Chapter 5

Prior-independent Approximation

In the last two chapters we discussed mechanism that performed well for a given Bayesian prior distribution. Assuming such a Bayesian prior is natural when deriving mechanisms for games of incomplete information as Bayes-Nash equilibrium requires the prior distribution to be common knowledge. For reasons to be discussed, it is desirable to relax this known prior assumption. The objective of prior-independent mechanism design is to identify a single mechanism that always has good performance, e.g., under any distributional assumption. A slightly relaxed objective would be to constrain the distributions to fall within some broad, natural class, e.g., i.i.d., regular distributions.

As is evident from our analysis of Bayesian optimal auctions, e.g., for profit maximization, for any auction that one might consider good, there is a value distribution for which another auction performs strictly better. This is obvious because optimal auctions for distinct distributions are generally distinct. While no auction is optimal for all value distributions, there may be a single auction that is approximately optimal across a wide range of distributions.

In this chapter we will take two approaches to prior-independent mechanisms. The first is a "resource" augmentation, a.k.a., bicriteria, approach. We will show that increasing competition, e.g., by recruiting more agents, and running the (prior-independent) surplus maximization mechanism sometimes earns more revenue more than the revenue-optimal mechanism would have without the increased competition. The second approach is to design mechanisms that do a little market analysis on the fly. We will show that for a large class of environments there is a single mechanism that approximates the revenue of the optimal mechanism.

5.1 Motivation

Since prior-independence is not without loss it is important to consider the motivation behind going from Bayesian optimal to prior-independent approximation mechanisms.

Remember why we adopted Bayesian optimality in the first place: we are considering a game of incomplete information and in games of incomplete information, in order to argue about strategic choice, we needed to formalize how players deal with uncertainty. In a

Stackelberg game, instead of moving simultaniously, players make actions in a prespecified order. We can view the mechanism designer as a player who moves first and the agents as players who (simultaniously) move second. To analize the Bayes-Nash equilibrium in such a Stackelberg game, the designer bases her strategy on the common prior. Without such prior knowledge, prediction the designer's strategy is ill posed.

Now consider from where the designer may have learned the prior-distributions. There are two most logical candidates. The first is from the designer's history in interacting with these or similar agents. The problem with this point of view is that the earlier agents may strategize so that information about their preferences is not exploited by the designer later. In fact, if a monopolist cannot commit not exploit the agents using information from prior interaction then the socially efficient (i.e., surplus maximizing) outcome is the only equilibrium. Its revenue can be far from the optimal revenue. This phenomenon is referred to as the *Coase Conjecture* (a theorem).

The second candidate is *market analysis*. The designer can hire a marketing firm to survey the market and provide distributional estimates of agent preferences. This mode of operation is quite reasonable in large markets. However, in large markets mechanism design is not such an interesting topic; each agent will have little impact on the others and therefore the designer may as well stick to posted-pricing mechanisms. Indeed, for commodity markets posted prices are standard in practice. Mechanisms on the other hand are most interesting in small, a.k.a., *thin*, markets. Contrast the large market for personal computers to the thin market for super computers. There may be five organizations in the world in the market for super computers. How would a designer optimize a mechanism for selling super computers? First, even if the agents' values do come from a distribution, the only way to sample the distribution is to interview the agents themselves. Second, even if we did interview the agents, the most data points we could obtain is five. This is hardly enough for statistical approaches to be able to estimate the distribution of agent values. This strongly motivates a question (which we will also answer in this chapter) that is closely related to prior-independent mechanism design: How many samples from a distribution are necessary to design a mechanism that can approximate the optimal mechanism for the distribution?

There are other reasons to consider prior-independent mechanism design besides the questionable origin of prior information. The most striking of which is the frequent inability of a designer to redesign a new mechanism for each scenario in which she wishes to run a mechanism. This is not just a concern, in many settings it is a principle. Consider the standard Internet routing protocol TCP/IP. This is the protocol responsible for sending emails, browsing web pages, steaming video, etc. Notice that the workloads for each of these tasks is quite different. Emails are small and can be delivered with several minutes delay without issue. Web pages are small, but must be delivered immediately. Comparably, video streaming requires a high responsiveness and a large bandwidth. There is not the flexibility to install new protocols in Internet routers each time a new network usage pattern arises. Instead, a good protocol, such as TCP/IP, should work pretty well in any setting, perhaps ones well beyond the imaginations of the original designers of the Internet.

The final motivation we will discuss for prior-independent mechanism design is that the solution of Bayesian optimal (or approximate) mechanisms is incomplete. It solves the problem of what a designer should do who knows the prior-distribution, but in many real situations a designer may not have such knowledge. Requiring the designer to acquire distribution information from "outside the system", therefore, does not completely solve the designer's problem.

5.2 "Resource" Augmentation

In this section we describe a classical result from auction theory which shows that a little more competition in a surplus maximizing mechanism revenue-dominates the profit maximizing mechanism without the increased competition. From an economic point of view this result questions the *exogenous-participation* assumption, i.e., that there a certain number of agents, say n, that will participate in the mechanism. If, for instance, agents only participate in the mechanism if their utility from doing so is large enough, i.e., with *endogenous participation*, then running an optimal mechanism may decrease participation and then result in a lower revenue than the surplus maximizing mechanism.

On the other hand, the suggestion of this result, that a little increasing competition can ensure good revenue, is inherently prior-independent. The designer does not need to know the prior distribution to market her service so as to attract more agent participation.

5.2.1 Single-item Auctions

The following theorem is due to Jeremy Bulow and Peter Klemperer and is known as the Bulow-Klemperer Theorem.

Theorem 5.1. *For i.i.d., regular distributions the expected revenue of the second-price auction on $n+1$ agents is at least the expected revenue of the optimal auction on n agents.*

Proof. First consider the following question. What is the optimal single-item auction for $n+1$ agents that always sells the item? The requirement to always sell the item means that, even if all virtual values are less than zero, a winner must still be selected. Clearly the optimal such auction is the one that assigns the item to the agent with the highest virtual value. Since the distribution is i.i.d. and regular, the agent with the highest virtual value is the agent with the highest value. Therefore, this optimal auction that always sells the item is the second-price auction.

Now consider an $n+1$ agent mechanism that we will call "mechanism B". Mechanism B runs the optimal auction on the first n agents and if this auction fails to sell the item, it gives the item away for free to the last agent. By definition, B's expected revenue is equal to the expected revenue of the optimal n-agent auction. It is, however, an $n+1$-agent auction that always sells. Therefore, its revenue is at most that of the optimal $n+1$-agent auction that always sells.

We conclude that the expected revenue of the second-price auction with $n+1$ agents is at least that of mechanism B which is equal to that of the optimal auction for n agents. □

This resource augmentation result provides the beginning of a prior-independent theory for mechanism design. For instance, we can easily obtain a prior-independent approximation result as a corollary to Theorem 5.1 and Theorem 5.2, below.

Theorem 5.2. *For i.i.d., regular, single-item environments the optimal* $(n-1)$*-agent auction is an* $\frac{n}{n-1}$*-approximation to the optimal* n*-agent auction revenue.*

Proof. See Exercise 5.1. □

Corollary 5.3. *For i.i.d., regular, single-item environments with* n *agents, the second-price auction is an* $\frac{n}{n-1}$*-approximation to the optimal auction revenue.*

5.2.2 Matroid Environments

Unfortunately, the "just add a single agent" result fails to generalize beyond single-item environments. Suppose instead that there are k identical units for sale. Is the $k+1$st-price auction (i.e., the one that sells to the k highest-valued agents at the $k+1$st value) revenue on $n+1$ agents at least that of the optimal k-unit auction on n agents? It is certainly not.

Consider the special case where $k = n$ and the values are distributed uniformly on $[0, 1]$. The expected revenue of the $n+1$st-price auction on $n+1$ agents is about one as there are n winners any the $n+1$st value is about $1/n$ in expectation. Of course the optimal auction will offer a price of $1/2$ and achieve an expected revenue of $n/4$.

It turns out that the resource augmentation result does extend, and in a very natural way, but we will have to recruit more than a single agent. For k-unit auctions we will have to recruit k additional agents. Notice that to extend the proof of Theorem 5.1 to the k-unit setting we can define the auction B to allocate optimally to the first n agents and then any remaining items can be given to the k additional agents. The desired conclusion results. In fact, this argument can be extended to matroids. Of course matroid set systems are generally asymmetric so we have to specific what kind of agents we are adding. This is formalized by the definition and theorem below.

Definition 5.4. *A* base *of a matroid is an independent set of maximal cardinality.*

Theorem 5.5. *For any i.i.d., regular, matroid environment the expected revenue of the surplus maximization mechanism is at least that of the optimal mechanism in the environment induced by removing the agents corresponding to any base of the matroid.*

Notice that by the augmentation property of matroids, all bases are the same size. Notice that the theorem implies the aforementioned result for k-unit auctions as any set of k agents forms a base.

5.3 Single-sample Mechanisms

While the assumption that it is possible to recruit an additional agent seems not to be too severe; once we have to recruit k new agents in k-unit auctions or a new base for matroid environments, the approach provided by Bulow-Klemperer Theorem seems less relevant. In this section we will show that a single additional agent is enough to obtain a good approximation to the optimal auction revenue. We will not, however, just add this agent to the market, instead we will use this agent for statistical purposes.

In the opening of this chapter we discussed the need to connect the size of the sample for market analysis with the size of the actual market. In this context, the assumption that the prior distribution is known is tantamount to assuming that an infinitely large sample is available for market analysis. In this section we show that this impossibly large sample market can be approximated by a single sample from the distribution.

Mechanism 5.1. *The* lazy single-sample mechanism *is the following:*

1. $(\mathbf{x}', \mathbf{p}') \leftarrow \mathrm{SM}(\mathbf{v})$,
2. *draw a single sample from the distribution* $r \sim F$,
3. $x_i = \begin{cases} x'_i & \text{if } v_i \geq r \\ 0 & \text{otherwise, and} \end{cases}$
4. $p_i = \max(r, p'_i)$,

where SM *denotes the surplus maximization mechanism.*

In comparison to the surplus maximization mechanism with reserve prices discussed in Chapter 4, where the reserve prices are used filter out low-valued agents out before finding the surplus maximizing set, in the lazy single-sample mechanism the reserve price filters out low-valued agents after finding the surplus maximizing set. In matroid environments, which include single- and k-unit auctions, the order in which the reserve price is imposed is irrelevant (i.e., the same outcome results), therefore, in such environments we will refer to the lazy single-sample mechanism as the *single-sample mechanism.*

5.3.1 The Geometric Interpretation

Consider a single-agent, single-item environment. The optimal auction in such an environment is simply to post the monopoly price as a take-it-or-leave-it offer. The single-sample mechanism in this context posts a random, from the distribution, price as a take-it-or-leave-it offer. We will give a geometric proof that shows that for regular distributions, the revenue from this random price is within a factor of two of the revenue from the (optimal) monopoly price.

This result can be viewed as the $n = 1$ special case of the Bulow-Klemperer Theorem, i.e., that the two-agent second-price auction obtains at least the (one-agent) monopoly revenue. In a two-agent second-price auction each agent is offered the a price equal to the value of the other, i.e., a random price from the distribution. Therefore, the two-agent second-price auction obtains twice the revenue of the single sample. The result showing that the single-sample revenue is at least half of the monopoly revenue then implies that the two-agent second-price auction obtains at least the (one-agent) monopoly revenue.

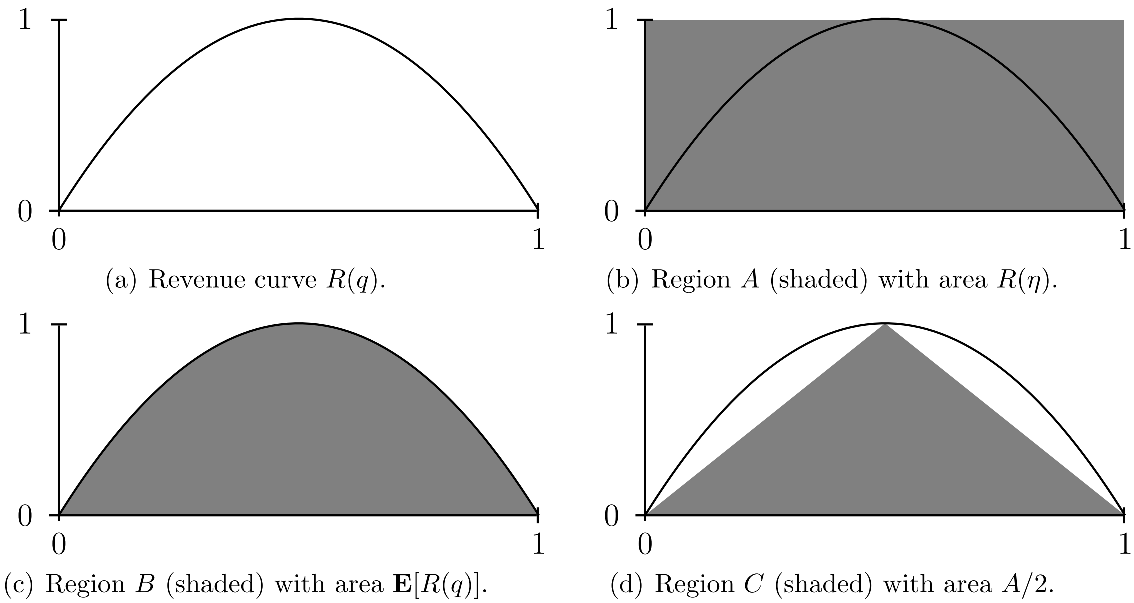

Figure 5.1: In the geometric proof of the that a random reserve is a 2-approximation to the optimal reserve, the areas of the shaded regions satisfy $A \geq B \geq C = A/2$.

Lemma 5.6. *For a single-agent with value drawn from regular distribution F, the revenue from a random take-it-or-leave-it offer $r \sim F$ is at least half the revenue of the (optimal) monopoly offer.*

Proof. Let $R(q)$ be the revenue curve for F in quantile space. Let η be the quantile corresponding to the monopoly price, i.e., $\eta = \operatorname{argmax}_q R(q)$. The expected revenue from such a price is $R(\eta)$. Recall that drawing a random value from the distribution F is equivalent to drawing a uniform quantile $q \sim U[0, 1]$. The expected revenue from such a random price is $\mathbf{E}_q[R(q)] = \int_0^1 R(q)\, dq$. In the Figure 5.1 the area of region A is $R(\eta)$. The area of region B is $\mathbf{E}_q[R(q)]$. Of course, the area of C is less than the area of B, by concavity of $R(\cdot)$, but at least half the area of A, by geometry. The lemma follows. $\square$

5.3.2 Random versus Monopoly Reserves

The geometric interpretation above is almost all that is necessary to show that the lazy single-sample mechanism is a good approximation to the optimal mechanism. We will show this in two pieces. First we will show that the lazy single-sample mechanism is a good approximation to the revenue of the surplus maximization mechanism with a lazy monopoly reserve. Then we argue that this lazy monopoly reserve mechanism is optimal or approximately optimal.

Theorem 5.7. *For any i.i.d., regular distribution, downward-closed environment, the revenue of the lazy single-sample mechanism is a 2-approximation to that of the lazy monopoly reserve mechanism.*

Proof. Let REF denote the lazy monopoly reserve mechanism and its revenue, and let APX denote the lazy single-sample mechanism and its revenue.

For $\mathbf{v}_{-i}$. We argue that the expected revenue from agent i in APX is at least half of that in REF. REF and APX are deterministic and dominant strategy IC, therefore in each mechanism agent i faces a critical value for winning. It will be useful to consider this critical value in quantile space, henceforth, "critical quantile". Let τ_i be the critical quantile of the surplus maximization mechanism (with no reserve). In APX, i's critical quantile is $\min(\tau_i, q)$ for $q \sim U[0, 1]$. In REF, i's critical quantile $\min(\tau_i, \eta)$, where η is the quantile of the monopoly price.

Now consider the induced revenue curve in APX from agent i with value $v_i \sim F$ as a function of q in the case where $\tau_i \leq \eta$ and where $\tau_i > \eta$ (Figure 5.2). We show, via the geometric interpretation, that in each case APX is a 2-approximation. Notice that if $q \leq \tau_i$ then APX's revenue from i is $R(q)$, otherwise it is $R(\tau_i)$. The REF revenue from i is $R(\min(\tau_i, \eta))$. By concavity of $R(\cdot)$ and geometry (Figure 5.2) the theorem follows. ☐

5.3.3 Single-sample versus Optimal

We have shown that lazy random reserve pricing is almost as good as lazy monopoly reserve pricing. We now connect lazy monopoly reserve pricing to the optimal mechanism to show that the lazy single-sample mechanism is a good approximation to the optimal mechanism.

As discussed above lazy monopoly reserve pricing is identical to (eager) monopoly reserve pricing in matroid settings. Also, Theorem 4.24 showed that monopoly reserve pricing is optimal. We conclude the following corollary. (Of course, k-unit auctions are a special case of matroid environments.)

Corollary 5.8. *For any i.i.d., regular, matroid environment, the single-sample mechanism is a 2-approximation to the optimal mechanism revenue.*

For downward-closed environments we have slightly more work to do. Theorem 4.20 showed that for monotone-hazard-rate distributions surplus maximization with (eager) monopoly reserves is a 2-approximation to the optimal mechanism. For downward-closed environments, eager and lazy reserve pricing are not identical. However, as suggested by Exercise 4.3 (for the eventual proof of Theorem 4.17), the revenue of lazy monopoly reserve pricing is an e-approximation to the optimal surplus. Clearly, then its revenue is an e-approximation to the optimal revenue. We can conclude the following.

Corollary 5.9. *For any i.i.d., monotone-hazard-rate, downward-closed environment, the lazy single-sample mechanism is a* $2e$*-approximation to the optimal mechanism revenue.*

The bound in the above corollary can be improved to a factor of four, but we will not discuss the details here.

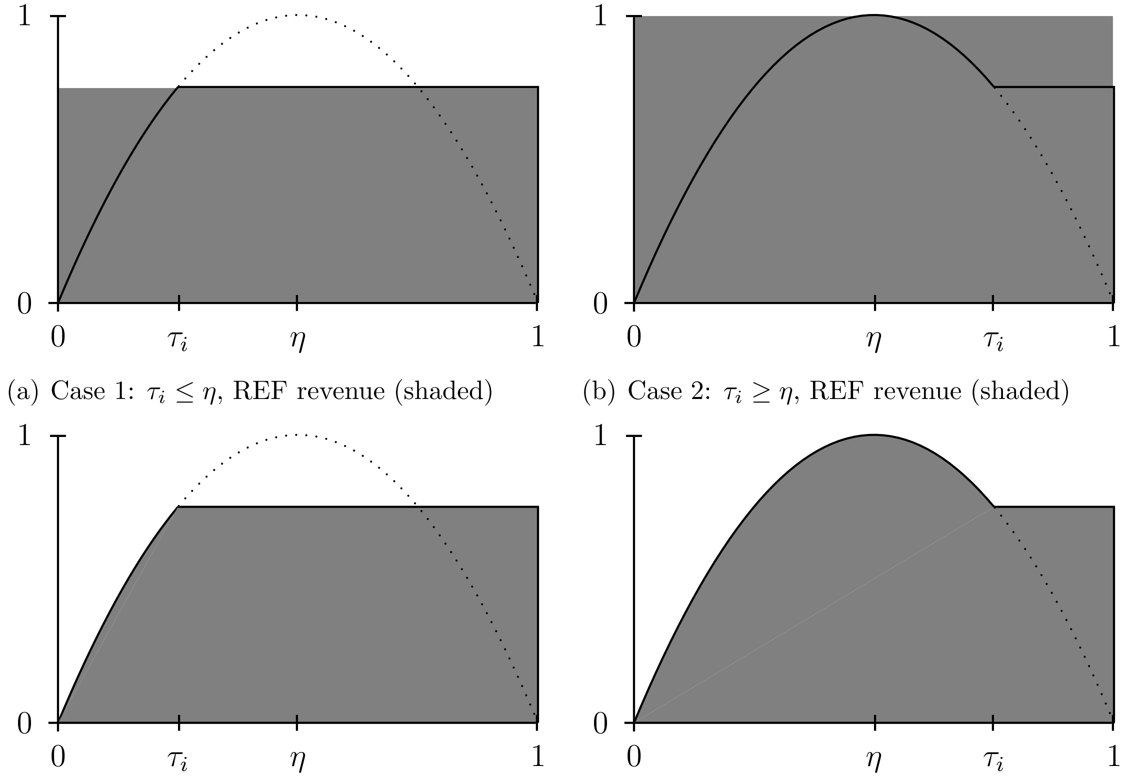

(a) Case 1: $\tau_i \leq \eta$, REF revenue (shaded)

(b) Case 2: $\tau_i \geq \eta$, REF revenue (shaded)

(c) Case 1: $\tau_i \leq \eta$, APX revenue (shaded)

(d) Case 2: $\tau_i \geq \eta$, APX revenue (shaded)

Figure 5.2: On the top is a geometric depiction of the payment (shaded area) of agent i in the lazy monopoly reserve (REF) mechanism; on the bottom is the same for the lazy single sample (APX) mechanism. On the left is the case where the critical quantile τ_i is less than the monopoly quantile η; on the right is the opposite case. The revenue curve is depicted with a dotted line, and the induced revenue curve given the critical quantile is depicted with a solid line.

5.4 Prior-independent Mechanisms

We now turn to mechanisms that are completely prior-independent. Unlike the mechanisms of the preceding section, these mechanisms will not require any distributional information in advance, not even a single sample from the distribution. We will, however, still assume that there is a distribution. We search for a single mechanism that has good expected performance for any distribution from a large class of distributions.

Definition 5.10. *A mechanism* APX *is a* prior-independent β-approximation *if*

$$\forall \mathbf{F}, \quad \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathrm{APX}(\mathbf{v})] \geq \tfrac{1}{\beta}\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathrm{REF}_{\mathbf{F}}(\mathbf{v})]$$

where $\mathrm{REF}_{\mathbf{F}}$ *is the optimal mechanism for distribution* $\mathbf{F}$*.*

The central idea behind the design of prior-independent mechanisms is that a small amount of market analysis can be done on-the-fly as the mechanism is being run; bids of some agents can be used for market analysis for other agents.

Consider the following k-unit auction:

0. Solicit bids.

1. Randomly reject an agent i^*.

2. Run the $k+1$st-price auction with reserve v_{i^*} on $\mathbf{v}_{-i^*}$.

This auction is clearly incentive compatible. Furthermore, it is easy to see that it is a $\frac{2n}{n-1}$-approximation for n agents with values drawn i.i.d. from a regular distribution. This follows from the fact that rejecting a random agent loses at most a $1/n$ faction of the optimal revenue and the previous single-sample result (Corollary 5.8). For $n \geq 2$ this mechanism guarantees a 4-approximation. The same approach can be applied to matroid and downward-closed environments as well; however, we will focus instead on a slightly more sophisticated approach.

5.4.1 Digital Good Environments

An important single-dimensional agent environment is that of a *digital good*. A digital good is one where there is little or no cost for duplication. The cost function for digital goods is $c(\mathbf{x}) = 0$ for all $\mathbf{x}$, or equivalently, all outcomes are feasible. Digital goods are the special case of k-unit auctions where $k = n$. Therefore the mechanism above obtains a $2n/(n-1)$-approximation.

There are a number of ways to improve this mechanism to remove the $n/(n-1)$ from the approximation factor. Two of the most natural are the following.

Definition 5.11.

- *The* pairing auction *arbitrarily pairs agents and runs the second-price auction on each pair (assuming n is even).*

- *The* circuit auction *orders the agents arbitrarily (e.g., lexicographically) and offers each agent a price equal to the value of the preceding agent in the order (the first agent is offered the last agent's value).*

The random pairing auction *and the* random circuit auction *are the variants where the implicit pairing or circuit is selected randomly.*

Theorem 5.12. *For i.i.d., regular, digital-good environments, any auction wherein each agent is offered the price of another random or arbitrary (but not value dependent) agent is a 2-approximation to the optimal auction revenue.*

The proof of this theorem follows directly from the geometric interpretation for the single-sample mechanism. Clearly, the pairing and circuit auctions satisfy the conditions of the above theorem. In conclusion, it is relatively easy, within a mechanism, to get samples from the distribution.

5.4.2 General Environments

We now adapt the results for digital goods to general environments. The main idea here is to replace the lazy single-sample reserve with a lazy circuit or pairing mechanism. Notice that in downward-closed environments we can view the lazy reserve pricing used with a surplus maximizing mechanism as a digital good auction. The surplus maximizing mechanism outputs a feasible outcome. Since all subsets of feasible outcomes are feasible, the induced environment is essentially one of a digital good.

Two deterministic DSIC mechanisms, $\mathcal{M}'$ and $\mathcal{M}''$, can be composed in many ways, perhaps the most natural is the following. Consider the critical values an agent i in each mechanism, τ_i' and τ_i'', respectively. Consider the composite mechanism $\mathcal{M}$ in which i's critical value is $\tau_i = \max(\tau_i', \tau_i'')$. Notice that the set of agents served by $\mathcal{M}$ is the intersection of those served by $\mathcal{M}'$ and $\mathcal{M}''$. This outcome is feasible by downward-closure and DSIC by its definition via critical values. Notice that the surplus maximization mechanism with lazy reserves is the composition, in this manner, of the surplus maximization mechanism with the mechanism that simply makes a take-it-or-leave-it offer of the reserve price to each agent.

Instead of composing the surplus maximization mechanism with the reserve pricing, we can compose it with either the pairing or circuit auctions. Both of the theorems below follow from analyses similar to that of the single-sample mechanism.

Definition 5.13.

- *The* pairing mechanism *is the composition of the surplus maximization mechanism with the (digital goods) pairing auction.*

- *The* circuit mechanism *is the composition of the surplus maximization mechanism with the (digital goods) circuit auction.*

Theorem 5.14. *For i.i.d., regular, matroid environments, the pairing and circuit mechanisms are 2-approximations to the optimal mechanism revenue.*

Theorem 5.15. *For i.i.d., monotone-hazard-rate, downward-closed environments, the pairing and circuit mechanisms are* 4*-approximations to the optimal mechanism revenue.*

Two issues remain undiscussed. First, our prior-independent mechanisms were derived from single-sample mechanisms. Clearly more on-the-fly samples can be used to obtain revenue that more closely approximates the lazy monopoly reserve pricing, and therefore, the optimal auction. Second, similarly more on-the-fly samples can be used to obtain prior-independent approximation mechanisms when distributions may be irregular. Both of these directions will be taken up during our discussion of prior-free mechanisms in Chapter 6.

Exercises

5.1 Prove Theorem 5.2: For i.i.d., regular, single-item environments the optimal $(n-1)$-agent auction is an $\frac{n}{n-1}$-approximation to the optimal n-agent auction revenue.

5.2 Suppose we are in a non-identical environment, i.e., agent i's value is drawn from independently from distribution F_i, and suppose the mechanism can draw one sample from each agent's distribution.

(a) Give a constant approximation mechanism for regular, matroid environments (and give the constant).

(b) Give a constant approximation mechanism for monotone-hazard-rate, downward-closed environments (and give the constant).

(c) Conclude with a two Bulow-Klemperer style theorems. Suppose you can recruit a competitor for each agent (i.e., from the same distribution but where the set system only allows an agent or her competitor to be served), then the surplus maximization mechanism with competitors obtains a constant fraction of the optimal revenue in the original environment. Give one theorem for regular, matroid environments and one for monotone-hazard-rate, downward-closed environments.

5.3 This chapter has been mostly concerned with the profit objective. Suppose we wished to have a single mechanism that obtained good surplus and good profit.

(a) Show that surplus maximization with monopoly reserves is not generally a constant approximation to the optimal social surplus in regular, single-item environments.

(b) Show that the lazy single sample mechanism is a constant approximation to the optimal social surplus in i.i.d., regular, matroid environments.

(c) Investigate the Pareto fontier between prior-independent approximation of surplus and revenue. I.e., if a mechanism is an α approximation to the optimal surplus and a β-approximation to the optimal revenue, plot it as point $(1/\alpha, 1/\beta)$ in the positive quadrant.

5.4 Suppose the agents are divided into k markets where the value of agents in the same market are identically distributed. Assume that the partitioning of agents into markets is known, but not the distributions of the markets. Assume there are at least two agents in each market. Unrelated to the markets, assume the environment has a downward-closed feasibility constraint.

(a) Give a prior-independent constant approximation to the revenue-optimal mechanism for regular, matroid environments.

(b) Give a prior-independent constant approximation to the revenue-optimal mechanism for monotone-hazard-rate, downward-closed environments.

Chapter Notes

The Coase conjecture, which states that a monopolist cannot sell early at a high price to high-valued consumers and late at a low price to low-valued consumers as the late price will compete with the high price, is due to Ronald Coase (1972).

The resource augmentation result that shows that recruiting one more agent to a single-item auction raises more revenue than setting the optimal reserve price is due to Bulow and Klemperer (1996). The proof of the Bulow-Klemperer Theorem that was presented in this text is due to René Kirkegaard (2006).

The single-sample mechanism and the geometric proof of the Bulow-Klemperer theorem is due to Dhangwatnotai et al. (2010). The pairing auction for digital good environments was proposed by Goldberg et al. (2001); however, in their, potentially irregular, environment it does not have good revenue guarantees.