

Envy, Truth, and Profit

Jason D. Hartline^{*}

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

Qiqi Yan[†]

Department of Computer Science
Stanford University
Stanford, CA, USA
qiqiyan@cs.stanford.edu

ABSTRACT

We consider profit maximizing (incentive compatible) mechanism design in general environments that include, e.g., position auctions (for selling advertisements on Internet search engines) and single-minded combinatorial auctions. We analyze optimal envy-free pricings in these settings, and give economic justification for using the optimal revenue of envy-free pricings as a benchmark for prior-free mechanism design and analysis. Moreover, we show that envy-free pricing has a simple nice structure and a strong connection to incentive compatible mechanism design, and we exploit this connection to design prior-free mechanisms with strong approximation guarantees.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General

General Terms

Economics, Theory, Algorithms

Keywords

envy-free pricing, truthful mechanisms, revenue maximization, optimal auction, prior-free

1. INTRODUCTION

Mechanism design theory provides guidelines for the design of economic systems that obtain good performance in the presence of selfish behavior on the part of the participants. A standard approach is to restrict attention to mechanisms where each participant has a dominant strategy of truthfully reporting her preference over possible outcomes. This restriction imposes a constraint, known as *incentive*

compatibility, on the mechanism that binds across all possible misreports of the agents. Reasoning about mechanisms, therefore, requires reasoning about the outcomes of the mechanism on all possible preferences, not just the true ones. A mechanism, therefore, may face a performance tradeoff across different possible preferences. For instance, there is such a performance tradeoff for the objective of profit maximization. Consequently, no mechanism obtains the maximum profit for all possible participant preferences.

In contrast to incentive compatibility, consider the “fairness” constraint of *envy-freedom*. The envy-free constraint requires that each participant prefers her own outcome to that of any other participant. Unlike incentive compatibility, envy-freedom is a constraint that binds point-wise on the true preferences. Consequently, for any performance metric, there is an envy-free outcome with the best performance.

As an example, suppose there are two identical units of an item for sale. Each agent wants one unit of the item, and her utility for the unit is her value for the unit minus the price she pays. There is one high-value agent with value 6 for a unit, and two low-value agents with value 4. It is envy-free to price each unit at 5, the high-value agent would opt to buy a unit, while others would not. It is not envy-free to price each unit at 2, as all three agents would want a unit, exceeding the supply.

Our goal is to design incentive compatible mechanisms that obtain good revenue point-wise, i.e., for any preferences of the participants. This contrasts from the standard *Bayesian* approach for optimal mechanism design in economics in which it is assumed that participant preferences are drawn from a known *prior* distribution. The knowledge about the distribution allows the designer to optimally trade-off revenue between different possible preferences of the participants. In contrast to the economics approach, our goal is to design good mechanisms without distributional assumptions, i.e., to be *prior free*.

Our approach to design good prior-free mechanisms is the following. We view envy-freedom (EF) as a relaxation of incentive compatibility (IC), characterize revenue-optimal envy-free pricings, relate the revenue of an IC mechanism to that of a corresponding EF pricing, and then use this connection between EF and IC revenues to design and analyze IC mechanisms.

We consider general single-parameter environments where each participant (a.k.a., agent) has a single value for receiving an abstract service and the designer has a feasibility constraint on which agents can be served simultaneously. In the previously mentioned multi-unit example, the abstract

^{*}Supported in part by NSF CAREER Award CCF-0846113.

[†]Supported by a Stanford Graduate Fellowship.

service is a unit of the item and the feasibility constraint requires we serve sets of cardinality at most two.

An environment is *downward-closed* if any subset of a feasible set is also feasible. The following special cases are of interest. In a *digital good environment* all allocations are feasible. In a *multi-unit environment* there is a fixed number k of identical units available and each agent desires one unit, i.e., the feasible sets are those with cardinality at most k . In a *matroid environments* the feasible allocations correspond to the independent sets of a matroid. Matroid environments include multi-unit and constrained matching environments as special cases.¹

We allow convex combinations of our basic feasibility environments. For instance, one can express *position auction environments*, as popularized by the auctioning advertising slots on Internet search engines, as a convex combination of multi-unit auctions. Position auction environments are specified by a decreasing sequence of probabilities for each position and a feasible outcome is a partial assignment of agents to positions. An agent assigned to a position receives a service with probability equal to the position’s associated probability.

For many interesting single-parameter environments, e.g., position auctions, the feasibility constraint is inherently symmetric. Given a feasible allocation, for any permutation of the names of the agents, the permuted assignment is still feasible. As an example of an asymmetric environment, consider single-minded combinatorial auctions, i.e., where each agent has a single value for receiving a particular bundle of goods (and no value for any other bundle). In such an environment, an agent is considered “served” if she receives her desired bundle. Therefore, the feasible subsets of agents are the ones for which the desired bundles are disjoint. For such asymmetric feasibility constraints, the notion of envy-freeness as a fairness condition is ill-defined. However, we can symmetrize any asymmetric environment by assuming that the agents’ roles with respect to the feasibility constraint are assigned by a random permutation. In the combinatorial auction example, this means that we start with a set of agent values and a set of desired bundles and the values are associated with the desired bundle via a random permutation. We refer to these environments as *permutation environments* and observe that the random permutation assumption on values is a natural prior-free analog of the i.i.d. assumption that is common in Bayesian mechanism design.

The first result of this paper is a characterization of revenue-optimal envy-free pricings. Importantly, this characterization mirrors the characterization of Bayesian revenue-optimal incentive compatible mechanisms of Myerson (for i.i.d. distributions). Optimal envy-free pricings are *ironed virtual surplus optimizers* in the sense that they are specified by a monotone non-decreasing function (known as the *ironed virtual valuation function*) and the agents served are the feasible set for which the total ironed virtual value is the highest. Furthermore, the ironed virtual valuation function that obtains this optimal revenue is the one that corresponds to the *empirical distribution* of the given agents’ values.

The above description of ironed virtual surplus optimizers

¹A constrained matching environment is one where each agent has a value for “good” items and a set of items that are considered good. The agent wants at most one item, and each item can be allocated to at most one agent.

specifies how to decide which agents should receive services. It does not explicitly specify the payments. In fact, incentive compatibility and envy-freedom generally require different payments. This is because incentive compatible payments are counteracting agents’ incentive to misreport while envy-free payments are counteracting agents’ envy (or desire to swap). As these are different constraints, they lead to different payments. Our second result shows that for any ironed virtual surplus optimizer these payments are, in settings of interest,² not too far apart, specifically: $EF \geq IC \geq EF/2$.

The first inequality above motivates the selection of the envy-free optimal revenue as a benchmark to which to compare prior-free mechanisms. Bayesian optimal IC mechanisms are ironed virtual surplus maximizers and their IC revenue is point-wise at most their EF revenue. Therefore, a prior-free mechanism that always approximates the optimal EF revenue is guaranteed to also approximate the revenue of any Bayesian optimal mechanism. Of course, the Bayesian optimal mechanism is the one that a designer would select if a prior were known, and so approximating it without any knowledge is a strong guarantee.

As a first example to illustrate the relevance of envy-freedom for prior-free mechanism design, we reduce the problem of approximating the optimal EF revenue for matroid permutation and position auction environments to multi-unit environments. This reduction is enabled by the fact that incentive compatibility is closed under convex combination and a structural equivalence between (1) the convex combinations of multi-unit auction environments, (2) position auction environments, and (3) matroid permutation environments. This implies that for any IC mechanism that is a β -approximation to the optimal EF revenue for multi-unit environments, there is an IC β -approximation mechanism for position auction environments and matroid permutation environments. In particular, results of [14, 8] imply a 6.5-approximation for these settings.

The second example we give considers general downward-closed environments such as single-minded combinatorial auctions. In such an environment we show that a prior-free incentive compatible mechanism based on random sampling approximates the optimal EF revenue (in permutation environments). The random sampling auction we consider works as follows. It first partitions the agents into a market set and a sample set. Then it calculates the ironed virtual value function for the empirical distribution of the sample and simulates ironed virtual value maximization with this function on the full set of agents. Finally, agents in the sample are rejected and agents in the market receive the allocation from the simulation and are charged the corresponding payments. To analyze this mechanism, we show that its EF revenue (as an ironed virtual surplus maximizer) approximates the optimal EF revenue, and consequently its IC revenue also approximates the optimal EF revenue.

Related work. Bayesian optimal mechanisms for single-parameter environments were characterized as ironed virtual surplus optimizers by Myerson [18]. The relationship between revenue curves and virtual valuations, a.k.a., marginal revenue, was identified by Bulow and Roberts [5]. We make heavy use of this machinery.

²The first inequality is known to hold for position auction and matroid permutation environments. The second inequality is known to hold under a technical condition.

This paper follows from a line of work that studies prior-free revenue properties of the random sampling auction of [12]. The tightest analysis of the random sampling auction for digital good environments is given by Alaei et al. [1]. For multi-unit environments, Hartline and Roughgarden [15] proposed a benchmark for prior-free analysis that is derived from Bayesian optimal auctions and discuss the desirability of a prior-free benchmark that implies approximation in the Bayesian setting. With this benchmark, Devanur and Hartline [8] extended the analysis from [1] to limited supply settings. In this context, the present paper further extends the benchmark of [15] and the analysis of [8] to general downward-closed environments.

It has been observed that position auctions are structurally equivalent to a convex combination of multi-unit auctions. For instance in Dughmi et al. [10], this equivalence was leveraged to show that the VCG mechanism with reserve gives the same revenue in these two environments. We extend this connection by doing the same thing for less well-behaved approximation mechanisms. We also make a similar connection between position auction and matroid permutation environments.

The prior-free mechanisms we discuss perform well even when there is no prior distribution; agent values can be adversarially chosen. If there is a prior distribution that is unknown it is possible to design good *prior-independent* mechanisms. Dhangwatnotai et al. [9] show that a mechanism based on a “single-sample” approach gives an 8-approximation for downward-closed environments with distributions that satisfy a standard *monotone hazard rate* condition, and a 2-approximation for structurally nicer matroid environments with 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 environments with nearly general i.i.d. valuation distributions. This is important as many distributions, e.g., bimodal, do not satisfy the regularity condition required by [9].

Connections between envy-free pricings and prior-free mechanism design have been made before (e.g. in [13, 3]). Much of the work on envy-free pricing has been focused on multi-dimensional agent preferences. Part of our characterization of envy-free pricings in single-parameter environments is a special case of Mu’alem’s “local efficiency” condition [17].

This paper uses envy-free pricings as a benchmark for gauging the performance of a prior-free mechanism. While it might be nice to have mechanisms that are both incentive compatible and envy-free, achieving both conditions together while obtaining a reasonable performance is impossible. This is discussed for profit maximization in digital good environments by Goldberg and Hartline [11], unrelated machine scheduling problems by Cohen et al. [6], and for welfare maximization in general combinatorial auctions by Ausubel and Milgrom [2]. In fact, Day and Milgrom [7] suggest that envy-freeness and related conditions (specifically the cooperative game theory condition of the *core*) are more important than incentive compatibility and when they are impossible to achieve together envy freedom should be chosen instead of incentive compatibility.

2. OPTIMAL ENVY-FREE PRICING

In this section we derive a theory of optimal envy-free pric-

ings in single-dimensional environments that mirrors that of Bayesian optimal (incentive compatible) mechanisms for i.i.d. prior distributions [18, 5].

There are $n \geq 2$ agents. Each agent i has a valuation v_i for receiving an abstract service. The *valuation profile* is $\mathbf{v} = (v_1, \dots, v_n)$. We assume that the agents are indexed in order of decreasing values, i.e., $v_i \geq v_{i+1}$. An agent i who is served with probability x_i and charged price p_i obtains utility $u_i = v_i x_i - p_i$. *Individual rationality* requires that u_i be non-negative.

We are allowed to serve certain feasible sets of agents as given by a set system. The set system is *downward-closed* in the sense if a set of agents is feasible, so is any of its subsets. The empty set is always feasible. We allow randomization in two senses (1) the set system constraint may be randomized (i.e., by convex combination over set systems) and (2) the set of agents served may be random (by convex combination over feasible sets). Notably, randomization in (1) is given by the environment and randomization in (2) is by our choice of outcome. We define an *allocation* as a vector $\mathbf{x} = (x_1, \dots, x_n) \in [0, 1]^n$ where x_i is the probability that agent i is served. An allocation is feasible if it is the characteristic vector induced by the process above. The environments permitted include digital good auctions, multi-unit auctions, position auction environments, matroid environments, and single-minded combinatorial auctions.

We further assume in this section that feasibility constraint imposed by the environment is symmetric, i.e., the set of feasible allocations is closed under permutation. Digital good, multi-unit auction, and position auction environments are all symmetric. Given any asymmetric environment its corresponding *permutation environment* is obtained by randomly permuting the agents with respect to the feasibility constraint. Of special interest for us will be downward closed permutation environments and matroid permutation environments. By definition these environments are symmetric.

DEFINITION 2.1 (ENVY FREEDOM). *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,*

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

We first characterize envy-free pricings in terms of the allocation. For a given allocation \mathbf{x} there may be several pricings \mathbf{p} for which the allocation is envy-free. Since our objective is profit maximization we will characterize the \mathbf{p} corresponding to \mathbf{x} that gives the highest total revenue. We omit the proof of this characterization as it is nearly identical to that of the analogous (and standard) characterization of incentive compatible mechanisms.

DEFINITION 2.2. *An allocation is swap monotone if the allocation probabilities have the same order as the valuations of the agents, i.e., $x_i \geq x_{i+1}$ for all i .*

LEMMA 2.1. *In symmetric environments, an allocation \mathbf{x} admits a non-negative and individual rational payment \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 which (\mathbf{x}, \mathbf{p}) is envy-free satisfy, for all i , (Figure 1)*

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




Figure 1: The solid curve depicts a swap monotone allocation as a function of the values (points). The shaded area corresponds to payment from Lemma 2.1 for agent i .

Importantly, the above characterization leaves us free to speak of the (maximum) envy-free revenue of any swap monotone allocation \mathbf{x} on values \mathbf{v} , which we denote by $\text{EF}^{\mathbf{x}}(\mathbf{v})$. For any \mathbf{v} and any symmetric environment we now solve for the envy-free optimal revenue, denoted by $\text{EFO}(\mathbf{v})$.




Figure 2: R and \bar{R} are the revenue curve and ironed revenue curve of the valuation profile $(6, 4, 4)$. The ironed virtual value of the high-value agent is 6, and the ironed virtual value of the two low-value agents are both $(12 - 6)/2 = 3$. E.g., the optimal EF revenue in the $k = 2$ unit environment is $\bar{R}(2) = 9$.

We will characterize the envy-free optimal revenue in terms of properties of the valuation profile \mathbf{v} . Given a valuation profile \mathbf{v} we denote the *revenue curve* by $R^{\mathbf{v}}(i) = i \cdot v_i$ for $i = \{1, \dots, n\}$ (recall v_i 's are indexed in decreasing order). For convenience we also let $R^{\mathbf{v}}(0) = R^{\mathbf{v}}(n+1) = 0$. The *ironed revenue curve*, denoted $\bar{R}^{\mathbf{v}}(i)$, is the minimum concave function that upper-bounds R . Likewise, define the *virtual valuation function* $\Phi^{\mathbf{v}}(v) = R^{\mathbf{v}}(i) - R^{\mathbf{v}}(i-1)$ and the *ironed virtual valuation function* $\bar{\Phi}^{\mathbf{v}}(v) = \bar{R}^{\mathbf{v}}(i) - \bar{R}^{\mathbf{v}}(i-1)$, where $i \in \{1, \dots, n+1\}$ is such that $v \in [v_i, v_{i-1})$. (We set $v_0 = \infty$ for notational convenience.) See Figure 2.

$R^{\mathbf{v}}(i)$ is the best envy-free revenue one can get from serving exactly i agents at the same price deterministically. Again consider the multi-unit auction example with two units of an item, one high-value agent with value 6, and two low-value agents with value 4. It is envy-free to serve one high-value agent and one low-value agent at price 4, achieving revenue $R(2) = 8$. Interestingly, this is not optimal. The following allocation and payments are also envy-free: serve the high-value agent with probability 1 at price 5, and serve a low-value agent chosen at random at price 4. Both units are always sold and the total revenue is $\bar{R}(2) = 9$. In what

follows we will derive that this revenue is optimal among all envy-free allocations.

LEMMA 2.2. *The (maximum) envy-free revenue of a swap monotone allocation \mathbf{x} satisfies:*

$$\text{EF}^{\mathbf{x}}(\mathbf{v}) = \sum_{i=1}^n R^{\mathbf{v}}(i) \cdot (x_i - x_{i+1}) = \sum_{i=1}^n \Phi^{\mathbf{v}}(v_i) \cdot x_i.$$

PROOF. The proof is by the following equalities:

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

□

An implication of the characterization of the envy-free revenue of a pricing as its *virtual surplus*, i.e., $\sum_i \Phi(v_i) x_i$, suggests that to maximize revenue, the allocation should maximize 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 symmetric environments, the allocation that maximizes *ironed* virtual surplus is both swap monotone and revenue optimal among all swap monotone allocations.

LEMMA 2.3. *In a symmetric environments, the allocation that maximizes ironed virtual surplus with ties broken randomly is swap monotone.*

PROOF. Suppose $\bar{\Phi}(v_i) > \bar{\Phi}(v_j)$ then $x_i > x_j$; otherwise, swapping x_i for x_j would have higher ironed virtual surplus. Suppose $\bar{\Phi}(v_i) = \bar{\Phi}(v_j)$, then $x_i = x_j$ because of random-tie-breaking and the symmetry of the environment. □

THEOREM 2.4. *In any symmetric environment with any valuation profile \mathbf{v} , the allocation \mathbf{x} that maximizes ironed virtual surplus w.r.t. $\bar{\Phi}^{\mathbf{v}}$ maximizes envy-free revenue among all swap-monotone allocations. I.e., $\text{EFO}(\mathbf{v}) = \text{EF}^{\mathbf{x}}(\mathbf{v})$.*

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

LEMMA 2.5. *For any swap-monotone allocation \mathbf{x} on valuation profile \mathbf{v} ,*

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

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

PROOF. To show the inequality, we have:

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

where we use the fact that $\bar{R}(i) \geq R(i)$ and $x_i \geq x_{i+1}$. Clearly the equality holds if and only if $x_i = x_{i+1}$ whenever $\bar{R}(i) > R(i)$. □

PROOF OF THEOREM 2.4. Consider \mathbf{x} that optimizes ironed virtual surplus with random tie breaking and also consider any other swap monotone \mathbf{x}' . Note that whenever $\bar{R}(i) > R(i)$, we have $\bar{\Phi}^{\mathbf{v}}(v_i) = \bar{\Phi}^{\mathbf{v}}(v_{i+1})$ for which random tie-breaking implies $x_i = x_{i+1}$. Therefore \mathbf{x} satisfies Lemma 2.5 with equality, whereas \mathbf{x}' satisfies it with inequality. Thus $\text{EF}^{\mathbf{x}}(\mathbf{v}) \geq \text{EF}^{\mathbf{x}'}(\mathbf{v})$ and \mathbf{x} is optimal. \square

As an example of this theorem, consider position auction environments with click probabilities $w_1 \geq w_2 \geq \dots \geq w_n$. An ironed virtual surplus maximizer assigns agents with higher ironed virtual values to slots with larger click probabilities, breaking ties randomly, ignoring agents with negative ironed virtual values. The ironed virtual surplus, and thus revenue, is $\sum_{i: \bar{\Phi}(v_i) \geq 0} \bar{\Phi}(v_i) \cdot w_i$, which can be read off the revenue curve, e.g., Figure 2.

3. MATROIDS, POSITION AUCTIONS, AND MULTI-UNIT AUCTIONS

In this section we consider matroid permutation, position auction, and multi-unit environments. We show that for both incentive compatible mechanism design and envy-free pricing, these environments are closely related. In fact, for either IC or EF, respectively, the optimal mechanisms across these environments are the same and approximation mechanisms give the same approximation factor. As an example, we will focus on approximating the optimal EF revenue with a prior-free mechanism. For reasons we motivate in Section 5, the EF revenue benchmark that we will approximate is $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(\mathbf{v}^{(2)})$ where $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \dots, v_n)$. Our solution will be via a two-step reduction: we reduce matroid permutation to position auction environments, which we then reduce to multi-unit environments.

Recall that in a multi-unit environment it is feasible to serve any set of agents of cardinality at most some given k . In position auction environments there are weights $w_1 \geq w_2 \geq \dots \geq w_n$ for positions and feasible outcomes are partial assignments of agents to positions. In matroid permutation environments there is a feasibility constraint given by independent sets of a matroid, but the roles of the agents are assigned by random permutation.

The property of these three settings that enables this reduction is that in each environment the greedy algorithm on ironed virtual values (with ties broken randomly) obtains the maximum ironed virtual surplus. The greedy algorithm works as follows: order the agents by ironed virtual value and serve each agent in this order if her ironed virtual value is positive and if doing so is feasible given the set of agents previously served. Notice that the only information needed to perform the surplus maximization is the ordering on the agents' ironed virtual values (but not their magnitudes).

DEFINITION 3.1. *The characteristic weights $w_1 \geq w_2 \geq \dots \geq w_n$ of a matroid environment are as follows: choose any valuation profile \mathbf{v} with all distinct values, assign the agents to elements in the matroid via a random permutation, run the greedy algorithm w.r.t. \mathbf{v} , and define w_i to be the probability that agent i is served.*

3.1 Reduction for Ironed Virtual Surplus Maximizers

We first show ironed virtual surplus optimization in the three environments is equivalent.

LEMMA 3.1. *The ironed virtual surplus maximizing assignment (and its virtual surplus) is equal in expectation in the following environments:*

1. a matroid permutation environment with characteristic weights \mathbf{w} ,
2. a position auction environment with weights \mathbf{w} ,
3. a convex combination of multi-unit environments where k units are available with probability $w_k - w_{k+1}$ for $k \in \{1, \dots, n\}$ and $w_{n+1} = 0$.

PROOF. Fix a tie-breaking rule, which induces an ordering on the agents. Consider the greedy algorithm 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 served in the convex combination setting. Taking expectation over all tie-breaking orders, agent i has the same probability of being served in the three settings. \square

The following corollary is immediate.

COROLLARY 3.2. *For any valuation profile \mathbf{v} and any weights \mathbf{w} , the envy-free optimal revenue is the same in each of the environments of Lemma 3.1.*

A basic fact about incentive compatibility is that it is closed under convex combination, i.e., a randomization over two incentive compatible mechanisms is incentive compatible: truth-telling is an optimal strategy in each, and so it remains an optimal strategy in the combination.

We now illustrate how to use Lemma 3.1 to show that an incentive compatible prior-free approximation mechanisms for multi-unit environments can be adapted to give the same approximation factor in position auction and matroid permutation environments. Consider the following incentive compatible mechanism.

DEFINITION 3.2. *The Random Sampling Empirical Myerson (RSEM) mechanism: (discussion of payments omitted)*

1. randomly partitions the population of agents $N = \{1, \dots, n\}$ into a market M and a sample S ,
2. calculates the ironed virtual surplus function $\bar{\Phi}^S$ for the sample S , and,
3. serves a feasible subset of M to maximize surplus with respect to $\bar{\Phi}^S$ and rejects all other agents.

LEMMA 3.3. [8] *In multi-unit auction environments, RSEM is a prior-free incentive compatible 50-approximation to the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$.*

Notice that this mechanism can easily be generalized to other downward-closed environments. It remains incentive compatible for these settings because it is essentially an ironed virtual surplus optimizer on the set M , and furthermore, it is incentive compatible even if the permutation that assigns agents to the set system is fixed. As a final corollary of Lemma 3.1, we can view its revenue in the matroid permutation or position auction environment as the analogous convex combination of its revenue in multi-unit auction environments.

COROLLARY 3.4. *In matroid permutation and position auction environments, RSEM is a prior-free incentive compatible 50-approximation to the envy-free revenue $\text{EFO}^{(2)}(\mathbf{v})$.*

3.2 General Reduction

The following prior-free approximations are essentially the best known for digital good and multi-unit environments. Notably, the mechanism from Corollary 3.7 below, is not based on ironed virtual surplus maximization and therefore Lemma 3.1 cannot be applied to a construct matroid permutation or position auction mechanism from it.

LEMMA 3.5. [14] *In the digital good environment, there is a prior-free incentive compatible 3.25-approximation to $\text{EFO}^{(2)}(\mathbf{v})$.*

LEMMA 3.6. [8] *If there is a prior-free incentive compatible mechanism for the digital good environment that β -approximates $\text{EFO}^{(2)}(\mathbf{v})$, then there is such a mechanism for multi-unit environments that 2β -approximates $\text{EFO}^{(2)}(\mathbf{v})$.*

COROLLARY 3.7. *In multi-unit environments, there is an incentive compatible prior-free 6.5-approximation to $\text{EFO}^{(2)}(\mathbf{v})$.*

We now show how to construct, from any multi-unit auction, a position auction and matroid permutation mechanism that has the exact same outcome (in expectation) as a convex combination of multi-unit auctions (as in Lemma 3.1). The challenge here is the distinct interfaces to the environment: in multi-unit auctions we are given a supply constraint k and we need to specify a set of at most k winners, whereas in position auctions, we are given weights and need to output a partial assignment of agents to positions.

DEFINITION 3.3 (POSITION AUCTION REDUCTION). *Given k -unit auction mechanisms for $k \in \{1, \dots, n\}$, we construct the following mechanism for the position auction environment with weights \mathbf{w} :*

1. *Introduce n dummy agents and n dummy positions into the system, indexed by $\{n+1, \dots, 2n\}$. Correspondingly, we pad weights \mathbf{w} and valuation profile \mathbf{v} with zeros such that they have dimension $2n$.*
2. *For each $k \in \{1, \dots, n\}$, simulate the k -unit auction on valuation profile \mathbf{v} , and give the unallocated leftover units to the dummy agents arbitrarily for free. Let the resulting allocation of all $2n$ agents be $\mathbf{x}^{(k)}$.*
3. *Calculate the probability that each agent is served in the convex combination: $x_i = \sum_{k=1}^n x_i^{(k)}(w_k - w_{k+1})$, for $i \in \{1, \dots, 2n\}$.*
4. *Solve for a set of permutation matrices $P_t \in \{0, 1\}^{2n \times 2n}$ and nonnegative weights r_t with $\sum_t r_t = 1$ such that $\sum_t r_t \cdot P_t \cdot \mathbf{w} = \mathbf{x}$.*
5. *With probability r_t , assign agents to positions according to the permutation specified by P_t .*
6. *Discard dummy agents and dummy position assignments.*

To justify step 4, one can verify that \mathbf{w} majorizes \mathbf{x} in the sense that $\sum_{i=1}^k w_i \geq \sum_{i=1}^k x_i$ for $k \in \{1, \dots, 2n\}$, with equality holding for $k = 2n$. Therefore by a theorem of Rado [19], the desired permutation matrices and weights exist. The following consequences are immediate.

LEMMA 3.8. *The resulting mechanism for position auction with weights \mathbf{w} obtained from the above reduction has the same allocation as the convex combination of k -unit auctions with $(w_k - w_{k+1})$'s as probabilities.*

LEMMA 3.9. *Given an incentive compatible multi-unit auction, the mechanism from the position auction reduction is also incentive compatible.*

DEFINITION 3.4 (MATROID PERMUTATION REDUCTION). *Given a position auction mechanism for weights \mathbf{w} , we construct the following mechanism for matroid permutation environment with characteristic weights \mathbf{w} :*

1. *We run the position auction and for $i = 1, \dots, n$, let j_i be the position assigned to agent i , or $j_i = \perp$ if i is not assigned a position.*
2. *Reject all agents i with $j_i = \perp$.*
3. *Run the greedy algorithm in the matroid permutation setting with agent i 's value reset to j_i .*

The following conclusions are immediate.

LEMMA 3.10. *The resulting mechanism for matroid permutation environment obtained from the above reduction has the same allocation as the position auction.*

LEMMA 3.11. *Given an incentive compatible position auction, the mechanism from the matroid permutation reduction is incentive compatible (in matroid permutation environments).*

THEOREM 3.12. *The factor β to which there is a prior-free incentive compatible approximation of $\text{EFO}^{(2)}(\mathbf{v})$ is the same for multi-unit, position auction, and matroid permutation environments.*

COROLLARY 3.13. *There is a prior-free incentive compatible 6.5-approximation to $\text{EFO}^{(2)}(\mathbf{v})$ in position auction and matroid permutation environments.*

There are two weakness in the reductions implied by Theorem 3.12 in comparison to those implied by Lemma 3.1. Recall that for the latter, ironed virtual surplus maximizations are via the greedy algorithm, and so the reductions were algorithmically trivial. In contrast, Theorem 3.12 requires knowledge of the characteristic weights to run the construction, which may be hard to compute. In addition the mechanism that results from the matroid permutation reduction is only incentive compatible if the agents are assigned to roles in the matroid via a random permutation as suggested in the model. In contrast, RSEM in matroid environments is incentive compatible without any random permutation (Corollary 3.4).

4. INCENTIVE COMPATIBILITY VERSUS ENVY FREEDOM

The major challenge in designing and analyzing incentive compatible mechanisms is that the incentive constraint binds across all possible misreports of the agents. We therefore view a mechanism as an allocation rule and payment rule pair where $\mathbf{x}(\mathbf{v})$ and $\mathbf{p}(\mathbf{v})$ denote the allocation and payments as a function of the agent values.

DEFINITION 4.1 (INCENTIVE COMPATIBILITY). *A mechanism is 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}) is obtained from \mathbf{v} with v_i replaced by z .

DEFINITION 4.2 (VALUE MONOTONICITY). *An allocation rule is value monotone if the probability that an agent is served is monotone non-decreasing in her value, i.e., $x_i(z, \mathbf{v}_{-i})$ is non-decreasing in z for all agents i .*

The following well-known theorem characterizes ex post IC mechanisms.

THEOREM 4.1. [18] *An allocation rule $\mathbf{x}(\cdot)$ admits a non-negative and individually rational payment rule $\mathbf{p}(\cdot)$ such that (\mathbf{x}, \mathbf{p}) is incentive compatible if and only if $\mathbf{x}(\cdot)$ is value monotone, and the uniquely determined payment rule is:*

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

Because the payments are uniquely determined by the allocation rule, for any allocation rule $\mathbf{x}(\cdot)$, we let $\text{IC}^{\mathbf{x}}(\mathbf{v})$ denote the IC revenue from running $\mathbf{x}(\cdot)$ over \mathbf{v} .

We now compare envy-free revenue to incentive compatible revenue for ironed virtual surplus optimizers in permutation environments, i.e., where agents are assigned to roles in the set system via a random permutation. We show that these quantities are often within a factor of two of each other.

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

LEMMA 4.2. *For downward-closed permutation environments, all valuations \mathbf{v} , and $\bar{\Phi}$, the ironed virtual valuation function corresponding to some \mathbf{v}' obtained from \mathbf{v} by setting a subset of agents' values to be 0, we have that $\text{IC}_i^{\bar{\Phi}}(\mathbf{v}) \geq \frac{1}{2} \text{EF}_i^{\bar{\Phi}}(\mathbf{v})$ for all i .*

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 4.1, $\text{IC}_i^{\bar{\Phi}}(\mathbf{v})$ 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 3, equals the area above the IC curve and below the horizontal dotted line. On the other hand, $\text{EF}_i^{\bar{\Phi}}(\mathbf{v})$ 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: $x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}) \geq \frac{1}{2} \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$. Note that $x_i(v_{j+1}, \mathbf{v}_{-i}) = x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$ as now agents i and $j+1$ have the same value, this is equivalent to $x_i(\mathbf{v}) + x_{j+1}(\mathbf{v}) \geq x_i(v_{j+1}, \mathbf{v}_{-i}) + x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$.

The last inequality 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} . To prove this, we fix the permutation that maps agents to roles of the set system, and show that the number of winning agents from i and $j+1$ can only be lower after

agent i decreases her value. There are two cases to verify: (1) Agent i wins after the decrease. Then before the decrease, agent i had higher value, and the optimal feasible set would be the same. (2) Agent $j+1$ wins and agent i loses after the decrease. Then before the decrease, at least one of agents i and $j+1$ would win. \square




Figure 3: 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 $x_i(z, \mathbf{v}_{-i})$.

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

LEMMA 4.3. *For matroid permutation environments, all valuations \mathbf{v} , and all ironed virtual valuation functions $\bar{\Phi}$, for all agent i , $\text{EF}_i^{\bar{\Phi}}(\mathbf{v}) \geq \text{IC}_i^{\bar{\Phi}}(\mathbf{v})$.*

PROOF. Recall that $\text{EF}_i^{\bar{\Phi}}(\mathbf{v}) = \sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$ and $\text{IC}_i^{\bar{\Phi}}(\mathbf{v}) = \int_0^{v_i} (x_i(\mathbf{v}) - x_i(z, \mathbf{v}_{-i})) dz$. By the monotonicity of $x_i(z, \mathbf{v}_{-i})$ in z , $\text{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is upper-bounded by $\sum_{j=i}^n (v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$. Recall that $x_i(v_{j+1}, \mathbf{v}_{-i}) = x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$. It suffices to prove that $x_{j+1}(\mathbf{v}) \leq x_{j+1}(v_{j+1}, \mathbf{v}_{-i})$. To see this, ironed virtual surplus maximizers are greedy algorithms in matroid permutation settings, and if agent i decreases her bid to v_{j+1} , agent $j+1$ is less likely to be blocked by i who was earlier in the greedy order, and is hence more likely to be allocated. \square

5. PRIOR-FREE MECHANISM DESIGN AND BENCHMARKS

As discussed previously, no incentive compatible mechanism obtains an optimal profit point-wise on all possible valuation profiles. Therefore, to obtain point-wise guarantees, the literature on prior-free mechanism design looks for the incentive compatible mechanism that minimizes, over valuation profiles, its worst-case ratio to a given performance benchmark. It is important to identify a good benchmark for such an analysis to be meaningful.

If the designer had a prior distribution over the agent valuations then she could design the mechanism that maximizes revenue in expectation over this distribution. This is the approach of Bayesian optimal mechanism design as characterized by Myerson [18] and refined by Bulow and Roberts [5]. Given a distribution F , virtual values and revenue curves can be derived. The optimal mechanism is the one that maximizes ironed virtual surplus.

THEOREM 5.1. [18] *When values are i.i.d. from distribution F the optimal mechanism, ICO^F , is the ironed virtual surplus optimizer for $\bar{\Phi}$ corresponding to F .*

If the agent values are indeed drawn from a prior distribution, but the designer is unaware of the distribution, then a reasonable objective might be to design a mechanism that is a good approximation to the optimal mechanism for any unknown distribution. This *prior-independent* objective is a relaxation of our prior-free objective.

One important criterion for a prior-free benchmark is that its approximation should imply prior-independent approximation: if a mechanism is a constant approximation to the benchmark, then for a relevant class of distributions, it should be a constant approximation to the Bayesian optimal mechanism under any distribution from the class.

For matroid permutation environments, Lemma 4.3 implies that for any values \mathbf{v} the optimal envy-free revenue $\text{EFO}(\mathbf{v})$ (which is at least the envy-free revenue of any ironed virtual surplus optimizer) is at least the incentive compatible revenue of any ironed virtual surplus optimizer. By Theorem 5.1, the Bayesian optimal mechanism is an ironed virtual surplus optimizer so $\text{EFO}(\mathbf{v})$ upper-bounds its revenue. Consequently, a prior-free β -approximation to EFO is also a prior-independent β -approximation for all distributions.

Unfortunately, even for simple the digital good environment it is not possible to obtain a prior-free constant approximation to EFO (see [12]). This impossibility arises because it is not possible to approximate the highest value v_1 . For essentially the same reason, it is not possible to design a prior-independent constant approximation for all distributions. We therefore restrict attention to the large family of distributions with tails that are not too irregular.

DEFINITION 5.1 (TAIL REGULARITY). *A distribution F is n -tail regular if in n -agent 1-unit environments, the expected revenue of the Vickrey auction is a 2-approximation to that of the Bayesian optimal mechanism.*

The definition of tail regularity is implied by Myerson’s regularity assumption via the Bulow-Klemperer Theorem [4]. The intuition for the definition is the following. For n -agent 1-unit environments, all the action happens in the tail of the distribution, i.e, values v for which $F(v) \approx 1 - 1/n$; therefore, irregularity of the rest of the distribution does not have much consequence on revenue. Tail regularity, then, restates the Bulow-Klemperer consequence, as a constraint on the tail of the distribution and leaves the rest unconstrained.

We now define the benchmark for prior-free mechanism design. Approximation of this benchmark guarantees prior-independent approximation of all n -tail-regular distributions.

DEFINITION 5.2. *The envy-free benchmark is $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(\mathbf{v}^{(2)})$ where $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \dots, v_n)$.*

THEOREM 5.2. *With any n -agent matroid permutation environment, any n -tail-regular distribution F , and any β -approximation mechanism to $\text{EFO}^{(2)}$, the expected revenue of the mechanism with valuations \mathbf{v} drawn i.i.d. from F is a 3β -approximation to the optimal mechanism for F .*

PROOF SKETCH. By the reduction from matroid permutation environments to multi-unit environments, we focus on showing the result for k -unit auctions. We use tail regularity to get a bound on the payment from the highest agent in terms of the Vickrey auction revenue, v_2 . Of course, $\text{EFO}^{(2)}(\mathbf{v})$ is at least v_2 , and so the payment from the highest agent is at most $2 \text{EFO}^{(2)}(\mathbf{v})$ (in expectation over i.i.d. draws

of \mathbf{v} from F). The second part of the argument involves bounding the total payments of agents $\{2, \dots, n\}$, point-wise from above, by $\text{EFO}^{(2)}(\mathbf{v})$. This is possible by Lemma 4.3 and detailed analysis of k -unit auction payments. \square

It is useful to compare the EFO benchmark to ones proposed in the literature that are based on the Vickrey-Clarke-Groves (VCG) mechanism with the best (for the particular valuation profile \mathbf{v}) reserve price (e.g., [16]). The VCG mechanism with a reserve price first rejects all agents whose values do not meet the reserve, then it serves the remaining agents to maximize the surplus (sum of values).

The VCG-with-reserve benchmark can be expressed as an ironed virtual surplus optimizer, and so by Lemma 4.3, in matroid permutation environments, EFO is no smaller. For a digital good, EFO and VCG-with-reserve are identical. For multi-unit auctions EFO is at most twice VCG-with-reserve [8]. For matroid permutation environments EFO can be (almost) a logarithmic factor larger than VCG-with-reserve. Therefore, the EFO -based benchmark results in stronger approximation guarantees.

LEMMA 5.3. *There exists a distribution F and n -agent matroid environment for which VCG with any reserve price is an $\Omega(\log n / \log \log n)$ -approximation to ICO^F .*

PROOF SKETCH. We construct a set system and an irregular distribution with a jagged revenue curve that has $m = \Omega(\frac{\log n}{\log \log n})$ deep “trenches”, such that the following are true. (1) Myerson’s mechanism gets about $1/m$ fraction of its revenue from each trench via ironing. (2) The VCG mechanism with reserve gets similar amount of good revenue from at most one of the trenches by setting an appropriate price, but only gets low revenue from the other trenches due to the lack of ironing. In total, VCG with reserve only gets about $1/m$ fraction of Myerson’s revenue. \square

6. DOWNWARD-CLOSED PERMUTATION ENVIRONMENTS

For downward-closed permutation settings, there are certain bizarre settings where the maximum envy-free revenue may not upper-bound the incentive compatible revenue, which means that approximating the EFO benchmark does not necessarily imply prior-independent approximation for i.i.d. distributions. (See full paper for proof.)

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

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

In this section, we will show that a variant of RSEM (recall Definition 3.2) approximates the envy-free benchmark by a constant factor.

DEFINITION 6.1 (RSEM’). *The variant RSEM' is identical to RSEM except Step 3 where instead it:*

- \mathcal{S}' finds the feasible subset W of N (the full set of agents) to maximize surplus with respect to $\bar{\Phi}^{\mathbf{S}}$ and serves agents in $M \cap W$ (the winners from the market M) only, rejecting all others.*

In the rest of this section we prove our main theorem.

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

Consider $\text{EF}^{\bar{\Phi}^S}(\mathbf{v})$, the envy-free revenue of Step 3' of RSEM'. Notice that $\bar{\Phi}^S$, the ironed virtual valuation function for S , is not the right ironed virtual valuation function for optimizing envy-free revenue on \mathbf{v} . The goal of this section is to understand the envy-free revenue that results from optimizing using the wrong ironed virtual valuation function. To do this we introduce two auxiliary revenue curves \tilde{R} and \hat{R} (and their corresponding valuation profiles). Intuitively, \hat{R} corresponds to the revenue we think we get when optimizing $\bar{\Phi}^S$ on \mathbf{v} , and \tilde{R} corresponds to the revenue curve we actually end up with.

DEFINITION 6.2 (EFFECTIVE REVENUE CURVE \tilde{R}).

For values \mathbf{v} and ironed virtual valuations $\bar{\Phi}^S$ for S : group agents with equal nonnegative $\bar{\Phi}^S$ values into consecutive classes $\{1, \dots, n_1\}, \{n_1 + 1, \dots, n_2\}, \dots, \{n_{t-1} + 1, \dots, n_t\}$ and define the effective revenue curve \tilde{R} from $R = R^{\mathbf{v}}$ by connecting the points $(0, 0), (n_1, R(n_1)), \dots, (n_t, R(n_t))$ and then extending horizontally to $(n, R(n_t))$, i.e., ironing the values in each class.




Figure 4: Effective Ironing

Figure 4 depicts an example of the effective revenue curve. The three rays from the origin, which correspond to values at which $\bar{\Phi}^S$ makes a piece-wise jump, divide the first orthant into four regions. For every region, every point $(i, R(i))$ in the region (which corresponds to value v_i) has the same $\bar{\Phi}^S$ value. In each region these points get “ironed”, and hence the line segment in \tilde{R} .

LEMMA 6.3. $\text{EF}^{\bar{\Phi}^S}(\mathbf{v}) = \sum_{i=1}^n \tilde{R}(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v}))$.

PROOF.

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

Here the first equality is by Lemma 2.5. To justify the second equality, note that whenever $\tilde{R}(i) \neq R(i)$, there are two cases: (1) i is in $\{n_{j-1} + 1, \dots, n_j - 1\}$ for some j , and so v_i and v_{i+1} have the same $\bar{\Phi}^S$ value, and hence $x_i^S(\mathbf{v}) = x_{i+1}^S(\mathbf{v})$. (2) i is bigger than n_t , and so v_i and v_{i+1} both have negative $\bar{\Phi}^S$ value, and hence $x_i^S(\mathbf{v}) = x_{i+1}^S(\mathbf{v}) = 0$. \square

For a set of agents S , 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. Note that $\mathbf{v} = \mathbf{v}_N$. Let R^S and \tilde{R}^S be the revenue curve and ironed revenue curve of the valuation profile \mathbf{v}_S respectively.

LEMMA 6.4. *For all $1 \leq i \leq n$, $\tilde{R}(i) \geq \tilde{R}^S(i)$.*




Figure 5: Revenue dominance

PROOF SKETCH. Figure 5 depicts the relationship between the revenue curves. Observe that revenue curve R dominates R^S in the sense that for every slope t , the intersection of the ray $y = tx$ with R is further away from the origin than its intersection with R^S . Transforming R and R^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 \tilde{R}^S is non-decreasing and concave, it follows that vertical dominance also holds, i.e., $\tilde{R}(i) \geq \tilde{R}^S(i)$ for all i . \square

DEFINITION 6.3 (PERCEIVED REVENUE CURVE \hat{R}).

The perceived revenue curve for $\bar{\Phi}^S$ on \mathbf{v} is given by $\hat{R}(i) = \sum_{j=1}^i \bar{\Phi}^S(v_j)$ for $i \in N$.

Let $\hat{\mathbf{v}}$ be the valuation profile corresponding to \hat{R} , i.e., $\hat{v}_i = \hat{R}(i)/i$, and let $\mathbf{x}^{\hat{\mathbf{v}}}$ be the ironed virtual surplus maximizer for $\bar{\Phi}^{\hat{\mathbf{v}}}$.

LEMMA 6.5. $x_i^S(\mathbf{v}) = x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})$.

PROOF. Compare running the ironed virtual surplus maximizer \mathbf{x}^S for $\bar{\Phi}^S$ on \mathbf{v} with running $\mathbf{x}^{\hat{\mathbf{v}}}$ for $\bar{\Phi}^{\hat{\mathbf{v}}}$ on $\hat{\mathbf{v}}$, 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 the lemma follows. \square

Recall that there are at least two agents. We will focus on the case that agent 1 is in M , and agent 2 is in S .

DEFINITION 6.4. *We say that the partitioning (S, M) of agents $N = \{1, \dots, n\}$ is balanced if $Y_i \leq 3i/4$ and double-side balanced if $i/4 \leq Y_i \leq 3i/4$ for all $i \in \{3, \dots, n\}$ and $Y_i = |\{1, 2, \dots, i\} \cap S|$.*

LEMMA 6.6. *Conditioning on $1 \in M, 2 \in S$, a random partitioning (S, M) of N is balanced with probability at least 0.8, and is double-side balanced with probability at least 0.6.*

LEMMA 6.7. *Given a balanced partitioning (S, M) , for every non-increasing sequence a_1, \dots, a_n of nonnegative reals and all $i \in N$, $\sum_{j \in M \cap \{1, \dots, i\}} a_j \geq \frac{1}{4} \sum_{j \in \{1, \dots, i\}} a_j$.*

LEMMA 6.8. *If (S, M) is a double-side balanced partitioning with $1 \in M, 2 \in S$, then $\tilde{R}^S(i) \geq \frac{1}{4} \hat{R}(i) \geq \frac{1}{4} \tilde{R}^S(i)$ for all $1 \leq i \leq n$.*

PROOF. For each i , $\widehat{R}(i) = \sum_{j=1}^i \phi_j$ and $\bar{R}^S(i)$ is the sum of the i largest ϕ_j values with $j \in S$. Therefore $\widehat{R} \geq \bar{R}^S(i)$. Since (S, M) is double-side balanced, applying Lemma 6.7, we also have that for all i , $\bar{R}^S(i) \geq \frac{1}{4}\widehat{R}(i)$. \square

Now we are ready to prove the following key lemma:

LEMMA 6.9. *For any downward-closed permutation environments, any valuation profile \mathbf{v} , and double-side balanced partitioning (S, M) , $\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{4}\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_S) = \frac{1}{4}\text{EFO}(\mathbf{v}_S)$.*

PROOF. Let $\mathbf{x}^{\hat{\mathbf{v}}}$ be short-hands for the ironed virtual surplus optimizers with ironed virtual valuation functions defined for \mathbf{v}_S and $\hat{\mathbf{v}}$, respectively. The proof is by the following inequalities:

$$\begin{aligned} \text{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) &= \sum_i \tilde{R}(i) \cdot (x_i^S(\mathbf{v}_N) - x_{i+1}^S(\mathbf{v}_N)) \\ &= \sum_i \tilde{R}(i) \cdot (x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) - x_{i+1}^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})) \\ &\geq \frac{1}{4} \cdot \sum_i \widehat{R}(i) \cdot (x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) - x_{i+1}^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})) \\ &\geq \frac{1}{4} \cdot \sum_i \widehat{R}(i) \cdot (x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)) \\ &\geq \frac{1}{4} \cdot \sum_i \bar{R}^S(i) \cdot (x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)). \end{aligned}$$

Here the first two equalities are guaranteed by our definitions of \tilde{R} and \widehat{R} . The first inequality is by Lemma 6.4 and Lemma 6.8, the second inequality is by the optimality of $\mathbf{x}^{\hat{\mathbf{v}}}$ for $\hat{\mathbf{v}}$, and the third inequality is by Lemma 6.8 again. \square

The following lemma, together with symmetry, will help us relate $\text{EFO}(\mathbf{v}_S)$ to $\text{EFO}(\mathbf{v}_N)$.

LEMMA 6.10. *For a partitioning (S, M) of N , we have that $\text{EFO}(\mathbf{v}_S) + \text{EFO}(\mathbf{v}_M) \geq \text{EFO}(\mathbf{v}_N)$.*

PROOF. $\text{EFO}(\mathbf{v}_N) = \text{EFO}(\mathbf{v}_{S \cup M})$ is the maximum revenue we can get from $S \cup M$ subject to the envy free constraints. Let agents in M contribute total revenue R to $\text{EFO}(\mathbf{v}_N)$. By setting the agents in S to have zero valuations to obtain valuation profile \mathbf{v}_S , we basically removed envy-free constraints between agents in S and agents in M . With less envy free constraints, the maximum envy-free revenue we can get from M , i.e., $\text{EFO}(\mathbf{v}_M)$, can only be larger. Similarly, the total revenue that S contributes to $\text{EFO}(\mathbf{v}_N)$ is at most $\text{EFO}(\mathbf{v}_S)$, and our lemma follows. \square

Now we can establish the performance guarantee for RSEM' .

PROOF OF THEOREM 6.2. We condition our analysis on that agent 1 is in M , and agent 2 is in S , which happens with probability $1/4$. Conditioning on this, by Lemma 6.6, the partitioning (S, M) is double-side balanced with probability 0.6, and by Lemma 6.10 and symmetry, $\text{EFO}(\mathbf{v}_S) \geq \frac{1}{2}\text{EFO}(\mathbf{v}_N)$ with probability 0.5. Both of these events happen with probability at least $1 - (1 - 0.6) - (1 - 0.5) = 0.1$.

We assume both events happen. For each i , by Lemma 4.2, we have $\text{IC}_i^{\mathbf{x}^S}(\mathbf{v}_N) \geq \frac{1}{2}\text{EF}_i^{\bar{\Phi}^S}(\mathbf{v}_N)$. Note that $\text{EF}_i^{\bar{\Phi}^S}(\mathbf{v}_N)$ is non-increasing in i . Because (S, M) is balanced, applying Lemma 6.7, we have $\sum_{i \in M} \text{EF}_i^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \sum_{i \in N} \frac{1}{4}\text{EF}_i^{\bar{\Phi}^S}(\mathbf{v}_N)$ and therefore by Lemma 4.2, $\sum_{i \in M} \text{IC}_i^{\mathbf{x}^S}(\mathbf{v}_N) \geq \frac{1}{8}\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_N)$. Together with that $\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{4}\text{EF}^{\bar{\Phi}^S}(\mathbf{v}_S) \geq \frac{1}{8}\text{EFO}(\mathbf{v}^{(2)})$ by Lemma 6.9 and Lemma 6.10, we have $\sum_{i \in M} \text{IC}_i^{\mathbf{x}^S}(\mathbf{v}_N) \geq \frac{1}{64}\text{EFO}(\mathbf{v}^{(2)})$, and our theorem follows by multiplying the ratio with the probabilities. \square

7. REFERENCES

- [1] S. Alaei, A. Malekian, and A. Srinivasan. On random sampling auctions for digital goods. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 187–196, 2009.
- [2] L. Ausubel and P. Milgrom. Ascending auctions with package bidding. In *Frontiers of Theoretical Economics*, volume 1, 2002.
- [3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *Proc. 46th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 605–614, 2005.
- [4] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.
- [5] J. Bulow and J. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90, 1989.
- [6] E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky. On the interplay between incentive compatibility and envy freeness. 2010.
- [7] R. Day and P. Milgrom. Core-selecting auctions. *International Journal of Game Theory*, 36(3-4):393–407, 2008.
- [8] N. Devanur and J. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, 2009.
- [9] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*, 2010.
- [10] S. Dughmi, T. Roughgarden, and M. Sundararajan. Revenue submodularity. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 243–252, 2009.
- [11] A. Goldberg and J. Hartline. Envy-free auctions for digital goods. In *Proc. 5th ACM Conf. on Electronic Commerce (EC)*. ACM Press, 2003.
- [12] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [13] 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 (SODA)*, 2005.
- [14] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proc. 7th ACM Conf. on Electronic Commerce (EC)*, pages 175–182, 2005.
- [15] J. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 75–84, 2008.
- [16] J. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, 2009.
- [17] A. Mu'alem. On multi-dimensional envy-free mechanisms. In *Conference on Algorithmic Decision Theory*, pages 120–131, 2009.
- [18] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [19] R. Rado. An equality. *J. London Math. Soc.*, 27:1–6, 1952.