Chapter 4

Bayesian Approximation

One of the most intriguing conclusions from the preceding chapter is that for i.i.d., regular, single-item environments the second-price auction with a reservation price is revenue optimal. This result is compelling as the solution it proposes is quite simple; therefore, making it easy to prescribe. Furthermore, reserve-price-based auctions are often employed in practice so this theory of optimal auctions is also descriptive. Unfortunately, i.i.d., regular, single-item environments are hardly representative of the scenarios in which we would like to design good mechanisms. Furthermore, if any of the assumptions are relaxed, reserve-price-based mechanisms are not optimal.

In this chapter we address this deficiency by showing that while reserve-price-based mechanisms are not optimal, they are approximately optimal in a wide range of environments. Furthermore, these approximately optimal mechanisms are more robust, less dependent on the details of the distribution, and provide more conceptual understanding than their optimal counterparts. The approximation factor obtained by most of these reserve-pricing mechanisms is two. Meaning, for the worst distributional assumption, their performance is within a factor two of the optimal mechanism. Of course, in any particular environment these mechanisms may perform better than their worst-case guarantee.

Distributional regularity, as implied by the concavity of the revenue curve, and independence will be instrumental in many of the approximation results obtained, as will two additional structural properties. First, the *monotone hazard rate* condition, a further restriction of regularity, is a property of a distribution that, intuitively, restricts how heavy the tails of the distribution are. An important consequence of the monotone hazard rate assumption is that the optimal revenue and optimal social surplus are within a factor of $e \approx 2.718$ of each other. Second, a *matroid set system* is one that is downward-closed and satisfies an additional "augmentation property." An important consequence of the matroid property is that the surplus maximizing allocation (subject to feasibility) is given by the *greedy-by-value* algorithm: sort the agents by value, then consider each agent in-turn and serve the agent if doing so is feasible.

4.1 Single-item Auctions

We start with single-item auctions and show that the second-price auction with suitably chosen agent-specific reserve prices is always a good approximation to the optimal mechanism.

Mechanism 4.1. *The* second-price auction with reserves $\mathbf{r} = (r_1, \ldots, r_n)$ *is:*

1. *reject each agent i with $v_i < r_i$,*
2. *allocate the item to the highest valued agent remaining (or none if none exists), and*
3. *charge the winner her critical price.*

4.1.1 Regular Distributions

Recall from Chapter 3 that when the agents values are i.i.d. from a regular distribution F (Definition 3.16) then the optimal auction is identically the second-price auction with reserves $\mathbf{r} = (\phi^{-1}(0), \ldots, \phi^{-1}(0))$ where $\phi(\cdot)$ is the virtual value function (Definition 3.14) for F. Further, if we just had a single agent with value $v \sim F$ we would offer her $\phi^{-1}(0)$ to maximize our revenue. This price is often referred to as the *monopoly price*.

Definition 4.1 (monopoly price)**.** *The* monopoly price*, denoted η, for a distribution F is the revenue-optimal price to offer an agent with value drawn from F, i.e., $\eta = \phi^{-1}(0)$.*

Notice that for asymmetric distributions, i.e., where $F_i \neq F_{i'}$, monopoly prices may differ for different agents. Furthermore, the second-price auction with monopoly reserves is not equivalent to the optimal auction when agent values are non-identically distributed. Instead the optimal auction carefully optimizes agents' virtual values at all points of their respective distributions. Therefore, the second-price auction with monopoly reserves has suboptimal revenue.

As an example consider a 2-agent single-item environment with agent 1's value from $U[0,1]$ and agent 2's value from $U[0,2]$. The virtual value functions are $\phi_1(v_1) = 2v_1 - 1$ and $\phi_2(v_2) = 2v_2 - 2$. We serve agent 1 whenever $\phi_1(v_1) > \max(\phi_2(v_2), 0)$, i.e., when $v_1 > \max(v_2 - 1/2, 1/2)$. This auction is not the second-price auction with reserves.

The main result of this section is enabled by the following consequence of distributional regularity. The virtual valuation function is monotone in value, therefore, the monopoly price is the boundary between positive virtual values and negative virtual values.

Fact 4.2. *Any agent whose value exceeds the monopoly price has non-negative virtual value.*

We will shortly show that the expected revenue of the second-price auction with monopoly reserves is close to the optimal revenue when the distributions are regular; however, before doing so, consider the following intuition. Either the monopoly-reserves auction and the optimal auction have the same winner or different winners. If they have the same winner then they have the same virtual surplus. By Fact 4.2, the monopoly-reserve auction always has non-negative virtual surplus, so the virtual surplus when both auctions have the same winner is a lower bound on its total virtual surplus and, thus, its revenue. If the two auctions have different winners then the optimal auction's winner is not the agent with the highest value. Of course this winner can pay at most her value, but the monopoly-reserves auction's winner pays at least the second highest value which must be least the value of the optimal auction's winner. Therefore, in this case the payment in the monopoly-reserves auction is higher than the payment in the optimal auction. We conclude that the revenue of the monopoly-reserves auction bounds the optimal revenue in both cases, therefore, it is a 2-approximation. Importantly, this analysis is driven by regularity and Fact 4.2.

Theorem 4.3. *For any regular, single-item environment the second-price auction with monopoly reserves gives a 2-approximation to the optimal expected revenue.*

Proof. Let REF denote the optimal auction and its expected revenue and APX denote the second-price auction with monopoly reserves and its expected revenue; our goal is to show that $\text{REF} \leq 2\,\text{APX}$. Let I be the winner of the optimal auction and J be the winner of the monopoly reserves auction. Notice that both auctions do not sell the item if and only if all virtual values are negative; in this situation define $I = J = 0$. I and J are random variables. With these definitions, $\text{REF} = \mathbf{E}[\phi_I(v_I)]$ and $\text{APX} = \mathbf{E}[\phi_J(v_J)]$.

We start by simply writing out the expected revenue of the optimal auction as its expected virtual surplus conditioned on $I = J$ and $I \neq J$.

$$\text{REF} = \underbrace{\mathbf{E}[\phi_I(v_I) \mid I = J]\,\mathbf{Pr}[I = J]}_{\text{REF}_=} + \underbrace{\mathbf{E}[\phi_I(v_I) \mid I \neq J]\,\mathbf{Pr}[I \neq J]}_{\text{REF}_{\neq}}.$$

We will prove the theorem by showing that both the terms on the right-hand side are bounded from above by APX. For the first term:

$$\begin{aligned}
\text{REF}_= &= \mathbf{E}[\phi_I(v_I) \mid I = J]\,\mathbf{Pr}[I = J] \\
&= \mathbf{E}[\phi_J(v_J) \mid I = J]\,\mathbf{Pr}[I = J] \\
&\leq \mathbf{E}[\phi_J(v_J) \mid I = J]\,\mathbf{Pr}[I = J] + \mathbf{E}[\phi_J(v_J) \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&= \text{APX}\,.
\end{aligned}$$

The inequality in the above calculation follows Fact 4.2 (i.e., regularity) as even when $I \neq J$ the virtual value of J must be non-negative. Therefore, the term added is non-negative. For the second term:

$$\begin{aligned}
\text{REF}_{\neq} &= \mathbf{E}[\phi_I(v_I) \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[v_I \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[p_J \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[p_J \mid I \neq J]\,\mathbf{Pr}[I \neq J] + \mathbf{E}[p_J \mid I = J]\,\mathbf{Pr}[I = J] \\
&= \text{APX}\,.
\end{aligned}$$

The first inequality in the above calculation follows because $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)} \leq v_i$ (since $\frac{1-F_i(v_i)}{f_i(v_i)}$ is always non-negative). The second inequality follows because, among agents that meet their reserve, J is the highest valued agent and I is a lower valued agent and therefore, in a second-price auction J's price is at least I's value. The third inequality follows because payments are non-negative so the term added is non-negative. □

This 2-approximation theorem is tight. We will give a distribution and show that there is an auction with expected revenue $2-\epsilon$ for any $\epsilon > 0$ but the revenue of the monopoly reserves auction is precisely one. The example that shows this separation is easiest to intuit for a distribution that is partly discrete, i.e., one that does not satisfy the continuity assumptions of the preceding chapter. It is of course possible to obtain the same result with continuous distributions.

A distribution that arises in many examples is the *equal-revenue distribution*. The equal-revenue distribution lies on the boundary between regularity and irregularity, i.e., it has constant virtual value. It is called the equal-revenue distribution because the same expected revenue is obtained by offering the agent any price in the distribution's support.

Definition 4.4. *The* equal-revenue distribution *has distribution function* $F(z) = 1 - 1/z$ *and density function* $f(z) = 1/z^2$. *Its support is* $[1, \infty)$.

Consider the asymmetric two-agent single-item auction setting where agent 1's value is deterministically 1 and agent 2's value is distributed according to a variant of the equal-revenue distribution. The monopoly price for the equal-revenue distribution is ill-defined because every price is optimal. Therefore, we slightly perturb the equal-revenue distribution for agent 2 so that her monopoly price is $\eta_2 = 1$. Clearly then, $\boldsymbol{\eta} = (1, 1)$ and the expected revenue of the second-price auction with monopoly reserve is one.

Of course, for this distribution it is easy to see how we can do much better. Offer agent 2 a high price h. If agent 2 rejects this price then offer agent 1 a price of 1. Notice that by the definition of the equal-revenue distribution, agent 2's expected payment is one, but still agent 2 rejects the offer with probability $1 - 1/h$ and the item can be sold to agent 1. The expected revenue of the mechanism is $h \cdot \frac{1}{h} + 1 \cdot (1 - \frac{1}{h}) = 2 - 1/h$. Choosing $\epsilon = 1/h$ gives the claimed result.

4.1.2 Irregular Distributions

Irregular distributions pose a challenge as where virtual valuations are not monotone, an agent whose value is above the monopoly price may yet have a negative virtual value. We first show that the regularity property is crucial to Theorem 4.3; without it the approximation factor of the second-price auction with monopoly reserves can be linear. We next make an aside to discuss *prophet inequalities* from *optimal stopping theory*. Finally, we use prophet inequalities to succinctly describe reserve prices for which the second-price auction is a 2-approximation. These approximation results rely critically on the assumption that the agents' values are independently distributed.

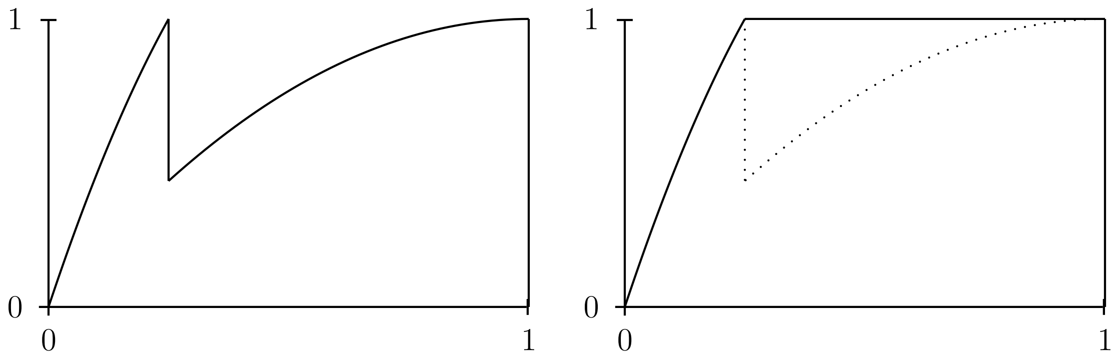

(a) Revenue curve $R(q)$. (b) Ironed Revenue curve $\bar{R}(q)$.

Figure 4.1: The revenue curve and ironed revenue curve for the Sydney opera house distribution for $n = 2$.

Lower-bound for Monopoly Reserves

The second-price auction with monopoly reserve prices is only a two approximation for regular distributions. The proof of Theorem 4.3 relied on regularity crucially when it assumed that the virtual valuation of the winning agent is always non-negative. We start our exploration of approximately optimal auctions for the irregular case with an example that shows that the second-price auction with monopoly reserves can be a linear factor from optimal even when the agents' values are identically distributed.

Definition 4.5. *The* Sydney opera house distribution *arises from drawing a random variable* $1+U[0,1-1/n^2]$ *with probability* $1-1/n^2$ *and* $n^2+U[0,1]$ *with probability* $1/n^2$. *Its revenue curve resembles the Sydney opera house (Figure 4.1).*

The Sydney opera house distribution is bimodal with $R(\cdot)$ maximized at $q = 1$ ($v = 1$) and $q = 1/n^2$ ($v = n^2$). Both give expected revenue of 1. For the purpose of discussion, consider the distribution perturbed sightly so that $\eta = 1$ is the unique monopoly price. The key property of this distribution is that, if there are n agents, the probability of exactly one high-valued agent (i.e., with value at least n^2) is about $1/n$ while the probability of two or more high-valued agents is about $1/(2n^2)$.

The expected revenue of the second-price auction with monopoly reserves is simply the expected second highest value (since the reserve price is never binding). If there is one or fewer high-valued agents then the second highest agent value at most 2. If there are two or more high-valued agents then the second highest agent value is about n^2. The expected revenue is thus about 2.5 (for large n).

To calculate the expected revenue of the optimal auction notice that low-valued agents are completely ironed (Figure 4.1(b)). Suppose there is one high-valued agent, say Alice, and the rest are low valued. If Alice bids a high value she wins. If she bids a low value she is placed in a lottery with all the other agents for a $1/n$ chance of winning. (Of course if she bids below 1 she always loses.) This allocation rule is depicted in Figure 4.2. Alice's payment

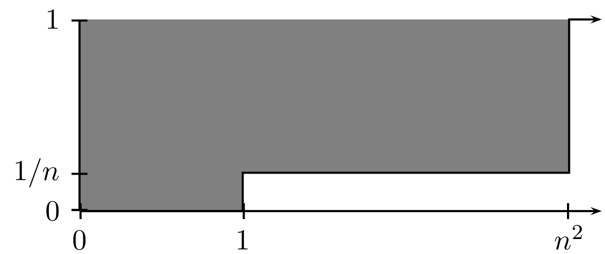

Figure 4.2: The optimal auction allocation rule (black line) and payment (area of gray region) for high-valued Alice when all other agents have low values.

(the area of the gray region in Figure 4.2) in this situation is $n^2 - (n^2 - 1)/n \approx n^2 - n$. There is one such high-valued agent with probability $1/n$ so the total expected revenue is about n.

The conclusion from this rough calculation is that the optimal auction's revenue can be a linear factor more than the second price auction with monopoly reserves.

Theorem 4.6. *There is an i.i.d., irregular distribution for which the second-price auction with monopoly reserves is a linear approximation to the optimal auction revenue.*

Prophet Inequalities

Consider the following scenario. A gambler faces a series of n games on each of n days. Game i has prize distributed according to F_i. The order of the games and distribution of the game prizes is fully known in advance to the gambler. On day i the gambler *realizes* the value $v_i \sim F_i$ of game i and must decide whether to keep this prize and *stop* or to return the prize and *continue* playing. In other words, the gambler is only allowed to keep one prize and must decide which prize to keep immediately on realizing the prize and before any other prizes are realized.

The gambler's optimal strategy can be calculated by *backwards induction*. On day n the gambler should stop with whatever prize is realized. This results in expected value $\mathbf{E}[v_n]$. On day $n-1$ the gambler should stop if the prize has greater value than $t_{n-1} = \mathbf{E}[v_n]$, the expected value of the prize from the last day. On day $n-2$ the gambler should stop with if the prize has greater value than t_{n-2}, the expected value of the strategy for the last two days. Proceeding in this manner the gambler can calculate a threshold t_i for each day where the optimal strategy is to stop with prize i if and only if $v_i \geq t_i$.

Of course, this optimal strategy suffers from many of the drawbacks of optimal strategies. It is complicated: it takes n numbers to describe it. It is not robust to small changes in the game, e.g., changing of the order of the games or making small changes to distribution i strictly above t_i. It does not allow for any intuitive understanding of the properties of good strategies. Finally, it does not generalize well to give solutions to other similar kinds of games.

Therefore, as we are predisposed to do in this text, we turn to approximation to give a crisper picture. A *threshold strategy* is given by a single t for acceptable prizes and an implicit tie-breaking rule which specifies which prize should be selected if there are multiple prizes above t. The implicit tie-breaking rule in the gambler's game is lexicographical: the gambler takes the first prize with value at least t. Threshold strategies are clearly suboptimal as even on day n if prize $v_n < t$ the gambler will not stop and will, therefore, receive no prize.

Theorem 4.7 (Prophet Inequality)**.** *There exists a threshold strategy such that the expected prize of the gambler is at least half the expected value of the maximum prize. Moreover, one such threshold strategy is the one where the probability that the gambler receives no prize is exactly 1/2. Moreover, this bound is invariant to the tie-breaking rule.*

The *prophet inequality* theorem is suggesting something quite strong. Most importantly it is saying that even though the gambler does not know the realizations of the prizes in advance, he can still do as well in comparison to a "prophet" who does. While this result implies that the optimal (backwards induction) strategy satisfies the same condition, such a implication was not at all clear from the original formulation of the optimal strategy. We can also observe that the result is driven by trading off the probability of not stopping and receiving no prize with the probability of stopping early with a suboptimal prize. The suggested threshold strategy is also quite robust. Notice that the order of the games makes no difference in the determination of the threshold. Furthermore, if the distribution above the threshold changes, nothing on the bound or suggested strategy is affected.

Implicit in definition of a threshold strategy is a tie-breaking rule for resolving which acceptable prize is selected when there is a tie, i.e., more than one prize above the threshold. In fact, the prophet inequality theorem, as is stated, is invariant to the tie-breaking rule. While some tie-breaking rules may bring more or less value to the gambler, the 2-approximation result still holds. This invariance of with respect to the tie-breaking rule means that the prophet inequality theorem has broad implications to other similar settings and in particular to auction design and posted pricing, as we will see later in this section.

Proof of Theorem 4.7. Let REF denote prophet and her expected prize, i.e., the expected maximum prize, $\mathbf{E}[\max_i v_i]$, and APX denote a gambler with threshold strategy t and her expected prize. Define $q_i = 1 - F_i(t)$ as the probability that $v_i \geq t$. Let $\chi = \prod_i (1 - q_i)$ be the probability that the gambler receives no prize. The proof follows in three steps. In terms of t and χ, we get an upper bound on the prophet, REF. Likewise, we get a lower bound on the gambler, APX. Finally, we plug in $\chi = 1/2$ to obtain the bound. If there is no t with $\chi = 1/2$, which is possible if the distributions F_i are not continuous, one of the t that corresponds to the smallest $\chi > 1/2$ or largest $\chi < 1/2$ suffices.

In the analysis below, the notation "$(v_i - t)^+$" is short-hand for "$\max(v_i - t, 0)$."

1. An upper bound on REF $= \mathbf{E}[\max_i v_i]$:

   Notice that regardless of whether there exists a $v_i \geq t$ or not, REF is at most $t + \max_i (v_i - t)^+$. Therefore,

$$\begin{aligned} \text{REF} &\leq t + \mathbf{E}\big[\max_i (v_i - t)^+\big] \\ &\leq t + \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+\big]. \end{aligned}$$

2. A lower bound on $\text{APX} = \mathbf{E}[\text{prize of gambler with threshold } t]$:

   Clearly, we get t with probability $1 - \chi$. Depending on which prize i is the earliest one that is greater than t we also get an additional $v_i - t$. It is easy to reason about the expectation of this quantity when there is exactly one such prize and much more difficult to do so when there are more than one. We will ignore the additional prize we get from the latter case and get a lower bound.

$$\begin{aligned}\text{APX} &\geq (1-\chi)t + \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+ \mid \text{other } v_j < t\big]\, \mathbf{Pr}[\text{other } v_j < t] \\ &\geq (1-\chi)t + \chi \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+ \mid \text{other } v_j < t\big] \\ &= (1-\chi)t + \chi \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+\big]\,.\end{aligned}$$

   The second inequality follows because $\mathbf{Pr}[\text{other } v_j < t] = \prod_{j \neq i}(1 - q_i) \geq \chi$. The final equality follows because the random variable v_i is independent of random variables v_j for $j \neq i$.

3. Plug in $\chi = 1/2$.

   From the upper and lower bounds calculated, if we can find a t such that $\chi = 1/2$ then $\text{APX} \geq \text{REF}/2$. Incidentally, as t increases $\sum_i \mathbf{E}[(v_i - t)^+]$ decreases; therefore, we can also solve for $t = \sum_i \mathbf{E}[(v_i - t)^+]$ to obtain same approximation result.

   Consider χ as a function of t denoted $\chi(t)$. For discontinuous distributions, e.g., ones with point-masses, $\chi(t)$ may be discontinuous. Therefore, there may be no t with $\chi(t) = 1/2$. Let $\chi_1 = \sup\{\chi(t) < 1/2\}$ and $\chi_2 = \inf\{\chi(t) > 1/2\}$. Notice that an arbitrarily small increase in threshold causes the jump from χ_1 to χ_2; let t^* be the limiting threshold for both these χs. Therefore, for $\chi \in \{\chi_1, \chi_2\}$ the lower-bound formula $\text{LB}(\chi) \geq (1-\chi)t^* + \chi \sum_i \mathbf{E}[(v_i - t^*)^+]$ which is linear in χ.

   We know that this function evaluated at $\chi = 1/2$ (which is not possible to implement) satisfies $\text{LB}(1/2) \geq \text{REF}/2$. Of course it is a linear function of χ so it is maximized on the end-points on which it is valid, namely χ_1 or χ_2. Therefore, one of $\chi \in \{\chi_1, \chi_2\}$ satisfies $\text{LB}(\chi) \geq \text{REF}/2$. If it is optimized at χ_1 then the threshold is exclusive, i.e., the gambler should accept the first prize in (t^*, ∞); if it is optimize at χ_2 then the threshold is exclusive, i.e., the gambler should accept the first prize in $[t^*, \infty)$.

□

Again the independence of the distributions of prizes is fundamental to the prophet inequality.

Uniform Virtual Prices

We return to our discussion of single-item auctions. Our goal in single-item auctions is to select the winner with the highest (positive) ironed virtual value. To draw a connection between the auction problem and the gambler's problem, we note that the gambler's problem in prize space is similar to the auctioneer's problem in ironed-virtual-value space. The gambler aims to maximize expected prize while the auctioneer aims to maximize expected virtual value. A uniform threshold in the gambler's prize space corresponds to a *uniform ironed virtual price* in ironed-virtual-value space. This strongly suggests that a uniform ironed virtual price would make good reserve prices in the second-price auction.

Definition 4.8. *A* uniform ironed virtual price *is* $\mathbf{p} = (p_1, \ldots, p_n)$ *such that* $\bar{\phi}_i(p_i) = \bar{\phi}_{i'}(p_{i'})$ *for all* i *and* i'.

Now compare the second-price auction with a uniform ironed virtual reserve price to the gambler's threshold strategy in the stopping game. The difference is the tie-breaking rule. The second-price auction breaks ties by value whereas the gambler's threshold strategy breaks ties by the ordering assumption on the games (i.e., lexicographically). Recall, though, that the tie-breaking rule was irrelevant for our analysis of the prophet inequality. We conclude the following theorem and corollary where, as in the prophet inequality, the uniform virtual price is selected so that the probability that the item remains sold is about 1/2.

Theorem 4.9. *For any independent, single-item environment the second-price auction with a uniform ironed virtual reserve price is a 2-approximation to the optimal auction revenue.*

It should be clear that what is driving this result is the specific choice of reserve prices and not explicit competition in the auction. Instead of running an auction imagine the agents arrived in any, perhaps worst-case, order and we made each in turn a take-it-or-leave-it offer of her reserve price? Such a *sequential posted pricing* mechanism is also a 2-approximation.

Theorem 4.10. *For any independent, single-item environment a sequential posted pricing of uniform ironed virtual prices is a 2-approximation to the optimal auction revenue.*

Proof. There may be several agents with values at least their posted price. Suppose that in such a situation the agent with the lowest price arrives first. The revenue under this assumption is certainly a lower bound on the revenue of any other ordering. Furthermore, the prophet inequality on virtual values with tie-breaking by "$-p_i$" guarantees a virtual surplus and, therefore, expected revenue that is a 2-approximation to the optimal expected revenue. □

In fact we already saw in Chapter 1 that posted pricing can be a $\frac{e}{e-1} \approx 1.58$ approximation to the optimal mechanism for social surplus for i.i.d. distributions (Theorem 1.6). This approximation factor also holds for revenue and i.i.d., regular distributions.

Corollary 4.11. *For any i.i.d., regular, single-item environment posting a uniform price is an* $\frac{e}{e-1}$ *approximation to the optimal revenue.*

We can also apply the prophet inequality in value space to argue, similarly to Theorem 4.10 that when the values are non-identically distributed posting a uniform price is is a 2-approximation to the optimal social surplus.

4.1.3 Anonymous Reserves

Thus far we have shown that simple reserve-price-based auctions approximate the optimal auction. Unfortunately, agent-specific reserve prices may be impractical for many scenarios, especially ones where agents could reasonably expect some degree of fairness of the auction protocol. Undoubtedly eBay faces such a constraint for the design of their auction. We therefore consider the extent to which an auction with an *anonymous reserve price*, i.e., the same for each agent, can approximate the revenue of the optimal, perhaps non-anonymous, auction.

We start by considering i.i.d., irregular distributions. For i.i.d., irregular distributions, the optimal auction is anonymous, but it is not a reserve-price-based auction. An immediate corollary of Theorem 4.9 is the following.

Corollary 4.12. *For any i.i.d., single-item environment, the second-price auction with an anonymous reserve is a 2-approximation to the optimal auction revenue.*

We now turn to the more challenging question of whether an anonymous reserve price will give a good revenue when the agents' values are not identically distributed. For instance, in the eBay auction the buyers are not identical. Some buyers have higher *ratings* and these ratings are public knowledge. The value distributions for agents with different ratings may generally be distinct. Therefore, the eBay auction may be suboptimal. Surely though, if the eBay auction was very far from optimal, eBay would have switched to a better auction. The theorem below justifies eBay sticking with the second-price auction with anonymous reserve.

Theorem 4.13. *For any independent, regular, single-item environment the second-price auction with an anonymous reserve is a 4-approximation to the optimal auction revenue.*

Proof. This proof can be obtained by extending the proof of Theorem 4.3 or by following a similar approach to the proof of the prophet inequality. We leave the details to Exercise 4.1. □

Note first that it is possible to prove Theorem 4.13 without considering the effect of competition between agents. Therefore, an anonymous price that satisfies the conditions of the theorem is the monopoly price for the distribution of the maximum value. Note second that the bound given in this theorem is not known to be tight. The two agent example with F_1, a point mass at one, and F_2, the equal-revenue distribution, shows that there is a distributional setting where approximation an factor of anonymous reserve pricing is least two.

We now turn to non-identical, irregular distributions. Here we show that anonymous reserve pricing cannot be better than a logarithmic approximation to the optimal (asymmetric) mechanism.

Theorem 4.14. *There is an n-agent, non-identical, irregular, single-item environment for which the second-price auction with an anonymous reserve is an $\Omega(\log n)$-approximation to the optimal auction revenue.*

Proof. The proof follows from analyzing the optimal revenue and the revenue of the second-price auction with any anonymous reserve on the following discrete distribution (which can, of course, be approximated by a continuous distribution). Agent i's value is drawn as:

$$v_i = \begin{cases} n^2/i & \text{w.p. } 1/n^2, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

The details of this analysis are left to Exercise 4.2. □

4.2 General Feasibility Settings

We now return to more general single-dimensional mechanism design problems, namely, those where the seller faces a combinatorial feasibility constraint. Feasibility constraints that are not downward closed will turn out to be exceptionally difficult and we will give no general approximation mechanisms for them. On the other hand, for regular, downward-closed environments, we show that the surplus maximization mechanism with monopoly reserves is often a 2-approximation. In particular, this result holds if we further restrict the distribution to those satisfying a "monotone hazard rate" condition. It also also holds if we instead restrict the feasible sets to those satisfying a natural augmentation property. These two results are driven by completely different phenomena.

Definition 4.15. *The surplus maximization mechanism with reserves* $\mathbf{r}$ *is:*

1. $\mathbf{v}' \leftarrow \{\text{agents with } v_i \geq r_i\}$.
2. $(\mathbf{x}, \mathbf{p}') \leftarrow \mathrm{SM}(\mathbf{v}')$.
3. *for all i:* $p_i \leftarrow \begin{cases} \max(r_i, p'_i) & \text{if } x_i = 1, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$

where SM *is the surplus maximization mechanism with no reserves.*

4.2.1 Monotone-hazard-rate Distributions (and Downward-closed Feasibility)

An important property of electronic devices, such as light bulbs or computer chips, is how long they will operate before failing. If we model the lifetime of such a device as a random variable then the failure rate, a.k.a., *hazard rate*, for the distribution at a certain point in time is the conditional probability (actually: density) that the device will fail in the next instant given that it has survived thus far. Device failure is naturally modeled by a distribution with a monotone hazard rate, i.e., the longer the device has been running the more likely it is to fail in the next instant. The uniform, normal, and exponential distributions all have monotone hazard rate. The equal-revenue distribution (Definition 4.4) does not.

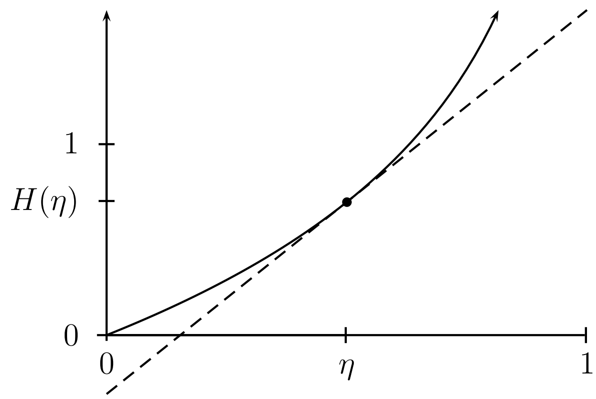

Figure 4.3: The cumulative hazard rate function (solid) for the uniform distribution is $H(v) = -\ln(1 - v)$ and it is lower bounded by its tangent (dashed) at $v = \eta = 1/2$.

Definition 4.16. *The* hazard rate *of distribution F (with density f) is $h(z) = \frac{f(z)}{1-F(z)}$. The distribution has* monotone hazard rate (MHR) *if $h(z)$ is monotone non-decreasing.*

Intuitively distributions with monotone hazard rate are not *heavy tailed*. In fact, the exponential distribution, with $F(z) = 1 - e^{-z}$, is the boundary between monotone hazard rate and non; its hazard rate is constant. Hazard rates are clearly important for optimal auctions as the definition of virtual valuations, expressed in terms of the hazard rate, is $\phi(v) = v - 1/h(v)$. An important property of monotone hazard rate distributions that will enable approximation by the surplus maximization mechanism with monopoly reserves is that, for MHR distributions, the optimal revenue is within a factor of $e \approx 2.718$ of the optimal surplus. We illustrate this with an example, then prove it for the case of a single agent. The proof of the general theorem, we will leave for Exercise 4.3.

Theorem 4.17. *For any downward-closed, monotone-hazard-rate environment, the optimal expected revenue is an e-approximation to the optimal expected surplus.*

To gain some intuition we will look at the exponential distribution. The expected value the exponential distribution (with rate one) is one. This can be calculated from the formula $\mathbf{E}[v] = \int_0^\infty (1 - F(z))\, dz$ with $F(z) = 1 - e^{-z}$. Since the exponential distribution has hazard rate $h(z) = 1$, the virtual valuation formula for the exponential distribution is $\phi(v) = v - 1$. The monopoly price is $\eta = \phi^{-1}(0) = 1$. The probability that the agent accepts the monopoly price is $1/e$ so its expected revenue is $1/e$. The ratio of the expected surplus to expected revenue is e as claimed.

Lemma 4.18. *For any monotone-hazard-rate distribution its expected value is at most e times more than the expected monopoly revenue.*

Proof. Let REF $= \mathbf{E}[v]$ be the expected value and APX $= \eta \cdot (1 - F(\eta))$ be the expected monopoly revenue. Let $H(v) = \int_0^v h(z)dz$ be the *cumulative hazard rate* of the distribution

F. We can write

$$1 - F(v) = e^{-H(v)}, \tag{4.1}$$

an identity that can be easily verified by differentiating $\log(1 - F(z))$. Recall of course that the expectation of $v \sim \mathbf{F}$ is $\int_0^\infty (1 - F(z))\, dz$. To get an upper bound on this expectation we need to upper bound $e^{-H(v)}$ or equivalently lower bound $H(v)$.

The main difficulty is that the lower bound must be tight for the exponential distribution where optimal expected value is exactly e times more than the expected monopoly revenue. Notice that for the exponential distribution the hazard rate is constant; therefore, the cumulative hazard rate is linear. This observation suggests that perhaps we can get a good lower bound on the cumulative hazard rate with a linear function.

Let $\eta = \phi^{-1}(0)$ be the monopoly price. Since $H(v)$ is a convex function (it is the integral of a monotone function). We can get a lower bound $H(v)$ by the line tangent to it at η. See Figure 4.3. I.e.,

$$\begin{aligned} H(v) &\geq H(\eta) + h(\eta)(v - \eta) \\ &= H(\eta) + \tfrac{v-\eta}{\eta}. \end{aligned} \tag{4.2}$$

The second part follows because $\eta = 1/h(\eta)$ by definition. Now we use this bound to calculate a bound on the expectation.

$$\begin{aligned} \text{REF} &= \int_0^\infty (1 - F(z))\, dz = \int_0^\infty e^{-H(z)}\, dz \\ &\leq \int_0^\infty e^{-H(\eta) - \frac{z-\eta}{\eta}}\, dz = e \cdot e^{-H(\eta)} \int_0^\infty e^{-\frac{z}{\eta}}\, dz \\ &= e \cdot e^{-H(\eta)} \cdot \eta = e \cdot (1 - F(\eta)) \cdot \eta = e \cdot \text{APX}\,. \end{aligned}$$

The first and last lines follow from (4.1); the inequality follows from (4.2). □

For non-monotone-hazard-rate distributions the separation between the optimal revenue and the optimal surplus can be arbitrarily large. To see this consider the equal-revenue distribution with $F(z) = 1 - 1/z$. The expected surplus is given by $\mathbf{E}[v] = 1 + \int_1^\infty \frac{1}{z}\, dz = 1 + [\log z]_1^\infty = \infty$. The expected monopoly revenue, of course, is one.

Shortly we will show that the surplus maximization mechanism with monopoly reserve prices is a 2-approximation to the optimal mechanism for MHR, downward-closed environments. This result derives from the intuition that revenue and surplus are close. The following lemma reformulates this intuition.

Lemma 4.19. *For any monotone-hazard-rate distribution F and $v \geq \eta$, $\phi(v) + \eta \geq v$.*

Proof. Since $\eta = \phi^{-1}(0)$ it solves $\eta = 1/h(\eta)$. By MHR, $v \geq \eta$ implies $h(v) \geq h(\eta)$. Therefore,

$$\phi(v) + \eta = v - 1/h(v) + 1/h(\eta) \geq v.$$ □

Theorem 4.20. *For any independent, monotone hazard rate, downward-closed environment the revenue of the surplus maximization mechanism with monopoly reserves is a 2-approximation to the optimal mechanism revenue.*

Proof. Let APX denote the surplus maximization mechanism with monopoly reserves (and its expected revenue) and let REF denote the revenue-optimal mechanism (and its expected revenue). We start with two bounds on APX and then add them.

$$\begin{aligned}\text{APX} &= \mathbf{E}[\text{APX's virtual surplus}]\,, \text{ and}\\ \text{APX} &\geq \mathbf{E}[\text{APX's winners' reserve prices}]\,.\end{aligned}$$

So, summing these two equations and letting $\mathbf{x}(\mathbf{v})$ denote the allocation rule of APX,

$$\begin{aligned}2 \cdot \text{APX} &\geq \mathbf{E}[\text{APX's winners' virtual values + reserve prices}]\\ &= \mathbf{E}\Big[\sum\nolimits_i (\phi_i(v_i) + \eta_i) x_i(\mathbf{v})\Big]\\ &\geq \mathbf{E}\Big[\sum\nolimits_i v_i x_i(\mathbf{v})\Big] = \mathbf{E}[\text{APX's surplus}]\\ &\geq \mathbf{E}[\text{REF's surplus}] \geq \mathbf{E}[\text{REF's revenue}] = \text{REF}\,.\end{aligned}$$

The second inequality follows from Lemma 4.19. By downward closure, neither REF nor APX sells to agents with negative virtual values. Of course, APX maximizes the surplus subject to not selling to agents with negative virtual values. Hence, the third inequality. The final inequality follows because the revenue of any mechanism is never more than its surplus. □

4.2.2 Matroid Feasibility (and Regular Distributions)

In Chapter 3 we saw that the second-price auction with the monopoly reserve was optimal for i.i.d., regular, single-item environments. In the first section of this chapter we showed that the second-price auction with monopoly reserves is a 2-approximation for regular, single-item environments. A very natural question to ask at this point is to what extent we can relax the single-item feasibility constraint and still preserve these results. Often the answer to such questions is *matroids*.

Definition 4.21. *A set system is $(E, \mathcal{I})$ where E is the* ground set *of elements and $\mathcal{I}$ is a set of feasible (a.k.a.,* independent*) subsets of E. A set system is a* matroid *if it satisfies:*

- downward closure*: subsets of independent sets are independent.*
- augmentation*: given two independent sets, there is always an element from the larger whose union with the smaller is independent.*

$$\forall I, J \in \mathcal{I},\ |J| < |I| \Rightarrow \exists e \in I \setminus J,\ \{e\} \cup J \in \mathcal{I}.$$

The augmentation property trivially implies that all maximal independent sets of a matroid have the same cardinality. This cardinality is known as the *rank* of the matroid. The most important theorem about matroids is that the *greedy-by-value* algorithm optimizes surplus. In fact, the most succinct proofs of many mechanism design results in matroid environments are obtained as consequences of the optimality of the greedy-by-value algorithm.

Algorithm 4.1. *A* greedy-by-value algorithm *is*

1. *Sort the agents in decreasing order by value.*

2. $\mathbf{x} \leftarrow \mathbf{0}$ *(the null assignment).*

3. *For each agent i (in this sorted order),*

   if $(\mathbf{x}_{-i}, 1)$ *is feasible,* $x_i \leftarrow 1$.

   (I.e., serve i if i can be served along side previously served agents.)

4. *Output* $\mathbf{x}$.

Theorem 4.22. *The greedy-by-value algorithm is selects the independent set with largest surplus for all valuation profiles if and only if feasible sets are a matroid.*

Proof. The "only if" direction follows from showing, by counter example, (a) downward-closure is necessary and (b) if the set system is downward-closed then the augmentation property is necessary; the "if" direction is as follows.

Let r be the *rank* of the matroid. Let $I = \{i_1, \ldots, i_r\}$ be the set of elements selected in the surplus maximizing assignment, and let $J = \{j_1, \ldots, j_r\}$ be the set of elements selected by greedy-by-value. The surplus from serving a subset S of the agents is $\sum_{i \in S} v_i$.

Assume for a contradiction that the surplus of set I is strictly more than the surplus of set J, i.e., greedy-by-value is not optimal. Assume the items of I and J are indexed in decreasing order. Therefore, there must exist a first index k such that $v_{i_k} > v_{j_k}$. Let $I_k = \{i_1, \ldots, i_k\}$ and let $J_{k-1} = \{j_1, \ldots, j_{k-1}\}$. Applying the augmentation property to sets I_k and J_{k-1} we see that there must exist some element $i \in I_k \setminus J_{k-1}$ such that $J_{k-1} \cup \{i\}$ is feasible. Of course, $v_i \geq v_{i_k} > v_{j_k}$ which means agent i was considered by greedy-by-value before it selected j_k. By downward-closure and feasibility of $J_{k-1} \cup \{i\}$, when i was considered by greedy-by-value it was feasible. By definition of the algorithm, i should have been added; this is a contradiction. □

The following matroids will be of interest.

- In a *k-uniform matroid* all subsets of cardinality at most k are independent. The 1-uniform matroid corresponds to a single-item auction; the k-uniform matroid corresponds to a k-unit auctions.

- In a *transversal matroid* the ground set is the set of vertices of part A of the bipartite graph $G = (A, B, E)$ (where vertices A are adjacent to vertices B via edges E) and independent sets are the subsets of A that can be simultaneously matched. E.g., if A is people, B is houses, and an edge from $a \in A$ to $b \in B$ suggests that b is acceptable to a; then the independent sets are subsets of people that can simultaneously be assigned acceptable houses with no two people assigned the same house. Notice that k-uniform matroids are the special case where $|B| = k$ and all houses are acceptable. Therefore, transversal matroids represent a generalization of k-unit auctions to a market environment where not all items are acceptable to every agent.

- In a *graphical matroid* the ground set is the set of edges in graph $G = (V, E)$ and independent sets are acyclic subgraphs (i.e., a *forest*). Maximal independent sets in a connected graph are spanning trees. The greedy-by-value algorithm for graphical matroids is known as *Kruskal's algorithm* and is studied in every introductory algorithms text.

It is important to be able to argue that a set system satisfies the augmentation property to verify that it is a matroid. As an example we show that acyclic subgraphs are indeed a matroid. For graph $G = (V, E)$ with vertex set $V = \{1, \ldots, n\}$ and edge set $E \subseteq V \times V$ the subgraph induced by edge set $E' \subseteq E$ is $G' = (V, E')$.

Lemma 4.23. *For graph $G = (V, E)$ with $\mathcal{I}$ the set of sets of edges for induced subgraphs that are acyclic, set system $(E, \mathcal{I})$ is a matroid.*

Proof. Downward closure is easy to argue: given an acyclic subgraph, removing edges cannot create cycles.

To show the augmentation property, consider the number of connected components of an acyclic subgraph $G' = (V, E')$ with $m' = |E'|$ edges. By induction the number of connected components is $n - m'$: when $m' = 0$ each vertex is its own connected component; the addition of any edge that does not create a cycle must connect two connected components thereby reducing the number of connected components by one.

Now consider two acyclic subgraphs given by edge sets $I, J \subseteq E$ satisfying the assumption of the augmentation property, i.e., that $|J| < |I|$. We conclude that the number of connected components of graph (V, J) is strictly more than that of (V, I) which is at least that of connected components of the graph $(V, I \cup J)$.

Consider adding edges $I \setminus J$ one at a time to J and let e be the first such edge that decreases the number of connected components. Then $(V, J \cup \{e\})$ is acyclic, as e connects two connected components of (V, J) and therefore does not create a cycle; and the augmentation property is satisfied. □

Since greedy by value is the optimal algorithm for matroid environments; the revenue-optimal mechanism for matroid environments is *greedy by ironed virtual value*. Of course, for i.i.d., regular distributions greedy by ironed virtual value is simply greedy by value with a reserve price of $\phi^{-1}(0) = \eta$. This is exactly the surplus maximization mechanism with reserve price η. This argument is implicitly taking advantage of the fact that the greedy-by-value algorithm is ordinal, i.e., only the relative order of values matters in determining the optimal feasible allocation.

Theorem 4.24. *For any i.i.d., regular, matroid environment, the surplus maximization mechanism with monopoly reserve price optimizes expected revenue.*

Proof. The optimal algorithm for maximizing virtual surplus (hence: the optimal mechanism) is greedy by virtual value with agents with negative virtual value discarded. In the regular case, i.e., when virtual values are monotone and identical, sorting by virtual values is the same as sorting by values and discarding negative virtual values is the same as discarding values less than the monopoly price. □

Of course, in matroid environments that are inherently asymmetric, the i.i.d. assumption is overly restrictive. It turns out that the surplus maximization mechanism with (agent-specific) monopoly reserves continues to be a good approximation even when the agents' values are non-identically distributed.

Theorem 4.25. *In regular, matroid environments the revenue of the surplus maximization mechanism with monopoly reserves is a 2-approximation to the optimal mechanism revenue.*

There are two very useful facts about the surplus maximization mechanism in matroid environments that enable the proof of Theorem 4.25. The first shows that the critical value (which determine agent payments) for an agent is the value of the agent's "best replacement." The second shows that the surplus maximization mechanism is point-wise revenue monotone, i.e., if the values of any subset of agents increases the revenue of the mechanism does not decrease. These properties are summarized by Lemma 4.29 and Lemma 4.28, below. The formal proofs of Theorem 4.25 and Lemma 4.28 are left for Exercises 4.4 and 4.5, respectively.

Definition 4.26. *If $I \cup \{i\} \in \mathcal{I}$ is surplus maximizing set containing i then the* best replacement *for i is $j = \operatorname{argmax}_{\{k \,:\, I \cup \{k\} \in \mathcal{I}\}} v_k$.*

Definition 4.27. *A mechanism is* revenue monotone *if for all valuation profiles $\mathbf{v} \geq \mathbf{v}'$ (i.e., for all i, $v_i \geq v'_i$), the revenue of the mechanism on $\mathbf{v}$ is no worse than its revenue on $\mathbf{v}'$.*

Lemma 4.28. *In matroid environments, the surplus maximization mechanism is revenue monotone.*

Lemma 4.29. *In matroid environments, the surplus maximization mechanism on valuation profile $\mathbf{v}$ has the critical values $\boldsymbol{\tau}$ satisfying, for each agent i, $\tau_i = v_j$ where j is the* best replacement *for i.*

Proof. The greedy-by-value algorithm is ordinal, therefore we can assume without loss of generality that the cumulative values of all subsets of agents are distinct. E.g., add a $U[0,\epsilon]$ random perturbation to each agent value, the event where two subsets sum to the same value has measure zero, and as $\epsilon \to 0$ the critical values for the perturbation approach the critical values for the original valuation profile, i.e., from equation (4.3).

To proceed with the proof, consider two alternative calculations of the critical value for player i. The first is from the proof of Lemma 3.6 where $\mathrm{OPT}(\mathbf{v}_{-i})$ and $\mathrm{OPT}_{-i}(\mathbf{v})$ are optimal surplus from agents other than i with i is not served and served, respectively.

$$\tau_i = \mathrm{OPT}(\mathbf{v}_{-i}) - \mathrm{OPT}_{-i}(\mathbf{v}). \tag{4.3}$$

The second is from the greedy algorithm. Sort all agents except i by value, then consider placing agent i at any position in this ordering. Clearly, when placed first i is served. Let j be the first agent after which i would not be served. Then,

$$\tau_i = v_j. \tag{4.4}$$

Now we compare these the two formulations of critical values given by equations (4.3) and (4.4). Notice that if i is ordered after j and this causes i to not be served, then j must be served as this is the only possible difference between i coming before or after j. Therefore, agent j must be served in the calculation of $\text{OPT}(\mathbf{v}_{-i})$. Let $J \cup \{j\}$ be the agents served in $\text{OPT}(\mathbf{v}_{-i})$ and let I be the agents served in $\text{OPT}_{-i}(\mathbf{v})$ (which does not include i). We can deduce (denoting by $v(S) = \sum_{k\in S} v_k$):

$$\begin{aligned} v_j &= \tau_i \\ &= \text{OPT}(\mathbf{v}_{-i}) - \text{OPT}_{-i}(\mathbf{v}) \\ &= v_j + v(J) - v(I). \end{aligned}$$

We conclude that $v(I) = v(J)$ which, by the assumption that the cumulative values of distinct subsets are distinct, implies that $I = J$. Meaning: j is a replacement for i; furthermore, by optimality of $J \cup \{j\}$ for $\text{OPT}(\mathbf{v}_{-i})$, j must be the best, i.e., highest valued, replacement. □

Exercises

4.1 Show that for any non-identical, regular distribution of agents, there exists a reserve price such that second-price auction with an anonymous reserve price obtains 4-approximation to the optimal single-item auction revenue.

4.2 Prove Theorem 4.14 by analyzing the revenue of the optimal auction and the second-price auction with any anonymous reserve when the agents values distributed as:

$$v_i = \begin{cases} n^2/i & \text{w.p. } 1/n^2\text{, and} \\ 0 & \text{otherwise.} \end{cases}$$

(a) Show that the expected revenue of the optimal auction is $\Omega(\log n)$.

(b) Show that for any anonymous reserve, the expected revenue of the second-price auction conditioned on exactly one agent having a non-zero value is $O(n)$.

(c) Show that for any anonymous reserve, the expected revenue of the second-price auction is $O(1)$.

(d) Combine the above three steps to prove the theorem.

4.3 Consider the following *surplus maximization mechanism with lazy monopoly reserves* where, intuitively, we run the the surplus maximization mechanism SM and then reject any winner i whose value is below her monopoly price η_i:

1. $(\mathbf{x}', \mathbf{p}') \leftarrow \mathrm{SM}(\mathbf{v})$,
2. $x_i = \begin{cases} x_i' & \text{if } v_i \geq \eta_i \\ 0 & \text{otherwise, and} \end{cases}$
3. $p_i = \max(\eta_i, p_i')$.

Prove that the revenue of this mechanism is an e-approximation to the optimal social surplus in any downward-closed, monotone-hazard-rate environment. Conclude Theorem 4.17 as a corollary.

4.4 Show that in regular, matroid environments the surplus maximization mechanism with monopoly reserves gives a 2-approximation to the optimal mechanism revenue, i.e., prove Theorem 4.25. Hint: this result can be proved using Lemmas 4.29 and 4.28 and a similar argument to the proof of Theorem 4.3.

4.5 A mechanism $\mathcal{M}$ is *revenue monotone* if for all pairs of valuation profiles $\mathbf{v}$ and $\mathbf{v}'$ such that for all i, $v_i \geq v_i'$, the revenue of $\mathcal{M}$ on $\mathbf{v}$ is at least its revenue on $\mathbf{v}'$. It is easy to see that the second-price auction is revenue monotone.

1. Give a single-parameter agent environment for which the surplus maximization mechanism (Mechanism 3.1) is not revenue monotone.
2. Prove that the surplus maximization mechanism is revenue monotone in matroid environments.

Chapter Notes

For non-identical, regular, single-item environments, the proof that the second-price auction with monopoly reserves is a 2-approximation is from Chawla et al. (2007). For the same environment, the second-price auction with anonymous reserve was shown to be a 4-approximation by Hartline and Roughgarden (2009).

The prophet inequality theorem was proven by Samuel-Cahn (1984) and the connection between prophet inequalities and mechanism design was first made by Taghi-Hajiaghayi et al. (2007). For irregular distributions and single-item auctions, the 2-approximation for the second-price auction with constant virtual reserves (and the related sequential posted pricing mechanism) was given by Chawla et al. (2010a).

Beyond single-item environments, Hartline and Roughgarden (2009) show that the surplus maximization mechanism with monopoly reserves is a 2-approximation to the optimal mechanism both for regular, matroid environments (generalizing the single-item auction proof of Chawla et al., 2007) and for monotone-hazard-rate, downward-closed environments.

The structural comparison between optimal surplus and optimal revenue for downward-closed, monotone-hazard-rate environments was given by Dhangwatnotai et al. (2010). The analysis of greedy-by-value under matroid feasibility was initiated by Joseph Kruskal (1956) and there are books written solely on the structural properties of matroids, see e.g., Welsh

(2010) or Oxley (2006). Mechanisms based on the greedy algorithm were first studied by Lehmann et al. (2002) where it was shown that even when these algorithms are not optimal, mechanisms derived from them are incentive compatible.

The first comprehensive study of surplus maximization in matroid environments was given by Talwar (2003); for instance, he proved critical values for matroid environments are given by the best replacement. The revenue monotonicity for matroid environments and non-monotonicity for non-matroids is discussed by Dughmi et al. (2009), Ausubel and Milgrom (2006), and Day and Milgrom (2007).