mathbf{E}[p_i(z, \mathbf{v}_{-i})] = \mathbf{E}\big[\Phi^F(z) x_i(z, \mathbf{v}_{-i})\big] = \mathbf{E}\big[\mathrm{R}^F(z) x_i'(z, \mathbf{v}_{-i})\big]\,,$$

where the expectation is over a random draw of z *from* F*, and* $x_i'(z, \mathbf{v}_{-i})$ *is the derivative of* x_i *with respect to* z*.*

Thus, the search for the Bayesian optimal auction can be rephrased as a search for the allocation rule that maximizes virtual surplus subject to monotonicity. Importantly, in the case that the virtual valuation function is monotone non-decreasing, pointwise optimization of virtual surplus results in a monotone allocation rule; therefore the optimal mechanism is simply the virtual surplus maximizer. In the general case, the ironed virtual valuation function, by the concavity of the ironed revenue curve, is monotone, therefore maximizing ironed virtual surplus always gives a monotone allocation rule and furthermore, though we omit the details, this gives the optimal mechanism.

Definition 2.1 (ironed virtual surplus optimizer) *For any feasibility constraint and ironed virtual valuation function* $\bar{\Phi}(\cdot)$*, the* ironed virtual surplus optimizer $\mathbf{x}^{\bar{\Phi}}$ *chooses the allocation* $\mathbf{x}$ *that maximizes total ironed virtual surplus,* $\sum_i \bar{\Phi}(v_i) x_i$*, subject to feasibility, with ties broken uniformly at random.*

Theorem 2.3 *[22] When values are drawn from distribution* F *the optimal mechanism,* ICO^F*, is the ironed virtual surplus optimizer for* $\bar{\Phi}^F$ *with the appropriate payment rule.*

3 Optimal Envy-free Pricing

In this section we derive a theory of optimal envy-free pricings in single-dimensional settings which mirrors that of Bayesian optimal (incentive-compatible) mechanisms. Proofs of the theorems herein are so similar to those in the optimal mechanism design literature that we defer them to Appendix A. Whereas incentive compatibility constrains the mechanism so that no agent would want to misreport her value, envy freedom constrains the outcome so that no agent would want to swap outcomes with another agent.

This swapping of outcomes only makes sense in settings where the agents are a priori symmetric. For general sets of feasible outcomes that are convex, downward-closed, and contain the origin, a

[1]Plug in the formula $q = 1 - F(v)$ to derive the familiar $\Phi^F(v) = v - \frac{1-F(v)}{f(v)}$ where f is the density function corresponding to distribution F.

natural way to symmetrize is by subjecting the agents' roles in the set system to a uniform random permutation. (This is a common intermediary between i.i.d. Bayesian settings and worst-case settings, e.g., the secretary problem [2]) Of course, the resulting feasibility constraint is symmetric, convex, downward-closed, and contains the origin.

An allocation $\mathbf{x}$ with payments $\mathbf{p}$ is *envy free* for valuation profile $\mathbf{v}$ if no agent prefers the outcome of another agent to her own. Formally,

$$\forall i, j,\ v_i x_i - p_i \geq v_i x_j - p_j.$$

Importantly, envy constraints bind point-wise for $\mathbf{x}$, $\mathbf{p}$, and $\mathbf{v}$. This contrasts sharply to IC constraints which constrain the functional form of the allocation rule, i.e., to be monotone.

We first characterize envy-free pricings. Notice that the maximum payment characterization are very similar in form to the payment characterization of incentive-compatible mechanisms.

An allocation is *swap monotone* if the allocation probabilities have the same order as the valuations of the agents.[2] I.e., for all i, j, $x_i \geq x_j$ whenever $v_i \geq v_j$. For convenience, our discussion of envy-free pricing will assume the agents are in descending order, i.e., $v_i \geq v_{i+1}$, and let $v_{n+1} = 0$.

Lemma 3.1 *In symmetric settings, an allocation* $\mathbf{x}$ *admits a non-negative and individual rational payment rule* $\mathbf{p}$ *such that* $(\mathbf{x}, \mathbf{p})$ *is envy-free if and only if* $\mathbf{x}$ *is swap monotone. If* $\mathbf{x}$ *is swap monotone, then the maximum payments for* $\mathbf{x}$ *satisfy, for all* i,

$$p_i = v_i x_i - \sum\nolimits_{k=i+1}^{n} x_k \cdot (v_{k-1} - v_k) = \sum\nolimits_{k=i}^{n} (v_k - v_{k+1}) \cdot (x_i - x_{k+1}).$$

Given a valuation profile $\mathbf{v}$ we denote the (empirical) revenue curve by $\mathrm{R}^{\mathbf{v}}(i) = i \cdot v_i$ for $i = \{1, \ldots, n\}$ (recall v_i's are indexed in decreasing order). For convenience we also let $\mathrm{R}^{\mathbf{v}}(0) = \mathrm{R}^{\mathbf{v}}(n+1) = 0$. The ironed revenue curve, denoted $\bar{\mathrm{R}}^{\mathbf{v}}(i)$, is the minimum concave function that upper-bounds R. Likewise, define the (empirical) virtual valuation function $\Phi^{\mathbf{v}}(v) = \mathrm{R}^{\mathbf{v}}(i) - \mathrm{R}^{\mathbf{v}}(i-1)$ and the (empirical) ironed virtual valuation function $\bar{\Phi}^{\mathbf{v}}(v) = \bar{\mathrm{R}}^{\mathbf{v}}(i) - \bar{\mathrm{R}}^{\mathbf{v}}(i-1)$, where $i \in \{1, \ldots, n+1\}$ is such that $v \in [v_i, v_{i-1})$. (We set $v_0 = \infty$ for notational convenience.) See Figure 1.

We now characterize the revenue of an envy-free pricing in terms of virtual values. Notice that the revenue characterization is very close in spirit to Myerson's characterization of revenue for Bayesian incentive compatible mechanisms. One difference is that our characterization is on the sum of all agent payments, whereas Myerson's characterization is per-agent.

Lemma 3.2 *The (maximum) envy-free revenue of a swap monotone allocation* $\mathbf{x}$ *satisfies:*

$$\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) = \sum\nolimits_{i=1}^{n} \Phi^{\mathbf{v}}(v_i) x_i = \sum\nolimits_{i=1}^{n} \mathrm{R}^{\mathbf{v}}(i)(x_i - x_{i+1}).$$

An implication of the characterization of maximum envy-free revenue as the (empirical) virtual surplus suggests that to optimize revenue, the allocation rule should optimize virtual surplus subject to swap monotonicity. In symmetric environments with monotone virtual valuation functions, the maximization of virtual surplus results in a swap monotone allocation. In general symmstric environments, the allocation that maximizes ironed virtual surplus is both swap monotone and revenue optimal among all swap monotone allocations.

Allocation rules of focus for this paper are ironed virtual surplus optimizers. We will abuse notation to let $\bar{\Phi}^{\mathbf{v}}$ denote $\mathbf{x}^{\bar{\Phi}^{\mathbf{v}}}$, the ironed virtual surplus optimizer with ironed virtual valuation function defined according to $\mathbf{v}$. For an allocation rule $\mathbf{x}(\cdot)$, we let $\mathrm{EF}^{\mathbf{x}}(\mathbf{v})$ to denote $\mathrm{EF}^{\mathbf{x}(\mathbf{v})}(\mathbf{v})$. Hence $\mathrm{EF}^{\bar{\Phi}^{\mathbf{v}}}(\cdot)$ is the envy-free revenue from using ironed virtual surplus optimizer $\bar{\Phi}^{\mathbf{v}}$, etc.

[2] This is the single-dimensional special case of the "local efficiency" condition of Mu'alem [21].

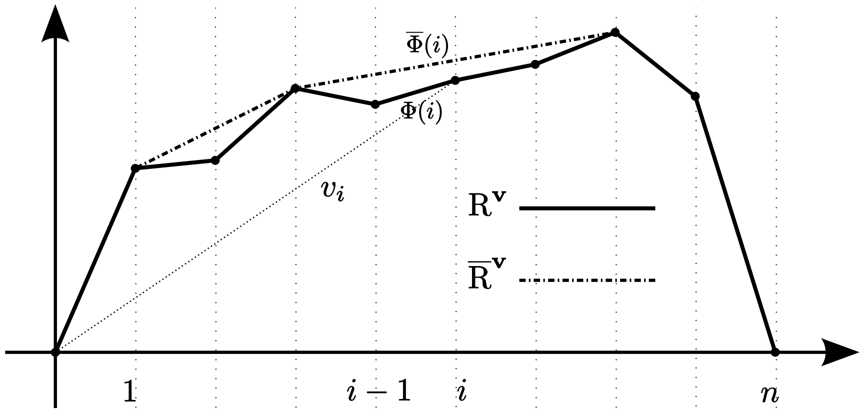

Figure 1: A depiction of the relationship between valuation profiles, revenue curves, and ironed revenue curves. In particular, v_i is equal to the slope of the line connecting $(i, \mathrm{R}^{\mathbf{v}}(i))$ with the origin. $\Phi(i)$ is equal to the "left slope" of $\mathrm{R}^{\mathbf{v}}(i)$ at i, i.e., the slope of the line segment connecting $(i-1, \mathrm{R}^{\mathbf{v}}(i-1))$ with $(i, \mathrm{R}^{\mathbf{v}}(i))$. Similarly $\bar{\Phi}(i)$ is the left slope $\bar{\mathrm{R}}$ at i, where $\bar{\mathrm{R}}^{\mathbf{v}}$ is the minimum concave function that upper-bounds of $\mathrm{R}^{\mathbf{v}}$.

Theorem 3.3 *For all valuation profiles* $\mathbf{v}$*, the ironed virtual surplus optimizer* $\bar{\Phi}^{\mathbf{v}}$ *computes an allocation that maximizes envy-free revenue among all swap-monotone allocations* $\mathbf{x}$*. I.e.,* $\mathrm{EFO}(\mathbf{v}) = \mathrm{EF}^{\bar{\Phi}^{\mathbf{v}}}(\mathbf{v}) \geq \mathrm{EF}^{\mathbf{x}}(\mathbf{v})$ *for all* $\mathbf{x}$*.*

This theorem is proved by a useful lemma that relates revenue to ironed virtual surplus.

Lemma 3.4 *For any swap-monotone* $\mathbf{x}$ *on valuation profile* $\mathbf{v}$*,*

$$\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) = \sum\nolimits_{i=1}^{n} \mathrm{R}^{\mathbf{v}}(i) \cdot (x_i - x_{i+1}) \leq \sum\nolimits_{i=1}^{n} \bar{\mathrm{R}}(i) \cdot (x_i - x_{i+1}) = \sum\nolimits_{i=1}^{n} \bar{\Phi}(i) \cdot x_i,$$

with equality holding if and only if $x_i = x_{i+1}$ *whenever* $\bar{\mathrm{R}}(i) > \mathrm{R}(i)$*.*

4 Matroids, Position Auctions, and Multi-unit Auctions

In this section we consider matroid permutation environments, position auctions, and multi-unit auctions. We show that for both incentive compatible mechanism design and envy-free pricing, these environments are closely related. In fact, for both IC and EF, the optimal mechanisms are the same and approximation mechanisms give the same approximation factor. In the interest of brevity, we will focus on approximating the optimal EF revenue with a prior free mechanism. Our solution to this will be by way of a two step reduction: we reduce matroid permutation environments to position auctions, which we reduce to multi-unit auctions.

Briefly, matroid permutation environments are ones with a feasibility constraint derived from the independent sets of a matroid set system. This matroid constraint is then randomly permuted with respect to the roles the agents play. For the purpose of our discussion, a matroid is a set system for which the greedy algorithm always selects the maximum feasible set of maximum value. Position auctions are given by non-negative and non-increasing probabilities $w_1, \ldots, w_n$ normalized with $w_1 = 1$. The auction assigns the agents to positions. Multi-unit auctions assign k identical items to n unit-demand agents. Multi-unit auctions are a special case of matroid environments, the k-uniform matroid, and a special case of position auctions with $w_i = 1$ for $i \leq k$ and zero otherwise.

The property of these three settings that enables this reduction is that (ironed virtual) surplus maximization is solved by the greedy algorithm (with ties broken randomly). Therefore the only information needed to perform the surplus maximization is a partial order on the agents.

A position auction with weights $\mathbf{w}$ is related to matroid permutation settings with *characteristic weights* $\mathbf{w}$, which we define as the following:

Definition 4.1 *In a matroid setting, choose any valuation profile* $\mathbf{v}$ *with all distinct values, assign the agents to elements in the matroid via a random permutation, and then run the greedy algorithm w.r.t.* $\mathbf{v}$*. Then for each* i*,* w_i *is the probability that agent* i *is serviced in this random process.*

Note that the characteristic weights are well-defined because the outcome of the greedy algorithm is fully determined by the relative ordering of the values.

4.1 Reductions

Here we show that the optimal IC mechanism and EF pricings in the three environments are essentially the same. Then we give a reduction that can be applied to an approximately optimal multi-unit auction to get an approximately optimal position auction, and then to the approximately optimal position auction to get an approximately optimal matroid permutation mechanism. The main challenge is in correctly reducing to approximation mechanisms that are not ironed virtual surplus optimizers.

Lemma 4.1 *Any ironed virtual surplus optimizer* $\bar{\Phi}$ *has the same allocation rule in the following three settings:*

1. *a matroid permutation setting with characteristic weights* $\mathbf{w}$*,*
2. *a position auction setting with weights* $\mathbf{w}$*,*
3. *a convex combination of* k*-unit auction settings where we run a* k*-unit auction with probability* $w_k - w_{k+1}$ *for* $k = 1, \ldots, n$*.*

Proof: Fix a tie-breaking rule, which induces an ordering on the agents. Then $\bar{\Phi}$ essentially runs greedy on the agents with non-negative $\bar{\Phi}$ values according to this ordering. The j-th agent with non-negative $\bar{\Phi}$ value in this ordering (1) gets allocated with probability w_j in the matroid permutation setting by definition of characteristic weights, (2) gets assigned to position j in the position auction and hence gets allocated with probability w_j, and, (3) gets allocated in k-unit auction for each $k \geq j$, and hence has probability $\sum_{k \geq j}(w_k - w_{k+1}) = w_j$ of being serviced in the convex combination setting. Taking expectation over all tie-breaking orders, agent i has the same probability of being serviced in the three settings. □

In particular, by plugging in $\bar{\Phi}^{\mathbf{v}}$ into Lemma 4.1, the empirical ironed virtual surplus optimizer, as a corollary, the optimal envy-free revenue for the three settings are equal.

Lemma 4.2 *(1) Given allocation rules* $\mathbf{x}^k$*'s for* k*-unit auctions, for* $k = 1 \ldots n$*, we can construct an allocation rule for position auction with weights* $\mathbf{w}$ *such that* $\mathbf{x}^P(\mathbf{v}) = \sum_{k=1}^{n}(w_k - w_{k+1}) \cdot \mathbf{x}^k(\mathbf{v})$ *for all* $\mathbf{v}$*.*

(2) Given an allocation rule $\mathbf{x}^P$ *for position auctions with weights* $\mathbf{w}$*, we can construct an allocation rule* $\mathbf{x}^M$ *for any matroid permutation setting with characteristic weights* $\mathbf{w}$ *such that* $\mathbf{x}^P(\mathbf{v}) = \mathbf{x}^M(\mathbf{v})$ *for all* $\mathbf{v}$*.*

Proof: To prove (1), we simulate a j-unit auction using $\mathbf{x}^k$ on the input $\mathbf{v}$ for each j and let $x_i^{(j)}$ be the probability that agent i is allocated in simulation j. Let $x_i = \sum_j x_i^{(j)}(w_j - w_{j+1})$ be the expected allocation to j in the convex combination setting. Reindex $\mathbf{x}$ in non-increasing order. Then $\mathbf{w}$ majorizes $\mathbf{x}$ in the sense that $\sum_i^k w_i \geq \sum_i^k x_i$, with equality holds for $k = m$. Therefore by a theorem of Rado [23] we can write $\mathbf{x} = S \cdot \mathbf{w}$ where S is a doubly stochastic matrix. Any doubly stochastic matrix is a convex combination of permutation matrices, so we can write $S = sum_t r_t P_t$

where $\sum_t r_t = 1$ and each P_t is a permutation matrix (Birkhoff–von Neumann Theorem). Finally, we pick a t with probability r_t and assign the agents to positions in the permutation specified by P_t. The resulting allocation is exactly the desired $\mathbf{x}^P$.

To prove (2), first, we run $\mathbf{x}^P$ and let $\mathbf{j}$ be the assignment where j_i is the position assigned to agent i, or $j_i = \perp$ if i is not assigned a slot. Reject all agents i with $j_i = \perp$. Now run the greedy matroid algorithm in the matroid permutation setting on an arbitrary valuation profile with distinct values and output its outcome. □

There are two important caveats that we will address when instantiating this reduction. First, as we are discussing incentive compatibility in the matroid permutation environments, it is important that the agents do not know their assignment to roles in the set system. If an agent does know their role in the set system, the generic reduction is not generally incentive compatible. It is however incentive compatible in the special case where the mechanism is an ironed virtual surplus optimizer. Second, we have assumed that in going from matroid permutation environments to position auctions that the characteristic weights were known. These may be challenging to compute exactly, making it difficulty to apply the reduction. However, in the case of ironed virtual surplus optimizers, the weights do not need to be computed as the partial ordering by the ironed virtual valuation functions specifies the order in which to consider agents in the matroid permutation environment.

4.2 Applying the Reduction

The above reduction allows us to extend all of the mechanisms for multi-unit auction, e.g., the digital goods literature, to position auctions and matroid permutation settings. We give two examples.

The optimal envy-free revenue upper-bounds the IC revenue of any ironed virtual surplus optimizer (see Lemma 5.2 in the Section 5). As these are precisely the mechanisms that are optimal in Bayesian settings, a prior-free mechanism that approximates the optimal envy-free revenue on all valuation profiles simultaneously approximates the Bayesian optimal revenue in the case the values are drawn from a distribution. This suggests that $\mathrm{EFO}(\mathbf{v})$ would be a good benchmark for prior-free approximation mechanisms, for more discussion, see [18]. For technical reasons we attempt to approximate $\mathrm{EFO}^{(2)}(\mathbf{v}) = \mathrm{EFO}(\mathbf{v}^{(2)})$ with $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \ldots, v_n)$ (see [14] for discussion).

For k-unit settings, the prior-free digital good auction literature gives approximations to the benchmark $\mathcal{F}^{(2)}(\mathbf{v}) = \max_{2 \le i \le k} i v_i$. This benchmark is a 2-approximation to $\mathrm{EFO}(\mathbf{v}^{(2)})$ [7]; therefore a β-approximation to $\mathcal{F}^{(2)}$ implies a 2β-approximation to $\mathrm{EFO}^{(2)}$. A mechanism of McGrew and Hartline [17] is a 3.25-approximation to $\mathcal{F}^{(2)}$, and hence a 6.5-approximation to $\mathrm{EFO}^{(2)}$. By the above reduction this mechanism can be used to give the same approximation to $\mathrm{EFO}^{(2)}$.

Theorem 4.3 *There is an IC mechanism $\mathcal{M}$ for position auctions such that $\mathrm{IC}^{\mathcal{M}}(\mathbf{v}) \ge \frac{1}{6.5} \cdot \mathrm{EFO}(\mathbf{v}^{(2)})$ for all $\mathbf{v}$. There is an IC mechanism $\mathcal{M}$ for matroid permutation settings such that $\mathrm{IC}^{\mathcal{M}}(\mathbf{v}) \ge \frac{1}{6.5} \cdot \mathrm{EFO}(\mathbf{v}^{(2)})$ for all $\mathbf{v}$.*

The Hartline-McGrew mechanism is not an ironed virtual surplus optimizer and as suggested by the discussion above (a) the induced matroid permutation mechanism is only IC if the agents do not know their roles in the set system, and (b) the characteristic weights may be difficult to compute. To address these deficiencies, consider the following mechanism.

Definition 4.2 (RSEM) *The* Random Sampling Empirical Myerson *(*RSEM*) mechanism randomly partitions the population of agents N into a market M and a sample S, i.e., each agent is in S independently with probability* 0.5*. It then uses the ironed virtual surplus optimizer for the sample, $\bar{\Phi}^{\mathbf{v}_S}$ or simply denoted by $\bar{\Phi}^S$, to choose a set $W \subseteq M$ of agents from the market that maximizes the total $\bar{\Phi}^S$ value, and then allocate to this set W.*

For k-unit auction settings, the IC revenue of RSEM is a 50-approximation to $\mathrm{EFO}(\mathbf{v}^{(2)})$, for all $\mathbf{v}$ (this is a corollary of the main theorem from [7]). Condition on a partitioning (S, M) of the agents, RSEM can be seen as applying the ironed virtual surplus optimizer $\bar{\Phi}^S$ to the valuation profile $(\mathbf{v}_M, \mathbf{0}_S)$ where agents in S are set to have zero value. By Lemma 4.1, the resulting allocation of RSEM for matroid is equivalent to a convex combination of the allocations in the k-unit auction settings, with multiplicative weights $(w_k - w_{k+1})$'s. As a result, by summing over all random partitioning, the IC revenue of RSEM for matroids is a convex combination of the IC revenue of RSEM for k-unit auctions. On the other hand, the EFO revenue on $\mathbf{v}^{(2)}$ of the matroid setting is also equal to the convex combination of the EFO revenue on $\mathbf{v}^{(2)}$ in the k-unit auctions with the same multiplicative weights. It follows that the revenue guarantee of *RSEM* for k-unit auctions extend to matroid permutation settings as well.

Theorem 4.4 *For matroid permutation settings,* $\mathrm{IC}^{\mathrm{RSEM}}(\mathbf{v}) \geq \frac{1}{50} \cdot \mathrm{EFO}(\mathbf{v}^{(2)})$ *for all* $\mathbf{v}$.

5 Incentive Compatibility versus Envy Freedom

In this section we compare envy-free revenue to incentive-compatible revenue for ironed virtual surplus optimizers in the permutation setting, where agents are assigned to roles in the set system via a random permutation. Here we assume that agents' values are ordered, i.e., $v_i \geq v_{i+1}$. For a mechanism $\mathcal{M}$, we let $\mathrm{IC}^{\mathcal{M}}(\mathbf{v})$ denote the IC revenue from running $\mathcal{M}$ over $\mathbf{v}$.

First we lower bound IC revenue by half of the maximum envy-free revenue under a technical condition. In the following we use $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ and $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$ to denote the IC and EF revenue from agent i by applying the ironed virtual surplus maximizer $\bar{\Phi}$, respectively.

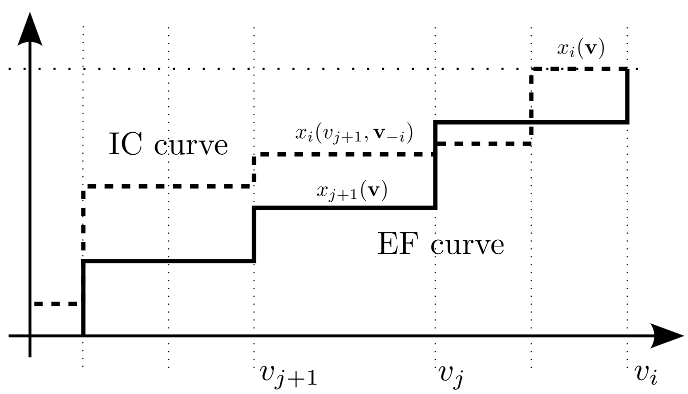

Figure 2: Depiction of EF allocation and IC allocation rule from which the payments for agent i are computed. The EF allocation curve maps each value in $[v_{j+1}, v_j)$ to $x_{j+1}(\mathbf{v})$, and the IC allocation curve maps each z to $x_i(z, \mathbf{v}_{-i})$.

Lemma 5.1 *For downward-closed permutation settings, all valuations* $\mathbf{v}$*, and all piece-wise constant ironed virtual valuation functions* $\bar{\Phi}$*, if for every maximal interval* $[l, r)$ *such that* $\bar{\Phi}(t)$ *is constant for* $t \in [l, r)$*, there is an agent* i *such that* $v_i = l$*, then for all* i*,* $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v}) \geq \frac{1}{2}\,\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$.

Proof: Let $\mathbf{x}(\cdot)$ denote the allocation rule of the ironed virtual surplus optimizer $\bar{\Phi}$. By the assumption of the lemma, for all j, $\bar{\Phi}(z)$ is constant for all $z \in [v_{j+1}, v_j)$, and hence the IC allocation rule in fact maps each $z \in [v_{j+1}, v_j)$ to $x_i(v_{j+1}, \mathbf{v}_{-i})$.

By Lemma 2.1, $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$ which, referring to Figure 2, equals the area above the IC curve and below the horizontal dotted line. On the other hand, $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$, which similarly corresponds to the area above the EF curve and below the horizontal dotted line. It suffices to prove that: $x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}) \geq \frac{1}{2} \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$, or equivalently, $x_i(\mathbf{v}) + x_{j+1}(\mathbf{v}) \geq 2 \cdot x_i(v_{j+1}, \mathbf{v}_{-i}) = 2 \cdot x_i(v_{j+1}, \mathbf{v}_{-j+1})$. This says that the total winning probability of agent i and $j+1$ can only decrease if agent i lowers her bid to v_{j+1}, which can be easily verified for ironed virtual surplus optimizers. □

In matroid permutation settings, envy-free revenue upper-bounds incentive-compatible revenue.

Lemma 5.2 *For matroid permutation settings, all valuations $\mathbf{v}$, and all ironed virtual valuation functions $\bar{\Phi}$, for all agent i, $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v}) \geq \mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$.*

Proof: Recall that $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$, and $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\int_0^{v_i}(x_i(\mathbf{v}) - x_i(z, \mathbf{v}_{-i}))dz$. By the monotonicity of $x_i(z, \mathbf{v}_{-i})$ in z, $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is upper-bounded by $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$. Also note that $x_i(v_{j+1}, \mathbf{v}_{-i}) = x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$ because agent i and agent $j+1$ have the same value. It suffices to prove $x_{j+1}(\mathbf{v}) \leq x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$. This says that agent $j+1$ is more likely to win if agent i decreases her bid to v_{j+1}, which is true for the greedy-based ironed virtual surplus optimizers in matroids. □

6 Downward-closed Permutation Settings

For downward-closed permutation settings, the maximum envy-free revenue may not upper-bound the incentive-compatible revenue, in certain bizarre cases, which means that approximating the EFO benchmark does not necessarily imply prior-free approximation in i.i.d. distributional settings.

Lemma 6.1 *There exists a downward-closed set system, and valuation profile $\mathbf{v}$, such that if $\bar{\Phi}$ is the ironed virtual valuation function of $\mathbf{v}$, then $\mathrm{IC}^{\bar{\Phi}}(\mathbf{v}) > \mathrm{EF}^{\bar{\Phi}}(\mathbf{v})$.*

These cases seem pathological, and even in these cases the envy-free revenue seems to be not too far below the IC revenue of the any an ironed virtual surplus optimizer. Therefore we believe it remains an interesting benchmark for approximation in downward closed settings.

In this section, we will show that a variant of RSEM approximates this maximum envy-free revenue benchmark by a constant factor. This variant of RSEM, called RSEM$'$, is similar to RSEM except that after partitioning into sample set S and the market set M, we choose a subset W of all agents that maximize total $\bar{\Phi}^S$ values, and then allocate only the agents in the intersection of W and M. Our main results for downward-closed settings is the following:

Theorem 6.2 *For downward-closed permutation settings, $\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v}) \geq \frac{1}{2560}\,\mathrm{EFO}(\mathbf{v}^{(2)})$ for all $\mathbf{v}$.*

The readers are referred to Section D for the detailed proof of Theorem 6.2. Here we only mention several important ingredients of the analysis. First of all, by our "IC vs EF" lemma (Lemma 5.1), the problem reduces to purely studying envy-free revenue. Second, a desirable property of envy-free revenue is that they can be related to revenue curves by Lemma 3.4, such that the problem of studying envy-free revenue, loosely speaking, reduces to studying revenue curves, and revenue curves have many nice combinatorial properties. Third, in the RSEM$'$ mechanism, we apply the ironed virtual surplus maximizer tailored for the sample set to the whole population, which is in some sense "ironing in the wrong way". To capture the effect of such sub-optimal ironing, we carefully construct "effective" revenue curves, and relate the envy-free revenue to such effective revenue curves. Finally, we invoke a double-sided version of the balanced sampling lemma of Feige et al. [10], and show that certain revenue curves are within a constant factor of each other, which imply the guarantee we need.

7 Conclusions

We have drawn a strong connection between what is possible via envy-free pricing and what is possible via prior-free incentive compatible mechanisms. Probably the most important open question motivated by this connection is whether the recent logarithmic approximation to envy-free makespan by Cohen et al. [6] implies that there is an incentive compatible mechanism that also achieves a logarithmic approximation, modulo, for instance, technical assumptions like one machine not being too important.

Another interesting direction is related to our general reduction from matroid permutation environments to multi-unit auctions, a.k.a., k-uniform matroids. In the matroid secretary problem of Babaioff et al. [2] (an online permutation environment) a constant approximation is possible for uniform matroids where as for general rank k matroids, the best algorithm only gives a logarithmic approximation (in the rank of the matroid). Perhaps our viewpoint of matroids in permutation settings as a convex combination of uniform matroids will help resolve questions in this area.

Finally, for downward closed settings the IC revenue of a virtual surplus maximizer may exceed that of the optimal envy-free pricing. This does not imply that the envy-free benchmark is not strong enough for downward closed settings. For instance, if we take expectations over valuation profiles drawn i.i.d. from any distribution, it may be that the optimal envy-free pricing is always at least the Bayesian optimal revenue. If this were true, then prior-free approximation of the envy-free benchmark implies simultaneous approximation of any Bayesian optimal mechanism.

References

[1] Saeed Alaei, Azarakhsh Malekian, and Aravind Srinivasan. On random sampling auctions for digital goods. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 187–196, 2009.

[2] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA07*, pages 434–443, 2007.

[3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *Proc. of the 46th IEEE Symp. on Foundations of Computer Science*, pages 605–614, 2005.

[4] J. Bulow and J. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90, 1989.

[5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[6] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. Envy-free makespan approximation: extended abstract. In *ACM Conference on Electronic Commerce*, pages 159–166, 2010.

[7] N. Devanur and J. D. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, 2009.

[8] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*, 2010.

[9] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, 2007.

[10] U. Feige, A. Flaxman, J. Hartline, and R. Kleinberg. On the Competitive Ratio of the Random Sampling Auction. In *Proc. 1st Workshop on Internet and Network Economics*, pages 878–886, 2005.

[11] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Generalized competitive auctions. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 72–81. ACM Press, 2002.

[12] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[13] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.

[14] Andrew Goldberg, Jason Hartline, Anna Karlin, Mike Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 2006.

[15] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[16] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proc. 16th ACM Symp. on Discrete Algorithms*, 2005.

[17] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 175–182, 2005.

[18] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 75–84, 2008.

[19] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, 2009.

[20] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce (EC)*, pages 96–102. ACM Press, 1999.

[21] Ahuva Mu'alem. On multi-dimensional envy-free mechanisms. In *Conference on Algorithmic Decision Theory*, pages 120–131, 2009.

[22] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[23] Rado R. An equality. *J. London Math. Soc.*, 27:1–6, 1952.

[24] Hal Varian. Position auctions. In *Working Paper, UC Berkeley*, 2006.

[25] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37, 1961.

A Envy-Free Pricing

A.1 Proof of Lemma 3.1

Proof: Suppose $\mathbf{x}$ admits $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is envy-free. By definition, $v_i x_i - p_i \geq v_i x_j - p_j$ and $v_j x_j - p_j \geq v_j x_i - p_i$. By summing these two inequalities and rearranging, $(x_i - x_j) \cdot (v_i - v_j) \geq 0$, and hence $\mathbf{x}$ is swap monotone.

Suppose $\mathbf{x}$ is swap monotone. Let $\mathbf{p}$ be given as in the lemma. We verify that $(\mathbf{x}, \mathbf{p})$ is envy-free. There are two cases: if $i \leq j$, we have:

$$\begin{aligned} p_i - p_j &= \sum_{k=i}^{j-1} v_k \cdot (x_k - x_{k+1}) \\ &\leq \sum_{k=i}^{j-1} v_i \cdot (x_k - x_{k+1}) \\ &= v_i \cdot \sum_{k=i}^{j-1} (x_k - x_{k+1}) \\ &= v_i \cdot (x_i - x_j), \end{aligned}$$

and if $i \geq j$, we have:

$$\begin{aligned} p_i - p_j &= -\sum_{k=j}^{i-1} v_k \cdot (x_k - x_{k+1}) \\ &\leq -\sum_{k=j}^{i-1} v_i \cdot (x_k - x_{k+1}) \\ &= -v_i \cdot \sum_{k=j}^{i-1} (x_k - x_{k+1}) \\ &= v_i \cdot (x_i - x_j). \end{aligned}$$

In particular, $p_i = p_i - p_{n+1} \leq v_i \cdot (x_i - x_{n+1}) \leq v_i$, and hence $p_i \leq v_i$.

To show that p_i is maximum, note that for any other envy-free pricing p_i, $p_i - p_j \leq v_i \cdot x_i - v_i \cdot x_j$, and hence:

$$\begin{aligned} p_i &= \sum_{j=i}^{n} (p_j - p_{j+1}) + p_{n+1} \\ &\leq \sum_{j=i}^{n} v_j \cdot (x_j - x_{j+1}) + v_{n+1} \\ &= p_i. \end{aligned}$$

$\square$

A.2 Proof of Lemma 3.2

Proof: The proof is by the following equalities:

$$\begin{aligned}
\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) &= \sum_{i=1}^{n}\sum_{j=i}^{n} v_j \cdot (x_j - x_{j+1}) \\
&= \sum_{i=1}^{n} i v_i \cdot (x_i - x_{i+1}) \\
&= \sum_{i=1}^{n} \mathrm{R}(i) \cdot (x_i - x_{i+1}) \\
&= \sum_{i=1}^{n} (\mathrm{R}(i) - \mathrm{R}(i-1)) \cdot x_i \\
&= \sum_{i=1}^{n} \Phi^{\mathbf{v}}(v_i) \cdot x_i
\end{aligned}$$

□

A.3 Proof of Theorem 3.3

Proof: Note that whenever $\bar{\mathrm{R}}(i) > R(i)$, we have $\bar{\Phi}^{\mathbf{v}}(v_i) = \bar{\Phi}^{\mathbf{v}}(v_{i+1})$. Since $\bar{\Phi}^{\mathbf{v}}$ does not distinguish between the i-th largest valuation and the $i+1$-th largest valuation, we have $x_i(\mathbf{v}) = x_{i+1}(\mathbf{v})$. By Lemma 3.4, we have that $\mathrm{EFO}(\mathbf{v}) = \sum_i \bar{\Phi}^{\mathbf{v}}(v_i) \cdot x_i(\mathbf{v})$.

Also by Lemma 3.4, $\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) \leq \sum_i \bar{\Phi}(v_i) \cdot x_i(\mathbf{v})$. Since $\mathbf{x}^{\mathbf{v}}$ maximizes $\sum_i x_i(\mathbf{v}) \cdot \bar{\Phi}^{\mathbf{v}}(v_i)$ over all $x_i(\mathbf{v})$, $\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) \leq \mathrm{EFO}(\mathbf{v})$.

□

A.4 Proof of Lemma 3.4

Proof: To show the inequality, we have:

$$\begin{aligned}
\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) &= \sum_{i=1}^{n} \mathrm{R}(i) \cdot (x_i - x_{i+1}) \\
&= \sum_{i=1}^{n} \bar{\mathrm{R}}(i) \cdot (x_i - x_{i+1}) \\
&\quad - \sum_{i=1}^{n} (\bar{\mathrm{R}}(i) - \mathrm{R}(i)) \cdot (x_i - x_{i+1}) \\
&\leq \sum_{i=1}^{n} \bar{\mathrm{R}}(i) \cdot (x_i - x_{i+1}) \\
&= \sum_{i=1}^{n} (\bar{\mathrm{R}}(i) - \bar{\mathrm{R}}(i-1)) \cdot x_i \\
&= \sum_{i=1}^{n} \bar{\Phi}^{\mathbf{v}}(v_i) \cdot x_i,
\end{aligned}$$

where we use the fact that $\bar{\mathrm{R}}(i) \geq \mathrm{R}(i)$ and $x_i \geq x_{i+1}$.

Clearly the equality holds if and only if $x_i = x_{i+1}$ whenever $\bar{\mathrm{R}}(i) > \mathrm{R}(i)$. □

B Proof of Lemma 6.1

Proof: Let there be $n+1$ agents. The "1 vs n" set system has two maximum feasible sets, one is a singleton set and the other one has size n. These two sets are disjoint. We define the valuation profile by specifying the virtual valuations. There are n "small" agents with virtual values $v+\epsilon, v+2\epsilon, \ldots, v+n\epsilon$ respectively, and one "big" agent with virtual value $nv + \frac{n(n+1)}{2}\epsilon - \epsilon^2$ for some small positive ϵ. The choice of the ϵ terms is such that for the sum of the virtual valuations of the first n agents to beat the big agent, no small agent can lower her virtual value to some other agent's virtual value. We will ignore ϵ terms from now on. Correspondingly, one can calculate the revenue curve, and then derive the valuations of the agents: the valuation of the big agent is nv, and the small agents have values $\frac{n+1}{2}v$, $\frac{n+2}{3}v, \ldots$, $\frac{2n}{n+1}v$, ignoring ϵ terms. The allocation rule is the ironed virtual surplus optimizer w.r.t. this valuation profile. Note that a reserve of $\frac{2n}{n+1}v$ is set because any value lower than this corresponds to a negative ironed virtual value.

Observe that every agent wins if and only if she is assigned to the size n set, which happens with probability $n/(n+1)$. Therefore the EF revenue is $\frac{2n}{n+1}v \cdot \frac{n}{n+1} \cdot (n+1) = \frac{2n^2}{n+1}v$. To calculate the IC revenue, with probability $n/(n+1)$, the big agent is assigned to the size n set, and every of the n winning agents pays the reserve $\frac{2n}{n+1}v$. Also with probability $1/(n+1)$, the big agent is assigned to the singleton set, and every agent has to pay her own value, which sums up to $\Theta(nv\log(n))$. Therefore the IC revenue is $\frac{2n}{n+1}v \cdot \frac{n}{n+1} \cdot n + \frac{1}{n+1}\Theta(nv\log(n))$, which is larger than EF revenue for sufficiently large n. □

C Benchmarks for Prior-free Mechanism Design

Prior-free mechanism design looks for the mechanism that minimizes, over valuation profiles, its worst case ratio to a given performance benchmark. For k-unit auctions, two benchmarks have been considered in the literature. Fiat et al. [11] proposed "$\mathcal{F}^{(2)}$" as "the optimal revenue from single price sale to between 2 and k agents" which has revenue $\max_{2 \leq i \leq k} \mathrm{R}^{\mathbf{v}}(i)$. Hartline and Roughgarden [18] proposed "$\mathcal{G}^{(2)}$" as "the supremum of Bayesian optimal mechanisms (with at least 2 winners)" which has revenue $\sup_F \mathrm{ICO}^F(\mathbf{v}^{(2)})$. Lemma 5.2 allows us to bound this below $\max_{2 \leq i \leq k} \bar{\mathrm{R}}^{\mathbf{v}}(i)$. Devanur and Hartline [7] shows that in fact these two benchmarks are within a factor of two of each other, for k unit auctions.

For more general settings Hartline and Roughgarden [19] showed that the Vickrey-Clarke-Groves (VCG) mechanism with a single reserve price can approximate the Bayesian optimal mechanism. This holds for downward closed set systems and i.i.d. distributions satisfying a standard "monotone hazard rate" condition and for matroid set systems and i.i.d. distributions satisfying a less restrictive "regularity" condition.[3] This motivates considering as a benchmark for prior-free mechanism design in general set systems "the revenue of VCG with optimal single reserve price", a generalization of $\mathcal{F}^{(2)}$. For instance, if a mechanism approximates this benchmark then for all monotone hazard rate distributions, the mechanism's expected revenue approximates the optimal mechanism's revenue. For matroid settings there is such a mechanism [19].

Unfortunately, for irregular distributions and matroid set systems VCG with a single reserve price may be very far from optimal. This is formalized by Lemma C.1, below. If our goal is a bench-

[3] A distribution is *regular* if $\mathrm{R}^F(\cdot)$ is concave.

mark to which approximation by a mechanism implies that for all distributions, the mechanism approximates the Bayesian optimal mechanism, then this benchmark is not good enough.

Lemma C.1 *For every sufficiently large n, there exists a distribution F and a matroid set system of n elements such that the expected revenue of VCG with any single reserve price is an $\Omega(\log n/\log\log n)$-approximation to the expected revenue of* ICO^F.

Lemma 5.2 shows that the revenue of the optimal envy-free pricing, EFO, is a benchmark for which prior-free approximation implies Bayesian approximation. This is summarized formally by the following theorem.

Theorem C.2 *For a symmetric matroid setting, if $\mathcal{M}$ is a β-approximation to* $\mathrm{EFO}(\mathbf{v})$ *for all* $\mathbf{v}$*, then for all F,*

$$\mathbf{E}_{\mathbf{v}\sim F}[\mathcal{M}(\mathbf{v})] \geq \tfrac{1}{\beta}\mathbf{E}_{\mathbf{v}\sim F}\left[\mathrm{ICO}^F(\mathbf{v})\right].$$

Unfortunately, for reasons discussed in [12], we can not approximate $\mathrm{EFO}(\mathbf{v})$ when the highest valuation is very large. Similar to the case of digital goods auctions [13], here we shall use the slightly weaker benchmark $\mathrm{EFO}(\mathbf{v}^{(2)})$ instead.

C.1 Proof of Lemma C.1

Proof: Fix some number m.

We first define the matroid in question. For each $k \in \{1,\ldots,m\}$, a type k graph contains two nodes, and m^{3k-1} parallel edges connecting these two nodes. The matroid is a graphic matroid where the graph contains m^{2m-2k} disjoint copies of type k graphs for each $k \in \{1,\ldots,m\}$. So total number of agents n is at most $m^{O(m)}$. Hence m is at least of order $\frac{\log n}{\log\log n}$.

Next we define the "sydney opera house distribution". The distribution F is such that the value is distributed according to uniform distribution $[m^{2k+1}-\epsilon, m^{2k+1}+\epsilon]$ with probability $\frac{1}{m^{3k}} - \frac{1}{m^{3k+3}}$ for $k \in \{0,\ldots,m-1\}$, and with probability $\frac{1}{m^{3k}}$ for $k = m$. Here we take ϵ to be some sufficiently negligible positive amount, and we will often omit ϵ related terms. So for each k the revenue function R at $\frac{1}{m^{3k+3}}$ has left limit $R(\frac{1}{m^{3k+3}}-) = \frac{m^{2k+3}}{m^{3k+3}} = \frac{1}{m^k}$, and right limit $R(\frac{1}{m^{3k+3}}+) = \frac{m^{2k+1}}{m^{3k+3}} = \frac{1}{m^{k+2}}$. Hence the ironed virtual valuation between quantile $\frac{1}{m^{3k+3}}$ to quantile $\frac{1}{m^{3k}}$ is $\frac{\frac{1}{m^{k-1}}-\frac{1}{m^k}}{\frac{1}{m^{3k}}-\frac{1}{m^{3k+3}}} \approx m^{2k+1}$. Note that the ironed virtual valuation is equal to valuation, ignoring minor terms.

To calculate the revenue of Myerson's auction, for a type k graph, there are m^{3k-1} agents. With probability at least $1-(1-\frac{1}{m^{3k}})^{m^{3k-1}} \approx \frac{1}{m}$, the highest agent is in quantile range $(0, \frac{1}{m^{3k}})$, with ironed virtual valuation at least m^{2k+1}. So the expected ironed virtual valuation from a type k graph is at least m^{2k}. Multiplied by the number of type k graphs, the total ironed virtual valuation, and hence expected revenue is at least $\sum_k m^{2k}\cdot m^{2m-2k} = m\cdot m^{2m}$.

To calculate the revenue of VCG with some reserve r, suppose w.l.o.g. $r \approx m^{2k'+1}$ for some k'. For a type k graphs with m^{3k-1} agents, the dominant amount of revenue is obtained from the following two cases:

1. When there are at least two agents with value at least $m^{2k+1}-\epsilon$ (i.e. in quantile $\frac{1}{m^{3k}}$), the lower of which has value at most $m^{2k+1}+\epsilon$. This happens with probability roughly $\frac{1}{m^2}$, and gives revenue m^{2k+1}. Therefore the expected revenue we get from this case is $\frac{1}{m^2}\cdot m^{2k+1}$, which multiplied by the number of type k graphs, is $O(m^{2m-1})$.

2. When $k = k'$, and there is at least one agent who beats the reserve m^{2k+1}. This happens with probability at most $\frac{1}{m}$. Therefore the expected revenue from this case is $m^{2k+1}\cdot\frac{1}{m}$, which multiplied by the number of type k graphs is $O(m^{2m})$.

Summing over all k, the total expected revenue of VCG with reserve r is at most $O(m^{2m}) + m \cdot O(m^{2m-1}) = O(m^{2m})$, which is less than that of Myerson's auction by a factor of $\Omega(m) = \Omega(\log n / \log\log n)$. $\square$

D Analysis of RSEM'

D.1 Effective Ironing

The RSEM' mechanism applies the ironed virtual surplus optimizer for the sample to the market, which is sub-optimal for the market. To capture the effect of applying a sub-optimal ironed virtual surplus optimizer, we introduce the notion of effective revenue curve.

Definition D.1 (effective revenue curve) *Let* $\mathbf{v}$ *be a valuation profile and* $\bar{\Phi}$ *be an ironed virtual valuation function. Group the agents with equal nonnegative* $\bar{\Phi}$ *values together into consecutive "equal priority" classes* $\{1, \ldots, n_1\}, \{n_1+1, \ldots, n_2\}, \ldots, \{n_{t-1}+1, \ldots, n_t\}$. *Let the effective revenue curve* $\widetilde{\mathrm{R}}$ *be obtained from* $\mathrm{R}^{\mathbf{v}}$ *by connecting the points of* $\mathrm{R}^{\mathbf{v}}$ *at* n_j*'s linearly (a.k.a. ironing) and extending horizontally beyond* n_t, *i.e.,* $\widetilde{\mathrm{R}}(i) = \mathrm{R}^{\mathbf{v}}(n_{j-1}) \cdot \frac{n_j - i}{n_j - n_{j-1}} + \mathrm{R}^{\mathbf{v}}(n_j) \cdot \frac{i - n_{j-1}}{n_j - n_{j-1}}$ *for* $n_{j-1} \le i \le n_j$ *and* $1 \le j \le t$ *(let* $n_0 = 0$*), and* $\widetilde{\mathrm{R}}(i) = \widetilde{\mathrm{R}}(n_t)$ *for all* $i \in \{n_t + 1, \ldots, n\}$.

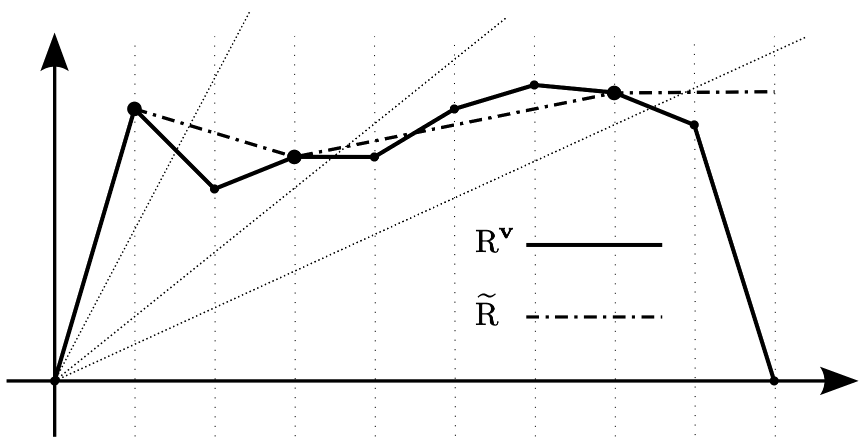

Figure 3: Effective revenue curve

Figure 3 depicts an example of the effective revenue curve. The three rays divide the first orthant into four regions. For every region, every point (i, y) in the region (which corresponds to value y/i) has the same $\bar{\Phi}$ value, i.e., $\bar{\Phi}(y/i)$ is the same. In particular, the points on the revenue curve $\mathrm{R}^{\mathbf{v}}$ in each region correspond to the same $\bar{\Phi}$ value, get "ironed". Our definition of effective revenue curve ensures the following:

Lemma D.1 *For all valuation profile* $\mathbf{v}$, $\mathrm{EF}^{\bar{\Phi}}(\mathbf{v}) = \sum_{i=1}^{n} \widetilde{\mathrm{R}}(i) \cdot (x_i^{\bar{\Phi}}(\mathbf{v}) - x_{i+1}^{\bar{\Phi}}(\mathbf{v}))$.

Proof:

$$\begin{aligned} \mathrm{EF}^{\bar{\Phi}}(\mathbf{v}) &= \sum_{i=1}^{n} \mathrm{R}^{\mathbf{v}}(i) \cdot (x_i^{\bar{\Phi}}(\mathbf{v}) - x_{i+1}^{\bar{\Phi}}(\mathbf{v})) \\ &= \sum_{i=1}^{n} \mathrm{R}^{\mathbf{v}^{\bar{\Phi}}}(i) \cdot (x_i^{\bar{\Phi}}(\mathbf{v}) - x_{i+1}^{\bar{\Phi}}(\mathbf{v})) \end{aligned}$$

Here the first equality is by Lemma 3.4. To justify the second equality, note that whenever $\mathrm{R}^{\mathbf{v}^{\bar{\Phi}}}(i) \neq \mathrm{R}^{\mathbf{v}}(i)$, there are two cases: (1) i is in $\{n_{j-1}+1, \ldots, n_j - 1\}$ for some j, and so v_i and

v_{i+1} have the same $\bar{\Phi}$ value, and hence $x_i^{\bar{\Phi}}(\mathbf{v}) = x_{i+1}^{\bar{\Phi}}(\mathbf{v})$. (2) i is bigger than n_t, and so v_i and v_{i+1} both have negative $\bar{\Phi}$ value, and hence $x_i^{\bar{\Phi}}(\mathbf{v}) = x_{i+1}^{\bar{\Phi}}(\mathbf{v}) = 0$.

□

D.2 Balanced Sampling and Revenue Curves

For an agent set S, we abuse notation to let $\mathbf{v}_S$ denote $(\mathbf{v}_S, \mathbf{0}_{N-S})$, i.e., the valuation profile (of n agents) obtained from $\mathbf{v}$ by decreasing the values of agents outside S to 0. A useful property that holds for envy-free revenue is "subadditivity", in the following sense.

Lemma D.2 *For a valuation profile* $\mathbf{v}$ *and two disjoint sets* A *and* B *of agents,* $\mathrm{EFO}(\mathbf{v}_A) + \mathrm{EFO}(\mathbf{v}_B) \geq \mathrm{EFO}(\mathbf{v}_{A\cup B})$.

Proof: Recall that $\mathrm{EFO}(\mathbf{v}_{A\cup B})$ is the maximum revenue we can get from $A \cup B$ subject to the envy free constraints. Let agents in B contribute total revenue R to $\mathrm{EFO}(\mathbf{v}_{A\cup B})$. By setting the agents in A to have zero valuations to obtain valuation profile $\mathbf{v}_B$, we basically removed envy-free constraints between agents in A and agents in B. With less envy free constraints, the maximum envy-free revenue we can get from B, i.e., $\mathrm{EFO}(\mathbf{v}_B)$, can only be larger. Similarly, the total revenue that A contributes to $\mathrm{EFO}(\mathbf{v}_{A\cup B})$ is at most $\mathrm{EFO}(\mathbf{v}_A)$, and our lemma follows. □

Let $\mathbf{v}_N^S$ be a short-hand for $(\mathbf{v}_N)^{\bar{\Phi}^S}$, which is the effective valuation profile obtained from ironing $\mathbf{v}$ using the ironed virtual valuation function $\bar{\Phi}^S$ for $\mathbf{v}_S$. The following lemma says that if we apply the ironed virtual valuation function $\bar{\Phi}^S$ for $\mathbf{v}_S$ to both the whole set N and samples S, then the effective revenue curve we get in the former case vertically dominates what we get in the latter case.

Lemma D.3 *For all* $1 \leq i \leq n$, $\mathrm{R}^{\mathbf{v}_N^S}(i) \geq \mathrm{R}^{\mathbf{v}_S^S}(i)$.

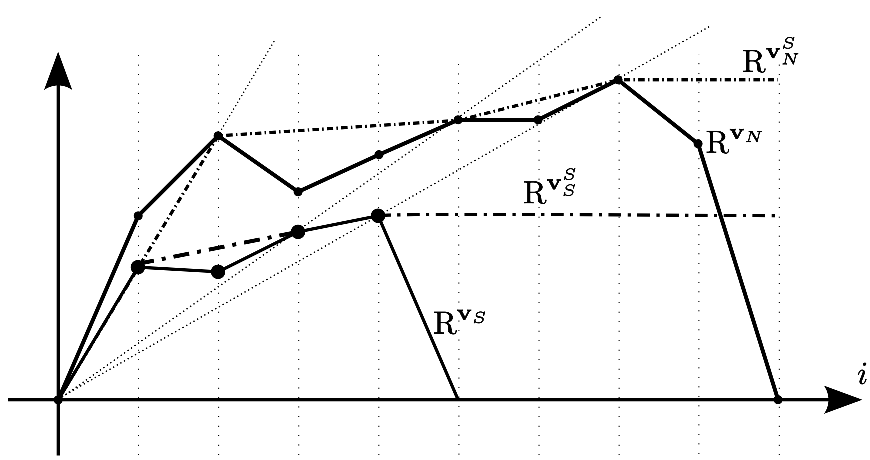

Figure 4: $\mathrm{R}^{\mathbf{v}_N^S}(i)$ dominates $\mathrm{R}^{\mathbf{v}_S^S}(i)$

Proof: The readers are referred to Figure 4 for an intuitive view of the relationship between the revenue curves, where revenue curves are piece-wise linearly interpolated between the integer points. Observe that revenue curve $\mathrm{R}^{\mathbf{v}_N}$ dominates $\mathrm{R}^{\mathbf{v}_S}$ in the sense that for every slope t, the intersection of the ray $y = tx$ with $\mathrm{R}^{\mathbf{v}_N}$ is farther away from the origin than its intersection with $\mathrm{R}^{\mathbf{v}_S}$. Transforming $\mathrm{R}^{\mathbf{v}_N}$ and $\mathrm{R}^{\mathbf{v}_S}$ to the effective revenue curves using the same ironed virtual valuation function $\bar{\Phi}^S$ do not change such dominance relationship, and moreover, because $\mathrm{R}^{\mathbf{v}_S^S}$ is non-decreasing and concave, it follows that vertical dominance also holds, i.e., $\mathrm{R}^{\mathbf{v}_N^S}(i) \geq \mathrm{R}^{\mathbf{v}_S^S}(i)$ for all i. $\square$

Definition D.2 *We say that the partitioning* (S, M) *of* $N = \{1, \ldots, n\}$ *is balanced if for all* $i \in \{1, \ldots, n\}$, $Y_i \leq 3i/4$ *where* $Y_i = |\{1, 2, \ldots, i\} \cap S|$.

We will focus on the case that agent 1 is in M, and agent 2 is in S. Recall that $\mathbf{v}_S$ denote $(\mathbf{v}_S, \mathbf{0}_{N-S})$.

Lemma D.4 *Conditioning on that* $1 \in M, 2 \in S$, *a random partitioning* (S, M) *of* N *is balanced with probability at least* 0.8.

Proof: Feige et al. [10] proved that conditioning on that $1 \in M$, (S, M) is balanced with probability at least 0.9. Let B denote the event that (S, M) is balanced. Clearly $0.9 \leq Pr[B \mid 1 \in M] = 0.5 \cdot Pr[B \mid 1 \in M, 2 \in S] + 0.5 \cdot Pr[B \mid 1 \in M, 2 \notin S]$. It follows that $Pr[B \mid 1 \in M, 2 \in S] \geq 0.8$. $\square$

The following consequence of balanced partitioning is useful, and easy to check.

Lemma D.5 *Given a balanced partitioning* (S, M), *for every non-increasing sequence* $a_1, a_2, \ldots, a_n$ *of nonnegative reals, for all* $i \in \{1, \ldots, n\}$, *we have* $\sum_{j \in M \cap \{1, \ldots, i\}} a_j \geq \frac{1}{4} \sum_{j \in \{1, \ldots, i\}} a_j$.

Definition D.3 *We say that the partitioning* (S, M) *of* N *is double-side balanced if for all* $i \in \{3, \ldots, n\}$, $i/4 \leq Y_i \leq 3i/4$ *where* $Y_i = |\{1, 2, \ldots, i\} \cap S|$.

Lemma D.6 *Conditioning on that* $1 \in M, 2 \in S$, *a random partitioning* (S, M) *of* N *is double-side balanced with probability at least* 0.6.

Proof: Lemma D.4 implies that with probability at least 0.8, for all $i \in \{3, \ldots, n\}$, $Y_i \leq 3i/4$. Symmetrically, with probability at least 0.8, for all $i \in \{3, \ldots, n\}$, $i/4 \leq Y_i$. The probability that both of these happen is at least $1 - (1 - 0.8) - (1 - 0.8) = 0.6$. $\square$

Let $\phi_i = \bar{\Phi}^S(v_i)$ for $i \in \{1, \ldots, n\}$ be the ironed virtual valuation values w.r.t. S of agents in N. For $i \in \{1, \ldots, n\}$, define $\widehat{\mathrm{R}}(i)$ as $\sum_{j=1}^{i} \phi_j$, and let $\widehat{\mathbf{v}}$ be the valuation profile corresponding to $\widehat{R}$, i.e. $\widehat{\mathbf{v}}_i = \widehat{R}(i)/i$. Compare running x^S on $\mathbf{v}_N$ with running $\mathbf{x}^{\widehat{\mathbf{v}}}$ on $\widehat{\mathbf{v}}$, the ironed virtual valuation of agent i in either case is equal to ϕ_i. Therefore these two ironed virtual surplus optimizers will choose the same allocation, and hence $x_i^S(\mathbf{v}_N) = x_i^{\widehat{\mathbf{v}}}(\widehat{\mathbf{v}})$.

Lemma D.7 *Given a double-side balanced partitioning* (S, M). *We have* $\mathrm{R}^{\mathbf{v}_S^S}(i) \geq \frac{1}{4} \mathrm{R}^{\widehat{\mathbf{v}}}(i) \geq \frac{1}{4} \mathrm{R}^{\mathbf{v}_S^S}(i)$ *for all* $1 \leq i \leq n$.

Proof: For each i, $\mathrm{R}^{\widehat{\mathbf{v}}}(i) = \sum_{j=1}^{i} \phi_j$ and $\mathrm{R}^{\mathbf{v}_S^S}(i)$ is the sum of the i largest ϕ_j values with $j \in S$. Clearly $\mathrm{R}^{\widehat{\mathbf{v}}}(i) \geq \mathrm{R}^{\mathbf{v}_S^S}(i)$. Since (S, M) is double-side balanced, applying Lemma D.5, we also have that for all i, $\mathrm{R}^{\mathbf{v}_S^S}(i) \geq \frac{1}{4} \mathrm{R}^{\widehat{\mathbf{v}}}(i)$. $\square$

D.3 Analysis of RSEM'

Lemma D.8 *For any downward-closed permutation setting, any valuation profile* $\mathbf{v}$, *and double-side balanced partitioning* (S, M), $\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{4} \mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_S)$.

Proof: Let $\mathbf{x}^S$ and $\mathbf{x}^{\widehat{\mathbf{v}}}$ be short-hands for $\mathbf{x}^{\bar{\Phi}^S}$ and $\mathbf{x}^{\bar{\Phi}^{\widehat{\mathbf{v}}}}$, respectively. I.e., they are the ironed virtual surplus optimizers with ironed virtual valuation functions defined for $\mathbf{v}_S$ and $\widehat{\mathbf{v}}$, respectively.

The proof is by the following inequalities:

$$\begin{aligned}
\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) &= \sum\nolimits_i \mathrm{R}^{\mathbf{v}_N^S}(i)\cdot(x_i^S(\mathbf{v}_N) - x_{i+1}^S(\mathbf{v}_N)) = \sum\nolimits_i \mathrm{R}^{\mathbf{v}_N^S}(i)\cdot(x_i^{\widehat{\mathbf{v}}}(\widehat{\mathbf{v}}) - x_{i+1}^{\widehat{\mathbf{v}}}(\widehat{\mathbf{v}})) \\
&\geq \tfrac{1}{4}\cdot\sum\nolimits_i \mathrm{R}^{\widehat{\mathbf{v}}}(i)\cdot(x_i^{\widehat{\mathbf{v}}}(\widehat{\mathbf{v}}) - x_{i+1}^{\widehat{\mathbf{v}}}(\widehat{\mathbf{v}})) \geq \tfrac{1}{4}\cdot\sum\nolimits_i \mathrm{R}^{\widehat{\mathbf{v}}}(i)\cdot(x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)) \\
&\geq \tfrac{1}{4}\cdot\sum\nolimits_i \mathrm{R}^{\mathbf{v}_S^S}(i)\cdot(x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)).
\end{aligned}$$

Here the first equality is by effective ironing (Lemma D.1), and the second equality is by that $x_i^S(\mathbf{v}_N) = x_i^{\widehat{\mathbf{v}}}(\widehat{\mathbf{v}})$. The first inequality is by Lemma D.3 and Lemma D.7, the second inequality is by the optimality of $\mathbf{x}^{\widehat{\mathbf{v}}}$ for $\widehat{\mathbf{v}}$, and the third inequality is by Lemma D.7 again. □

Now we can establish the performance guarantee for RSEM' in the downward-closed case.

Proof: (of Theorem 6.2) With probability 1/4, agent 1 is in M, and agent 2 is in S. Conditioning on this, by Lemma D.4, the partitioning (S, M) is double-side balanced with probability 0.6, and by Lemma D.2 and symmetry, $\mathrm{EFO}(\mathbf{v}_S) \geq \frac{1}{2}\mathrm{EFO}(\mathbf{v}_N)$ with probability 0.5. Both of these events happen with probability at least $1-(1-0.6)-(1-0.5) = 0.1$. Assume both events happen. As in the matroid case, we have $\sum_{i\in M} \mathrm{IC}_i^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{8}\mathrm{EF}^{\mathbf{x}^S}(\mathbf{v}_N)$. Together with that $\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{4}\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_S) \geq \frac{1}{8}\mathrm{EFO}(\mathbf{v}^{(2)})$, we have $\sum_{i\in M}\mathrm{IC}_i^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{64}\mathrm{EFO}(\mathbf{v}_N)$, and our theorem follows by multiplying the ratio with the probabilities. □