

7

Prior-free Mechanisms

In [Chapter 3](#) we derived optimal mechanisms for social surplus and profit. For social surplus, the externality pricing mechanism ([Definition 3.2.1](#), page [60](#)) is optimal pointwise on all valuation profiles. For profit, the virtual surplus maximization mechanism ([Definition 3.3.2](#) page [67](#)) is optimal in expectation for values drawn from the given distribution. The difference between the statement of these results is significant: for social surplus there is a pointwise optimal mechanism whereas optimal mechanisms for expected profit are parameterized by the distribution from which values are drawn. The goal of this chapter is to design mechanisms that obtain approximately optimal profit pointwise on all valuation profiles.

As an example, consider a digital good environment with $n = 100$ agents. Consider first the valuation profile where agent i has value $v_i = i$ for all i . How much revenue could a mechanism hope to obtain in such an environment? For example, this valuation profile seems similar to the uniform distribution on $[0, 100]$ for which the Bayesian optimal mechanism would post a price of 50 and obtain an expected revenue of $2500 = 50 \times 50$. Consider second the valuation profile where all agents have value one. This valuation profile seems similar to a pointmass distribution where the Bayesian optimal mechanism post a price of one for a revenue of 100. Can we come up with one mechanism that on the first profile obtains revenue close to 2500 and on the second profile obtains revenue close to 100? Moreover, what is an appropriate target revenue in general and is there an auction that approximates this target? These are the questions we address in this chapter.

The main difficulty in prior-free mechanism design for non-trivial objectives like profit (or, e.g., social surplus with a balanced budget, see [Section 3.5](#)) is that there is no pointwise, i.e., for all valuation pro-

files, optimal mechanism. Recall that incentive constraints in mechanism design bind across all valuation profiles. For example, the payment of an agent depends on the what the mechanism would have done had the agent possessed a lower value (Theorem 2.5.1). Therefore, mechanisms for the profit objective must trade off performance on one input for another. In Chapter 3 this tradeoff was optimized in expectation with respect to the prior distribution from which the agents' values are drawn; without a prior another method for navigating this tradeoff is needed.

This challenge can be resolved with approximation by comparing the performance of a mechanism to an economically meaningful prior-free benchmark. A mechanism approximates a prior-free benchmark if, for all valuation profiles, the mechanism's performance approximates the benchmark performance. A benchmark is economically meaningful if, for a large class of distributions, the expected value of the benchmark is at least the expected performance of the Bayesian optimal mechanism. If a mechanism approximates an economically meaningful benchmark then, as a corollary, the mechanism is also a prior-independent approximation (as defined in Chapter 5). Notice that this approach gives a purely prior-free design and analysis framework, but still requires returning to the Bayesian setting for economic justification of the benchmark.

A final concern is the equilibrium concept. Recall from Chapter 2 that we introduced the common prior assumption (Definition 2.3.2 page 28) so that strategic choice in games of incomplete information is well defined. Recall also that most of the optimal and approximately optimal mechanisms that we discussed in previous chapters were dominant strategy incentive compatible. In this chapter we resolve the issue of strategic choice absent a common prior by requiring that the designed mechanisms satisfy this stronger dominant-strategy incentive-compatibility condition.

The chapter begins by formalizing the framework for design and analysis of prior-free mechanisms via an economically meaningful prior-free benchmark. This framework is instantiated first in the structurally simple environment of a digital good and then subsequently generalized to environments with richer structure. The prior-free mechanism discussed will all be based a natural market analysis metaphor.

7.1 The Framework for Prior-free Mechanism Design

A main challenge for prior-free mechanism design is in identifying an economically meaningful method for evaluating a mechanism's performance. While the prior-independent mechanisms (of [Chapter 5](#)) can be compared to the optimal mechanism for the unknown distribution; absent a prior, there is no optimal mechanism with which to compare. This challenge can be resolved by decomposing the prior-independent analysis into two steps. Fix a large, relevant class of prior distributions. In the first step, a prior-free benchmark is identified and normalized so that for all distributions in the class the expected benchmark is at least the Bayesian optimal performance. In the second step an auction is constructed and proven to approximate the benchmark pointwise on all valuation profiles. These steps combine to imply a prior-independent approximation and are formalized below.

Definition 7.1.1. A prior-free benchmark maps valuation profiles to target performances. A prior-free benchmark (APX) is normalized for a class of distributions if for all distributions in the class, in expectation the benchmark is at least the performance of the Bayesian optimal mechanism (REF_F) for the distribution. I.e., for all F in the class, $\mathbf{E}_{\mathbf{v} \sim F}[\text{APX}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v} \sim F}[\text{REF}_F(\mathbf{v})]$.

Definition 7.1.2. A mechanism (APX) is a prior-free β approximation to prior-free benchmark (REF) if for all valuation profiles, its performance is at least a β fraction of the benchmark. I.e., for all \mathbf{v} , $\text{APX}(\mathbf{v}) \geq 1/\beta \text{REF}(\mathbf{v})$.

Proposition 7.1.1. For any prior-free mechanism, class of distributions, and prior-free benchmark, if the benchmark is normalized for the class of distributions and the mechanism is a prior-free β approximation

Chapter 7: Topics Covered.

- prior-free benchmarks,
- envy-free optimal pricings,
- random sampling auctions,
- profit extraction as a decision problem for mechanism design, and
- stochastic analysis of random walks and the gamblers ruin.

to the benchmark, then the mechanism is a prior-independent β approximation for the class of distributions.

We can distinguish good prior-free benchmarks from bad prior-free benchmarks by how much they overestimate the performance. (Note: a normalized prior-free benchmark never underestimates performance.) The extent to which a prior-free benchmark overestimates performance can be quantified by again considering the benchmark relative to a class of prior distributions. As the benchmark is normalized, for any distribution the expected benchmark exceeds the expected performance of the Bayesian optimal mechanism. Of course, the performance of any mechanism is no better than the Bayesian optimal mechanism for the distribution; therefore, the extent to which the Bayesian optimal mechanism approximates the benchmark gives a lower bound on the prior-free approximation of any mechanism to the benchmark. This is formalized in the following definition and proposition.

Definition 7.1.3. *The resolution γ of a prior-free benchmark (REF) is the largest ratio of the benchmark to the performance of the Bayesian optimal mechanism (APX_F) for any prior-distribution F . I.e., γ satisfies $\mathbf{E}_{\mathbf{v} \sim F}[\text{APX}_F(\mathbf{v})] \geq 1/\gamma \mathbf{E}_{\mathbf{v} \sim F}[\text{REF}(\mathbf{v})]$ for all F .*

Proposition 7.1.2. *For any class of distributions and any prior-free benchmark, the prior-free approximation β of any mechanism is at least the benchmark's resolution γ .*

This prior-free design and analysis framework turns the question of approximation into one of optimization. There is some mechanism that obtains the optimal prior-free approximation relative to the benchmark. In most of the cases we will discuss in this chapter the optimal mechanism has an approximation factor that matches the resolution of the benchmark.

Definition 7.1.4. *The optimal prior-free approximation β^* for a prior-free benchmark (REF) satisfies*

$$\beta^* = \min_{\text{APX}} \max_{\mathbf{v}} \frac{\text{REF}(\mathbf{v})}{\text{APX}(\mathbf{v})}$$

where APX ranges over all dominant strategy incentive compatible mechanisms for the given environment.

In summary, we need a normalized benchmark so that its approximation has economic meaning, and we need a benchmark with fine

resolution as its resolution lower bounds the best prior-free approximation. Intuitively, a benchmark with finer resolution will be better for distinguishing good mechanisms from bad mechanisms. A first and fundamental task in prior-free mechanism design is to identify a benchmark with fine resolution.

Example: Prior-free Monopoly Pricing

We conclude this section by instantiating the framework for design and analysis of prior-free mechanisms for the single-agent monopoly pricing problem. This is the problem of selling a single item to a single agent to maximize revenue. When the agent's value is drawn from a known distribution F , the seller's optimal mechanism, is to post the monopoly price $\hat{v}^* = \operatorname{argmax} \hat{v}(1 - F(\hat{v}))$ for the distribution (see [Section 3.3.3](#)).

Consider the class of distributions over a single agent's value with support $[1, h]$. The surplus gives a normalized prior-free benchmark and is defined by the identity function. Notice that (a) for any distribution the expected value of the benchmark exceeds the monopoly revenue and (b) and this inequality is tight for pointmass distributions. The latter observation implies that the surplus is the smallest normalized benchmark (hence, it obtains the finest resolution).

We approach the problem of analyzing the resolution of a benchmark in tandem with its optimal prior-free approximation. First, we give a lower bound on the resolution by considering the expected benchmark on the distribution for which all mechanisms attain the same performance. For the revenue objective, this distribution is the equal-revenue distribution. Second, we give a mechanism with a prior-free approximation factor that matches the lower bound. As, by [Proposition 7.1.2](#), the optimal prior-free approximation factor is at least the resolution this upper bound implies that the lower bound on resolution is tight.

Lemma 7.1.3. *For single-agent environments, the class of distributions with support $[0, h]$, and the objective of revenue, the surplus benchmark has resolution γ at least $1 + \ln h$.*

Proof. Consider the equal revenue distribution (truncated to the range $[1, h]$ with a pointmass at h with probability $1/h$). The monopoly revenue for the equal-revenue distribution is one and the expected surplus (and therefore the expected benchmark) is $1 + \ln h$ (as also calculated in [Example 4.6.1](#), page [148](#)); therefore, the resolution γ of the benchmark is at least $1 + \ln h$. \square

Now consider the purely prior-free question of posting a price to obtain a revenue that approximates the surplus benchmark. It should be clear that no deterministic price \hat{v} will do: if $\hat{v} > 1$ the prior-free approximation is infinite for value $v = 1$, and if $\hat{v} = 1$ then the prior-free approximation is h for value $v = h$. On the other hand, picking a randomized price uniformly from the powers of two on the $[1, h]$ interval gives a logarithmic approximation to the surplus. For such a randomized pricing, with probability $1/\log h$ the power of two immediately below v is posted and when this happens the revenue is at least half of the surplus benchmark. This approach and analysis can be tightened to give an approximation ratio that exactly matches the resolution of the benchmark.

Lemma 7.1.4. *For values in the interval $[1, h]$ there is a prior-free distribution over posted prices with revenue that is a $1 + \ln h$ approximation to the surplus benchmark.*

Proof. Consider the distribution over prices G with cumulative distribution function $G(z) = 1 + \ln z / 1 + \ln h$ and a pointmass at one with probability $1/1 + \ln h$. For any particular value $v \in [1, h]$, the expected revenue from a random price drawn from G is $v/1 + \ln h$. \square

Theorem 7.1.5. *For single-agent environments, values in $[1, h]$, and the objective of revenue, the resolution of the surplus benchmark and the optimal prior-free approximation are $1 + \ln h$.*

Notice that the resolution of the surplus benchmark, which is optimal among all normalized benchmarks, is not constant. In particular, it grows logarithmically with h and, when the agent's value is not bounded within some interval $[1, h]$, it is infinite. We will address this deficiency in the subsequent section where a benchmark with constant resolution and prior-free mechanisms with constant approximation ratios are derived (for $n \geq 2$ agents).

7.2 The Digital-good Environment

Our foray into prior-free mechanism design begins with the simplistic digital-good environment. In a digital-good environment any subset of agents can be simultaneously served. The absence of a feasibility constraint will enable us to focus directly on the main challenge of prior-free mechanism design which is in overcoming the lack of a prior.

We begin by deriving a benchmark with constant resolution. This

benchmark is based on a theory of envy-free pricing and we will refer to it as the *envy-free benchmark*. The resolution of the envy-free benchmark and the prior-free optimal approximation are 2.42 (in the limit with n). In the remainder of the section, we will focus on the design of simple mechanisms that approximate this envy-free benchmark (but are not optimal). First, we show that anonymous deterministic auctions cannot give good prior-free approximation. Second, we describe two approaches for designing randomized prior-free auctions for digital goods. The first auction is based on a straightforward market analysis metaphor: use a random sample of the agents to estimate the distribution of values, and run the optimal auction for the estimated distribution on the remaining agents. The approximation ratio of this auction is upper bounded by 4.68. The second auction is based on a standard algorithmic design paradigm: reduction to the a “decision version” of the problem. It gives a four approximation. These mechanisms are randomizations over deterministic dominant strategy incentive compatible (DSIC) mechanisms, the characterization of which is restated from [Corollary 2.10.4](#) as follows.

Theorem 7.2.1. *A direct, deterministic mechanism (\mathbf{x}, \mathbf{p}) is DSIC if and only if for all i and \mathbf{v} ,*

- (i) (step-function) $\mathbf{x}_i(\mathbf{v}_i, \mathbf{v}_{-i})$ steps from 0 to 1 at some $\hat{v}_i(\mathbf{v}_{-i})$, and
- (ii) (critical value) $\mathbf{p}_i(\mathbf{v}_i, \mathbf{v}_{-i}) = \begin{cases} \hat{v}_i(\mathbf{v}_{-i}) & \text{if } \mathbf{x}_i(\mathbf{v}_i, \mathbf{v}_{-i}) = 1 \\ 0 & \text{otherwise} \end{cases} + \mathbf{p}_i(0, \mathbf{v}_{-i})$.

7.2.1 The Envy-free Benchmark

Consider the following definition of and motivation for the envy-free benchmark. Recall that, when the agents’ values are drawn from an i.i.d. distribution, the Bayesian optimal digital-good auction would simply post the monopoly price for the distribution as a take-it-or-leave-it offer independently to each agent. For such a posted pricing, the agents with values above the monopoly price would choose to purchase the item and the agents with values below the monopoly price would not. As each agent selects her preferred outcome, this outcome is *envy free* no agent is envious of the outcome obtained by any other agent.

Without a prior, the monopoly price is not well defined. Instead, the *empirical monopoly price* for valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ is the monopoly price of the empirical distribution; it is calculated as $v_{(i^*)}$ with

$i^* = \operatorname{argmax}_i i v_{(i)}$ and $v_{(i)}$ denoting the i th highest value in \mathbf{v} . It is easy to see that the empirical monopoly revenue $\max_i i v_{(i)}$ is an upper bound on the revenue that would be obtained by monopoly pricing if there were a known prior distribution on values. While it is not incentive compatible to inspect the valuation profile, calculate the empirical monopoly price $v_{(i^*)}$, and offer it to each agent; it is envy free. Furthermore, as we will see subsequently, for digital good environments empirical monopoly pricing gives the *envy-free optimal revenue* which we will denote by $\text{EFO}(\mathbf{v})$.

For the class of i.i.d. distributions, the envy-free optimal revenue is a normalized benchmark. Unfortunately, the resolution of the envy-free optimal revenue, as a benchmark, is super constant. When there is $n = 1$ agent the optimal envy-free revenue is the surplus and, from the discussion of the monopoly pricing problem in the preceding section, its resolution is $1 + \ln h$ for values in $[1, h]$ and unbounded in general. The only thing, however, preventing $\text{EFO}(\mathbf{v}) = \max_i i v_{(i)}$ from being a good benchmark is the case where the maximum is obtained at $i^* = 1$ by selling to the highest value agent at her value. This discussion motivates the definition of an envy-free benchmark that explicitly excludes the $i^* = 1$ case.

Definition 7.2.1. *The envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$ for digital goods is the optimal revenue from posting a uniform price that is bought by two or more agents. I.e., $\text{EFO}^{(2)}(\mathbf{v}) = \max_{i \geq 2} i v_{(i)}$.*

Our discussion will distinguish between the envy-free optimal revenue, EFO , and the envy-free benchmark, $\text{EFO}^{(2)}$. The difference between them is that the latter excludes the possibility of selling to just the highest-valued agent. While the envy-free optimal revenue (as a benchmark) is normalized for all i.i.d. distributions, the envy-free benchmark is not. The envy-free benchmark is, however, normalized for a large class of distributions; we omit a precise characterization of this class, though subsequently in [Section 7.3](#), we show that it includes all i.i.d. regular distributions on $n \geq 2$ agents.

Analysis of resolution of the envy-free benchmark is difficult because we must quantify over all distributions. We follow the same high-level approach as for bounding the benchmark resolution in the monopoly pricing problem. First, we analyze the ratio between the expected benchmark and the Bayesian optimal auction revenue for the equal revenue distribution to get a lower bound on the resolution. Second, we observe that an auction exists with prior-free approximation that matches this resolution. [Proposition 7.1.2](#) which states that any prior-free approxi-

mation is an upper bound on the resolution, implies that the resolution and optimal prior-free approximation are equal. The following theorem summarizes this analysis.

Theorem 7.2.2. *In digital good environments, the resolution and optimal prior-free approximation of the envy-free benchmark are equal. For $n = 2, 3,$ and $4,$ the resolution and optimal prior-free approximation are $2, 13/6 \approx 2.17,$ and $2^{15}/96 \approx 2.24,$ respectively; in the limit with n it is $2.42.$*

We give the complete two-step proof of the $n = 2$ special case of Theorem 7.2.2: [Lemma 7.2.3](#) proves a lower bound of two on the resolution, and [Lemma 7.2.4](#) proves an upper bound of two on the optimal prior-free approximation. [Proposition 7.1.2](#) then, implies the equality. The generalization of this proof to $n \geq 3$ agents is technical and the subsequent discussion will treat it only at a high level.

Lemma 7.2.3. *For two-agent digital-good environments, the resolution of the envy-free benchmark is at least two.*

Proof. We give a lower bound on the resolution by comparing the expected envy-free benchmark (REF) to the expected revenue of the Bayesian optimal auction (APX) for the equal revenue distribution.

Recall that, on the equal revenue distribution, all mechanisms that allocate with probability 1 to an agent with infinite value are Bayesian optimal. Thus, the second-price auction is Bayesian optimal and $\text{APX} = \mathbf{E}_{\mathbf{v}}[v_{(2)}]$. In the case that $n = 2,$ the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$ simplifies to $2v_{(2)},$ i.e., twice the revenue of the second-price auction, and $\text{REF} = 2\mathbf{E}_{\mathbf{v}}[v_{(2)}]$. Thus, the resolution of the envy-free benchmark is at least $\text{REF}/\text{APX} = 2.$ \square

The generalization of [Lemma 7.2.3](#) to $n > 2$ follows same proof structure. The expected revenue of the second-price auction, which is Bayesian optimal for the equal-revenue distribution and any number n of agents, is $\text{APX} = n.$ The main difficulty of the analysis is in calculating the expectation of the benchmark REF. This calculation is complicated because it becomes the maximum of many terms. E.g., for $n = 3$ agents, $\text{EFO}^{(2)}(\mathbf{v}) = \max(2v_{(2)}, 3v_{(3)}).$ Nonetheless, for general n its expectation can be calculated exactly; in the limit with n it is $\text{REF} \approx 2.42n.$

Lemma 7.2.4. *For two-agent digital-good environments, the second-price auction is a prior-free two approximation of the envy-free benchmark.*

Proof. For $n = 2$ agents the the envy-free benchmark is $2v_{(2)}$ which is twice the revenue of the second-price auction. Therefore, the second-price auction is a prior-free two approximation to the envy-free benchmark. \square

The generalization of [Lemma 7.2.4](#) beyond $n = 2$ agents is technical and does not give a natural auction. For example, the $n = 3$ agent optimal auction offers each agent a price drawn from a probability distribution with a pointmass at each of the other two agents' values and continuous density at prices strictly higher than these values. The probabilities depend on the ratio of the two other agents' values. For larger $n \geq 4$ no closed-form expression is known; though, the prior-free optimal auction can be seen to match the lower bound on the resolution by a brute-force construction. This prior-free optimal auction suffers from the main drawback of optimal mechanisms: it is quite complicated. In the next sections, we will derive simple mechanisms that approximate the prior-free optimal digital-good auction.

7.2.2 Deterministic Auctions

The main idea that enables approximation of the envy-free benchmark is that when selecting a price to offer agent i we can use statistics from the values of all other agents as given by their reports \mathbf{v}_{-i} . This motivates the following mechanism which differs from empirical monopoly pricing in that the price to agent i is from the empirical distribution for \mathbf{v}_{-i} not \mathbf{v} .

Definition 7.2.2. *The deterministic optimal price auction offers each agent i the take-it-or-leave-it price of \hat{v}_i set as the monopoly price for the profile of other agent values \mathbf{v}_{-i} .*

The deterministic optimal price auction is dominant strategy incentive compatible. It is possible to show that the auction is a prior-independent constant approximation (cf. [Chapter 5](#)); however it is not a prior-free approximation. In fact, this deficiency of the deterministic optimal price auction is one that is fundamental to all anonymous (a.k.a., symmetric) deterministic auctions.

Example 7.2.1. *Consider the valuation profile with ten high-valued agents, with value ten, and 90 low-valued agents, with value one. What does the auction do on such a valuation profile? The offer to a high-valued agent is $\hat{v}_h = 1$, as \mathbf{v}_{-h} consists of 90 low-valued agents and 9*

high-valued agents. The revenue from the high price is 90; while the revenue from the low price is 99. The offer to a low-valued agent is $\hat{v}_1 = 10$, as \mathbf{v}_{-1} consists of 89 low-valued agents and 10 high-valued agents. The revenue from the high price is 100; while the revenue from the low price is 99. With these offers all high-valued agents will win and pay one, while all low-valued agents will lose. The total revenue of ten is far from the envy-free benchmark revenue of $\text{EFO}^{(2)}(\mathbf{v}) = 100$.

Theorem 7.2.5. *No n -agent anonymous deterministic dominant-strategy incentive-compatible digital-good auction is better than an n approximation to the envy-free benchmark.*

Proof. Consider valuation profiles \mathbf{v} with values $v_i \in \{1, h\}$. Let $n_h(\mathbf{v})$ and $n_1(\mathbf{v})$ denote the number of h values and 1 values in \mathbf{v} , respectively. By [Theorem 7.2.1](#) any deterministic and dominant strategy incentive compatible auction APX has a critical value at which each agent is served. That APX is anonymous implies that the critical value for agent i , as a function of the reports of other agents, is independent of the index i and only a function of $n_h(\mathbf{v}_{-i})$ and $n_1(\mathbf{v}_{-i})$. Thus, we can let $\hat{v}(n_h, n_1)$ represent the offer price of APX for any agent i when we plug in $n_h = n_h(\mathbf{v}_{-i})$ and $n_1 = n_1(\mathbf{v}_{-i})$. Finally, we assume that $\hat{v}(n_h, n_1) \in \{1, h\}$ as this restriction cannot hurt the auction revenue on the valuation profiles we are considering.

We assume for a contradiction that the auction is a good approximation and proceed in three steps.

- (i) Observe that for any deterministic auction that is a good approximation, it must be that for all m , $\hat{v}(m, 0) = h$. Otherwise, on the $n = m + 1$ agent all h 's input, the auction only achieves revenue n while the envy-free benchmark is hn . Thus, the auction would be at most an $h \geq n$ approximation on profiles with $h \geq n$.
- (ii) Likewise, observe that for any deterministic auction that is a good approximation, it must be that for all m , $\hat{v}(0, m) = 1$. Otherwise, on the $n = m + 1$ agent all 1's input, the auction achieves no revenue and is clearly not an approximation of the envy-free benchmark n .
- (iii) We now identify a bad valuation profile for the auction. Take m sufficiently large and consider $\hat{v}(k, m - k)$ as a function of k . As we have argued for $k = 0$, $\hat{v}(k, m - k) = 1$. Consider increasing k until $\hat{v}(k, m - k) = h$. This must occur since at $k = m$ we have $\hat{v}(k, m - k) = h$. Let $k^* = \min\{k : \hat{v}(k, m - k) = h\} > 1$ be this

transition point. Now consider an $n = m + 1$ agent valuation profile with $\mathbf{n}_h(\mathbf{v}) = k^*$ and $\mathbf{n}_1(\mathbf{v}) = m - k^* + 1$.

- For low-valued agents: $\hat{v}(\mathbf{n}_h(\mathbf{v}_{-1}), \mathbf{n}_1(\mathbf{v}_{-1})) = \hat{v}(k^*, m - k^*) = h$. Thus, all low-valued agents are rejected and contribute nothing to the auction revenue.
- For high-valued agents: $\hat{v}(\mathbf{n}_h(\mathbf{v}_{-h}), \mathbf{n}_1(\mathbf{v}_{-h})) = \hat{v}(k^* - 1, m - k^* + 1) = 1$. Thus, all high-valued agents are offered a price of one which they accept. Thus, the contribution to the auction revenue from such agents is $1 \times \mathbf{n}_h(\mathbf{v}) = k^*$.

Thus, the total auction revenue for this valuation profile is $\text{APX} = k^*$.

- (iv) For $h = n$, the envy-free benchmark on this valuation profile is $\text{REF} = nk^*$. There are two cases. If $k^* = 1$ then the benchmark is n (from selling to all agents at price 1); of course, for $k^* = 1$ then $n = nk^*$. If $k^* \geq 2$ the benchmark is also nk^* (from selling to the k^* high-valued agents at price $h = n$).

In conclusion, we have identified a valuation profile where the auction revenue is $\text{APX} = k^*$ and the envy-free benchmark is $\text{REF} = nk^*$; the auction is at best a prior-free n approximation. \square

[Theorem 7.2.5](#) implies that either randomization or asymmetry is necessary to obtain good prior-free mechanisms. While either approach will permit the design of good mechanisms, all deterministic asymmetric auctions known to date are based on derandomizations of randomized auctions. This text will discuss only these randomized auctions.

7.2.3 The Random Sampling Auction

We now discuss a prior-free auction based on a natural market-analysis metaphor. Notice that the problem with the deterministic optimal price auction in the preceding section is that it may simultaneously offer high-valued agents a low price and low-valued agents a high price. Of course, either of these prices would have been good if it were offered consistently to all agents. One approach for combating this lack of coordination is to coordinate using random sampling. The idea is roughly to partition the agents into a market and sample and then use the sample to estimate a good price and then offer that price to the agents in the market. With a random partition we expect a fair share of high- and low-valued agents to be in both the market and the sample; therefore, a price that is good for the sample should also be good for the market.

Definition 7.2.3. *The random sampling (optimal price) auction works as follows:*

- (i) *randomly partition the agents into sample S and market M (by flipping a fair coin for each agent),*
- (ii) *compute (empirical) monopoly prices \hat{v}_S^* and \hat{v}_M^* for S and M respectively, and*
- (iii) *offer \hat{v}_S^* to M and \hat{v}_M^* to S .*

We first, and easily, observe that the random sampling auction is dominant strategy incentive compatible.

Theorem 7.2.6. *The random sampling auction is dominant strategy incentive compatible.*

Proof. Fix a randomized partition of the agents into a market and sample. For this partitioning, each agent faces a critical value that is a function only of other agent reports. [Theorem 7.2.1](#) then implies that the auction for this partitioning is dominant strategy incentive compatible. Of course, if it is dominant strategy for any fixed partitioning it is certainly dominant strategy in expectation over the random partitioning. \square

The following example, as a warm up exercise, demonstrates that the random sampling auction is not better than a four approximation to the envy-free benchmark.

Example 7.2.2. *Consider the 2-agent input $\mathbf{v} = (1.1, 1)$ for which the envy-free benchmark is $\text{EFO}^{(2)}(\mathbf{v}) = 2$. To calculate the random-sampling auction's revenue on this input, notice that these two agents are in the same partition with probability $1/2$ and in different partitions with probability $1/2$. In the former case, the auction's revenue is zero. In the latter case it is the lower value, i.e., one. The auction's expected revenue is therefore $1/2$, which is a four approximation to the benchmark.¹*

Theorem 7.2.7. *For digital good environments and all valuation profiles, the random sampling auction is at most a 4.68 approximation to the envy-free benchmark.*

¹ It is natural to think this example could be improved if the auction were to partition half of the agents into the market and half into the sample. However in worst case, this improved partitioning cannot help. Pad the valuation profile with agents who have zero value for the item and then observe that the same analysis on this padded valuation profile gives a lower bound of four on the auction's approximation ratio.

This theorem is technical and it is generally believed that the bound it provides is loose and the random sampling auction is in fact a worst-case four approximation. Below we will prove the weaker claim that it is at worst a 15 approximation. This weaker claim highlights the main techniques involved in proving that variants and generalizations of the random sampling auction are constant approximations.

Lemma 7.2.8. *For digital good environments and all valuation profiles, the random sampling auction is at most a 15 approximation to the envy-free benchmark.*

Proof. Assume without loss of generality that the highest-valued agent is in the market M . This terminology comes from the fact that if the highest agent value is sufficiently large then all agents in other partition (in this case S) will be rejected; the role of S is then only as a sample for statistical analysis. There are two main steps in the proof. Step (i) is to show that the optimal envy-free revenue from the sample $\text{EFO}(\mathbf{v}_S)$ is close to the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$. Step (ii) is to show that the revenue from price \hat{v}_S^* on M is close to the envy-free optimal revenue from the sample which is, recall, the revenue from price \hat{v}_S^* on S .

We will use the following definitions. First sort the agents by value so that v_i is the i th largest valued agent. Define y_i as an indicator variable for the event that $i \in S$ (the sample). Notice that $\mathbf{E}[y_i] = 1/2$ except for $i = 1$; $y_1 = 0$ by the assumption that the highest valued agent is in the market. Define $Y_i = \sum_{j \leq i} y_j$ as the number of the i highest-valued agents who are in the sample. Let $\text{EFO}^{(2)}(\mathbf{v}) = i^* \hat{v}^*$ where i^* is the number of winners in the benchmark and $\hat{v}^* = v_{i^*}$ is the benchmark price.

- (i) With good probability, the optimal envy-free revenue for the sample, $\text{EFO}(\mathbf{v}_S)$, is close to the envy-free benchmark, $\text{EFO}^{(2)}(\mathbf{v})$.

Define \mathcal{B} as the event that the sample contains at least half of the i^* highest-valued agents, i.e., $Y_{i^*} \geq i^*/2$. Of course the envy-free optimal revenue for the sample is at least the revenue from posting price \hat{v}^* (which is envy-free), i.e., $\text{EFO}(\mathbf{v}_S) \geq Y_{i^*} \hat{v}^*$. Event \mathcal{B} then implies that $Y_{i^*} \hat{v}^* \geq 1/2 i^* \hat{v}^*$, or equivalently $\text{EFO}(\mathbf{v}_S) \geq 1/2 \text{EFO}^{(2)}(\mathbf{v})$.

We now show that $\Pr[\mathcal{B}] = 1/2$ when i^* is even. Recall that the highest valued agent is always in the market. Therefore there are $i^* - 1$ (an odd number) of agents which we partition between the market and the sample. One partition receives at least $i^*/2$ of these and half the time it is the sample; therefore, $\Pr[\mathcal{B}] = 1/2$.

When i^* is odd $\Pr[\mathcal{B}] < 1/2$, and a slightly more complicated argument is needed to complete the proof. A sketch of the argument is as follows. Define \mathcal{C} as the event that $Y_{i^*} \geq i^{*-1}/2$. When this event occurs, by a similar analysis as in the even case, $\text{EFO}(\mathbf{v}_S) \geq 1/2(1 - 1/i^*) \text{EFO}^{(2)}(\mathbf{v})$. The implied bound is worse than the analogous bound for the even case by an $1 - 1/i^*$ factor. The probability that the event \mathcal{C} holds improves over event \mathcal{B} , however, and this improvement more than compensates for the loss. Notice that strictly more of the top $i^* - 1$ agents are in the sample or market with equal probability but event \mathcal{C} also occurs when the numbers are equal. Thus, $\Pr[\mathcal{C}] > 1/2 = \Pr[\mathcal{B}]$. The intuition that these bounds combine to improve over the even case, above, is that the probability that the $i^* - 1$ top agents are split evenly grows as $\Theta(\sqrt{1/i^*})$ and the loss from the event providing a weaker bound grows as $\Theta(1/i^*)$.

- (ii) With good probability, the revenue from price \hat{v}_S^* on M is close to $\text{EFO}(\mathbf{v}_S)$.

Define \mathcal{E} as the event that for all indices i that the market contains at least a third as many of the i highest-valued agents as the sample, i.e., $\forall i, i - Y_i \geq 1/3 Y_i$. Notice that the left hand side of this equation is the number of agents with value at least v_i in the market, while the right hand side is a third of the number of such agents in the sample. Importantly, this event implies that the partitioning of agents is not too imbalanced in favor of the sample. We refer to this event as the *balanced sample* event; though, note that it is only a one-directional balanced condition.

Let the envy-free optimal revenue for the sample be $\text{EFO}(\mathbf{v}_S) = Y_{i_S^*} \hat{v}_S^*$ where i_S^* is the index of the agent whose value is used as its price, $\hat{v}_S^* = v_{i_S^*}$ is its price, and $Y_{i_S^*}$ is its number of winners. The revenue of the random sampling auction is equal to $(i_S^* - Y_{i_S^*}) \hat{v}_S^*$. Under the balanced sample condition this is lower bounded by $1/3 Y_{i_S^*} \hat{v}_S^* = 1/3 \text{EFO}(\mathbf{v}_S)$.

Subsequently, we will prove a *balanced sampling lemma* (Lemma 7.2.9) that shows that $\Pr[\mathcal{E}] \geq 0.9$.

We combine the two steps, above, as follows. If both good events \mathcal{C} and \mathcal{B} hold, then the expected revenue of random sampling auction is at least $1/6 \text{EFO}^{(2)}(\mathbf{v})$. By the union bound, the probability of this good fortune is $\Pr[\mathcal{C} \wedge \mathcal{B}] \geq 1 - \Pr[\neg\mathcal{C}] - \Pr[\neg\mathcal{B}] \geq 0.4$ ² We conclude that

² We denote the event that \mathcal{E} does not occur by $\neg\mathcal{E}$, which should be read as “not \mathcal{E} .” The probabilities of any event \mathcal{E} and its complement $\neg\mathcal{E}$ satisfy

the random sampling auction is a $15 = 6 \times 1/0.4$ approximation to the envy-free benchmark. \square

Lemma 7.2.9 (Balanced Sampling). *For $y_1 = 0$, y_i for $i \geq 2$ an indicator variable for a independent fair coin flipping to heads, and sum $Y_i = \sum_{j \leq i} y_j$,*

$$\Pr[\forall i, (i - Y_i) \geq 1/3 Y_i] \geq 0.9.$$

Proof. We relate the condition of the lemma to the *probability of ruin* in a *random walk* on the integers. Notice that $(i - Y_i) \geq 1/3 Y_i$ if and only if, for integers i and Y_i , $3i - 4Y_i + 1 > 0$. So let $Z_i = 3i - 4Y_i + 1$ and view Z_i as the position, in step i , of a random walk on the integers. Since $Y_1 = y_1 = 0$ this random walk starts at $Z_i = 4$. Notice that at step i the random walk is in position Z_i , so at step $i + 1$ we have

$$Z_{i+1} = \begin{cases} Z_i - 1 & \text{if } y_{i+1} = 1, \text{ and} \\ Z_i + 3 & \text{if } y_{i+1} = 0; \end{cases}$$

i.e., the random walk either takes three steps forward or one step back. We wish to calculate the probability that this random walk never touches zero. This type of calculation is known as a *probability of ruin* analysis in reference to a gambler's fate when playing a game with such a payoff structure.

Let r_k denote the probability of ruin from position k . This is the probability that the random walk eventually takes k steps backwards. Clearly $r_0 = 1$, as at position $k = 0$ we are already ruined, and $r_k = r_1^k$, as taking k steps back is equivalent to stepping back k times. By the definition of the random walk, we have the recurrence,

$$r_k = 1/2 (r_{k-1} + r_{k+3}).$$

Plugging in the above identities for $k = 1$ we have,

$$r_1 = 1/2 (1 + r_1^4).$$

This is a quartic equation that can be solved, e.g., by *Ferretti's formula* (though we omit the details). Since our random walk starts at $Z_1 = 4$ we

$\Pr[\neg \mathcal{E}] = 1 - \Pr[\mathcal{E}]$. A typical approach for bounding the probability of the conjunction (i.e., the “and”) of two events is by the disjunction (i.e., the “or”) of their negations, i.e., $\Pr[\mathcal{E} \wedge \mathcal{B}] = 1 - \Pr[\neg(\mathcal{E} \wedge \mathcal{B})] = 1 - \Pr[\neg \mathcal{E} \vee \neg \mathcal{B}]$. The *union bound* states that the probability of the disjunction of two events is at most the sum of the probabilities of each event. (This bound is tight for disjoint events, while for events that may simultaneously occur, it double counts the probability of outcomes that satisfy both events.) Thus, $\Pr[\mathcal{E} \wedge \mathcal{B}] \geq 1 - \Pr[\neg \mathcal{E}] - \Pr[\neg \mathcal{B}]$.

calculate $r_4 = r_1^4 \leq 0.1$, meaning that the probability that the balanced sampling condition is satisfied is at least 0.9. \square

The proof of [Theorem 7.2.7](#) follows a very similar structure to that of [Lemma 7.2.8](#). The main additional idea is that, instead of fixing the level of imbalance to be tolerated, it is a random variable. In [Lemma 7.2.8](#) the imbalance is fixed to $1/3$. Notice that the performance bound constructed in the lemma scales linearly with the imbalance. Thus, the expected bound can be factored into the expected imbalance times the performance bound for an imbalance of one.

7.2.4 Decision Problems for Mechanism Design

Decision problems play a central role in computational complexity and algorithm design. Where as an optimization problem is to find the optimal solution to a problem, a *decision problem* is to decide whether or not there exists a solution that meets a given objective criterion. While it is clear that decision problems are no harder to solve than optimization problems, the opposite is also true, for instance, we can search for the optimal objective value of any feasible solution by making repeated calls to an algorithm that solves the decision problem. This search is single-dimensional and can be effectively solved, e.g., by *binary search*. In this section we develop a similar theory for prior-free mechanism design.

Profit Extraction

For profit maximization in mechanism design, recall, there is no point-wise optimal mechanism. Therefore, we define the mechanism design decision problem in terms of the envy-free optimal revenue EFO. The decision problem for profit target Π is to design a mechanism that obtains profit at least Π on any valuation profile \mathbf{v} with $\text{EFO}(\mathbf{v}) \geq \Pi$. We call a mechanism that solves the decision problem a *profit extractor*.

Definition 7.2.4. *The digital good profit extractor for target Π and valuation profile \mathbf{v} finds the largest k such that $v_{(k)} \geq 1/k \Pi$, sells to the top k agents at price $1/k \Pi$, and rejects all other agents. If no such k exists, it rejects all agents.*

Lemma 7.2.10. *The digital good profit extractor is dominant strategy incentive compatible.*

Proof. Consider the following ascending auction. See if all agents can evenly split the target Π . If some agents cannot afford to pay their fair

share, reject them. Repeat with the remaining agents. Notice that the number of remaining agents in this process is decreasing, and thus, the fair share of each remaining agent is increasing. Therefore, each agent faces an ascending price until she drops out. It is a dominant strategy for her to drop out when the ascending price exceeds her value (c.f. the single-item ascending-price auction of [Definition 1.1.5](#), page [6](#)).

The outcome selected by this ascending auction is identical to that of the profit extractor. Therefore, we can interpret the profit extractor as the revelation principle ([Theorem 2.10.1](#)) applied to the ascending auction. The dominant strategy equilibrium of the ascending auction, then, implies that the profit extractor is dominant strategy incentive compatible. \square

Lemma 7.2.11. *For all valuation profiles \mathbf{v} , the digital good profit extractor for target Π obtains revenue Π if $\text{EFO}(\mathbf{v}) \geq \Pi$ and zero otherwise.*

Proof. Recall, $\text{EFO}(\mathbf{v}) = i^* v_{(i^*)}$. If $\Pi \leq \text{EFO}(\mathbf{v})$ then there exists a k such that $v_{(k)} \geq 1/k \Pi$, e.g., $k = i^*$. In this case its revenue is exactly Π . On the other hand, if $\Pi > \text{EFO}(\mathbf{v}) = \max_k k v_{(k)}$ then there is no such k for which $v_{(k)} \geq 1/k \Pi$ and the mechanism has no winners and no revenue. \square

Approximate Reduction to Decision Problem

We now employ random sampling to approximately reduce the mechanism design problem of optimizing profit to the decision problem. The key observation in this reduction is an analogy. Notice that given a single agent with value v , if we offer this agent a threshold \hat{v} the agent buys and pays \hat{v} if and only if $v \geq \hat{v}$. Analogously a profit extractor with target Π on a subset S of the agents obtains revenue Π if and only if $\text{EFO}(\mathbf{v}_S) \geq \Pi$. We can thus view the subset S of agents like a single “meta agent” with value $\text{EFO}(\mathbf{v}_S)$. The idea then is to randomly partition the agents into two parts, treat each part as a meta agent, and run the second-price auction on these two meta agents. The last step is accomplished by attempting to profit extract the envy-free optimal revenue for one part from the other part, and vice versa.

Definition 7.2.5. *The random sampling profit extraction auction works as follows:*

- (i) *randomly partition the agents into S and M (by flipping a fair coin for each agent),*

- (ii) Calculate $\Pi_M = \text{EFO}(\mathbf{v}_M)$ and $\Pi_S = \text{EFO}(\mathbf{v}_S)$, the benchmark profit for each part.
- (iii) Profit extract Π_S from M and Π_M from S .

Notice that the intuition from the analogy to the second-price auction implies that the revenue of the random sampling profit extraction auction is exactly the minimum of Π_M and Π_S . Since the profit extractor is dominant strategy incentive compatible, so is the random sampling profit extraction auction.

Lemma 7.2.12. *The random sampling profit extraction auction is dominant strategy incentive compatible.*

Before we prove that the random sampling profit extraction auction is a four approximation to the envy-free benchmark, we give a simple proof of a lemma that will be important in the analysis.

Lemma 7.2.13. *With $k \geq 2$ fair coin flips, the expected minimum of the number of heads or tails is at least $1/4 k$.*

Proof. Let W_i be a random variable for the minimum number of heads or tails in the first i coin flips. The following calculations are elementary:

$$\begin{aligned}\mathbf{E}[W_1] &= 0, \\ \mathbf{E}[W_2] &= 1/2, \text{ and} \\ \mathbf{E}[W_3] &= 3/4.\end{aligned}$$

We now obtain a general bound on $\mathbf{E}[W_i]$ for $i > 3$. Let $w_i = W_i - W_{i-1}$ representing the change to the minimum number of heads or tails when we flip the i th coin. Notice that linearity of expectation implies that $\mathbf{E}[W_i] = \sum_{j=1}^i \mathbf{E}[w_j]$. Thus, it will suffice to calculate $\mathbf{E}[w_i]$ for all i . We consider this calculation in three cases:

Case 1 (i even): This implies that $i - 1$ is odd, and prior to flipping the i th coin it was not the case that there was a tie. Assume without loss of generality that there were more tails than heads. Now when we flip the i th coin, there is probability $1/2$ that it is heads and we increase the minimum by one; otherwise, we get tails have no increase to the minimum. Thus, $\mathbf{E}[w_i] = 1/2$.

Case 2 (i odd): Here we use the crude bound that $\mathbf{E}[w_i] \geq 0$. Note that this is actually the best we can claim in worst case since $i - 1$ is even so before flipping the i th coin it could be that there is a tie. If this were the case then regardless of the i th coin flip,

$w_i = 0$ and the minimum number of heads or tails would be unchanged.

Case 3 ($i = 3$): This is a special case of Case 2; however we can get a better bound using the calculations of $\mathbf{E}[W_2] = 1/2$ and $\mathbf{E}[W_3] = 3/4$ above to deduce that $\mathbf{E}[w_3] = \mathbf{E}[W_3] - \mathbf{E}[W_2] = 1/4$.

Finally we are ready to calculate a lower bound on $\mathbf{E}[W_k]$.

$$\begin{aligned} \mathbf{E}[W_k] &= \sum_{i=1}^k \mathbf{E}[w_i] \\ &\geq 0 + 1/2 + 1/4 + 1/2 + 0 + 1/2 + 0 + 1/2 + \dots \\ &= 1/4 + 1/2 \lfloor k/2 \rfloor \\ &\geq 1/4 k. \end{aligned} \quad \square$$

Theorem 7.2.14. *For digital good environments and all valuation profiles, the random sampling profit extraction auction is a four approximation to the envy-free benchmark.*

Proof. For valuation profile \mathbf{v} , let REF be the envy-free benchmark and its revenue and APX be the random sampling profit extraction auction and its expected revenue. From the analogy to the second-price auction on meta-agents, the expected revenue of the auction is $\text{APX} = \mathbf{E}[\min(\Pi_M, \Pi_S)]$ (where the expectation is taken over the randomized of the partitioning of agents).

Assume that the envy-free benchmark sells to $i^* \geq 2$ agents at price \hat{v}^* , i.e., $\text{REF} = i^* \hat{v}^*$. Of the i^* winners in REF, let i_M^* be the number of them that are in M and i_S^* the number that are in S . Since there are i_M^* agents in M above price \hat{v}^* , then $\Pi_M \geq i_M^* \hat{v}^*$. Likewise, $\Pi_S \geq i_S^* \hat{v}^*$.

$$\frac{\text{APX}}{\text{REF}} = \frac{\mathbf{E}[\min(\Pi_M, \Pi_S)]}{i^* \hat{v}^*} \geq \frac{\mathbf{E}[\min(i_M^* \hat{v}^*, i_S^* \hat{v}^*)]}{i^* \hat{v}^*} = \frac{\mathbf{E}[\min(i_M^*, i_S^*)]}{i^*} \geq \frac{1}{4}.$$

The last inequality follows by applying Lemma 7.2.13 when we consider $i^* \geq 2$ coins and heads as putting an agent in S and a tails as putting the agent in M .

This bound is tight by an adaptation of the analysis of Example 7.2.2 from which we concluded that the random sampling optimal price auction is at best a four approximation. \square

One question that should seem pertinent at this point is whether partitioning into two groups is optimal. We could alternatively partition into three parts and run a three-agent auction on the benchmark revenue of these parts. Of course, the same could be said for partitioning

into ℓ parts for any ℓ . In fact, the optimal partitioning comes from $\ell = 3$, though we omit the proof and full definition of the mechanism.

Theorem 7.2.15. *For digital good environments and all valuation profiles, the random three-partitioning profit extraction auction is a 3.25 approximation to the envy-free benchmark.*

7.3 The Envy-free Benchmark

The first step in generalizing the framework for prior-free approximation from the preceding sections to richer environments is to generalize the envy-free benchmark. In this section we consider envy-free optimal pricing in general environments. We will give characterizations of envy-free pricings and envy-free optimal pricings that mirror those of incentive compatibility. These characterizations will promote the viewpoint that envy freedom is a relaxation of incentive compatibility that admits point-wise optimization. The section will conclude with the general definition and discussion of the envy-free benchmark.

Definition 7.3.1. *For valuation profile \mathbf{v} , an outcome with allocation \mathbf{x} and payments \mathbf{p} is envy free if no agent prefers the outcome of another agent to her own, i.e.,*

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

Example 7.3.1. *As a running example for this section consider an $n = 90$ agent, $k = 20$ unit environment with a valuation profile \mathbf{v} that consists of ten high-valued agents each with value ten and 80 low-valued agents each with value two. The following three pricings are envy free (and feasible for the environment).*

- (i) *Post a price of ten. Serve the ten high-valued agents at this price, reject the low-valued agents. This pricing is envy free: the high-valued agents weakly prefer buying and the low-valued agents prefer not buying. The total revenue is $100 = 10 \times 10$.*
- (ii) *Post a price of two. Serve the ten high-valued agents and ten of the low-valued agents at this price. This pricing is envy free: the high-valued agents prefer buying and the low-valued agents are indifferent between buying and not buying. The total revenue is $40 = 20 \times 2$.*
- (iii) *Post a price of nine to buy the item with certainty and a price of $1/4$ to buy the item with probability $1/8$ (equivalently, a probability $1/8$*

chance to buy at price of two). Serve the ten high-valued agents with the certainty outcome, and serve the 80 low-valued agents with the probabilistic outcome. By an elementary analysis this pricing is envy free: the high-valued agents weakly prefer to buy the certainty outcome and the low-valued agents weakly prefer to buy the probabilistic outcome (over nothing). The total revenue is $110 = 10 \times 9 + 80 \times 1/4$.

7.3.1 Envy-free Pricing

The definition of envy freedom can be contrasted to definition of incentive compatibility as given by the revelation principle and the defining inequality of Bayes-Nash equilibrium (Proposition 2.4.1, page 31). Importantly, incentive compatibility constrains the outcome an agent would receive upon a unilateral misreport where as envy freedom constrains the outcome she would receive upon swapping with another agent. The similarity of envy freedom and incentive compatibility enables an analogous characterization (cf. Section 2.5, page 31) and optimization (cf. Section 3.3, page 61) of envy-free pricings. However, unlike the incentive-compatibility constraints, envy-freedom constraints bind pointwise on the given valuation profile; therefore, there is always a pointwise optimal envy-free outcome.

Theorem 7.3.1. *For valuation profile \mathbf{v} (sorted with $v_1 \geq \dots \geq v_n$), an outcome (\mathbf{x}, \mathbf{p}) is envy free if and only if*

- (monotonicity) $x_1 \geq \dots \geq x_n$.
- (payment correspondence) there exists a p_0 and monotone function $y(\cdot)$ with $y(v_i) = x_i$ such that for all i

$$p_i = v_i x_i - \int_0^{v_i} y(z) dz + p_0,$$

where usually $p_0 = 0$.

Notice that the envy-free payments are not pinned down precisely by the allocation; instead, there is a range of appropriate payments. As these payment can be interpreted as the “area above the curve $y(\cdot)$,” the maximum payments are given when $y(\cdot)$ is the smallest monotone function consistent with the allocation. Given our objective of revenue maximization, for any monotone allocation rule, we focus on the maximum envy-free payments. These maximum envy-free payments are thus




Figure 7.1. The allocation is depicted as points (v_j, x_j) for each agent j . The envy-free payment of agent i is depicted as the total shaded area. The j th term in the sum of equation (7.3.1) is the dark shaded rectangle. The effective allocation rule y from Theorem 7.3.1 is the stair function depicted by a solid line.

given by the following formula and depicted in Figure 7.1:

$$p_i = \sum_{j \geq i}^n v_j (x_j - x_{j+1}), \quad (7.3.1)$$

again, with \mathbf{v} sorted as $v_1 \geq \dots \geq v_n$.

Proof of Theorem 7.3.1 We prove the theorem for the maximum envy-free payments as specified by (7.3.1) and leave the general payment correspondence as an exercise.

Monotonicity and the payment identity of equation (7.3.1) imply envy freedom: Suppose \mathbf{x} is monotone. Let \mathbf{p} be given as by equation (7.3.1). We verify that (\mathbf{x}, \mathbf{p}) is envy-free. There are two cases: if $i \leq j$, we have:

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

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

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

Each equation above can be rearranged to give the definition of envy freedom.

Envy freedom implies monotonicity: 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, $(x_i - x_j)(v_i - v_j) \geq 0$, and hence \mathbf{x} is monotone.

The maximum envy-free prices satisfy the payment identity of equation (7.3.1). Agent i does not envy $i+1$ so $v_i x_i - p_i \geq v_i x_{i+1} - p_{i+1}$,

or rearranging: $p_i \leq v_i(x_i - x_{i+1}) + p_{i+1}$. Given p_{i+1} the maximum p_i satisfies this inequality with equality. Letting $p_n = v_n x_n$ (the maximum individually rational payment) and induction gives the payment identity: $p_i = \sum_{j=i}^n v_j(x_j - x_{j+1})$. \square

7.3.2 Envy-free Optimal Revenue

Definition 7.3.2. *Given any symmetric environment and valuation profile \mathbf{v} , the envy-free optimal revenue, denoted $\text{EFO}(\mathbf{v})$, is the maximum revenue attained by a feasible envy-free outcome.*

In [Section 7.2](#) we discussed the envy-free optimal revenue for digital good environments and observed that it can be viewed as the revenue from the monopoly pricing of the *empirical distribution* for the valuation profile. The empirical distribution for a valuation profile \mathbf{v} is the discrete distribution with probability $1/n$ at value v_i .

Consider envy-free optimal pricing in multi-unit environments where, unlike digital goods, there is a constraint on the number of agents that can be simultaneously served (see [Example 7.3.1](#)). Recall that for irregular multi-unit auction environments the Bayesian optimal auction is not just the second-price auction with the monopoly reserve (in particular, it may iron). For these environments the envy-free optimal pricing also may iron. In particular, it corresponds to a virtual value optimization for virtual values given by the empirical distribution. Below we define the empirical revenue and empirical marginal revenue from which the envy-free optimal revenue can be calculated (cf. [Definition 3.3.8](#), [Definition 3.4.1](#) and [Definition 3.4.4](#) in [Section 3.3](#) page [61](#)).

Definition 7.3.3. *For valuation profile \mathbf{v} sorted as $v_1 \geq \dots \geq v_n$, the empirical price-posting revenues are $\mathbf{P} = (P_0, \dots, P_n)$ with $P_0 = 0$ and $P_i = i v_i$ for all $i \in \{1, \dots, n\}$. The empirical price-posting revenue curve is the piece-wise linear function connecting the points $(0, P_0), \dots, (n, P_n)$. The empirical revenue curve is the smallest concave function that upper bounds the empirical price-posting revenue curve; i.e., the empirical revenue curve is given by ironing the empirical price-posting revenue curve. The empirical revenues are $\mathbf{R} = (R_0, \dots, R_n)$ with R_i obtained by evaluating the empirical revenue curve at i . Empirical marginal revenues and empirical marginal price-posting revenues are given by the left slope of their respective empirical revenue curves, or equivalently, as $P'_i = P_i - P_{i-1}$ and $R'_i = R_i - R_{i-1}$.*




Figure 7.2. The empirical revenue and empirical price-posting revenue curves corresponding to the $n = 90$ agent valuation profile with ten high-valued agents and 80 low-valued agents (Example 7.3.1). The three envy-free pricings of the example are depicted as P_{10} , P_{20} , and R_{20} .

Example 7.3.2. *The empirical marginal revenues for Example 7.3.1 ($n = 90$ agents, ten with value ten and 80 with value two). The empirical revenues and price-posting revenues for this valuation profile are given in Figure 7.3.2. The empirical marginal revenues are:*

$$R'_i = \begin{cases} 10 & i \in \{1, \dots, 10\}, \text{ and} \\ 1 & i \in \{11, \dots, 90\}. \end{cases}$$

Analogously to the Bayesian optimal incentive compatible auction, the envy-free optimal pricing is a virtual value maximizer for virtual values defined by the empirical marginal revenue. The proofs of Theorem 7.3.2 and Corollary 7.3.3 below, are essentially the same as the proofs of Theorem 3.3.9 and Corollary 3.3.12.

Theorem 7.3.2. *The maximal envy-free revenue for monotone allocation \mathbf{x} is*

$$\sum_i P'_i x_i = \sum_i P_i (x_i - x_{i+1}) \leq \sum_i R_i (x_i - x_{i+1}) = \sum_i R'_i x_i$$

with equality if and only if $R_i \neq P_i \Rightarrow x_i = x_{i+1}$.

Corollary 7.3.3. *In symmetric environments, with virtual values defined as the empirical marginal revenues, virtual surplus maximization*

(with random tie-breaking) gives the envy-free outcome with the maximum revenue.

Proof of [Theorem 7.3.2](#) The inner inequality holds by the following sequence of inequalities:

$$\begin{aligned} \sum_{i=1}^n P_i &= \sum_{i=1}^n \sum_{j=i}^n v_j (x_j - x_{j+1}) \\ &= \sum_{i=1}^n i v_i (x_i - x_{i+1}) = \sum_{i=1}^n P_i (x_i - x_{i+1}) \\ &= \sum_{i=1}^n R_i (x_i - x_{i+1}) - \sum_{i=1}^n (R_i - P_i) (x_i - x_{i+1}) \\ &\leq \sum_{i=1}^n R_i (x_i - x_{i+1}), \end{aligned}$$

where the final inequality follows from the facts that $R_i \geq P_i$ and $x_i \geq x_{i+1}$. Clearly the inequality holds with equality if and only if $x_i = x_{i+1}$ whenever $R_i > P_i$.

The outer equalities hold by collecting like terms in the summation as follows,

$$\sum_{i=1}^n P_i (x_i - x_{i+1}) = \sum_{i=1}^n (P_i - P_{i-1}) x_i = \sum_{i=1}^n P'_i x_i,$$

with the analogous equations relating R_i and R'_i . \square

7.3.3 Envy freedom versus Incentive Compatibility

Optimal envy-free pricing and Bayesian optimal mechanisms are structurally similar; they are both virtual value maximizers. In this section we observe that their optimal revenues are also similar.

An *empirical virtual value function* can be defined from a valuation profile \mathbf{v} with empirical marginal revenues \mathbf{R}' as follows (recall $v_{n+1} = 0$):

$$\varphi(\mathbf{v}) = \begin{cases} R'_{i+1} & \text{if } \mathbf{v} \in [v_{i+1}, v_i) \text{ for some } i \in \{1, \dots, n\}, \text{ and} \\ \mathbf{v} & \text{otherwise.} \end{cases} \quad (7.3.2)$$

This definition is true to the geometric revenue curve interpretation where the value \mathbf{v} can be represented as a diagonal line from the origin with slope \mathbf{v} and the marginal revenue for \mathbf{v} is the left slope of the revenue curve at its intersection with this line.

For any virtual value function, symmetric environment, and valuation profile; virtual surplus maximization gives an allocation that is monotone, i.e., $v_i > v_j \Rightarrow x_i \geq x_j$, as well as an allocation rule that is monotone, i.e., $\mathbf{z} > \mathbf{z}^\dagger \Rightarrow x_i(\mathbf{z}) \geq x_i(\mathbf{z}^\dagger)$. From this allocation and allocation

rule the incentive-compatible and envy-free revenues can be calculated and compared. Recall that the maximal envy-free payment of agent i for this allocation comes from equation (7.3.1) whereas the payment of the incentive compatible mechanism with this allocation rule comes from Corollary 2.10.3. These payments are related but distinct.

Example 7.3.3. Compare the envy-free revenue and incentive-compatible revenue corresponding to Example 7.3.1 ($k = 20$ units, $n = 90$ agents, ten with value ten, and 80 with value two). The virtual value function from equation (7.3.2) is:

$$\varphi(v) = \begin{cases} -180 & v < 2, \\ 1 & v \in [2, 10), \text{ and} \\ v & v \in [10, \infty). \end{cases}$$

We now calculate the revenue of the incentive compatible mechanism that serves the 20 agents with the highest virtual value. In the virtual-surplus-maximizing auction, on the valuation profile \mathbf{v} , the high-valued agents win with probability one and the low-valued agents win with probability $1/8$ (as there are ten remaining units to be allocated randomly among 80 low-valued agents). To calculate payments we must calculate the allocation rule for both high- and low-valued agents. Low-valued agents, by misreporting a high value, win with probability one. The allocation rule for low-valued agents is depicted in Figure 7.3(a). High-valued agents, by misreporting a low value, on the other hand, win with probability $11/81$. Such a misreport leaves only nine high-value-reporting agents and so there are 11 remaining units to allocate randomly to the 81 low-value-reporting agents. The allocation rule for high-valued agents is given in Figure 7.3(b). Payments can be determined from the allocation rules: a high-valued agent pays about 8.9 and a low-valued agent (in expectation) pays $1/4$. The total incentive compatible revenue is about 109. Notice that this revenue is only slightly below the envy-free optimal revenue of 110.

The revenue calculation in Example 7.3.3 is complicated by the fact that when a high-valued agent reports truthfully there are ten remaining units to allocate to the 80 low-valued agents; whereas when misreporting a low value, there are 11 remaining units to allocate to 81 low-value reporting agents. Importantly: the allocation rule for high-valued agents and low-valued agents are not the same (compare Figure 7.3(a) and Figure 7.3(b)). The envy-free payments for both high- and low-valued agents, on the other hand, are calculated from the same “allocation

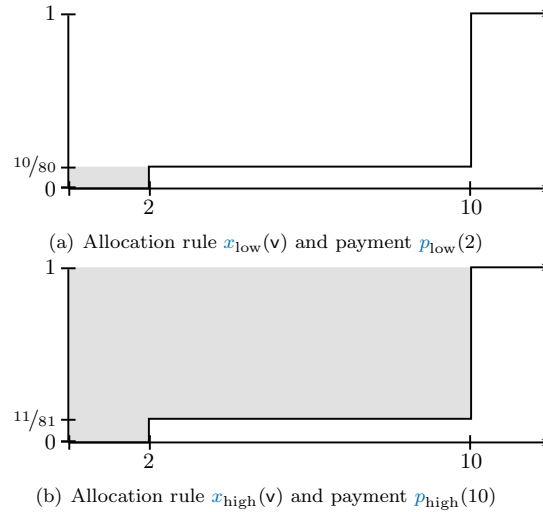


Figure 7.3. The allocation rules for high- and low-valued agents induced by the valuation profile and mechanism with virtual values given in Example 7.3.3. The incentive compatible payments are given by the area of the shaded region.

rule" (denoted as $y(\cdot)$ in [Theorem 7.3.1](#)) which is, in fact, identical to the incentive-compatible allocation rule of the low-value agents ([Figure 7.3\(a\)](#)). Thus, the envy-free revenue can be viewed as a relaxation of the incentive-compatible revenue that is simpler and, therefore, more analytically tractable.

We now formalize the fact that the envy-free optimal revenue is an economically meaningful benchmark by showing that it is pointwise normalized (which implies that it is normalized for any i.i.d. distribution).

Theorem 7.3.4. *For multi-unit environments and any virtual value function $\varphi(\cdot)$, the envy-free revenue of virtual surplus maximization is at least its incentive-compatible revenue.*

Proof. We show that the envy-free payment of agent i is at least her incentive-compatible payment. In particular if we let $x_i(\mathbf{v})$ be the allocation rule of the virtual surplus optimizer, then for $\mathbf{z} \leq \mathbf{v}_i$, $x_i(\mathbf{z}, \mathbf{v}_{-i})$ (as a function of \mathbf{z}) is at most the smallest $y(\mathbf{z})$ that satisfies the conditions of [Theorem 7.3.1](#). Since the incentive-compatible and envy-free payments,

respectively, correspond to the area “above the allocation curve” this inequality implies the desired payment inequality.

Since $x_i(\mathbf{z}, \mathbf{v}_{-i})$ is monotone, we only evaluate it at $v_j \leq v_i$ and show that $x_i(v_j, \mathbf{v}_{-i}) \geq x_j(\mathbf{v})$. In particular,

$$x_i(v_j, \mathbf{v}_{-i}) = x_j(v_j, \mathbf{v}_{-i}) \geq x_j(\mathbf{v}).$$

The equality above comes from the symmetry of the environment and the fact that agent i and j have the same value in profile (v_j, \mathbf{v}_{-i}) . The inequality comes from greedy maximization with random tie breaking for multi-unit auctions: when agent i reduces her bid to tie agent j 's value v_j the probability that j receives a unit does not decrease as agent i is only less competitive. \square

We will see later that this theorem generalizes beyond multi-unit environments (see [Section 7.5](#)). In particular, the only properties of multi-unit environments that we employed in the proof were symmetry and that the greedy algorithm is optimal.

7.3.4 Permutation Environments

Envy freedom is less natural in asymmetric environments such as those given by matroid or downward-closed feasibility constraints. To extend the envy-free benchmark to asymmetric environments we assume a symmetry imposing prior-free analog of the (standard) Bayesian assumption that the agents' value distribution is independent and identically distributed. Specifically, the valuation profile can be arbitrary, but the roles the agents play with respect to the environment (e.g., feasibility constraint or cost function) are assigned by random permutation.

Definition 7.3.4. *Given an environment, specified by cost function $c(\cdot)$, the permutation environment is the environment with the identities of the agents uniformly permuted. I.e., for permutation π drawn uniformly at random from all permutations, the permutation environment has cost function $c(\pi(\cdot))$.*

Our goal is a prior-free analysis framework for which approximation implies prior-independent approximation in i.i.d. environments. Of course the expected revenue of the optimal auction in an i.i.d. environment is unaffected by a random permutation of the identities of the agents. Therefore, with respect to the goal of obtaining a prior-independent corollary from a prior-free analysis (by [Proposition 7.1.1](#)),

it is without loss to assume a permutation environment. Importantly, while a matroid or downward-closed environment may be asymmetric, a matroid permutation or downward-closed permutation environment is inherently symmetric. This symmetry admits a meaningful study of envy freedom.

The environments considered heretofore have been given deterministically, e.g., by a cost function or set system (Chapter 3, Section 3.1). A generalization of this model would be to allow randomized environments. We view a randomized environment as a probability distribution over deterministic environments, i.e., as a convex combination. For the purpose of incentives and performance, we will view mechanism design in randomized environments as follows. First, the agents report their preferences; second, the designer's cost function (or feasibility constraint) is realized; and third, the mechanism for the realized cost function is run on the reported preferences. The performance in such probabilistic environment is measured in expectation over both the randomization in the mechanism and the environment. Agents act before the set system is realized and therefore from their perspective the game they are playing in is the composition of the randomized environment with the (potentially randomized) mechanism.

An example of such a probabilistic environment comes from *display advertising*. Banner advertisements on web pages are often sold by auction. Of course the number of visitors to the web page is not precisely known at the time the advertisers bid; instead, this number can be reasonably modeled as a random variable. Therefore, the environment is a convex combination of multi-unit auctions where the supply k is randomized.

7.3.5 The Envy-free Benchmark

We are now ready to formally define the envy-free benchmark. To do so we must address the potential asymmetry in the environment and the technicality that the envy-free revenue itself may have unbounded resolution (recall the discussion above Definition 7.2.1 on 216). Finally, we must give economic justification to the benchmark by showing that it is normalized.

Definition 7.3.5. *For any environment and valuation profile \mathbf{v} , the envy-free benchmark, denoted $\text{EFO}^{(2)}(\mathbf{v})$, is the optimal envy-free revenue in the permutation environment on the valuation profile $\mathbf{v}^{(2)}$ where*

the highest value $v_{(1)}$ is replaced with twice the second highest value $2v_{(2)}$, i.e., $\mathbf{v}^{(2)} = (2v_{(2)}, v_{(2)}, v_{(3)}, \dots, v_{(n)})$.

Theorem 7.3.5. *For i.i.d., regular, $n \geq 2$ agent, multi-unit environments, the envy-free benchmark is normalized.*

Proof. Recall that for i.i.d. regular distributions F , the n -agent k -unit Bayesian optimal auction REF_F is the $k + 1$ st-price auction with the monopoly reserve \hat{v}^* for the distribution. We will show a stronger claim than the statement of the lemma. The *anonymous-reserve benchmark* APX for valuation profile \mathbf{v} is the revenue of the k -unit auction with the best reserve price for the valuation profile $\mathbf{v}^{(2)} = (2v_{(2)}, v_{(2)}, \dots, v_{(n)})$. For $k = 1$, $\text{APX}(\mathbf{v}) = 2v_{(2)}$ and in general $\text{APX}(\mathbf{v}) = \max_{2 \leq i \leq k} i v_{(i)}$. The outcome of the anonymous-reserve benchmark is envy-free for $\mathbf{v}^{(2)}$; therefore, it lower bounds the envy-free optimal revenue for $\mathbf{v}^{(2)}$; and therefore, the normalization of the anonymous-reserve benchmark implies normalization of the envy-free benchmark.

We first argue the $n = 2$ agent, $k = 2$ unit special case (a.k.a., the two-agent digital good environment). Fix an i.i.d. regular distribution F over valuation profiles. We show that the expected anonymous-reserve benchmark (APX) is at least the performance of the Bayesian optimal mechanism (REF_F) for the distribution, i.e., that $\mathbf{E}_{\mathbf{v} \sim F}[\text{APX}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v} \sim F}[\text{REF}_F(\mathbf{v})]$.

Recall [Theorem 5.2.1](#) (also [Lemma 5.3.2](#)) which states that for i.i.d. regular distributions that the revenue of the two-agent second-price auction exceeds that of the single-agent monopoly pricing. Thus, twice the second-price revenue exceeds twice the monopoly pricing revenue. For $n = k = 2$, the former is equal to the expected anonymous-reserve benchmark and the latter is equal to the expected Bayesian optimal revenue.

Under the regularity assumption, the normalization of the anonymous-reserve benchmark for the two-agent digital good environment implies its normalization for multi-unit environments with general $n \geq 2$ agents and k units. To show this extension, consider any $k \geq 2$ and any $n \geq 2$ and condition on the third-highest value $v_{(3)}$. The following argument shows that $\mathbf{E}[\text{APX}(\mathbf{v}) \mid v_{(3)}] \geq \mathbf{E}[\text{REF}_F(\mathbf{v}) \mid v_{(3)}]$ for $v_{(3)} < \hat{v}^*$ and $v_{(3)} \geq \hat{v}^*$ considered as separate cases.

When the third-highest value $v_{(3)}$ is less than the monopoly price \hat{v}^* , then all agents except for the top two are rejected. The conditional distribution on of the two highest valued agents is regular (the conditioning only truncates and scales the revenue curve; therefore, its convexity is

preserved), moreover, the remaining feasibility constraint is that of a digital good. Hence, the normalization for the two-agent digital good environment implies normalization for this conditional environment.

When the third-highest value $v_{(3)}$ is more than the monopoly price \hat{v}^* , then the Bayesian optimal auction REF_F on \mathbf{v} sells at least two units at a uniform price and the empirical anonymous-reserve revenue from selling the same number of units is pointwise no smaller. Thus, the desired bound holds pointwise.

Now consider the final case of $k = 1$ unit, $n \geq 2$ agent environments. We will reduce normalization of the anonymous-reserve benchmark in this environment to that of the $k = 2$ unit environment. The benchmark in the two environments is the same: the one-unit benchmark is $2v_{(2)}$; the two-unit benchmark is $2v_{(2)}$. The Bayesian optimal revenue is only greater for two units than with one unit. Therefore, normalization for two units implies normalization for one unit. \square

It is evident from this proof that the anonymous-reserve benchmark is also normalized for multi-unit environments. We will prefer to use the envy-free benchmark for three reasons. First, the envy-free benchmark remains normalized for a larger class of distributions which admit a large degree of irregularity (though not arbitrary irregular distributions). Second, the envy-free benchmark is easier to work with as it is structurally a virtual surplus optimization. Third, for position environments discussed subsequently, the envy-free benchmark is linear in the position weights, while the anonymous reserve benchmark is not. This linearity will be important for our analysis.

7.4 Multi-unit Environments

In this section we will discuss two approaches for multi-unit environments. In the first, we will give an approximate reduction to digital good environments. This reduction will give a $\beta + 1$ approximation mechanism for multi-unit environments from any β approximation mechanism for digital goods. Applied to the prior-free optimal digital good auction, a 2.42 approximation, this approach yields a multi-unit 3.42 approximation. The second approach will be to directly generalize the random sampling optimal price auction to multi-unit environments. This generalization randomly partitions the agents into a market and sample, calculates the empirical distribution of the sample, and then runs optimal

multi-unit auction on the market according to the empirical distribution for the sample.

7.4.1 Reduction to Digital Goods

Our first approach is an approximate reduction. For i.i.d. irregular single-item environments, [Corollary 4.4.3](#) shows that the second-price auction with anonymous reserve is a two approximation to the optimal auction. I.e., the loss in performance from not ironing when the distribution is irregular is at most a factor of two. In fact, this result extends to multi-unit environments (as does the prophet inequality from which it is proved) and the approximation factor only improves. Given the close connection between envy-free optimal outcomes and Bayesian optimal auctions, it should be unsurprising that this result translates between the two models.

Consider the revenue of the surplus maximization mechanism with the best (ex post) anonymous reserve price. For instance, for the k -unit environment and valuation profile \mathbf{v} , this revenue is $\max_{i \leq k} i v_{(i)}$. It is impossible to approximate this revenue with a prior-free mechanism so, as we did for the envy-free benchmark, we exclude the case that it sells to only the highest-valued agent at her value. Therefore, for k -unit environments the *anonymous-reserve benchmark* is $\max_{2 \leq i \leq k} i v_{(i)}$ for $k > 2$ (and $2v_{(2)}$ for $k = 1$), i.e., it is the optimal anonymous reserve revenue for the valuation profile $\mathbf{v}^{(2)} = (2v_{(2)}, v_{(2)}, \dots, v_{(n)})$. Notice that for digital goods, i.e., $k = n$, the anonymous-reserve benchmark is equal to the envy-free benchmark. Of course, an anonymous reserve is envy free so the envy-free benchmark is at least the anonymous-reserve benchmark.

We now give an approximate reduction from multi-unit environments to digital-good environments in two steps. We first show that the envy-free benchmark is at most twice the anonymous-reserve benchmark in multi-unit environments. We then show an approximation preserving reduction from multi-unit to digital-good environments with respect to the anonymous-reserve benchmark. In the next section a more sophisticated approach that attains a better bound is given.

Theorem 7.4.1. *For any valuation profile, in multi-unit environments, the envy-free benchmark (resp. optimal revenue) is at most the sum of the anonymous-reserve benchmark (resp. optimal revenue) and $k + 1$ st-price auction revenue, which is at most twice the anonymous-reserve benchmark (resp. optimal revenue).*




Figure 7.4. Depiction of ironed revenue curve \mathbf{R} for the geometric proof of Theorem 7.4.1. The solid piece-wise linear curve is \mathbf{R} , the convex hull of \mathbf{P} , and contains the line-segment connecting point $A = (i, P_i)$ and point $D = (j, P_j)$. The envy-free benchmark is achieved at point $C = (k, R_k)$. The parallel dashed lines have slope $v_{(j)}$, the other dashed line has slope $v_{(i)}$.

Proof. We prove the statement with respect to the optimal revenues and any valuation profile \mathbf{v} and then apply the theorem to the valuation profile $\mathbf{v}^{(2)}$ to obtain the statement with respect to the benchmarks.

If the envy-free optimal revenue sells fewer than k units or the revenue curve is not ironed at k then the anonymous-reserve revenue equals the envy-free revenue and the theorem trivially holds. Otherwise, assume that the envy-free optimal revenue sells all k units and irons between index $i < k$ and $j > k$ (see Figure 9.4). In terms of empirical revenue curves (Definition 7.3.3), the envy-free optimal revenue for \mathbf{v} is REF = $R_k = C$. Note that the AC line has slope R'_k , i.e., $C = A + (k - i) R'_k$. The line from the origin to D has slope $v_{(j)}$. By geometry $v_{(j)} > R'_k$. Thus, extending a line from $A = R_i$ with slope $v_{(j)}$ to point $B = A + (k - i) v_{(j)}$ satisfies $B > C$.

The anonymous-reserve revenue exceeds the $k + 1$ st-price auction revenue; thus, twice the anonymous-reserve revenue exceeds the sum of the anonymous-reserve revenue and the $k + 1$ st-price auction revenue. The anonymous-reserve revenue is at least R_i and the $k + 1$ st-price revenue

is $k v_{(k+1)} \geq (k-i)v_{(j)}$ (as $j \geq k+1$ and $i \geq 1$); thus the sum of their revenues exceeds $R_i + (k-i)v_{(j)} = B > C = \text{REF}$. The theorem follows. \square

Theorem 7.4.1 reduces the problem of approximating the envy-free benchmark to that of approximating the anonymous-reserve benchmark. There is a general construction for converting a digital good auction \mathcal{M} into a limited supply auction and if \mathcal{M} is a β approximation to the anonymous-reserve benchmark (which is identical to the envy-free benchmark for digital goods) then so is the resulting multi-unit auction.

Definition 7.4.1. *The $k \geq 2$ unit restriction \mathcal{M}_k of digital good auction \mathcal{M} is the following:*

- (i) *Simulate the $k+1$ st-price auction (i.e., the k highest valued agents win and pay $v_{(k+1)}$).*
- (ii) *Simulate \mathcal{M} on the k winners $v_{(1)}, \dots, v_{(k)}$.*
- (iii) *Serve the winners from the second simulation and charge them the higher of their prices in the two simulations.*

The 1-unit restriction \mathcal{M}_1 is the second-price auction.

Implicit in this definition is a new notion of mechanism composition (cf. [Chapter 5](#) [Section 5.4.2](#)). It is easy to see that this mechanism composition is dominant strategy incentive compatible. In general such a composition is DSIC whenever no winner of the first mechanism can manipulate her value to change the set of winners while simultaneously remaining a winner (see [Exercise 7.3](#)); mechanisms that satisfy this property are said to be *non-bossy*.

Theorem 7.4.2. *If \mathcal{M} is a β approximation to the envy-free benchmark in digital good environments then its multi-unit restriction \mathcal{M}_k is a 2β approximation in multi-unit environments.*

Proof. For 1-unit environments, the second-price auction (with revenue $v_{(2)}$) is a 2-approximation to the 1-unit envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v}) = 2v_{(2)}$. For $k \geq 2$ unit environments, the k -unit restriction is a β approximation to the envy-free benchmark restricted to the k highest-valued agents. This benchmark is equal to the anonymous-reserve benchmark on the full set of agents. This benchmark, by [Theorem 7.4.1](#) is at least half the envy-free benchmark on the full set of agents. Thus, the k -unit restriction is a 2β approximation to the envy-free benchmark. \square

This theorem can be applied to any digital good auction; for instance, from Theorem [7.2.2](#) we can conclude that there is a multi-unit auction that is a 4.84 approximation to the envy-free benchmark.

7.4.2 Combination of Benchmarks and Auctions

[Theorem 7.4.1](#) which shows that the envy-free benchmark is bounded by the sum of the anonymous-reserve benchmark and the $k + 1$ st-price auction revenue, can be employed to construct a $\beta + 1$ approximation for multi-unit environments from a β approximation for digital goods. Applied to the prior-free optimal auction for digital goods, this yields an multi-unit 3.42 approximation. The approach is to view the envy-free benchmark as the sum of two benchmarks, design prior-free mechanisms for each benchmark, and then consider an appropriate convex combination of the two mechanisms to optimize the approximation with respect to the original benchmark. This approach provides two conclusions. First, it gives a modular approach to prior-free mechanism design. Second, it suggests that, even in pursuit of prior-free approximation with respect to the economically motivated envy-free benchmark, it may be useful to understand prior-free approximation for other benchmarks.

Definition 7.4.2. For benchmark $\mathcal{G}(\mathbf{v}) = \mathcal{G}_A(\mathbf{v}) + \mathcal{G}_B(\mathbf{v})$, mechanism \mathcal{M}_A giving a prior-free β_A approximation to benchmark \mathcal{G}_A , and mechanism \mathcal{M}_B giving a prior-free β_B approximation to benchmark \mathcal{G}_B ; the prior-free combination \mathcal{M} runs \mathcal{M}_A with probability $\beta_A/\beta_A+\beta_B$ and \mathcal{M}_B with probability $\beta_B/\beta_A+\beta_B$.

Theorem 7.4.3. With respect to [Definition 7.4.2](#), the prior-free combination \mathcal{M} is a prior-free $\beta = \beta_A + \beta_B$ approximation to benchmark $\mathcal{G}(\mathbf{v}) = \mathcal{G}_A(\mathbf{v}) + \mathcal{G}_B(\mathbf{v})$.

Proof.

$$\begin{aligned} \mathcal{M}(\mathbf{v}) &= \frac{\beta_A}{\beta_A + \beta_B} \mathcal{M}_A(\mathbf{v}) + \frac{\beta_B}{\beta_A + \beta_B} \mathcal{M}_B(\mathbf{v}) \\ &\geq \frac{\beta_A}{\beta_A + \beta_B} \frac{\mathcal{G}_A(\mathbf{v})}{\beta_A} + \frac{\beta_B}{\beta_A + \beta_B} \frac{\mathcal{G}_B(\mathbf{v})}{\beta_B} = \frac{\mathcal{G}_A(\mathbf{v}) + \mathcal{G}_B(\mathbf{v})}{\beta_A + \beta_B}. \quad \square \end{aligned}$$

As described above, the k -unit restriction \mathcal{M}_k of β approximate digital good auction \mathcal{M} is an β approximation to the multi-unit anonymous reserve benchmark. The $k + 1$ st-price auction is, obviously, a one approximation to the $k + 1$ st-price auction revenue. Thus, we can invoke [Theorem 7.4.1](#) and [Theorem 7.4.3](#) to obtain the following corollary. The

best multi-unit auction via this construction is attained by instantiating the reduction with the prior-free optimal digital good auction which is a 2.42 approximation ([Theorem 7.2.2](#)).

Corollary 7.4.4. *For prior-free β -approximate digital-good auction \mathcal{M} , the prior-free combination of the multi-unit restriction \mathcal{M}_k (with probability $\beta/\beta+1$) and the $k+1$ st-price auction (with probability $1/\beta+1$) is a prior-free $\beta+1$ approximation to the envy-free benchmark in multi-unit environments. For the prior-free optimal digital-good auction, this multi-unit auction is a prior-free 3.42 approximation.*

7.4.3 The Random Sampling Auction

An alternative approach to the multi-unit auction problem is to directly generalize the random sampling optimal price auction. Intuitively, the random sampling auction partitions the agents into a market and sample and then runs the optimal auction for the empirical distribution of the sample on the market. For digital goods the optimal auction for the empirical distribution sample is just the to post the empirical monopoly price. For multi-unit environments, the optimal auction irons when the empirical distribution of the sample is irregular.

Definition 7.4.3. *The random sampling (virtual surplus maximization) auction for the k -unit environment*

- (i) *randomly partitions the agents into market M and sample S by assigning the highest-valued agent to M and flipping a fair coin for all other agents,*
- (ii) *computes virtual valuation function φ_S for the empirical distribution of S , and*
- (iii) *maximizes virtual surplus of selling at most k units to S with respect to φ_M .*

The proof of the following theorem can be derived similarly to the proof of [Lemma 7.2.8](#) (page [222](#)); we omit the details.

Theorem 7.4.5. *For multi-unit environments and all valuation profiles, the random sampling auction is a constant approximation to the envy-free benchmark.*

The random sampling auction shares some good properties with optimal mechanisms. The first is that the mechanism on the market is a virtual-surplus optimization. I.e., it sorts the agents in the market by

virtual value and allocates to the agents greedily in that order. This property is useful for two reasons. First, in environments where the supply k of units is unknown in advance, the mechanism can be implemented *incrementally*. Each unit of supply is allocated in to the agent remaining in the market with the highest virtual valuation. Second, as we will see in the next section, it can be applied without specialization to matroid permutation and position environments.

7.5 Matroid Permutation and Position Environments

Position environments are important as they model auctions for selling advertisements on Internet search engines such as Google and Microsoft's Bing. In these auctions agents bid for positions with higher positions being better. The feasibility constraint imposed by position auctions is a priori symmetric.

Definition 7.5.1. *A position environment is one with n agents, n positions, each position j described by weight w_j . An auction assigns each position j to an agent i which corresponds to setting $x_i = w_j$. Positions are usually assumed to be ordered in non-increasing order, i.e., $w_j \geq w_{j+1}$. (Often w_1 is normalized to one.)*

Position auctions correspond to advertising on Internet search engines as follows. Upon each search to the search engine, *organic search results* appear on the left-hand side and *sponsored search results*, a.k.a., advertisements, appear on the right-hand side of the search results page. Advertiser i receives a revenue of v_i in expectation each time her ad is clicked (e.g., if the searcher buys the advertiser's product) and if her ad is shown in position j it receives *click-through rate* w_j , i.e., the probability that the searcher clicks on the ad is w_j . If the ad is not clicked on the advertiser receives no revenue. Searchers are more likely to click on the top slots than the bottom slots, hence $w_j \geq w_{j+1}$. An advertiser i shown in slot j receives value $v_i w_j$. Understandably, this model of Internet search advertising omits many details of the environment; nonetheless, it has proven to be quite relevant.

We now show that mechanism design for matroid permutation environments can be reduced to that for position environments which can be reduced to that for multi-unit environments. These reductions follow, essentially, because each of these environments are *ordinal*, i.e., because

surplus is maximized by the greedy algorithm. The greedy algorithm does not compare magnitudes of the values of agents, it only considers their relative order. This intuition is summarized by the following definition.

Definition 7.5.2. *The characteristic weights \mathbf{w} for a matroid are defined as follows: Set $v_i = n - i + 1$, for all i , and consider the surplus maximizing allocation when agents are assigned roles in the set system via random permutation and then the maximum feasible set is calculated, e.g., via the greedy algorithm. Let w_i be the probability of serving agent i , i.e., the i th highest-valued agent.*

To see why the characteristic weights are important, notice that since the greedy algorithm is optimal for matroids, the cardinal values of the agents do not matter, just the sorted order. Therefore, e.g., when maximizing virtual value, w_i is the probability of serving the agent with the i th highest virtual value.

Theorem 7.5.1. *The problem of revenue maximization (or approximation) in matroid permutation environments reduces to the problem of revenue maximization (or approximation) in position environments.*

Proof. We show two things. First, we show that for any matroid permutation environment with characteristic weights \mathbf{w} , the position environment with weights \mathbf{w} has the same optimal expected revenue. Second, for any such environments any position auction can be converted into an auction for the matroid permutation environment that achieves the same expected revenue as the position auction in the position environment given by the characteristic weights of the matroid. These two results imply that any Bayesian, prior-independent, or prior-free approximation results for position auctions extend to matroid permutation environments.

- (i) Revenue optimal auctions are virtual surplus optimizers. Let \mathbf{w} be the characteristic weights for the given matroid environment. By the definition of \mathbf{w} , the optimal auctions for both the matroid permutation and position environments serve the agent with the j th highest positive virtual value with probability w_j . (In both environments agents with negative virtual values are discarded.) Expected revenue equals expected virtual surplus; therefore, the optimal expected revenues in the two environments are the same.
- (ii) Consider the following matroid permutation mechanism which is based

on the position auction with weights \mathbf{w} . The input is \mathbf{v} . First, simulate the position auction 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 environment on input $\mathbf{v}_i^\dagger = n - j_i + 1$ and output its outcome.

Notice that any agent i is allocated in the matroid permutation setting with probability equal to the expected weight of the position it is assigned in the position auction. Therefore the two mechanisms have the exact same allocation rule (and therefore, the exact same expected revenue). \square

We are now going to reduce the design of position auctions to that of multi-unit auctions. This reduction implies that the prior-free approximation factor for multi-unit environments extends to matroid permutation and position environments. Furthermore, the mechanism that gives this approximation can be derived from the multi-unit auction.

Theorem 7.5.2. *The problem of revenue maximization (or approximation) in position auctions reduces to the problem of revenue maximization (or approximation) in k -unit auctions.*

Proof. This proof follows the same high-level argument as the proof of Theorem [7.5.1](#)

Let $\mathbf{w}'_j = \mathbf{w}_j - \mathbf{w}_{j+1}$ be the difference between successive position weights. Recall that without loss of generality $\mathbf{w}_1 = 1$ so $\sum_j \mathbf{w}'_j = 1$ and \mathbf{w}' can be interpreted as a probability measure over $\{1, \dots, n\}$.

- (i) The expected revenue of an optimal position auction is equal to the expected revenue of the convex combination of optimal j -unit auctions under measure \mathbf{w}' . In the optimal position auction and the optimal auction for the above convex combination of multi-unit auctions the agent with the j th highest positive virtual value is served with probability w_j . (In both settings agents with negative virtual values are discarded.) Therefore, the expected revenues in the two environments are the same.
- (ii) Now consider the following position auction which is based on a multi-unit auction. Simulate a j -unit auction on the input \mathbf{v} for each $j \in \{1, \dots, n\}$ and let $x_i^{(j)}$ be the (potentially random) indicator for whether agent i is allocated in simulation j . Let $x_i = \sum_j x_i^{(j)} \mathbf{w}'_j$ be the expected allocation to j in the convex combination of multi-unit

auctions given by measure \mathbf{w}' . The vector of position weights \mathbf{w} majorizes the allocation vector \mathbf{x} in the sense that $\sum_i^k w_i \geq \sum_i^k x_i$ (and with equality for $k = n$). Therefore we can write $\mathbf{x} = S\mathbf{w}$ where S is a doubly stochastic matrix. Any doubly stochastic matrix is a convex combination of permutation matrices, so we can write $S = \sum_\ell \rho_\ell P_\ell$ where $\sum_\ell \rho_\ell = 1$ and each P_ℓ is a permutation matrix (Birkhoff–von Neumann Theorem). Finally, we pick an ℓ with probability ρ_ℓ and assign the agents to positions according to the permutation matrix P_ℓ . The resulting allocation is exactly the desired \mathbf{x} .

Let β be the worst case, over number of units k , approximation factor of the multi-unit auction in the Bayesian, prior-independent, or prior-free sense. The position auction constructed is at worst a β approximation in the same sense. \square

We conclude that matroid permutation auctions reduce to position auctions which reduce to multi-unit auctions. But multi-unit environments are the simplest of matroid permutation environments, i.e., the uniform matroid (Section 4.6.1 page 137), where even the fact that the agents are permuted is irrelevant because uniform matroids are inherently symmetric. Therefore, from the perspective of optimization and approximation all of these problems are equivalent.

It is important to note, however, that this reduction may not preserve non-objective aspects of the mechanism. For instance, we have discussed that anonymous reserve pricing is a two approximation to virtual surplus maximization in multi-unit environments (e.g., Corollary 4.4.3 and Theorem 7.4.1). The reduction from matroid permutation and position environments does not imply that surplus maximization with an anonymous reserve gives a two approximation in these more general environments. This is because in the multi-unit two approximation via an anonymous reserve, the reserve is tailored to k , the number of units. Therefore, constructing a position auction or matroid mechanism would require simulating the multi-unit auction with potentially distinct reserve prices for each multi-unit supply constraint; the resulting mechanism will not generally be an anonymous-reserve mechanism.

In fact, for i.i.d., irregular, position and matroid permutation environments the surplus maximization mechanism with anonymous reserve is not generally a constant approximation to the optimal mechanism. The approximation factor via the anonymous reserve in these environments is $\Omega(\log n / \log \log n)$, i.e., there exists a distribution and matroid permutation and position environments such that the anonymous-reserve

mechanism has expected revenue that is a $\Theta(\log n / \log \log n)$ multiplicative factor from the optimal mechanism revenue (Exercise 7.4). The same inapproximation result holds with comparison between the anonymous-reserve and envy-free benchmarks.

Theorem 7.5.3. *There exists an i.i.d. distribution (resp. valuation profile), a matroid permutation environment, and position environment such that the (optimal) anonymous-reserve mechanism (resp. benchmark) is a $\Theta(\log n / \log \log n)$ approximation the Bayesian optimal mechanism (resp. envy-free benchmark).*

Implicit in the above discussion (and reductions) is the assumption that the characteristic weights for a matroid permutation setting can be calculated, or fundamentally, that the weights in the position auction are precisely known. Notice that in our application of position auctions to advertising on Internet search engines the position weights were the likelihood of a click for an advertisement in each position. These weights can be estimated but are not known exactly. The general reduction from matroid permutation and position auctions to multi-unit auctions requires foreknowledge of these weights.

Recall from the discussion of the multi-unit random sampling auction (Definition 7.4.3) that, as a virtual surplus maximizer, it does not require foreknowledge of the supply k of units. Closer inspection of the reductions of Theorem 7.5.2 reveals that if the given multi-unit auction is a virtual surplus maximizer then the weights do not need to be known to calculate the appropriate allocation. Simply maximize the virtual surplus for the realized environment.

In the definition of permutation environments, it is assumed that the agents are unaware of their roles in the set system, i.e., the agents' incentives are taken in expectation over the random permutation. A mechanism that is incentive compatible in this permutation model may not generally be incentive compatible if agents do know their roles. Therefore, matroid permutation auctions that result from the above reductions are not generally incentive compatible without the uniform random permutation. Of course the random sampling auction is a virtual surplus maximizer for the market and virtual surplus maximizers are dominant strategy incentive compatible (Theorem 3.3.11). Thus, the reduction applied to the random sampling auction is incentive compatible even if the permutation is known.

Corollary 7.5.4. *For any matroid environment and valuation profile,*

the random sampling auction is dominant strategy incentive compatible and when the values are randomly permuted, its expected revenue is a β approximation to the envy-free benchmark where β is its approximation factor for multi-unit environments.

7.6 Downward-closed Permutation Environments

In this section we consider downward-closed permutation environments. In multi-unit, position, and matroid permutation environments, virtual surplus maximization is ordinal, i.e., it depends on the relative order of the virtual values and not their magnitudes. In contrast, the main difficulty of more general downward-closed environments is that virtual surplus maximization is not generally ordinal. Nonetheless, variants of the random sampling (virtual surplus maximization) and the random sampling profit extraction auctions give constant approximations to the envy-free benchmark in downward-closed environments. We will describe only the latter result, which can be viewed as transforming the non-ordinal environment into an ordinal one.

The first step in this construction is to generalize the notion of a profit extractor (from [Section 7.2.4](#)). Our approach to profit extraction in downward-closed permutation environments will be the following. The true (and unknown) valuation profile is \mathbf{v} . Suppose we knew a profile \mathbf{v}^\dagger that was a coordinate-wise lower bound on \mathbf{v} , i.e., $v_{(i)} \geq v_{(i)}^\dagger$ for all i (short-hand notation: $\mathbf{v} \geq \mathbf{v}^\dagger$). A natural goal with this side-knowledge would be to design an incentive compatible mechanism that obtains at least the envy-free optimal revenue for \mathbf{v}^\dagger . We refer to mechanism that obtains this revenue, in expectation over the random permutation and whenever the coordinate-wise lower-bound assumption holds, as a profit extractor.

Definition 7.6.1. *The downward-closed profit extractor for \mathbf{v}^\dagger is the following:*

- (i) Sort \mathbf{v} and \mathbf{v}^\dagger in decreasing order.
- (ii) Reject all agents if there exists an i with $v_i < v_i^\dagger$.
- (iii) Calculate the empirical virtual values φ^\dagger for \mathbf{v}^\dagger .
- (iv) For all i , assign the i th highest-valued agent the i th highest virtual value φ_i^\dagger .
- (v) Serve the agents to maximize the virtual surplus.

Theorem 7.6.1. *For any downward-closed environment and valuation profiles \mathbf{v} and \mathbf{v}^\dagger , the downward-closed profit extractor for \mathbf{v}^\dagger is dominant strategy incentive compatible and if $\mathbf{v} \geq \mathbf{v}^\dagger$ then its expected revenue under a random permutation is at least the envy-free optimal revenue for \mathbf{v}^\dagger .*

Proof. See Exercise 7.5 □

To make use of this profit extractor we need to find a \mathbf{v}^\dagger that satisfies the assumption of the theorem and that is non-manipulable. The idea is to use biased random sampling. In particular, if the agents are partitioned into a sample with probability $p < 1/2$ and market with probability $1 - p$, then there is a high probability the valuation profile of the sample is a coordinate-wise lower bound on that of the market. Furthermore, we will show that even conditioned on this event, the expected optimal envy-free revenue of the sample approximates the envy-free benchmark. The approximate optimality of the mechanism follows.

Definition 7.6.2. *The biased (random) sampling profit extraction mechanism for downward-closed environments (with parameters $p \in (0, 1/2)$ and $\ell \in \{0, 1, 2, \dots\}$) is:*

- (i) *Assign the top ℓ agents to the market M .*
- (ii) *Randomly partition the remaining agents into S (with probability p) and M (with probability $1 - p$).*
- (iii) *Reject agents in S .*
- (iv) *Run the downward-closed profit extractor for \mathbf{v}_S on M* ³

Lemma 7.6.2. *The biased sampling profit extraction mechanism is dominant strategy incentive compatible.*

Proof. Fix any outcome of the n coins. Each agent i faces a critical value. Pretend the agent is in the market, and simulate the rest of the auction. The profit extractor is deterministic and dominant strategy incentive compatible; thus by Theorem 7.2.1 it induces a critical value \hat{v}_i . Now consider i 's coin. If the coin puts i in the market then she is offered critical value \hat{v}_i ; if the coin puts i in the sample, then she is offered $\max(v_{(\ell+1)}, \hat{v}_i)$, i.e., she wins only if she is in the top ℓ and would win in the profit extractor. □

³ The payments of the top ℓ agents are adjusted as follows. Flipping a biased coin for each such agent, but if she ends up in the sample (with probability p), she can buy her way into the market by agreeing to pay at least $v_{(\ell+1)}$. In such a case, her final payment is the maximum of her payment in the profit extraction mechanism and $v_{(\ell+1)}$.

The following lemma, which is key to the analysis, shows that the probability that $\mathbf{v}_M \geq \mathbf{v}_S$ in the biased sampling profit extraction auction is at least $1 - (p/1-p)^{\ell+1}$.

Lemma 7.6.3. *The probability of ruin of a biased random walk on the integers⁴ that steps back with probability $p < 1/2$, steps forward with probability $1 - p$, and starts from position one; is exactly $p/1-p$. If it starts at position k the probability of ruin is $(p/1-p)^k$.*

Proof. The proof is similar to that of Lemma 7.2.9. See Exercise 7.6. \square

The remainder of this section follows the the approach of prior-free combination developed in Section 7.4.2. Lemma 7.6.5 will bound the envy-free benchmark by the sum of two benchmarks, the envy-free benchmark restricted to the two highest-valued agents and the envy-free optimal revenue excluding these two agents. Lemma 7.6.6 will show that the second-price auction (to serve at most one agent) is a two approximation to the first benchmark and Lemma 7.6.7 will show that a biased sampling profit extraction auction is a 4.51 approximation to the second benchmark. We will conclude by Theorem 7.4.3 that the prior-free combination (Definition 7.4.2) of the two auctions is a 6.51 approximation to the envy-free benchmark.

Theorem 7.6.4. *In downward-closed permutation environments, the prior-free combination of the second-price auction with a biased sampling profit extraction auction is a 6.51 approximation to the envy-free benchmark.*

Lemma 7.6.5. *For any valuation profile \mathbf{v} , the envy-free optimal revenue for a subset S of agents is a subadditive function S . In particular, $\text{EFO}(\mathbf{v}) \leq \text{EFO}(\mathbf{v}_1, \mathbf{v}_2) + \text{EFO}(\mathbf{v}_{-1,2})$.*

Proof. Observe for disjoint sets A and B of agents,

$$\begin{aligned} \text{EFO}(\mathbf{v}_{A \cup B}) &= \text{EFO}_A(\mathbf{v}_{A \cup B}) + \text{EFO}_B(\mathbf{v}_{A \cup B}) \\ &\leq \text{EFO}(\mathbf{v}_A) + \text{EFO}(\mathbf{v}_B). \end{aligned}$$

The first line follows by definition where $\text{EFO}_A(\mathbf{v}_{A \cup B})$ denotes the contribution to the envy-free optimal revenue of $A \cup B$ from the agents in A , likewise for B . Of course, the envy-free optimal outcome for $A \cup B$ is envy free with respect to subset A . However, if we are only to consider

⁴ Recall, the probability of ruin of a random walk is the probability that it ever reaches position zero.

envy-freedom constraints of A , then this outcome for $A \cup B$ is not necessarily optimal. Thus, $\text{EFO}_A(\mathbf{v}_{A \cup B}) \leq \text{EFO}(\mathbf{v}_A)$; likewise for B ; and the second line follows. The left- and right-hand side of this equation give the definition of subadditivity. \square

Lemma 7.6.6. *For any downward-closed environment, the second-price auction is a 2-approximation to the envy-free benchmark restricted to the two highest-valued agents.*

Proof. Assume all singleton sets are feasible with respect to the downward-closed environment and the two highest valued agents have values $v_1 \geq v_2$. The second-price auction, which always only serves a single agent, is feasible and its revenue is v_2 . For the valuation profile $\mathbf{v}^\dagger = (2v_2, v_2)$, the revenues are $\mathbf{R}^\dagger = (2v_2, 2v_2)$ and the marginal revenues are $(2v_2, 0)$. Thus, the envy-free optimal revenue is obtained by only serving the first agent at a price of $2v_2$. \square

Lemma 7.6.7. *For any downward-closed permutation environment and any valuation profile, the biased sampling profit extraction auction with $p = .29$ and $\ell = 2$ is a 4.51 approximation to the envy-free optimal revenue on the valuation profile without the two highest-valued agents.*

Proof. Index the two highest-valued agents by 1 and 2. Let $\text{REF}(\mathbf{v}) = \text{EFO}(\mathbf{v}_{-1,2})$ be the envy-free optimal revenue on the valuation profile without the two highest valued agents, and $\text{APX}(\mathbf{v})$ be the expected revenue of the biased sampling profit extraction mechanism. We have,

$$\begin{aligned} \text{APX}(\mathbf{v}) &\geq \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathbf{v}_M \geq \mathbf{v}_S] \Pr[\mathbf{v}_M \geq \mathbf{v}_S] \\ &= \mathbf{E}[\text{EFO}(\mathbf{v}_S)] - \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathbf{v}_M \not\geq \mathbf{v}_S] \Pr[\mathbf{v}_M \not\geq \mathbf{v}_S]. \\ &\geq p \text{EFO}(\mathbf{v}_{-1,2}) - \text{EFO}(\mathbf{v}_{-1,2}) \Pr[\mathbf{v}_M \not\geq \mathbf{v}_S] \\ &\geq \left(p - \left(\frac{p}{1-p}\right)^3\right) \text{REF}(\mathbf{v}). \end{aligned}$$

The first line is by the definition of the mechanism and [Theorem 7.6.1](#). The second line is by the definition of conditional expectation. The first and second part of the third line are by subadditivity ([Lemma 7.6.5](#)) and monotonicity of the envy-free optimal revenue, respectively. The last line is from $\Pr[\mathbf{v}_M \not\geq \mathbf{v}_S] \leq (p/1-p)^3$ as guaranteed by [Lemma 7.6.3](#) for a random walk starting at position $\ell + 1 = 3$.

The expression $p - (p/1-p)^3$ is maximized at $p \approx 0.29$ giving an approximation of about 4.51 with respect to $\text{REF}(\mathbf{v})$. \square

Exercises

- 7.1 Complete the prior-free analysis framework for the objective of residual surplus in a two-agent single-item environment. The residual surplus is the sum of the values of the winners less any payments made.
- Identify a normalized benchmark.
 - Identify a distribution for which all auctions have the same residual surplus.
 - Give a lower bound on the resolution of your benchmark.
 - Give an upper bound on the prior-free optimal approximation with respect to your benchmark.
- Ideally, your lower bound on resolution should match your upper bound on prior-free optimal approximation.
- 7.2 Consider the design of prior-free incentive-compatible mechanisms with revenue that approximates the (optimal) social-surplus benchmark, i.e., $\text{OPT}(\mathbf{v})$, when all values are known to be in a bounded interval $[1, h]$. For downward-closed environments, give a $\Theta(\log h)$ approximation mechanism.
- 7.3 Consider a generalization of the mechanism composition from the construction of the multi-unit variant of a digital good auction, i.e., where the $k + 1$ st-price auction and the given digital good auction are composed ([Definition 7.4.1](#)). Two dominant strategy incentive compatible mechanisms A and B can be composed as follows: Simulate mechanism A ; run mechanism B on the winners of mechanism A ; and charge the winners of B the maximum of their critical values for A and B . A deterministic mechanism is *non-bossy* if there are no two values for any agent i such that the sets of winners of the mechanism are distinct but contain i .
- Show that the composite mechanism is dominant strategy incentive compatible when mechanism A is non-bossy.
 - Show that the surplus maximization mechanism in any single-dimensional environment is non-bossy.
- 7.4 Prove the envy-free variant of [Theorem 7.5.3](#) i.e., that there exists a valuation profile and a position environment for which the anonymous-reserve benchmark is a $\Omega(\log n / \log \log n)$ approximation to the envy-free benchmark.
- 7.5 Show that for any downward-closed environment and valuation profiles \mathbf{v} and \mathbf{v}^\dagger , the downward-closed profit extractor for \mathbf{v}^\dagger is

dominant strategy incentive compatible and if $\mathbf{v} \geq \mathbf{v}^\dagger$ then its expected revenue under random permutation is at least the envy-free optimal revenue for \mathbf{v}^\dagger . I.e., prove Theorem 7.6.1

- 7.6 Prove Lemma 7.6.3: The probability of ruin of a biased random walk on the integers; that steps back with probability $p < 1/2$, steps forward with probability $1 - p$, and starts from position one; is exactly $p/1-p$. If it starts at position k the probability of ruin is $(p/1-p)^k$.

Chapter Notes

The prior-free auctions for digital good environments were first studied by Goldberg et al. (2001) where the deterministic impossibility theorem and the random sampling optimal price auction were given. The random sampling optimal price auction was shown to be a constant approximation by Goldberg et al. (2006). The proof that the random sampling auction is a prior-free 15 approximation is from Feige et al. (2005); the bound was improved to 4.68 by Alaei et al. (2009). The profit extraction mechanism and the random sampling profit extraction mechanism were given by Fiat et al. (2002). The extension of this auction to three partitions was studied by Hartline and McGrew (2005).

The lower-bound on the approximation factor of prior-free auctions for digital goods of 2.42 was given by Goldberg et al. (2004); this bound was proven to be tight by Chen et al. (2014b). For the special cases of $n = 2$ and $n = 3$ agents the form of the optimal auction is known. For $n = 2$, Fiat et al. (2002) showed that the second-price auction is optimal and its approximation ratio is $\beta^* = 2$. For $n = 3$, Hartline and McGrew (2005) identified the optimal three-agent auction and showed that its approximation ratio is $\beta^* = 13/6 \approx 2.17$.

The formal prior-free design and analysis framework for digital good auctions was given by Goldberg et al. (2006). This framework was refined for general symmetric auction problems and grounded in the theory of Bayesian optimal auctions by Hartline and Roughgarden (2008). The connection between prior-free mechanism design and envy-freedom was given by Devanur et al. (2015) (originally as Hartline and Yan, 2011).

The 2-approximate reduction from multi-unit to digital-good environments combines results from Fiat et al. (2002) and Devanur et al. (2015). The improved reduction via “prior-free combination” that gives a multi-

unit $\beta + 1$ approximation from a digital-good β approximation is from [Chen et al. \(2014a\)](#).

Analysis of the random sampling auction for limited supply, position, matroid permutation, and downward-closed permutation environments was given by [Devanur et al. \(2015\)](#) (originally as [Devanur and Hartline, 2009](#)). For multi-unit auctions they prove the random sampling auction is a 9.6 approximation to the envy-free benchmark (i.e., [Theorem 7.4.5](#)) by extending the analysis of [Alaei et al. \(2009\)](#). They prove the equivalence between distributions over multi-unit environments, position environments, and matroid permutation environments which allows the 9.6 approximation bound for multi-unit environments to extend. For downward-closed permutation environments they give a variant of the random sampling auction that is a prior-free 189 approximation.

The downward-closed profit extractor is from [Ha and Hartline \(2011\)](#). [Devanur et al. \(2013\)](#) study the random sampling profit extraction auction, similar to the one described in this chapter, and show that it is a 7.5 approximation in downward-closed permutation environments. (They also give a variant of the auction for the case that the agents have a common budget.) The biased sampling profit extraction auction ([Definition 7.6.2](#)) and its analysis ([Theorem 7.6.4](#)) are from [Chen et al. \(2014a\)](#).

This chapter omitted discussion of a very useful technique for designing prior-free mechanisms using a “consensus mechanism” on statistically robust characteristics of the input. In this vein the consensus estimates profit extraction mechanism from [Goldberg and Hartline \(2003\)](#) obtains a 3.39 approximation for digital goods. This approach is also central in obtaining an asymmetric deterministic auction that gives a good approximation ([Aggarwal et al., 2005](#)). [Ha and Hartline \(2011\)](#) extend the consensus approach to downward-closed permutation environments.

This chapter omitted asymptotic analysis of the random sampling auction which is given [Balcan et al. \(2008\)](#). This analysis allows agents to be distinguished by publicly observable attributes and agents with distinct attributes may receive distinct prices.

The statistical methods of this chapter foreshadow a literature on sample complexity in mechanism design, e.g., [Cole and Roughgarden \(2014\)](#). Results from this literature identify the number of samples from the value distribution, as a function of ϵ , are sufficient to obtain a $(1 + \epsilon)$ -approximation to the Bayesian optimal auction.