

Chapter 6

Prior-free Mechanisms

The big challenge that separates mechanism design from (non-game-theoretic) optimization is that the incentive constraints in mechanism design bind across all valuation profiles. E.g., the payment of an agent depends on the what the mechanism does when the agent has a lower value (Theorem 2.7). Therefore, where optimization gives an outcome that is good point-wise (i.e., for any input), mechanism design gives a mechanism for all of type-space that must trade-off performance on one input for another.

In the last chapter we gave mechanisms that made this trade-off obliviously to the actual distribution. The resulting mechanisms were prior-independent and approximated the optimal mechanism for the implicit distribution. Furthermore, the described mechanisms were dominant-strategy incentive compatible, meaning, agent also need not know the distribution to act. This lack of distributional requirement for both the agent and the designer suggests that there must be a completely prior-free theory of mechanism design.

Intuitively, the class of good prior-free mechanisms should be smaller than the class of good prior-independent mechanisms. The prior-independent mechanism can rely on there being a distribution where as the prior-free mechanism cannot. Therefore, we demand from our prior-free design and analysis framework, that prior-free approximation implies prior-independent approximation. Indeed, up to constance factors, the results of this chapter subsume the results of the previous chapter.

A main challenge in considering a formal framework in which to design and analyze prior-free mechanisms is in identifying a meaningful benchmark against which to evaluate a mechanism's performance. For instance, it was natural to compare the prior-independent mechanisms of the previous chapter to the (Bayesian) optimal mechanism for the implicit, and unspecified, distribution. We define the meaningfulness of a benchmark by the implications of its approximation. A mechanism is a prior-free approximation to a given benchmark if the mechanisms performance on any valuation profile always approximates the benchmark performance. The benchmark is economically meaningful if, as desired by the previous paragraph, its approximation implies prior-independent approximation.

In this chapter we introduce the envy-free optimal revenue benchmark. An outcome, i.e., allocation and payments (\mathbf{x}, \mathbf{p}) , is envy-free if no agent prefers to swap outcome (allocation and payment) with another agent. Notice that the envy-freedom constraint binds point-wise

on valuation profiles; therefore, for any objective and valuation profile there is an envy-free outcome that is optimal. Envy-freeness can be viewed as a relaxation incentive compatibility, a view point that can be made precise in many environments, e.g., as the envy-free optimal revenue dominates the revenue of any (Bayesian) optimal mechanism. Thus, the focus of the chapter is on designing prior-free approximation mechanisms for this benchmark in general downward-closed environments.

6.1 The Digital Good Environment

Recall the digital good environment wherein all allocations are feasible. Given an i.i.d. distribution, the optimal mechanism would post the monopoly price as a take-it-or-leave-it offer to each agent. Of course, agents with values above the monopoly price would choose to purchase the item, while, agents with values below the monopoly price would not. This outcome is inherently envy-free as each agent was permitted to choose from among the two possible outcomes: either take item at the monopoly price, or take nothing and pay nothing.

Without a prior the monopoly price is not well defined; however, on inspection of the valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ it is easy to obtain an upper bound on the revenue of any monopoly pricing as $\max_i iv_{(i)}$. While it is not incentive compatible to inspect the valuation profile and offer the revenue maximizing price to each agent, it is envy free. Furthermore, though we do not argue it here, it gives the envy-free optimal revenue, denoted $\text{EFO}(\mathbf{v})$. Clearly, any mechanism that approximates this envy-free optimal revenue would also approximate the (Bayesian) optimal auction for any i.i.d. distribution. Therefore, this envy-free benchmark is economically meaningful.

Unfortunately, there is no prior-free constant approximation to this benchmark. In particular, when there is $n = 1$ agent the optimal envy-free revenue is the surplus, while we know that, even if the distribution on values is known (cf., Chapter 4, Section 4.2.1), the optimal surplus and revenue can be separated by more than a constant. For instance, if the agent's value is known to fall within the range $[1, h]$ then the best approximation factor is $1 + \ln h$ (See Exercise 6.1). Clearly, if nothing is known about the range of values then no finite approximation is possible.

In fact the only thing preventing $\max_i iv_{(i)}$ from being a good benchmark is the case where the maximization is obtained at $i = 1$ by selling to the highest value agent at her value. We therefore slightly alter the benchmark to exclude this scenario. The *envy-free (optimal) benchmark* for digital goods is $\text{EFO}^{(2)}(\mathbf{v}) = \max_{i \geq 2} iv_{(i)}$.

We now consider approximating this benchmark. In the remainder of this section we will show that deterministic auctions cannot give good prior-free approximation. We will then describe two approaches for designing 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, run the optimal auction for the estimated distribution on the remaining agents. The resulting auction is known to be a 4.68-approximation. The second auction is based on a standard algorithmic design paradigm: reduction to the “decision version” of the problem. The resulting auction is known to be a 4-approximation.

Finally, we describe a method for proving lower bounds on the approximation factor of any prior-free auction; no auction is better than a 2.42-approximation.

6.1.1 Deterministic Auctions

The main idea that enables approximation of the envy-free benchmark is that when figuring out a price to offer agent i we can use statistics from the values of all other agents \mathbf{v}_{-i} . This motivates the following mechanism.

Mechanism 6.1. *The deterministic optimal price auction offers each agent i the take-it-or-leave-it price of τ_i equal to the monopoly price for \mathbf{v}_{-i} .*

It is possible to show that the deterministic optimal price auction is a prior-independent constant approximation; however it is not a prior-free approximation. For example, 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 $\tau_h = 1$, 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 $\tau_1 = 10$, 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. Clearly with these offers all high-valued agents will win and pay one, while all low-valued agents will lose. The total revenue is ten, a far cry from the envy-free benchmark revenue of $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(\mathbf{v}) = 100$. In fact, this deficiency of the deterministic optimal price auction is one that is fundamental to all *anonymous* (a.k.a., symmetric) deterministic auctions.

Theorem 6.1. *No anonymous, deterministic digital good auction is better than an n -approximation to the envy-free benchmark.*

Proof. We consider only valuation profiles with values $v_i \in \{1, h\}$. Let $n_h(\mathbf{v})$ and $n_1(\mathbf{v})$ represent the number of h values and 1 values in \mathbf{v} , respectively. That an auction \mathcal{A} is anonymous implies that the critical value for agent i as a function of the reports of other agents is independent of i and only a function of $n_h(\mathbf{v}_{-i})$ and $n_1(\mathbf{v}_{-i})$. Thus, we can let $\tau(n_h, n_1)$ represent the offer price of \mathcal{A} 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 $\tau(n_h, n_1) \in \{1, h\}$ as this restriction cannot hurt the auction profit on the valuation profiles we are considering.

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

1. Observe that for any auction that is a good approximation, it must be that for all m , $\tau(m, 0) = h$. Otherwise, on the all h 's input, the auction only achieves profit n while the envy-free benchmark is hn . Thus, the auction would be at most an h -approximation which is not constant.

2. Likewise, observe that for any auction that is a good approximation, it must be that for all m , $\tau(0, m) = 1$. Otherwise, on the all 1's input, the auction achieves no profit and is clearly not an approximation of the envy-free benchmark n .
3. For the final argument, consider taking m sufficiently large and looking at $\tau(k, m - k)$. As we have argued for $k = 0$, $\tau(k, m - k) = 1$. Consider increasing k until $\tau(k, m - k) = h$. This must occur since $\tau(k, m - k) = h$ when $k = m$. Let $k^* = \min\{k : \tau(k, m - k) = h\} \geq 1$ be this transition point. Now consider an $n = m + 1$ agent valuation profile with $n_h(\mathbf{v}) = k^*$ and $n_1(\mathbf{v}) = m - k^* + 1$. Consider separately the offer prices to high- and low-valued agents:
 - For low-valued agents: $\tau(n_h(\mathbf{v}_{-1}), n_1(\mathbf{v}_{-1})) = \tau(k^*, m - k^*) = h$. Thus, all low-valued agents are rejected and contribute nothing to the auction profit.
 - For high-valued agents: $\tau(n_h(\mathbf{v}_{-h}), n_1(\mathbf{v}_{-h})) = \tau(k^* - 1, m - k^* + 1) = 1$. Thus, all high-valued agents are offered a price of one which they accept. Thus, the contribution to the auction profit from such agents is $1 \times n_h(\mathbf{v}) = k^*$.

Set $h = n$. 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 $k^* > 1$ the benchmark is also nk^* (from selling to the k^* high-valued agents at price n). Therefore, the auction profit k^* is at-best an n -approximation. \square

6.1.2 Random Sampling

The conclusion from the preceding discussion is that either randomization or asymmetry is necessary to obtain prior-free approximations. While either approach will permit the design of good mechanisms, all deterministic asymmetric auctions known to date are based on derandomizations of randomized auctions. In this text we will discuss only these randomized mechanisms.

Notice that the problem with the deterministic optimal price auction is that it sometimes offers high-valued agents a low price and low-valued agents a high price. Either of these prices would have been good if only it offered consistently to all agents. The first idea to combat 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.

Mechanism 6.2. *The random sampling (optimal price)*

1. *randomly partitions the agents into S' and S'' (by flipping a fair coin for each agent),*
2. *computes (empirical) monopoly prices η' and η'' for S' and S'' respectively, and*
3. *offers η' to S'' and η'' to S' .*

As a warm-up exercise for analyzing this random sampling auction we observe that its not better than a 4-approximation to the envy-free benchmark. 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 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 profit is therefor $1/2$, which is a 4-approximation to the benchmark.

Theorem 6.2. *For digital good environments and all valuation profiles, the random sampling auction is at least a 4.68-approximation to the envy-free benchmark.*

This theorem is involved and it is generally believed that the bound it provides is loose and the random sampling auction is in fact a worst-case 4-approximation. Below we will prove the weaker claim that it is at worst at 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 6.3. *For all valuation profiles, the random sampling auction is at least a 15-approximation to the envy-free benchmark.*

Proof. Assume without loss of generality that $v_{(1)} \in S'$ and call S' the market; call S'' the sample. This terminology comes from the fact that if $v_{(1)}$ is much bigger than all other agent values then all agents in 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 1 is to show that $\text{EFO}(\mathbf{v}_{S''})$ is close to $\text{EFO}^{(2)}(\mathbf{v})$. Step 2 is to show that the revenue from price η'' on S' is close to $\text{EFO}(\mathbf{v}_{S''})$, i.e., the revenue from price η'' 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 X_i is an indicator variable for the event that $i \in S''$ (the sample). Notice that $\mathbf{E}[X_i] = 1/2$ except for $i = 1$; $X_1 = 0$ by our assumption that the highest valued agent is in the market. Define $S_i = \sum_{j < i} X_j$. Let k be the number of winners in $\text{EFO}(\mathbf{v})$, i.e., $k = \text{argmax}_i i v_i$.

1. With good probability, the optimal revenue for the sample, $\text{EFO}(\mathbf{v}_{S''})$, is close to the benchmark, $\text{EFO}^{(2)}(\mathbf{v})$.

Define the event \mathcal{B} that $S_k \geq k/2$. Of course $\text{EFO}(\mathbf{v}_{S''}) \geq S_k v_k$ as the former is the optimal single price revenue on S'' and the latter is the revenue from S'' with price v_k . Event \mathcal{B} implies that $S_k v_k \geq k v_k / 2 = \text{EFO}^{(2)}(\mathbf{v}) / 2$, and thus, $\text{EFO}(\mathbf{v}_{S''}) \geq \text{EFO}^{(2)}(\mathbf{v}) / 2$.

We now show that $\Pr[\mathcal{B}] = 1/2$ when k is even. Recall that the highest valued agent is always in the market. Therefore there are $k - 1$ (an odd number) of agents which we partition between the market and the sample. One partition receives at least $k/2$ of these and half the time it is the sample; therefore, $\Pr[\mathcal{B}] = 1/2$. When k is odd $\Pr[\mathcal{B}] < 1/2$, and a slightly more complicated argument is needed to complete the proof. We omit the details.

2. With good probability, the revenue from price η'' on S' is close to $\text{EFO}(\mathbf{v}_{S''})$.

Define the event \mathcal{E} that “ $\forall i, (i - S_i) \geq S_i/3$.” 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. I.e., 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 k'' be index of the agent whose value is the monopoly price for the sample, i.e., $v_{k''} = \eta''$ and $\text{EFO}(\mathbf{v}_{S''}) = S_{k''}v_{k''}$. The profit of the random sampling auction is equal to $(k'' - S_{k''})v_{k''}$. Under the balanced sample condition this is lower bounded by $S_{k''}v_{k''}/3 = \text{EFO}(\mathbf{v}_{S''})/3$.

We defer to later the proof of a *balanced sampling lemma* (Lemma 6.4) that shows that $\Pr[\mathcal{E}] \geq .9$.

Finally, we combine these two pieces. If both good events \mathcal{E} and \mathcal{B} hold, then the expected revenue of random sampling auction is at least $\text{EFO}^{(2)}(\mathbf{v})/6$. By the union bound, the probability of this good fortune is $\Pr[\mathcal{E} \wedge \mathcal{B}] = 1 - \Pr[\neg\mathcal{E}] - \Pr[\neg\mathcal{B}] \geq 0.4$. We conclude that the random sampling auction is a 15-approximation to the envy-free benchmark. \square

Lemma 6.4 (Balanced Sampling). *For $X_1 = 0$, X_i for $i \geq 1$ an indicator variable for a independent fair coin flipping to heads, and sum $S_i = \sum_{j \leq i} X_j$,*

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

Proof. We relate the condition to the *probability of ruin* in a *random walk* on the integers. Notice that $(i - S_i) \geq S_i/3$ if and only if, for integers i and S_i , $3i - 4S_i + 1 > 0$. So let $Z_i = 3i - 4S_i + 1$ and view Z_i as the position, in step i , of a random walk on the integers. Since $S_1 = 0$ this random walk starts at $Z_1 = 4$. Notice that at step i in the random walk with $Z_i = k$ then at step $i + 1$ we have

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

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 the *probability of ruin* in analogy 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 $r_0 = 1$ (at $k = 0$ we are already ruined) and $r_k = r_1^k$ (taking k steps back is equal to stepping back k times). By the definition of the random walk, we have the recurrence

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

Plugging in the above identities,

$$r_1 = \frac{1}{2}(1 + r_1^4).$$

This is a quartic equation that can be solved, e.g., by *Ferarri's formula*. Since our random walk starts at $Z_1 = 4$ we calculate $r_4 = r_1^4 \leq 0.1$, meaning that the success probability for the random walk satisfying the balanced sampling condition is at least 0.9. \square

6.1.3 Decision Problems

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, often times the opposite is also true. For instance, with binary search and repeated calls to an algorithm that solves the decision problem, the optimal solution can be found. In this section we develop a similar theory for mechanism design.

Profit extraction

For profit maximization in mechanism design, recall, there is no absolutely optimal mechanism. Therefore, we define the mechanism design decision problem in terms of the aforementioned profit benchmark EFO. The decision problem for EFO and profit target R to design a mechanism that obtains profit at least R on any input \mathbf{v} with $\text{EFO}(\mathbf{v}) \geq R$. We call the mechanism that solves the decision problem a *profit extractor*.

Definition 6.5. *The digital good profit extractor for target R and valuation profile \mathbf{v} finds the largest k such that $v_{(k)} \geq R/k$, sells to the top k agents at price R/k , and rejects all other agents. If no such set exists, it rejects all agents.*

Lemma 6.6. *The digital good profit extractor is dominant strategy incentive compatible.*

Proof. Consider the following indirect mechanism. See if all agents can evenly split the target R . If some agents cannot afford to pay their fair share, reject them. Repeat with the remaining agents. Notice that as the number of agents in this process is decreasing, the fair share that each agent faces is increasing. Therefore, any agent rejected for inability to pay their fair share could not afford any of the future prices considered in the mechanism either. Thus, the incentives are identical to that of the English auction. An agent wishes to drop out when the increasing price surpasses her value.

The digital good profit extractor is obtained by applying the revelation principle to the above ascending price mechanism. \square

Lemma 6.7. *For all valuation profiles \mathbf{v} , the digital good profit extractor for target R obtains revenue R if $R \leq \text{EFO}(\mathbf{v})$ and zero otherwise.*

Proof. $\text{EFO}(\mathbf{v}) = kv_{(k)}$ for some k . If $R \leq \text{EFO}(\mathbf{v})$ then $R/k \leq v_{(k)}$. The digital good profit extractor may yet find a larger k that satisfies the same property, however, it can certainly find some k . On the other hand, if $R > \text{EFO}(\mathbf{v}) = \max_k kv_{(k)}$ then there is no such k for which $R/k \leq v_{(k)}$ and the mechanism has no winners and no revenue. \square

Approximate Reduction to Decision Problem

We now use 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 t the agent buys and pays t if and only if $v \geq t$. Analogously a profit extractor with target R obtains revenue R on \mathbf{v} if and only if $\text{EFO}(\mathbf{v}) \geq R$. The idea then is to randomly partition the agents and use profit extraction to run the second-price auction on the benchmark profit from each partition.

Definition 6.8. *The random sampling profit extraction auction works as follows:*

1. *Randomly partition the agents by flipping a fair coin for each agents and assigning her to S' or S'' .*
2. *Calculate $R' = \text{EFO}(\mathbf{v}_{S'})$ and $R'' = \text{EFO}(\mathbf{v}_{S''})$, the benchmark profit for each part.*
3. *Profit extract R'' from S' and R' 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 R' and R'' . Since the profit extractor is dominant strategy incentive compatible, so is the random sampling profit extraction auction.

Lemma 6.9. *The random sampling profit extraction auction is dominant strategy incentive compatible.*

Before we prove that the auction is a 4-approximation to to the envy-free benchmark, we give a simple proof of a lemma that will be important in the analysis.

Lemma 6.10. *Flip $k \geq 2$ fair coins, then*

$$\mathbf{E}[\min\{\#heads, \#tails\}] \geq \frac{k}{4}.$$

Proof. Let M_i be a random variable for the $\min\{\#heads, \#tails\}$ after only i coin flips. We make the following basic calculations (verify these as an exercise):

- $\mathbf{E}[M_1] = 0.$
- $\mathbf{E}[M_2] = 1/2.$
- $\mathbf{E}[M_3] = 3/4.$

We now obtain a general bound on $\mathbf{E}[M_i]$ for $i > 3$. Let $X_i = M_i - M_{i-1}$ representing the change to $\min\{\#\text{heads}, \#\text{tails}\}$ after flipping one more coin. Notice that linearity of expectation implies that $\mathbf{E}[M_k] = \sum_{i=1}^k \mathbf{E}[X_i]$. Thus, it would be enough to calculate $\mathbf{E}[X_i]$ for all i . We consider this in two 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, i.e., $\#\text{heads} \neq \#\text{tails}$. Assume without loss of generality that $\#\text{heads} < \#\text{tails}$. 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, $\mathbf{E}[X_i] = 1/2$.

Case 2 (i odd): Here we use the crude bound that $\mathbf{E}[X_i] \geq 0$. Note that this is actually the best we can claim in worst case since $i - 1$ is even and it could have been that $\#\text{heads} = \#\text{tails}$ in the previous round. If this were the case then regardless of the i th coin flip, $X_i = 0$ and the minimum of $\#\text{heads}$ and $\#\text{tails}$ 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 $\mathbf{E}[M_2] = 1/2$ and $\mathbf{E}[M_3] = 3/4$ above to deduce that $\mathbf{E}[X_3] = \mathbf{E}[M_3] - \mathbf{E}[M_2] = 1/4$.

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

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

Theorem 6.11. *For digital good environments and all valuation profiles, the revenue of the random sampling profit extraction auction is a 4-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 aforementioned analogy, the expected revenue of the auction is $\text{APX} = \mathbf{E}[\min(R', R'')]$ (where the expectation is taken over the randomized of the partitioning of agents).

Assume that envy-free benchmark sells to $k \geq 2$ agents at price p , i.e., $\text{REF} = kp$. Of the k winners in REF, let k' be the number of them that are in S' and k'' the number that are in S'' . Since there are k' agents in S' at price p , then $R' \geq k'p$. Likewise, $R'' \geq k''p$.

$$\begin{aligned} \frac{\text{APX}}{\text{REF}} &= \frac{\mathbf{E}[\min(R', R'')]}{kp} \\ &\geq \frac{\mathbf{E}[\min(k'p, k''p)]}{kp} \\ &= \frac{\mathbf{E}[\min(k', k'')]}{k} \\ &\geq \frac{1}{4}. \end{aligned}$$

The last inequality follows from applying Lemma 6.10 when we consider $k \geq 2$ coins and heads as putting an agent in S' and a tails as putting the agent in S'' .

This bound is tight as is evident from the same example from which we concluded that the random sampling optimal price auction is at best a 4-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 k parts for any k . In fact, the optimal partitioning comes from $k = 3$, though we omit the proof and full definition of the mechanism.

Theorem 6.12. *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.*

6.1.4 Lower bounds

We have discussed three auctions with known approximation factors 4.62, 4, and 3.25. What is the best approximation factor possible? This question, of course, turns our framework of approximation into one of optimality.

Definition 6.13. *The prior-free optimal auction for a envy-free benchmark $\text{EFO}^{(2)}$ is*

$$\operatorname{argmin}_{\mathcal{A}} \max_{\mathbf{v}} \frac{\text{EFO}^{(2)}(\mathbf{v})}{\mathcal{A}(\mathbf{v})}.$$

Unfortunately, this optimal auction suffers from the main drawback of optimal mechanisms. In general it is quite complicated. The auctions described heretofore can be viewed as simple approximations to this potentially complex optimal auction.

For the special case of $n = 2$, however, the prior-free optimal auction is simple. In this case, the envy-free benchmark is $\text{EFO}^{(2)}(\mathbf{v}) = 2v_{(2)}$. Recall that the revenue of the second-price auction is $v_{(2)}$. Therefore, in this special case, the second-price auction is a 2-approximation to the benchmark. Is this the best possible or is there some better approximation factor possible by a more complicated auction? In fact, it is the best possible.

Lemma 6.14. *For any auction, there is a $n = 2$ agent valuation profile such that the auction is at best a 2-approximation to the envy-free benchmark.*

Proof. The proof follows a simple structure that is useful for proving lower bounds for this type of problem. First, we consider values drawn from a random distribution. Second, we argue that for any auction \mathcal{A} and \mathbf{v} i.i.d. from F , $\mathbf{E}_{\mathbf{v}}[\mathcal{A}(\mathbf{v})] \leq \mathbf{E}_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v})]/2$. By the definition of expectation this implies that there exists a valuation profile \mathbf{v}^* such that $\mathcal{A}(\mathbf{v}^*) \leq \text{EFO}^{(2)}(\mathbf{v}^*)/2$ (as otherwise the expected values could not satisfy this condition).

We choose a distribution to make the analysis of $\mathbf{E}_{\mathbf{v}}[\mathcal{A}(\mathbf{v})]$ simple. This is important because we have to analyze it for all auctions \mathcal{A} . The idea is to choose the distribution for \mathbf{v} such that all auctions obtain the same expected profit. The distribution that satisfies this condition is the equal-revenue distribution (Definition 4.4), i.e., $F(z) = 1 - 1/z$. Note that whatever price $\tau_i \geq 1$ that \mathcal{A} offers agent i , the expected payment made by agent

i is $\tau_i \times \Pr[v_i \geq \tau_i] = 1$. Thus, for $n = 2$ agents the expected profit of the auction is $\mathbf{E}_{\mathbf{v}}[\mathcal{A}(\mathbf{v})] = n = 2$.

We must now calculate $\mathbf{E}_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v})]$. $\text{EFO}^{(2)}(\mathbf{v}) = \max_{i \geq 2} i v_{(i)}$ where $v_{(i)}$ is the i th highest valuation. In the case that $n = 2$, this simplifies to $\text{EFO}^{(2)}(\mathbf{v}) = 2v_{(2)} = 2 \min(v_1, v_2)$. We recall that a non-negative random variable X has $\mathbf{E}[X] = \int_0^\infty \Pr[X \geq z] dz$ and calculate $\Pr[\text{EFO}^{(2)}(\mathbf{v}) > z]$.

$$\begin{aligned} \Pr_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v}) > z] &= \Pr_{\mathbf{v}}[v_1 \geq z/2 \wedge v_2 \geq z/2] \\ &= \Pr_{\mathbf{v}}[v_1 \geq z/2] \Pr_{\mathbf{v}}[v_2 \geq z/2] \\ &= 4/z^2. \end{aligned}$$

Note that this equation is only valid for $z \geq 2$. Of course for $z < 2$, $\Pr[\text{EFO}^{(2)}(\mathbf{v}) \geq z] = 1$.

$$\begin{aligned} \mathbf{E}_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v})] &= \int_0^\infty \Pr[\text{EFO}^{(2)}(\mathbf{v}) \geq z] dz \\ &= 2 + \int_2^\infty \frac{4}{z^2} dz = 4. \end{aligned}$$

Thus we see that for this distribution and any auction \mathcal{A} , $\mathbf{E}_{\mathbf{v}}[\mathcal{A}(\mathbf{v})] = 2$ and $\mathbf{E}_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v})] = 4$. Thus, the inequality $\mathbf{E}_{\mathbf{v}}[\mathcal{A}(\mathbf{v})] \leq \mathbf{E}_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v})] / 2$ holds and there must exist some input \mathbf{v}^* such that $\mathcal{A}(\mathbf{v}^*) \leq \text{EFO}^{(2)}(\mathbf{v}^*)/2$. \square

For $n > 2$ the same proof schema gives lower bounds on the approximation factor of the prior-free optimal auction. The main difficulty of the $n > 2$ case is in calculating the expectation of the benchmark. This 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, its expectation can be calculated exactly.

For any auction, there is a $n = 2$ agent valuation profile such that the auction is at best a 2-approximation to the envy-free benchmark.

Theorem 6.15. *For any auction, there is a valuation profile such that the auction is at best a 2.42-approximation to the envy-free benchmark. Furthermore, for special cases of $n = 2, 3,$ and 4 agents the lower bound on approximation factors are exactly $2, 13/6,$ and $96/215,$ respectively.*

It is known that there is a $13/6$ -approximation for $n = 3$ agents. It is not known whether the $96/215$ bound is tight for $n = 4$.

6.2 The Envy-free Benchmark

To generalize beyond digital good environments we must be formal about the envy-free benchmark. First, is the envy-free benchmark meaningful in multi-unit, matroid, or downward-closed environments? For instance, informally we would like prior-free approximation of the benchmark to imply prior-independent approximation for any i.i.d. prior. Second, is it analytically tractable, i.e., is there an easy to interpret description of envy-free optimal pricings? Both of these issues are important.

Recall that in the digital goods example the envy-free benchmark is the revenue from the monopoly pricing of the empirical distribution given by the valuation profile. This seems like a reasonable benchmark as the Bayesian optimal auction for digital goods is the monopoly pricing (for the real distribution). Recall that for irregular multi-unit auction environments the optimal auction is not just the second-price auction with the monopoly reserve (in particular, it may iron). For these environments the envy-free benchmark is also more complex.

Up to this point, we have assumed that the environment is 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., 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 is randomized.

Definition 6.16. *Given an environment, specified by cost function $c(\cdot)$, the permutation environment is the convex combination of the environment with the identities of the agents 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 random permutations. Therefore, with respect to our goal, 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 permits a meaningful study of envy-freedom.

Definition 6.17. 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.,

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

The definition of envy freedom should be contrasted to the Bayes-Nash equilibrium condition given by Fact 2.6. Importantly, Bayes-Nash equilibrium 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. However, unlike the incentive-compatibility constraints, no-envy constraints bind point-wise on the given valuation profile; therefore, there is always a point-wise optimal envy-free outcome. The similarity of envy freedom and incentive compatibility enables virtually identical characterization and optimization of envy free outcomes (cf. Theorem 2.7).

Theorem 6.18. For valuation profile \mathbf{v} (sorted with $v_1 \geq v_2 \geq \dots \geq v_n$), an outcome (\mathbf{x}, \mathbf{p}) is envy free if and only if

- (monotonicity) $x_1 \geq x_2 \geq \dots \geq x_n$.
- (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. Given our objective of profit maximization, for any monotone allocation rule, we focus on the largest envy-free payments. As this 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. Formulaically this revenue can be calculated as:

$$p_i = \sum_{j \geq i}^n v_j (x_j - x_{j+1}). \quad (6.1)$$

We can define the revenue curve, marginal revenue, virtual values, and their ironed equivalents that correspond to envy-free revenue (cf. Definitions 3.11, 3.14, and 3.23). In fact, these terms are exactly those that govern the Bayesian optimal revenue for the *empirical distribution*. The empirical distribution for a valuation profile is the distribution with mass i/n above value $v_{(i)}$.

For envy-free revenue, the index of an agent (in the sorted order) plays the same role as quantile in the analogous definitions of Bayesian optimal mechanisms in Chapter 3 (cf.




Figure 6.1: The revenue curves corresponding to 10 high-valued agents and 80 low-valued agents. Depicted on $R(\cdot)$ are the revenues of the second-price auctions with reserves 10 and 2 with $k = 20$ units. Depicted on $\bar{R}(\cdot)$ is the envy-free optimal revenue with $k = 20$ units.

Definition 3.9). For these natural definitions, the optimal envy-free outcome in any symmetric environment is the ironed virtual surplus optimizer (cf. Corollary 3.27). In particular for permutation environments the desired allocation can be calculated as follows: first, calculate ironed virtual values from values; second, realize the random permutation; and third, serve the subset of agents to maximize the ironed virtual surplus.

Definition 6.19. For the i th index, the revenue is $R_i = iv_i$; the virtual value is $\phi_i = R_i - R_{i-1}$; the ironed revenue is denoted by \bar{R}_i and given by the evaluating at i the smallest concave function that upper bounds the point set given by $\{0, 0\} \cup \{(i, R_i) : i \in [n]\}$; and the ironed virtual value is $\bar{\phi}_i = \bar{R}_i - \bar{R}_{i-1}$.

Theorem 6.20. The maximal envy-free revenue for monotone allocation \mathbf{x} is

$$\sum_i \phi_i x_i = \sum_i R_i(x_i - x_{i+1}) \leq \sum_i \bar{\phi}_i x_i = \sum_i \bar{R}_i(x_i - x_{i+1})$$

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

Theorem 6.21. In symmetric environments, ironed virtual surplus maximization (with random tie-breaking) gives the envy-free outcome with the maximum profit.

In a symmetric environment, ironed 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., $z > z' \Rightarrow x_i(z) \geq x_i(z')$. The maximal envy-free payment of agent i for this allocation comes from equation (6.1) whereas the payment of the incentive compatible mechanism with this allocation rule comes from Corollary 2.17. These payments are related but distinct.

For example, consider a $k = 20$ unit environment and 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 revenue curve for this valuation profile is given in Figure 6.1(a). Both selling ten




Figure 6.2: The allocation rules for high- and low-valued agents induced by the mechanism with virtual values given in the text on the valuation profile given in the text. The payments are given by the area of the shaded region.

units at price ten and 20 units at price two are envy-free. The envy-free optimal revenue, however, is given by selling to the high-valued agents with probability one and price nine and selling to the low-valued agents with probability $1/8$ and price two. It is easy to verify that this outcome is envy-free and that its total revenue is 110. The ironed revenue curve for this valuation profile is given in Figure 6.1(b). The ironed virtual values are given by the following function:

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

We now calculate the revenue of the incentive compatible mechanism that serves the 20 agents with the highest ironed virtual value. In the virtual-surplus-maximizing auction, on the valuation profile \mathbf{v} (with ten high-valued agents and 80 low-valued agents), 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 6.2(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 6.2(b). Payments can be read from the allocation rules: a high-valued agent pays about 8.9 and a low-valued agent (in expectation) pays $1/4$. The total revenue from ten of each is about 109. Notice that this revenue is only slightly below the envy-free optimal revenue.

The revenue calculation above was 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 Figures 6.2(a) and 6.2(b)). The envy-free optimal revenue can be viewed as an approximation of the incentive-compatible revenue that is more analytically tractable.

We now formalize the fact that the envy-free revenue is an economically meaningful benchmark. The theorem below implies that, in matroid permutation environments, prior-free approximation of the benchmark implies prior-independent approximation.

Theorem 6.22. *For any matroid permutation environment and any ironed virtual value function $\bar{\phi}(\cdot)$, ironed virtual surplus maximization’s envy-free revenue 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 ironed virtual surplus optimizer in the permutation environment, then for $z \leq v_i$, $x_i(\mathbf{v}_{-i}, z)$ (as a function of z) is at most the smallest $y(z)$ that satisfies the conditions of Theorem 6.18. Since the incentive-compatible and envy-free payments, respectively, correspond to the area “above the curve” this inequality implies the desired payment inequality.

Since $x_i(\mathbf{v}_{-i}, z)$ is monotone, we only evaluate it at $v_j \leq v_i$ and show that $x_i(\mathbf{v}_{-i}, v_j) \geq x_j(\mathbf{v})$. This can be seen by the following sequence of inequalities.

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

The equality above comes from the symmetry of the environment and the fact that agent i and j have the same value in profile (\mathbf{v}_{-i}, v_j) . The inequality comes from the matroid assumption and the fact that the greedy algorithm is optimal (Theorem 4.22): when agent i reduces her bid to v_j , agent j is less likely to be blocked by i . \square

We are now ready to formally define the envy-free benchmark. Notice that the envy-free benchmark is well defined in all environments not just symmetric environments. For instance, when we wish to compare a mechanisms performance to the envy-free benchmark, it is not necessary for the environment to be symmetric.

Definition 6.23. *Given any environment, let $\text{EFO}(\mathbf{v})$ denote the maximum revenue attained by an envy-free outcome in the corresponding permutation environment.*

Definition 6.24. *Let $\mathbf{v}^{(2)} = (v_{(2)}, v_{(2)}, \dots, v_{(n)})$ be the valuation profile with $v_{(1)}$ replaced with a duplicate of $v_{(2)}$ and define $\text{EFO}^{(2)}(\mathbf{v}) = \text{EFO}(\mathbf{v}^{(2)})$.*

6.3 Multi-unit Environments

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 lose a factor of two

in the approximation ratio, i.e., it will derive a 2β -approximation for multi-unit environments from any β -approximation for digital goods. 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 two part, calculates ironed virtual valuation functions for the empirical distribution of each part, and then runs optimal $k/2$ -unit auction on each part using the ironed virtual valuation function from the opposite part.

Our first approach is an approximate reduction. For i.i.d., irregular, single-item environments Corollary 4.12 shows that the second-price auction with anonymous reserve is a 2-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 $\max_{i \leq k} iv_{(i)}$. 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 $\max_{2 \leq i \leq k} iv_{(i)}$. 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.

Theorem 6.25. *For any valuation profile, in multi-unit environments, the envy-free benchmark is at most twice the anonymous-reserve benchmark.*

Proof. Assume without loss of generality that the envy-free optimal revenue is derived from selling all k units. In terms of revenue curves (Definition 6.19), the envy-free optimal revenue for \mathbf{v} is $\text{REF} = \max_{i \leq k} \bar{R}_i$ whereas the anonymous-reserve revenue is $\text{APX} = \max_{i \leq k} R_i$.

Assume without loss of generality that the envy-free optimal revenue sells all k units and irons between index $i < k$ and $j > k$ as depicted in Figure 6.3. We will use a short hand notation and refer to the value of a point as the value of its y -coordinate. Accordingly, $C = \text{REF} = \bar{R}_k$, $A = R_i = iv_{(i)}$, $E = R_j = jv_{(j)}$, and $D = \frac{k}{j}R_j = kv_{(j)}$. Note that $v_{(k)} \geq v_{(j)}$ so $R_k \geq D$.

By definition the anonymous-reserve revenue satisfies $\text{APX} = \max_i R_i$ so $A \leq \text{APX}$ and $D \leq \text{APX}$ so $A + D \leq 2 \text{APX}$. But, line segment AB is certainly longer than line segment CD so $\text{REF} = C \leq A + D \leq 2 \text{APX}$.

Finally, this inequality holds for any \mathbf{v} therefore it also holds for $\mathbf{v}^{(2)}$. □




Figure 6.3: Depiction of ironed revenue curve $\bar{\mathbf{R}}$ for the pictorial proof of Theorem 6.25. The solid piece-wise linear curve is $\bar{\mathbf{R}}$, the convex hull of \mathbf{R} , and contains the line-segment connecting point $A = (i, R_i)$ and point $E = (j, R_j)$. The envy-free benchmark is achieved at point $D = (k, \bar{R}_k)$. The dashed lines have slope $v_{(i)}$ and $v_{(j)}$.

Theorem 6.25 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{A} into a limited supply auction and if \mathcal{A} 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.

Mechanism 6.3. *The k -unit variant \mathcal{A}_k of digital good auction \mathcal{A} is the following:*

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

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 (Exercise 6.3). The proof of the following theorem is immediate.

Theorem 6.26. *If \mathcal{A} is a β -approximation in digital good environments then its multi-unit variant \mathcal{A}_k is a 2β -approximation in multi-unit environments (with respect to the envy-free benchmark).*

We can of course apply this theorem to any digital good auction; for instance, from Theorem 6.12 we can conclude the following corollary.

Corollary 6.27. *There is an multi-unit auction that is a 6.5-approximation to the envy-free benchmark.*

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 two parts and then derives the optimal auction for each part and runs that auction on the opposite part. For digital goods the optimal auction for each part is just the to post the monopoly price for the valuation profile. Of course, multi-unit environments the optimal auction, e.g., for irregular distributions, may iron.

Mechanism 6.4. *The random sampling (ironed virtual surplus maximization) auction for the k -unit environment*

1. *randomly partitions the agents into S' and S'' (by flipping a fair coin for each agent),*
2. *computes ironed virtual valuation functions $\bar{\phi}'$ and $\bar{\phi}''$ for the empirical distributions of S' and S'' respectively, and*
3. *maximizes ironed virtual surplus on S'' with respect to $\bar{\phi}'$ and S' with respect to $\bar{\phi}''$ with $k/2$ -units each.*

If k is odd the last unit is allocated with probability $1/2$ to each part.

The proof of the following theorem can be derived similarly to the proof of Lemma 6.3; we omit the details.

Theorem 6.28. *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 each part is an ironed-virtual-surplus optimization. I.e., in each part it sorts the agents by ironed virtual surplus 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 to alternating partitions to the agent remaining with the highest ironed virtual valuation. Second, as we will see in the next section, it can be applied without specialization to matroid permutation and position environments.

6.4 Matroid Permutation and Position Environments

Position environments are important as they model auctions for selling advertisements on Internet search engines such as Google, Yahoo!, and 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 6.29. A position environment is one with n agents, m 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 their ad is clicked (e.g., if the searcher buys the advertisers product) and if their 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 $v_i w_j$. Though this model of Internet advertising leaves out many details of the environment, it captures many others.

We now show that mechanism design for matroid permutation environments can be reduced to position auctions which can be reduced to k -unit auctions. The main intuition that underlies this reduction is provided by the following definition.

Definition 6.30. 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., by definition, 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 ironed virtual value, w_i is the probability of serving the agent with the i th highest ironed virtual value.

Theorem 6.31. 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 matroid permutation auction that achieves the same approximation factor to the optimal mechanism. These two results imply that any Bayesian, prior-independent, or prior-free approximation results for position auctions extend to matroid permutation environments.

1. Revenue optimal auctions are ironed 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 ironed virtual value with probability w_j . (In both environments agents with negative ironed virtual values are discarded.) Expected revenue equals expected virtual surplus; therefore, the optimal expected revenues in the two environments are the same.

2. 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 $v'_i = n - j_i + 1$ 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 position auctions to single-item multi-unit auctions. This reduction implies that the approximation factor of a given multi-unit auction in an i.i.d. distributions can immediately be extended to matroid permutation and position environments. Furthermore, the mechanism that gives this approximation can be derived from the multi-unit auction.

Theorem 6.32. *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 6.31.

Let $d_j = w_j - w_{j+1}$ be the difference between successive weights. Recall that without loss of generality $w_1 = 1$ so \mathbf{d} gives a probability measure over $[m]$.

1. 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{d} . 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 ironed virtual value is served with probability w_j . (In both settings agents with negative ironed virtual values are discarded.) Therefore, the expected revenues in the two environments are the same.
2. 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 $j \in [m]$ and let $x_i^{(j)}$ be the (potentially random) indicator for whether agent i is allocated in simulation j . Let $x_i = \sum_j x_i^{(j)} d_j$ be the expected allocation to j in the convex combination of multi-unit auctions given by measure \mathbf{d} . Reindex \mathbf{x} in non-increasing order. Then \mathbf{w} majorizes \mathbf{x} in the sense that $\sum_i^k w_i \geq \sum_i^k x_i$ (and with equality for $k = m$). 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 in the permutation specified by 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, 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 2-approximation to ironed virtual surplus maximization in multi-unit environments (e.g., Corollary 4.12 and Theorem 6.25). The reduction from matroid permutation and position environments does not imply that surplus maximization with an anonymous reserve gives a 2-approximation in these more general environments. This is because in the multi-unit 2-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 various supply constraints and reserve prices; the resulting mechanism would not 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 matroid permutation and position environments, and distribution 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. We leave this result as an exercise with the hint that the distribution that gives this result is a generalization of the Sydney opera house distribution (Definition 4.5). The same inapproximation result holds with comparison between the anonymous-reserve and envy-free benchmarks.

Theorem 6.33. *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) 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 (Mechanism 6.4) that, as an ironed virtual surplus maximizer, it does not require foreknowledge of the supply k of units. Closer inspection of the reductions of Theorem 6.32 reveals that if the given multi-unit auction is an ironed virtual surplus maximizer then the weights do not need to be known to calculate the appropriate allocation. Simply maximize the ironed virtual surplus.

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 permutation. Of course the random sampling auction is a ironed virtual surplus maximizer and ironed virtual surplus maximizers are dominant strategy incentive compatible (Theorem 3.25).

Corollary 6.34. *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.*

6.5 Downward-closed Permutation Environments

In multi-unit, position, and matroid permutation environments, ironed virtual surplus maximization is ordinal, i.e., it depends on the relative order of the ironed virtual values and not their magnitudes. In contrast, the main difficulty of downward-closed environments is that ironed virtual surplus maximization is not ordinal. Nonetheless, for downward-closed environments variants of the random sampling (ironed virtual surplus maximization) and the random sampling profit extraction auctions give constant approximations to the envy-free benchmark. We will describe only the latter result.

Our approach to profit extraction in general downward-closed environments will be the following. The true (and unknown) valuation profile is \mathbf{v} . Suppose we knew a profile \mathbf{v}' that was a coordinate-wise lower bound on \mathbf{v} , i.e., $v_{(i)} \geq v'_{(i)}$ for all i (short-hand notation: $\mathbf{v} \geq \mathbf{v}'$). A natural goal with this side-knowledge would be to obtain the optimal envy-free revenue for \mathbf{v}' . A mechanism that obtains this revenue (in expectation over the random permutation) whenever the coordinate-wise lower-bound assumption holds is a profit extractor.

Mechanism 6.5. *The downward-closed profit extractor for \mathbf{v}' is the following:*

1. *Reject all agents if there exists an i with $v_{(i)} < v'_{(i)}$.*
2. *Calculate the ironed virtual values $\bar{\phi}'$ for \mathbf{v}' .*
3. *For all i , assign the i th highest-valued agent the i th highest ironed virtual value $\bar{\phi}'_{(i)}$.*

4. Serve the agents to maximize the ironed virtual surplus.

Theorem 6.35. *For any downward-closed environment and valuation profiles \mathbf{v} and \mathbf{v}' , the downward-closed profit extractor for \mathbf{v}' is dominant strategy incentive compatible and if $\mathbf{v} \geq \mathbf{v}'$ then its expected revenue under a random permutation is at least the envy-free optimal revenue for \mathbf{v}' .*

Proof. See Exercise 6.5. □

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

Mechanism 6.6. *The biased (random) sampling profit extraction mechanism for downward-closed environments (with parameter $p < 1/2$) is:*

1. Randomly partition the agents into S (with probability p) and M (with probability $1-p$).
2. Reject agents in S .
3. Run the downward-closed profit extractor for \mathbf{v}_S on M .

The main lemma that enables the proof that this biased sampling profit extraction mechanism performs well is very similar to Lemma 6.4.

Lemma 6.36. *For $X_1 = 0$, X_i for $i \geq 1$ an indicator variable for a independent biased coin flipping to heads with probability $p < 1/2$ (tails otherwise), and sum $S_i = \sum_{j \leq i} X_j$,*

$$\Pr[\forall i, S_i \leq (i - S_i)] = 1 - \left(\frac{p}{1-p}\right)^2.$$

Proof. See Exercise 6.6. □

Theorem 6.37. *For any downward-closed environment and any valuation profile, the biased sampling profit extraction auction with $p \approx .21$ is dominant strategy incentive compatible and its expected revenue under a random permutation is a 18.2-approximation to the envy-free benchmark.*

Proof. We define the event \mathcal{B} that $\mathbf{v}_M \geq \mathbf{v}_S$ and the event \mathcal{C} that the highest-valued agent (a.k.a., agent 1) is in the market. Lemma 6.36 implies that $\mathbf{E}[\mathcal{B} \mid \mathcal{C}] = 1 - p^2/(1-p)^2$. Of course, $\Pr[\mathcal{C}] = 1 - p$.

The expected revenue of the biased sampling profit extraction mechanism is, by the definition of conditional expectation,

$$\begin{aligned} \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{C} \wedge \mathcal{B}] \Pr[\mathcal{C} \wedge \mathcal{B}] &= \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{C}] \Pr[\mathcal{C}] \\ &\quad - \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{C} \wedge \neg\mathcal{B}] \Pr[\mathcal{C} \wedge \neg\mathcal{B}]. \end{aligned}$$

We now bound the terms on the right hand side in terms of $\text{EFO}(\mathbf{v}_{-1})$, the envy-free optimal revenue on the valuation profile without the highest-valued agent. For the first term,

$$\begin{aligned} \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{C}] \Pr[\mathcal{C}] &\geq p \text{EFO}(\mathbf{v}_{-1}) \Pr[\mathcal{C}] \\ &= p(1-p) \text{EFO}(\mathbf{v}_{-1}). \end{aligned}$$

To see the inequality: Event \mathcal{C} means that agent 1 is in M , the remaining valuation profile is \mathbf{v}_{-1} . Envy-free revenue is super-additive so the expectation of the envy-free optimal revenue is super-linear. For the second term,

$$\begin{aligned} \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{C} \wedge \neg\mathcal{B}] \Pr[\mathcal{C} \wedge \neg\mathcal{B}] &\leq \mathbf{E}[\text{EFO}(\mathbf{v}_{-1})] \Pr[\neg\mathcal{B} \mid \mathcal{C}] \Pr[\mathcal{C}] \\ &= \frac{p^2}{1-p} \mathbf{E}[\text{EFO}(\mathbf{v}_{-1})]. \end{aligned}$$

The above inequality follows from the coarse upper bound that $\text{EFO}(\mathbf{v}_S) \leq \text{EFO}(\mathbf{v}_{-1})$ under event \mathcal{C} . Combining the bounds above we get:

$$\mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathcal{C} \wedge \mathcal{B}] \Pr[\mathcal{C} \wedge \mathcal{B}] \geq \left(p(1-p) - \frac{p^2}{1-p} \right) \text{EFO}(\mathbf{v}_{-1}).$$

Optimizing for p and using the inequality that $\text{EFO}(\mathbf{v}_{-1}) \geq \text{EFO}^{(2)}(\mathbf{v})/2$ (Exercise 6.7) we get the desired bound in the theorem. \square

Exercises

6.1 Consider the following single-agent prior-free pricing game. There is a value $v \in [1, h]$. If you offer a price $p \leq v$ you get p otherwise you get zero.

- (a) Design a randomized pricing strategy to minimize the ratio of the value to the revenue.
- (b) Prove that your randomized pricing strategy is optimal. Hint: Use the lower-bounding technique for digital-goods auctions from class.
- (c) Discuss the connection between your above results and the claim from class that it is impossible for a digital-goods auction to approximate the envy-free benchmark $\text{EFO}(\mathbf{v}) = \max_i i v_{(i)}$.

6.2 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.

6.3 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 (Mechanism 6.3). 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) Show that the composite mechanism is dominant strategy incentive compatible if the set of winners of A is non-manipulable in the following sense. There are no two values for an agent i such that the sets of winners in A are distinct but contain i .

(b) Show that the set of winners in the surplus maximization mechanism in matroid environments is non-manipulable.

6.4 Prove the envy-free variant of Theorem 6.33, 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.

6.5 Show that for any downward-closed environment and valuation profiles \mathbf{v} and \mathbf{v}' , the downward-closed profit extractor for \mathbf{v}' is dominant strategy incentive compatible and if $\mathbf{v} \geq \mathbf{v}'$ then its expected revenue under random permutation is at least the envy-free optimal revenue for \mathbf{v}' . I.e., prove Theorem 6.35.

6.6 Prove Lemma 6.36: For $X_1 = 0$, X_i for $i \geq 1$ an indicator variable for a independent biased coin flipping to heads with probability $p < 1/2$ (tails otherwise), and sum $S_i = \sum_{j \leq i} X_j$,

$$\Pr[\forall i, S_i \leq (i - S_i)] = 1 - \left(\frac{p}{1-p}\right)^2.$$

6.7 Given a valuation profile \mathbf{v} in sorted order, i.e., $v_1 \geq v_2 \geq \dots \geq v_n$, and any (single-dimensional) downward-closed permutation environment, show that the envy-free optimal revenue for $\mathbf{v}^{(2)} = (v_2, v_2, \dots, v_n)$ and $\mathbf{v}_{-1} = (v_2, v_3, \dots, v_n, 0)$ are within a factor of two of each other.

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 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). Profit extraction and the random sampling profit extraction mechanism were given by Fiat et al. (2002). The extension of this auction to three partitions can be found in Hartline and McGrew (2005). The downward-closed profit extractor is from Ha and Hartline (2011).

The lower-bound on the approximation factor of prior-free auctions for digital goods of 2.42 was given by Goldberg et al. (2004). It is conjectured that this lower bound is tight for general n -agent environment; however, optimal prior-free auctions have not been identified for $n \geq 4$. The second-price auction is optimal for $n = 2$ and its approximation ratio is $\beta = 2$. The optimal three-agent auction can be found in Hartline and McGrew (2005), its approximation ratio is $\beta = 13/6 \approx 2.17$.

This chapter omitted 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 a tractable 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 formal prior-free design and analysis framework for digital good auctions was given by Fiat et al. (2002). 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 Hartline and Yan (2011).

Analysis of the random sampling auction for limited supply (i.e., k -unit auctions) was given by Devanur and Hartline (2009). This result implies prior-free approximation results for matroid permutation and position environments. This result is enabled by the equivalence between position auctions and convex combinations of k -unit auctions (for each k) that is described by Dughmi et al. (2009) and an equivalence between matroid permutation and position environments by Hartline and Yan (2011). Generalizations that give prior-free auctions for downward-closed permutation environments are given by Hartline and Yan (2011).

