Chapter 5

Bayesian Approximation

One of the most intriguing conclusions from the preceding chapter on Bayesian profit maximization is that for i.i.d. regular distributions the optimal single-item auction is the second-price auction with a reservation price. 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 settings 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 settings. These approximately optimal mechanisms are more robust, less dependent on the details of the setting, and provide more conceptual understanding than their optimal counterparts. The approximation factor of these mechanisms is 2 under worst-case distributional assumptions. Of course, in any particular setting they may perform better than their worst-case guarantee.

5.1 Single-item Auctions

We start with single-item auctions and show that the second-price auction with suitably chosen reserve prices is always a good approximation to the optimal mechanism. The mechanism is the following and is parameterized by agent-specific reserve prices $\mathbf{r} = (r_1, \ldots, r_n)$.

Mechanism 5.1 *The* second-price auction with reserves $\mathbf{r}$, $\mathrm{Vic}_{\mathbf{r}}$ *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 their critical price.*

5.1.1 Regular Distributions

Recall that when the agents values are i.i.d. from a regular distribution F then the Myerson auction is identically the second-price auction with reserves $\mathbf{r} = (\phi^{-1}(0), \ldots, \phi^{-1}(0))$. Notice that if we just had a single agent with value $v \sim F$ we would offer them $\phi^{-1}(0)$ to maximize our revenue. This price is often referred to as the *monopoly price*.

Definition 5.1 The *monopoly price* for $v \sim F$ is $r^* = \phi^{-1}(0)$.

Notice that in asymmetric settings where $F_i \neq F_{i'}$, monopoly prices may differ for different agents. Notice that the second-price auction with monopoly reserves, $\text{Vic}_{\mathbf{r}^*}$, is not equivalent to the Myerson auction when agent values are non-identically distributed. The Myerson auction carefully optimizes agents' virtual values at all points of their respective distributions whereas the second-price auction with reserves only considers the inverse virtual values of zero. Therefore, the second-price auction with monopoly reserves has suboptimal revenue. The following theorem shows that its expected revenue is never too far from the optimal revenue.

Theorem 5.2 *For any regular product distribution* $\mathbf{F}$*, the second-price auction with monopoly reserves,* $\text{Vic}_{\mathbf{r}^*}$*, gives a 2-approximation to the optimal auction revenue.*

Before giving the proof of this theorem consider the following intuition. Either $\text{Vic}_{\mathbf{r}^*}$ and Myerson have the same winner or different winners. If they have the same winner then they have the same virtual surplus. If they have different winners then Myerson's winner is not the agent with the highest value. Of course Myerson's winner can pay at most their value, but $\text{Vic}_{\mathbf{r}^*}$'s winner pays at least the second highest value which is at least the value of Myerson's winner. Therefore, in this case the payment in $\text{Vic}_{\mathbf{r}^*}$ is higher than the payment of Myerson. An important observation that is glossed over in this informal argument but necessary for a formal proof is that for regular distributions $\text{Vic}_{\mathbf{r}^*}$ never sells to an agent with negative virtual value.

Proof: (of theorem) Let I be the winner of Myerson and J be the winner of $\text{Vic}_{\mathbf{r}^*}$. Notice Myerson and $\text{Vic}_{\mathbf{r}^*}$ both 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.

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

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

We will prove the theorem by showing that both the terms on the right-hand side are bounded from above by $\mathbf{E}[\text{Vic}_{\mathbf{r}^*}]$.

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

The inequality in the above calculation follows because regularity of the distribution implies that the virtual value of the winner of $\text{Vic}_{\mathbf{r}^*}$ is always non-negative. Therefore, this added term is non-negative.

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

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 J is the highest valued agent and I is a lower valued agent and therefore, 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 $\text{Vic}_{\mathbf{r}^*}$'s expected revenue is precisely 1. The example that shows this separation is easiest to intuit in a partly discrete setting, i.e., one that does not satisfy the continuity assumptions of the preceding section. 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 5.3 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 $r_2^* = 1$. Clearly then, $\mathbf{r}^* = (1, 1)$ and $\mathbf{E}[\text{Vic}_{\mathbf{r}^*}] = 1$.

Of course, it is easy to see how we can do much better than $\text{Vic}_{\mathbf{r}^*}$. First 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 accepts the offer with probability $1/h$. The expected revenue of this mechanism is $h \cdot \frac{1}{h} + 1 \cdot (1 - \frac{1}{h}) = 2 - 1/h$. Choosing $\epsilon = 1/h$ gives the desired result.

5.1.2 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 5.2; without it the approximation

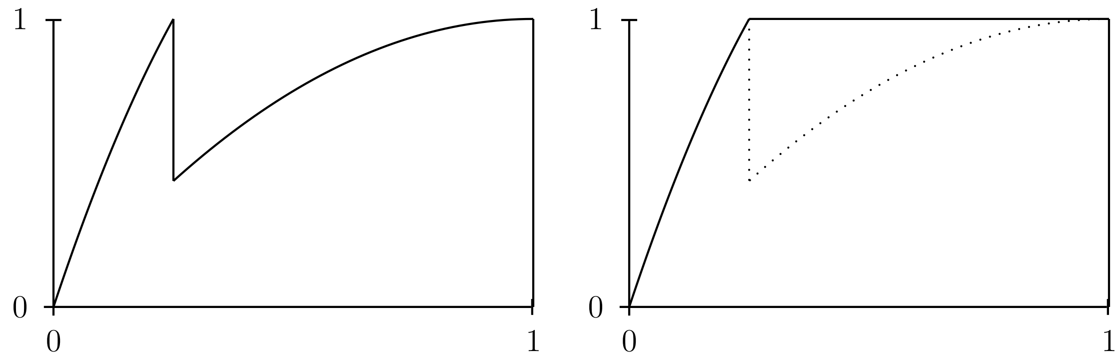

(a) Revenue curve $R(q)$. (b) Ironed Revenue curve $\bar{R}(q)$.

Figure 5.1: The revenue curve and ironed revenue curve for the Sydney opera house distribution for $n = 2$.

factor of the second price auction with monopoly reserves can be linear. We next make an aside 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.

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 5.2 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 5.4 The *Sydney opera house distribution* is 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 5.1).

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 $r^* = 1$ is the unique monopoly price. The key property of the distribution, however, is that if there are n agents, the probability of a high-valued agent (i.e., with value at least n^2) is about $1/n$ where as the probability of two 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 Myerson auction notice that low-valued agents are completely ironed (Figure 5.1(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 5.2. Alice's payment (the area of the gray region in Figure 5.2) 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.

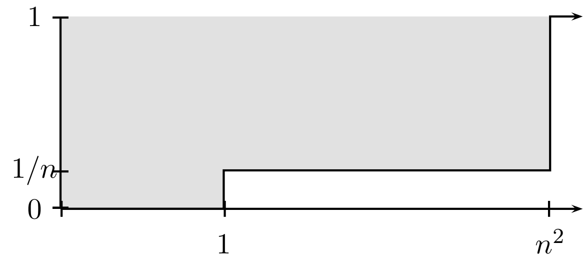

Figure 5.2: The optimal auction allocation rule (black line) and payment (area of gray region) for high-valued Alice when all other agents have low values.

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

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 ascertained from *backwards induction*. On day n the gambler should stop with whatever prize is realized. This results in some expected value. On day $n-1$ the gambler should stop with any price t_{n-1} with value at least the expected value of the prize from the last day. On day $n-2$ the gambler should stop with any prize with greater value than the expected value of the strategy thus-far calculated. 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 and requires the gambler to accept the first prize i with $v_i \geq t$. 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 5.5 (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.*

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 as 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.

We will refer to the setting where there are more than one prize above the threshold as a *tie*. There is an implicit tie-breaking strategy given by the definition of the game: ties are broken in favor of the earliest game. The aforementioned invariance of the performance bound to the order of the games then suggests that the bound might invariant on the tie-breaking rule entirely. This is indeed the case as is evident from the proof of theorem. This invariance of with respect to the tie-breaking rule means that 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.

We conclude with the proof of the theorem.

Proof: (of theorem) 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 expected maximum prize. Likewise, we get a lower bound on the algorithm performance. 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 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 $\text{OPT} = \mathbf{E}[\max_i v_i]$.

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

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

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

   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.

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

   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 $\mathcal{A} \geq \mathrm{OPT}/2$. Incidentally, as t increases $\sum_i \mathbf{E}[(v_i - t)^+]$ decreases; therefore, we can also set $t = \sum_i \mathbf{E}[(v_i - t)^+]$ for the same result.

   For discontinuous distributions, e.g., ones with point-masses, χ may by a discontinuous function of t. There may be no t with $\chi = 1/2$. Let $\chi_1 = \sup\{\chi < 1/2\}$ and $\chi_2 = \inf\{\chi > 1/2\}$. 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, as a function of χ, t is constant and we get the same lower bound formula $\mathrm{LB}(\chi) \geq (1-\chi)t + \chi \sum_i \mathbf{E}[(v_i - t)^+]$ which is linear in χ.

   We know that this function evaluated at $\chi = 1/2$ (which is not possible to implement) satisfies $\mathrm{LB}(1/2) \geq \mathrm{OPT}/2$. Of course this is a linear function 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 $\mathrm{LB}(\chi) \geq \mathrm{OPT}/2$. We conclude that either the threshold that corresponds to χ_1 or χ_2 is a 2-approximation.

■

Constant 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 constant threshold in the gambler in prize space corresponds to a constant *ironed virtual price* in ironed-virtual-value space. This suggests strongly that constant ironed virtual prices would make good reserve prices in the second-price auction.

Definition 5.6 A *constant ironed virtual price* is $\mathbf{p}$ 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 constant ironed virtual reserve price with the gambler's threshold strategy. 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 constant virtual price is selected so that the probability that the item remains sold is about 1/2.

Theorem 5.7 *For any product distribution* $\mathbf{F}$*, the second price auction with a constant ironed virtual reserve price is a 2-approximation to the optimal revenue.*

Corollary 5.8 *For any i.i.d. distribution* $\mathbf{F}$*, the second price auction with an anonymous reserve is a 2-approximation to the optimal 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. So 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 their reserve price? Such a *sequential posted pricing* mechanism is certainly also a 2-approximation.

Theorem 5.9 *For any distribution* $\mathbf{F}$*, a sequential posted pricing with constant ironed virtual prices is a 2-approximation to revenue.*

5.1.3 Anonymous Reserves

The reserve prices discussed thus far, except for the i.i.d. regular case where the monopoly reserve price is optimal and the i.i.d. irregular case where an anonymous reserve price is a 2-approximation. The question of whether anonymous reserve prices with the second price auction are good approximations in non-identical settings is natural and important.

For instance, the eBay auction is essentially a second-price auction with an anonymous reserve. Of course, the buyers are not identical. Some buyers have higher *ratings* and these ratings are public knowledge. The distributions values of 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 may have switched to a better auction. The theorem below justifies eBay sticking with the second-price auction with anonymous reserve.

Theorem 5.10 *For regular product distribution* $\mathbf{F}$*, the second price auction with an anonymous reserve is a 3-approximation to the optimal revenue.*

The proof of this theorem is a non-trivial extension of basic prophet inequality proof that we do not give here. It should be noted that the proof does not make use of competition between agents, i.e., the tie-breaking rule. Therefore, an anonymous price that satisfies the conditions of the theorem is the monopoly price for the distribution of the maximum value. Finally, the bound given in this theorem is not known to be tight. The two agent example with F_1, a point mass at 1, and F_2, the equal-revenue distribution, shows that the worst-case distributional setting has approximation factor at least 2.

5.2 General Feasibility Settings

We now return to more general kinds of single-dimensional mechanism design problems, namely, those where the seller faces a combinatorial feasibility constraint. Settings that are not downward closed will turn out to be exceptionally difficult and we will give no general approximation mechanisms for them. For downward-closed settings and regular distributions, we show that the VCG 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 results are driven by completely different phenomena.

Definition 5.11 The VCG mechanism with reserves $\mathbf{r}$ is:

1. $\mathbf{v}' \leftarrow \{$ agents with $v_i \geq r_i\}$.
2. $(\mathbf{x}, \mathbf{p}') \leftarrow \mathrm{VCG}(\mathbf{v}')$.
3. for all i: $p_i \leftarrow \max(r_i, p'_i)$.

5.2.1 Monotone-hazard-rate Distributions and Downward-closed Settings

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 does not.

Definition 5.12 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 $\phi(v) = v - 1/h(v)$. The important property of monotone hazard rate distributions that will enable approximation by VCG with monopoly reserves is that, for MHR distributions, the optimal revenue is within a factor of e 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 as an exercise.

Theorem 5.13 *For any monotone-hazard-rate product distribution* $\mathbf{F}$ *and any downward-closed setting, the optimal expected revenue (e.g., from the Myerson mechanism) is an* e*-approximation to the optimal expected surplus (e.g., from the VCG mechanism).*

The expected value the exponential distribution is 1. This can be calculated from the formula $\mathbf{E}[v] = \int_0^\infty (1 - F(z))\, dz$. As previously stated, the exponential distribution has hazard rate 1. Therefore, the virtual valuation formula for the exponential distribution is $\phi(v) = v - 1$. The monopoly price is $r^* = \phi^{-1}(0) = 1$. The probability that the agent accepts the monopoly price is e^{-1} so its expected revenue is e^{-1}. The ratio of the expected surplus to expected revenue is e as claimed.

Lemma 5.14 *For any monotone-hazard-rate product distribution* F *the expected value is at most* e *times more than the expected monopoly revenue.*

Proof: Let $H(v) = \int_0^v h(z)dz$ be the *cumulative hazard rate*. We can write $1 - F(v) = e^{-H(v)}$, an identity that can be easily verified by differentiating both sides. 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 expected surplus is exactly equal to expected revenue. Notice that for the exponential distribution that the hazard rate is constant; therefore, the cumulative hazard rate is linear. This suggests that perhaps we can get a good lower bound on the cumulative hazard rate with a linear function.

Let $r^* = \phi^{-1}(0)$ 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 r^*. I.e.,

$$
\begin{aligned}
H(v) &\geq H(r^*) + h(r^*)(v - r^*) \\
&= H(r^*) + \tfrac{v - r^*}{r^*}
\end{aligned}
$$

The second part follows because $r^* = 1/h(r^*)$ by definition. Now we use this bound on the expectation.

$$
\begin{aligned}
\mathbf{E}[v] &= \int_0^\infty (1 - F(z))dz \\
&= \int_0^\infty e^{-H(z)}dz \\
&\leq \int_0^\infty e^{-H(r^*) - \frac{z - r^*}{r^*}}dz \\
&= e^{-H(r^*)} \int_0^\infty e^{-\frac{z}{r^*}+1}dz \\
&= e^{-H(r^*)} \cdot r^* \cdot e.
\end{aligned}
$$

This final line completes the lemma as the optimal revenue is $e^{-H(r^*)} \cdot r^*$. ■

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 revenue, of course, is 1.

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

Lemma 5.15 *For any monotone-hazard-rate distribution F and $v \geq r^*$, $\phi(v) + r^* \geq v$.*

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

$$
\begin{aligned}
\phi(v) + r^* &= v + 1/h(v) + 1/h(r^*) \\
&\geq v.
\end{aligned}
$$

■

Theorem 5.16 *For any downward closed setting and any monotone-hazard-rate product distribution $\mathbf{F}$, VCG with monopoly reserves is a 2-approximation to the optimal revenue.*

Proof: We start with two equations, then add them.

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

So, summing these two equations,

$$\begin{aligned} 2\mathbf{E}[\text{VCG revenue}] &\geq \mathbf{E}[\text{VCG winners' virtual values + reserve prices}] \\ &= \mathbf{E}\Big[\sum\nolimits_i (\phi_i(v_i) + r_i^*) x_i(\mathbf{v})\Big] \\ &\geq \mathbf{E}\Big[\sum\nolimits_i v_i x_i(\mathbf{v})\Big] \\ &= \mathbf{E}[\text{VCG surplus}] \\ &\geq \mathbf{E}[\text{Mye surplus}] \\ &\geq \mathbf{E}[\text{Mye revenue}]\,. \end{aligned}$$

The third line follows from Lemma 5.15. The fourth line is by definition. The fifth line follows from downward closure. By downward closure Myerson never sells to agents with negative virtual values; VCG with monopoly reserves does not either. Of course, VCG with monopoly reserves maximizes the surplus subject to not selling to agents with negative virtual values. Hence, the inequality. The final line follows because the revenue of any mechanism is never more than its surplus. ■

5.2.2 Regular Distributions and Matroid Settings

In the first section of this chapter we showed that the second-price auction with monopoly reserves is a 2-approximation for regular product distributions. A very natural question to ask at this point is for what general single-dimensional agent settings such result holds. Often the answer to such questions is *matroids*.

Matroids

Definition 5.17 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.

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

It should be noted that all maximal independent sets of a matroid have the same cardinality. This follows directly from the augmentation property. The most important theorem about matroids, one that we will not prove here, is that the *greedy-by-value* algorithm is optimal. In particular. Assume each element had a corresponding value and we wished to find an independent set with the largest cumulative value. The optimal algorithm would sort the elements by value and then starting with the largest, select for inclusion each element if its union with the previously selected elements is independent. (Recall Algorithm 3.1.)

Theorem 5.18 *The greedy-by-value algorithm is selects the independent set with largest cumulative value if and only if feasible sets are a matroid.*

The reader will recall from Chapter 3 that greedy by value is not optimal for single-minded combinatorial auction settings; indeed, these do not satisfy the augmentation property required of matroids.

The following matroid settings 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 *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 course.

- In a *transversal matroid* the ground set is the set of vertices are 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. 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 setting not all items are acceptable to every agent.

Optimal Mechanisms and Approximations

Since greedy by value is the optimal algorithm for matroid settings; the optimal mechanism for matroid settings 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 $\phi^{-1}(0) = r^*$. This is exactly VCG with reserve price r^*.

Theorem 5.19 *For any i.i.d. regular distribution* $\mathbf{F}$ *and any matroid feasibility setting, VCG with monopoly reserve price is optimal.*

Of course, in matriod setting which may be inherently asymmetric, the i.i.d. assumption is overly restrictive. The following theorem can be proved analogously to the same result for single-item settings (i.e., Theorem 5.2).

Theorem 5.20 *For any regular product distribution* $\mathbf{F}$ *and any matroid feasibility setting, VCG with monopoly reserves,* $\mathrm{VCG}_{\mathbf{r}^*}$*, gives a 2-approximation to the optimal revenue.*

5.3 Notes

For non-identical, regular distributions and single-item auctions, the proof that the second-price auction with monopoly reserves is a 2-approximation is from Chawla, Hartline, and Kleinberg [9]. For the same setting, the second-price auction with anonymous reserve was shown to be a 4-approximation by Hartline and Roughgarden [15]. This theorem has since been improved to give a 3-approximation via a prophet-inequality-like proof.

The prophet inequality theorem was proven by Samuel-Cahn [20] and the connection between prophet inequalities and mechanism design questions was first made by Taghi-Hajiaghayi, Kleinberg, and Sandholm [21]. 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, Hartline, Malec, and Sivan [10].

Beyond single-item settings, Hartline and Roughgarden [15] show that VCG with monopoly reserves is a 2-approximation to the optimal mechanism both for regular distributions and matroid feasibility constraints (generalizing the single-item auction proof of [9]) and for monotone-hazard-rate distributions and downward-closed feasibility constraints. The result for monotone-hazard-rate distributions is obtained by a connection between optimal surplus and optimal revenue that was formalized by Dhangwatnotai, Roughgarden, and Yan [12].