

5

Prior-independent Approximation

In the last two chapters we discussed mechanism that performed well for a given Bayesian prior distribution. Assuming the existence of such a Bayesian prior is natural when deriving mechanisms for games of incomplete information as the Bayes-Nash equilibrium concept requires a prior distribution that is common knowledge. In this chapter we will relax the assumption the designer has knowledge of the prior distribution and is able to tune the parameters of her mechanism with it. The goal of *prior-independent* mechanism design is to identify a single mechanism that has good performance for all distributions in a large family of relevant distributions, e.g., the family of 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 for one prior, there is another prior for which another auction performs strictly better. This consequence is obvious because optimal auctions for distinct distributions are generally distinct. Nonetheless, while no single 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 mechanism design. The first approach considers “resource” augmentation. We will show that in some environments the (prior-independent) externality pricing mechanism with increasing competition, e.g., by recruiting more agents, earns more revenue than the revenue-optimal mechanism without the increased competition. The second approach is to design mechanisms that do a little market analysis on the fly. Via this second approach, 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 its motivation; however, before doing so recall the original justification for the common prior assumption (see [Section 2.3](#)). Auctions and mechanisms are games of incomplete information and in such games, in order to argue about strategic choice, we needed to formalize how players deal with uncertainty. We did this by assuming a Bayesian prior. In a *Stackelberg game*, instead of moving simultaneously, players make actions in a prespecified order. We can view mechanism design as a two stage Stackelberg game where the designer is a player who moves first and the agents are players who (simultaneously) move second. To analyze the Bayes-Nash equilibrium in such a Stackelberg game, the designer bases her strategy on the common prior. Without such prior knowledge, the problem of predicting the designer's strategy is ill posed. Thus, in so far as the theory of mechanism design should describe (or predict) the outcome of a game, within the standard equilibrium concept for games of incomplete information, a prior assumption is necessary.

As discussed in [Chapter 1](#) in addition to being descriptive, the theory of mechanism design should be prescriptive. It should suggest to a designer how to solve a given mechanism design problem that she may confront. If the designer does not have prior information, then she cannot directly employ the suggestions of Bayesian mechanism design. The Bayesian theory of mechanism design is, thus, incomplete in so far as it would require the designer to acquire distribution information from "outside the system." In contrast, a prior-independent mechanism is required to solve both information acquisition and incentive problems and, therefore, must insure that losses due to inaccuracies in information acquisition and the interplay between information acquisition and incentives are properly accounted for.

It is important to consider the incentives of information acquisition within the mechanism design problem; even if the designer has knowledge of a prior distribution, it may be problematic to employ this knowledge

Chapter 5: Topics Covered.

- the Bulow-Klemperer Theorem,
- the single-sample mechanism, and
- mechanism composition.

in a mechanism. Suppose the designer obtained her prior knowledge from previous market experience. A problem with designing the mechanism with this knowledge is that the earlier agents may strategize so that information about their preferences is not exploited by the designer later. For example, a monopolist who cannot commit not to lower the sale price in the future cannot sell at a high price now (see [Exercise 5.1](#)).

It is similarly important to consider the losses due to inaccuracies in information acquisition within the mechanism design problem. To learn the prior a designer could perform a *market analysis*, for example, by hiring 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 automobiles to the thin market for spacecrafts. There may be five organizations in the world in the market for spacecrafts; how would a designer optimize a mechanism for selling them? First, even if the agents' values do come from a distribution, the only way to sample from the distribution is to interview the agents themselves. Second, even if we did interview the agents, we could obtain at most five data points. This sample size is hardly enough for statistical approaches to be able to estimate the distribution of agent values. A motivating question this perspective raises, and one that is closely tied to prior-independent mechanism design, is: How many samples from a distribution are sufficient for the design of an approximately optimal mechanism?

There are other reasons to consider prior-independent mechanism design besides the questionable origin of prior information. Perhaps 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 IP. This protocol is responsible for sending emails, browsing web pages, streaming 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 permits high latency but requires continuous bandwidth. It would be impractical to install new protocols in Internet

routers each time a new network usage pattern arises. Instead, a protocol for computer networks, such as IP, should work pretty well in any setting, even ones well beyond the imaginations of the original designers.

5.2 “Resource” Augmentation

In this section we describe a classical result from auction theory that shows that a little more competition in a surplus maximizing mechanism revenue dominates the revenue 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 when their utility from doing so is large enough, i.e., with *endogenous participation*, then running an optimal mechanism may decrease participation and thus result in a lower revenue than the surplus maximizing mechanism.

On the other hand, the suggestion of this result, that slightly 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 Environments

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

Theorem 5.2.1. *For i.i.d. regular single-item environments, the expected revenue of the second-price auction with $n + 1$ agents is at least that of the optimal auction with 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 that the item always be sold implies that, even if all virtual values are negative, 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 (cf. [Corollary 3.3.5](#)). 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 LB that runs the optimal

auction on agents $1, \dots, n$ and if this auction fails to sell the item, it gives the item away for free to agent $n + 1$. Obviously, LB'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 the item. Therefore, its revenue is a lower bound on 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 LB which is equal to that of the optimal auction for n agents. \square

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 5.2.1 and Theorem 5.2.2 below.

Theorem 5.2.2. *For i.i.d. single-item environments the optimal $(n-1)$ -agent auction is an $n/n-1$ approximation to the optimal n -agent auction.*

Proof. See Exercise 5.2 \square

Corollary 5.2.3. *For i.i.d. regular single-item environments with $n \geq 2$ agents, the second-price auction is an $n/n-1$ approximation to the optimal auction revenue.*

5.2.2 Multi-unit and Matroid Environments

Unfortunately, the “just add a single agent” result fails to generalize beyond single-item environments. Consider a multi-unit environment; is the revenue of the $(k + 1)$ st-price auction (i.e., the one that sells a unit to each of the k highest-valued agents at the $(k + 1)$ st highest value) with $n + 1$ agents at least that of the optimal k -unit auction with n agents? No.

Example 5.2.1. *For large n consider an n -unit environment and agents with uniformly distributed values on $[0, 1]$. With $n + 1$ agents, the expected revenue of the $(n + 1)$ st-price auction on $n + 1$ agents is about one as there are n winners and the $(n + 1)$ st value is $1/n+2 \approx 1/n$ in expectation¹. On the other hand, the optimal auction with n agents will post a price of $1/2$ to each agent and achieve an expected revenue of $n/4$.*

The resource augmentation result does extend, and in a very natural

¹ In expectation, uniform random variables evenly divide the interval they are over.

way, but more than a single agent must be recruited. For k -unit environments we have to recruit k additional agents. Notice that to extend the proof of [Theorem 5.2.1](#) to a k -unit environment we can define the auction LB for $n+k$ agents to run the optimal n -agent auction on agents $1, \dots, n$ and to give any remaining units to agents $n+1, \dots, n+k$. The desired conclusion follows. In fact, this argument can be extended to matroid environments. Of course matroid set systems are generally asymmetric, so we have to be specific as to the role with respect to the feasibility constraint of the added agents. The result is more intuitive when stated in terms of removing agents from the optimal mechanism instead of adding agents to the (surplus maximizing) externality pricing mechanism, though the consequence is analogous. Recall from [Section 4.6](#) that a base of a matroid is a feasible set of maximal cardinality.

Theorem 5.2.4. *For any i.i.d. regular matroid environment the expected revenue of the externality pricing mechanism [\(Definition 3.2.1\)](#) is at least that of the optimal mechanism in the environment obtained by removing any set of agents that corresponds to a base of the matroid.*

Recall that by the augmentation property of matroids, all bases are the same size. Notice that the theorem implies the aforementioned result for k -unit environments as any set of k agents forms a base of the k -uniform matroid. Similarly, for transversal matroids, which model constrained matching markets, recruiting a new base requires one additional agent for each of the items.

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 environments or a new base for matroid environments, the approach seems less actionable. 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 market analysis.

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 can be approximated by a single sample from the distribution.

Definition 5.3.1. *The externality pricing mechanism with lazy reserves $\hat{\mathbf{v}}$ is the following:*

- (i) *simulate the externality pricing mechanism on the bids,*

$$(\mathbf{x}^\dagger, \mathbf{p}^\dagger) \leftarrow \text{EP}(\mathbf{v}),$$

- (ii) *serve the winners of the simulation who exceed their reserve prices,*

$$x_i = \begin{cases} x_i^\dagger & \text{if } v_i \geq \hat{v}_i \\ 0 & \text{otherwise, and} \end{cases}$$

- (iii) *charge the winners (with $x_i = 1$) their critical values $p_i = \max(\hat{v}_i, p_i^\dagger)$,*

where EP denotes the externality pricing mechanism.

Definition 5.3.2. *The lazy single-sample-reserve mechanism is the externality pricing mechanism with lazy reserves $\hat{\mathbf{v}} = (\hat{v}, \dots, \hat{v})$ for $\hat{v} \sim F$. The lazy monopoly-reserve mechanism sets $\hat{\mathbf{v}} = \hat{\mathbf{v}}^*$.*

Proposition 5.3.1. *The externality pricing mechanism with lazy reserves is dominant strategy incentive compatible.*

In comparison to the externality pricing mechanism with reserve prices discussed in Chapter 4 where the reserve prices are used filter out low-valued agents before finding the surplus maximizing set (i.e., eagerly), lazy reserve prices filter out low-valued agents after finding the surplus maximizing set. It is relatively easy to find examples of downward-closed environments for which the order in which the reserve is applied affects the outcome (see Exercise 5.3). On the other hand, matroid environments, which include single-item and multi-unit environments, are distinct in that the order in which an anonymous reserve price is imposed does not change the auction outcome. Thus, for i.i.d. matroid environments we will not specify the order, i.e., lazy versus eager, of the reserve pricing.

5.3.1 The Geometric Interpretation

Consider a single-agent environment. The optimal auction in such an environment is simply to post the monopoly price as a take-it-or-leave-it offer. In comparison, the single-sample-reserve mechanism would post a

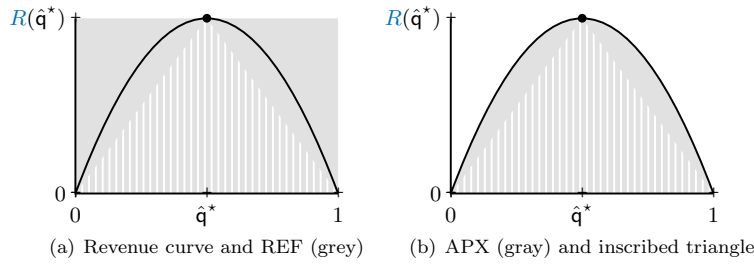


Figure 5.1. The revenue curve (black line) for the uniform distribution is depicted. REF is the area of the rectangle (gray); by geometry the area of the inscribed triangle (white striped) is $1/2$ REF. APX is the area under the revenue curve (gray); by convexity it is lower bounded by the area of the inscribed triangle (white striped). Thus, $\text{REF} \geq \text{APX} \geq 1/2 \text{REF}$.

random price that is drawn from the same distribution as the agent's value is drawn. We will give a geometric proof that shows that for regular distributions, the revenue from posting such a random price is within a factor of two of that of the (optimal) monopoly price.

This statement can be viewed as the $n = 1$ special case of the Theorem 5.2.1, 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 a single sample reserve. 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.

Lemma 5.3.2. *For a regular single-agent environment, posting a random price from the agent's value distribution obtains at least half the revenue from posting the (optimal) monopoly price.*

Proof. Let $R(\cdot)$ be the agent's revenue curve. Let \hat{q}^* be the quantile corresponding to the monopoly price, i.e., $\hat{q}^* = \operatorname{argmax}_{\hat{q}} R(\hat{q})$. The expected revenue from (optimal) monopoly pricing is $\text{REF} = R(\hat{q}^*)$; this revenue is represented in Figure 5.1(a) by the area of the rectangle (gray) of width one and height $R(\hat{q}^*)$. Recall that drawing a random value from the distribution is equivalent to drawing a uniform quantile. The expected revenue from the corresponding random price is $\text{APX} = \mathbf{E}_{\hat{q}}[R(\hat{q})] = \int_0^1 R(\hat{q}) d\hat{q}$; this revenue is depicted in Figure 5.1(b)

by the region below the revenue curve (grey). This region is convex because the revenue curve is concave; therefore, by geometry it contains an inscribed triangle with vertices corresponding to 0, \hat{q}^* , and 1 on the revenue curve (Figure 5.1 white striped). This triangle has width one, height $\text{REF} = R(\hat{q}^*)$, and therefore its area is equal to $1/2 \text{REF}$. Thus, $\text{APX} \geq 1/2 \text{REF}$. \square

Example 5.3.1. For the uniform distribution where $R(\hat{q}) = \hat{q} - \hat{q}^2$, the quantities in the proof of Lemma 5.3.2 can be easily calculated:

$$\begin{aligned} \text{REF} &= R(\hat{q}^*) = 1/4 \\ &\geq \text{APX} = \mathbf{E}_{\hat{q} \sim U[0,1]}[R(\hat{q})] = 1/6 \\ &\geq 1/2 \text{REF} = 1/8. \end{aligned}$$

5.3.2 Monopoly versus Single-sample Reserves

The geometric interpretation above is almost all that is necessary to show that the lazy single-sample-reserve mechanism is a good approximation to the optimal mechanism. We will show the result in two steps. First we will show that the lazy single-sample-reserve mechanism is a good approximation to the lazy monopoly-reserve mechanism. Then we argue that this lazy monopoly-reserve mechanism is approximately optimal.

Theorem 5.3.3. For i.i.d. regular downward-closed environments, the expected revenue of the lazy single-sample-reserve mechanism is at least half of that of the lazy monopoly-reserve mechanism.

Proof. With the values \mathbf{v}_{-i} of the other agents fixed, we will argue the stronger result that the contribution to the expected revenue from any agent i (Alice) in the lazy single-sample-reserve mechanism is at least half of that in the lazy monopoly-reserve mechanism (in expectation over her value and the sampled reserve). Let REF denote the lazy monopoly-reserve mechanism and Alice's contribution to its revenue, and let APX denote the lazy single-sample-reserve mechanism and her contribution to its revenue (again, both for fixed \mathbf{v}_{-i}).

Denote the monopoly quantile by \hat{q}^* , denote the critical quantile for Alice in the externality pricing mechanism with no reserve by \hat{q}_i^{EP} , and denote the quantile of a lazy reserve by \hat{q} . Alice's wins in the externality pricing mechanism with this lazy reserve when her quantile is below $\min(\hat{q}, \hat{q}_i^{\text{EP}})$. For a fixed \hat{q}_i^{EP} , the revenue from Alice, in expectation over her own quantile and as a function of the lazy reserve quantile \hat{q} ,

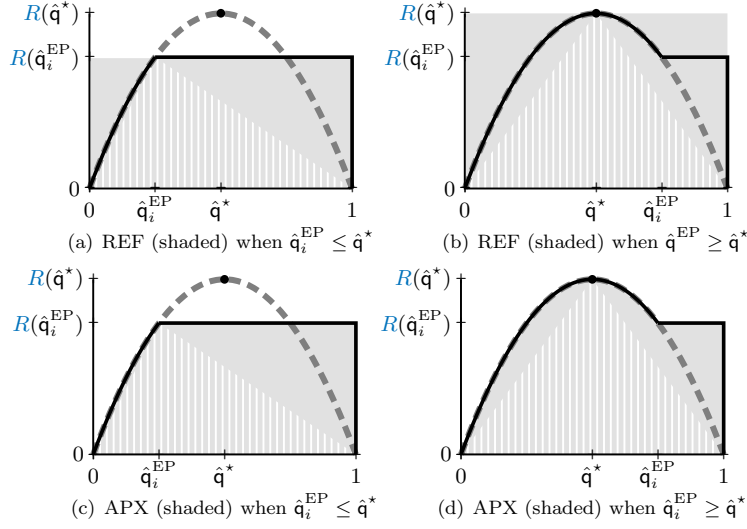


Figure 5.2. In each diagram, the revenue curve $R(\cdot)$ (thick, dashed, grey line) of the uniform distribution and the induced revenue curve $R^\dagger(\cdot) = R(\max(\cdot, \hat{q}_i^{\text{EP}}))$ (thin, solid, black line). On the left is the case that $\hat{q}_i^{\text{EP}} \leq \hat{q}^*$; on the right is the case that $\hat{q}_i^{\text{EP}} \geq \hat{q}^*$. On the top the revenue of REF is shaded grey; on the bottom the revenue of APX is shaded in gray. The inscribed triangles (white striped) have area $1/2$ REF. Both on the left and on the right $\text{REF} \geq \text{APX} \geq 1/2 \text{REF}$.

induces the revenue curve $R^\dagger(\hat{q}) = R(\min(\hat{q}, \hat{q}_i^{\text{EP}}))$. [Figure 5.2](#) depicts Alice's original revenue curve $R(\cdot)$ and this induced revenue curve $R^\dagger(\cdot)$ in the cases that $\hat{q}_i^{\text{EP}} \leq \hat{q}^*$ and $\hat{q}_i^{\text{EP}} \geq \hat{q}^*$.

Alice's expected payment in the lazy monopoly-reserve mechanism is $\text{REF} = R^\dagger(\hat{q}^*)$ which is geometrically the maximum height of the revenue curve R^\dagger ; and her expected payment in the lazy single-sample-reserve mechanism, where $\hat{q} \sim U[0, 1]$, is $\text{APX} = \mathbf{E}_{\hat{q}}[R^\dagger(\hat{q})]$. We conclude with the same geometric argument as in [Lemma 5.3.2](#) that relates REF to a rectangle, APX to the area under the induced revenue curve, and $1/2 \text{REF}$ to the area of an inscribed triangle (see [Figure 5.2](#)). \square

5.3.3 Optimal versus Lazy Single-sample-reserve Mechanism

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

For i.i.d. matroid environments, as discussed previously, lazy monopoly reserve pricing is identical to (eager) monopoly reserve pricing. Moreover, the externality pricing mechanism with the monopoly reserve is revenue optimal (Proposition 4.6.1). We conclude the following corollary. Recall that matroid environments include multi-unit environments as a special case.

Corollary 5.3.4. *For any i.i.d. regular matroid environment, the revenue of the single-sample-reserve mechanism is a two approximation to that of the revenue-optimal mechanism.*

Theorem 4.7.4 shows that for monotone-hazard-rate distributions the externality pricing mechanism with (eager) monopoly reserves is a two approximation to the optimal mechanism; however, as in downward-closed environments, where eager and lazy reserve pricing are not identical (see Exercise 5.3), we have slightly more work to do. Recall Theorem 4.7.1 which states that for MHR distributions the optimal revenue and the optimal surplus are within an ϵ factor of each other. One way to prove this theorem is, in fact, by showing that the revenue of the externality pricing mechanism with lazy monopoly reserve prices is an ϵ approximation to the optimal social surplus and hence so is the optimal revenue (see Exercise 4.25). Combining this observation with Theorem 5.3.3 it is evident that the lazy single-sample-reserve mechanism is a 2ϵ approximation. The approximation bound can be improved to four via a more careful analysis (Exercise 5.7).

Theorem 5.3.5. *For any i.i.d. monotone-hazard-rate downward-closed environment, the revenue of the lazy single-sample-reserve mechanism is a four approximation to that of the revenue-optimal mechanism.*

Proof. See Exercise 5.7

□

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, not even a single sample from the distribution. We will, however, still assume that there is a distribution.

Definition 5.4.1. *A mechanism APX is a prior-independent β approximation if*

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

where $\text{REF}_{\mathbf{F}}$ is the (prior-dependent) optimal mechanism for distribution \mathbf{F} and “ $\forall \mathbf{F}$ ” quantifies over all distributions in a given family.

The central idea behind the design of prior-independent mechanisms is that a small amount of market analysis can be done while the mechanism is being run. For example, the bids of some agents can be used as a market analysis to calculate the prices to be offered to other agents.

Consider the following k -unit auction:

- (i) Solicit bids,
- (ii) randomly reject an agent j , and
- (iii) run the $(k + 1)$ st-price auction with reserve v_j on \mathbf{v}_{-j} .

This auction is clearly incentive compatible. Furthermore, it is easy to see that it is a $2^{n/n-1}$ approximation for $n \geq 2$ 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 $1/n$ fraction of the optimal revenue (Theorem 5.2.2), and from the previous single-sample-reserve result (Corollary 5.3.4). This approximation bound is clearly worst for $n = 2$ where it guarantees a four approximation. The same approach can be applied to matroid and downward-closed environments as well; instead, we will discuss a slightly more sophisticated approach.

5.4.1 Digital Good Environments

An important single-dimensional agent environment is that of a *digital good*, i.e., one where there is little or no cost for duplication. In terms of single-dimensional environments for mechanism design, the cost function for digital goods is $c(\mathbf{x}) = 0$ for all \mathbf{x} ; in other words, all outcomes are feasible. Digital goods can also be viewed as the special case of k -unit auctions where $k = n$. Therefore the mechanism above obtains a $2^{n/n-1}$ approximation.

There are a number of approaches for improving this mechanism to remove the $n/n-1$ from the approximation factor. The following two approaches are natural.

Definition 5.4.2. *For digital-good environments,*

- *the (digital good) pairing auction arbitrarily pairs agents and runs the second-price auction on each pair (assuming n is even), and*
- *the (digital good) 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 pairing or circuit is selected randomly.

Theorem 5.4.1. *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 two approximation to the optimal auction revenue.*

The proof of this theorem follows directly from the geometric analysis of single-sample pricing ([Lemma 5.3.2](#)). Clearly, the pairing and circuit auctions satisfy the conditions of the above theorem. In conclusion, in i.i.d. environments it is relatively easy to obtain samples from the distribution while running a mechanism.

5.4.2 General Environments

We now adapt the results for digital goods to general environments. Consider the externality pricing mechanism with a lazy reserve price. First, the surplus maximizing set is found. Second, the agents that do not meet the reserve are rejected. We can view this second step as a digital good auction as, once we have selected a surplus maximizing feasible set, downward closure requires that any subset is feasible. The main idea of this section is to replace the lazy reserve part of the single-sample mechanism with any approximately optimal digital good auction (e.g., the circuit or pairing auction).

Consider [the following definition of mechanism composition](#) (cf. Exercise 5.10). Notice that the mechanisms we have been discussing can all be interpreted as calculating a critical value for each agent, serving each agent whose value exceeds her critical value, and charging each served

agent her critical value. In fact, by [Corollary 2.10.4](#) any randomization over deterministic dominant strategy incentive compatible mechanisms admits such an interpretation.

Definition 5.4.3. *The parallel composite \mathcal{M} of two (randomizations over) deterministic DSIC mechanisms, \mathcal{M}^\dagger and \mathcal{M}^\ddagger is as follows:*

- (i) Calculate the critical values \hat{v}^\dagger and \hat{v}^\ddagger of \mathcal{M}^\dagger and \mathcal{M}^\ddagger , respectively.
- (ii) The critical values of \mathcal{M} are $\hat{v}_i = \max(\hat{v}_i^\dagger, \hat{v}_i^\ddagger)$ for each agent i .
- (iii) Allocation and payments are $x_i = x_i^\dagger x_i^\ddagger$ and $p_i = \hat{v}_i x_i$ for all i , respectively.

Notice that in the parallel composite, \mathcal{M} , the set of agents served is the intersection of those served by \mathcal{M}^\dagger and \mathcal{M}^\ddagger . By downward closure, then, the outcome of the composition is feasible as long as the outcome of one of \mathcal{M}^\dagger or \mathcal{M}^\ddagger is feasible. The mechanism is dominant strategy incentive compatible by its definition via critical values and [Corollary 2.10.4](#).

Proposition 5.4.2. *The parallel composite of two (randomizations over) deterministic dominant strategy incentive compatible mechanisms is dominant strategy incentive compatible and, if one of the mechanisms is feasible, feasible.*

Notice that the externality pricing mechanism with a lazy reserve price is the composition, in the manner above, of the externality pricing mechanism with a (digital good) uniform posted pricing. Consider composing the externality pricing mechanism with either the pairing or circuit auctions. Both of the theorems below follow from analyses similar to that of the single-sample-reserve mechanism.

Definition 5.4.4. *For downward-closed environments,*

- *the pairing mechanism is the parallel composite of the externality pricing mechanism with the (digital goods) pairing auction, and*
- *the circuit mechanism is the parallel composite of the externality pricing mechanism with the (digital goods) circuit auction.*

Theorem 5.4.3. *For i.i.d. regular matroid environments, the revenues of the pairing and circuit mechanisms are two approximations to the optimal mechanism revenue.*

Theorem 5.4.4. *For i.i.d. monotone-hazard-rate downward-closed environments, the revenues of the pairing and circuit mechanisms are four approximations to the optimal mechanism revenue.*

The results presented in this chapter are representative of the techniques for the design and analysis of prior-independent approximation mechanisms; however, a number of extensions are possible. If we use more than one samples from the distribution, bounds for regular distributions can be improved and bounds for irregular distributions can be obtained. Both of these directions will be taken up during our discussion of prior-free mechanisms in Chapter 7. Finally, the i.i.d. assumption can be relaxed, either by assuming that agents are partitioned by demographic (see Exercise 5.11) or by an ordering assumption.

Exercises

- 5.1 Consider the sale of a magazine subscription over two periods to a single agent who has a linear uniform additive value for each period's issue of the magazine. Her value v is drawn from a regular distribution F and if x_1 , x_2 , p_1 , and p_2 denote her allocation and payments in each period then her utility is $v(x_1 + x_2) - p_1 - p_2$. In each period, the designer publishes her mechanism and then the agent bids for receiving that period's issue of the magazine.
- Suppose that the designer can commit to the mechanism to be used in period two before the agent bids in period one, describe the revenue optimal mechanisms and the equilibrium behavior of the agent.
 - Suppose that the designer cannot commit to the mechanism to be used in period two before the agent bids in period one, describe the revenue optimal mechanisms and the equilibrium behavior of the agent. Assume that the designer deterministically posts a price in period 1 and the agent follows a threshold strategy, i.e., the agent accepts the price if her value is at least some threshold and otherwise she rejects the price.
 - Compare the revenues from the previous steps for the uniform distribution.
- 5.2 Prove Theorem 5.2.2 For i.i.d. single-item environments the optimal auction with $n - 1$ agents auction is an $n/n-1$ approximation to the optimal auction with n agents.
- 5.3 Consider the externality pricing mechanism with an anonymous reserve that is either lazy or eager.
- Find a valuation profile, downward-closed feasibility constraint,

- and anonymous reserve price such that different outcomes result from lazy and eager reserve pricing.
- (b) Prove that for anonymous reserve pricing in matroid environments, lazy and eager reserve pricing give the same outcome.
- 5.4 Consider a regular single-agent environment. Show that posting the median price from the agent's value distribution obtains at least half the revenue from posting the monopoly price. The median price for an agent with inverse demand function $v(\cdot)$ is $\hat{v} = v(1/2)$.
- 5.5 In [Example 5.3.1](#) it is apparent that the approximation bound of a sample reserve to the monopoly reserve for a uniform distribution is $3/2$. Use this bound to derive better bounds for the lazy single-sample-reserve mechanism versus the lazy monopoly-reserve mechanism. In particular, show that if the single-agent approximation of sample reserve to monopoly reserves is β then the same bound holds in general for the lazy single-sample-reserve and lazy monopoly reserve mechanism.
- 5.6 Consider the externality pricing mechanism with lazy monopoly reserve prices in downward-closed monotone-hazard-rate environments.
- (a) Show that in a single-agent environment, that its expected surplus is at most twice its expected revenue.
- (b) Show that in a downward-closed environment, that its expected surplus is at most twice its expected revenue.
- 5.7 Prove [Theorem 5.3.5](#): *For any i.i.d. monotone-hazard-rate downward-closed environment, the revenue of the lazy single-sample-reserve mechanism is a four approximation to that of the revenue-optimal mechanism.*
- 5.8 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).
- 5.9 This chapter has been mostly concerned with the profit objective. Suppose we wished to have a single mechanism that obtains good surplus and good profit.

- (a) Show that the externality pricing mechanism 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 frontier 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.
- 5.10 Define the *sequential composite* \mathcal{M} of two mechanism \mathcal{M}^\dagger and \mathcal{M}^\ddagger as first simulating \mathcal{M}^\dagger , second simulating \mathcal{M}^\ddagger on the winners of \mathcal{M}^\dagger , and serving the agents served by the second mechanism at the maximum of their prices in the two mechanisms.
- (a) Give an example of deterministic DSIC mechanisms \mathcal{M}^\dagger and \mathcal{M}^\ddagger such that the sequential composite \mathcal{M} is not DSIC.
 - (b) Show that if \mathcal{M}^\dagger is the externality pricing mechanism (and \mathcal{M}^\ddagger is any randomization over DSIC mechanisms) then the composition is DSIC.
 - (c) Describe a property of the externality pricing mechanism as \mathcal{M}^\dagger that enables the incentive compatibility of the sequential composite \mathcal{M} .
- 5.11 Suppose the agents are divided into k markets where the value of agents in the same market are identically distributed, e.g., by demographic. 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 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 \(2006\)](#). A generalization of the Bulow-Klemperer Theorem to non-identical distributions was given by [Hartline and Roughgarden \(2009\)](#).

The single-sample mechanism and the geometric proof of the Bulow-Klemperer theorem are due to [Dhangwatnotai et al. \(2010\)](#). They also considered a relaxation of the i.i.d. assumption where there is a known partitioning of the agents into markets, e.g., by demographic or zip code, where there are at least two agents in each market. The pairing auction for digital good environments was proposed by [Goldberg et al. \(2001\)](#); however, in the possibly irregular environments that they considered it does not have good revenue guarantees.

A line of work has investigated optimal single-sample and prior-independent mechanisms for i.i.d. regular distributions. Notably, [Fu et al. \(2015\)](#) proves that (a) the optimal single-agent mechanism with a single sample does not simply offer the agent the sample and (b) the optimal two-agent single-item auction is not the second-price auction. In both cases the optimal mechanism is strictly better than a two-approximation.