Chapter 6

Prior-free 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-free mechanism design is in identifying a single mechanism that always has good performance, e.g., under any distributional assumption.

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 distributional setting for which another auction performs strictly better. This is obvious because the optimal auctions for distinct distributional settings are distinct. This lack of an absolute optimal mechanism requires relaxing our objective. Therefore, instead of absolute optimal mechanisms we will look for relative optimal mechanisms. The prior-free optimal mechanism is then the one with the smallest approximation factor to the Bayesian optimal mechanism.

In this chapter we will take three approaches to prior-free mechanisms. The first is a resource augmentation, a.k.a., bicriteria, approach. We will show that in some cases increasing competition, e.g., by recruiting more agents, and running the VCG mechanism increases revenue more than running the (optimal) Myerson mechanism. The second is a average-case approach. We will show that for a large class of distributions and settings there is a single mechanism that approximates the revenue of the Bayesian optimal mechanism. The third and final approach is worst case. We will define a prior-free performance benchmark and derive mechanisms that approximate this benchmark on any worst-case valuation profile. The benchmark is so selected to imply that such an approximation mechanism simultaneously approximates all Bayesian optimal mechanisms.

6.1 Motivation

Since prior-freedom is not without loss it is important to consider the motivation behind going from Bayesian optimal (or approximations) to prior-free optimal (or approximations) mechanisms.

First 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 agents deal with uncertainty. In the sense that our single-shot games actually modeled static behavior in a repeated scenario it made sense to assume that agents are able to best-respond to a distribution over their opponents' preferences. If we view mechanism design as a *Stackleberg game*, i.e., one were players move in a prespecified order, where the designer is a player who moves first, if we are to stick to the standard Bayes-Nash equilibrium concept then to reason about the designer's strategy we must assume the designer knows the distribution.

Now consider from where the designer may have learned the prior-distributions. There are two most logical candidates. The first is, as alluded to above, 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 automobiles to the thin market for space crafts. There may be five organizations in the world in the market for space crafts. How would a designer optimize a mechanism for selling space crafts? 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 related to prior-free mechanism design which is 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-free mechanism design besides the questionable origin of prior information. The most striking is the frequent inability of a designer to redesign a new mechanism for each scenario they wish to run a mechanism in. 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-free 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.

6.2 "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 is often better than 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 partition and harm its revenue which could result in a worse 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-free. The designer does not need to know the prior distribution to market their service so as to attract more agent participation.

6.2.1 Single-item Auctions

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

Theorem 6.1 *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 it 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. ■

6.2.2 Matroid Settings

Unfortunately the "just add a single agent" result fails to generalize beyond single-item auctions. Suppose instead that there are k identical units for sale. Is the k-Vickrey auction (i.e., the one that sells to the k highest-valued agents at the $k+1$st value) 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-Vickrey on $n+1$ agents is about 1 as there are n winners any the $n+1$st value is about $1/n$. Of course the optimal auction will offer a price of $1/2$ and achieve an expected revenue of about $n/2$.

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 6.1 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 6.2 A *base* of a matroid is an independent set of maximal cardinality.

Theorem 6.3 *For i.i.d., regular distributions in matroid settings, the expected revenue of the VCG mechanism on n agents plus an additional base is at least the expected revenue of the optimal auction on the original n agents.*

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.

6.3 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 a k new agents in k-unit auctions or a new base for matroid settings, 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.

Definition 6.4 The *lazy single-sample mechanism* LSS is the following:

1. Simulate VCG($\mathbf{v}$).
2. Draw a single sample from the distribution $r \sim F$.
3. Reject any winners i of VCG with $v_i < r$.

In comparison to the VCG mechanism with reserve prices discussed in Chapter 5, where the reserve prices are used filter out low-valued agents out before running VCG, in the lazy single-sample mechanism the reserve price filters out low-valued agents after running VCG. In matroid settings, 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 contexts we will therefore refer to the lazy single-sample mechanism as the *single-sample mechanism*.

6.3.1 The Geometric Interpretation

Consider the one agent, one item setting. The optimal auction in such a setting 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 for regular distributions, the revenue from this random price is within a factor of two of the revenue from the (optimal) monopoly price.

This can be viewed as a proof of the $n = 1$ special case of the Bulow-Klemperer Theorem. The Bulow-Klemberer theorem in this case says that the 2-agent second-price auction achieves more revenue than the 1-agent optimal auction. Of course, the in a 2-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 2-agent second-price auction obtains twice the revenue of the single sample. Since the single-sample revenue is at least half of the optimal revenue, the 2-agent second-price auction obtains at least the revenue of the optimal 1-agent auction.

Lemma 6.5 *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 (optimal) monopoly offer.*

Proof: Let $R(q)$ be the revenue curve for F in quantile space. Let q^* be the quantile corresponding to the monopoly price, i.e., $q^* = \text{argmax}_q R(q)$. The expected revenue from such a price is $R(q^*)$. Recall that drawing a random value from the distribution F is equivalent to drawing a uniform quantile $q \sim U[0, 1]$. The revenue from such a random price is $R(q)$. In

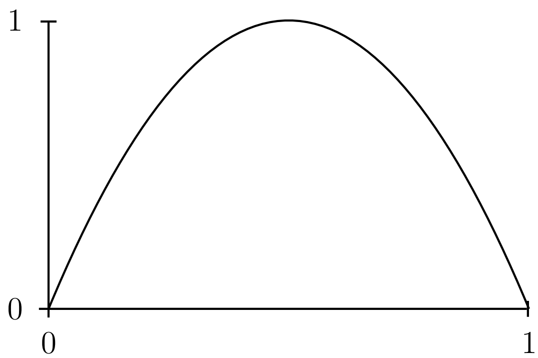

(a) Revenue curve $R(q)$.

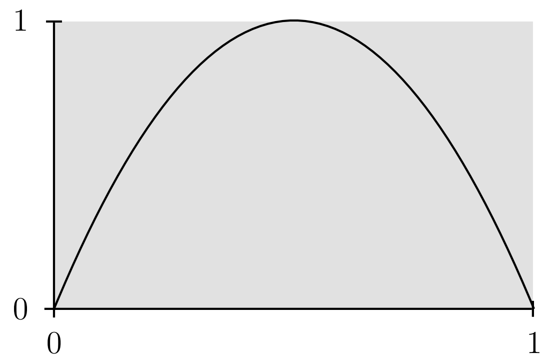

(b) Region A (shaded) with area $R(q^*)$.

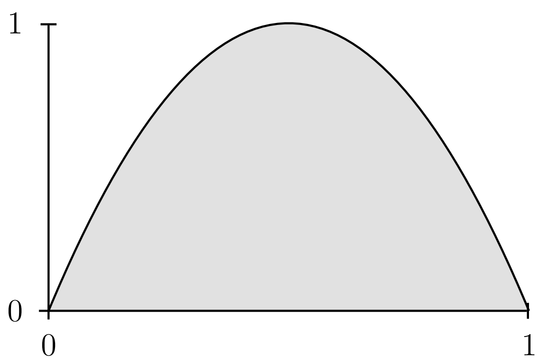

(c) Region B (shaded) with area $\mathbf{E}[R(q)]$.

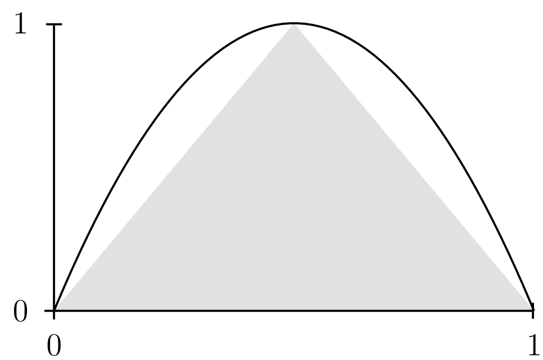

(d) Region C (shaded) with area $A/2$.

Figure 6.1: 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 $A \geq B \geq C = A/2$.

the Figure 6.1 the area of region A is $R(q^*)$. 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. ■

6.3.2 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 VCG with a lazy monopoly reserve. Then we argue that VCG with a lazy monopoly reserve is optimal or near optimal.

Theorem 6.6 *For any i.i.d., regular distribution and any downward-closed setting, the expected revenue of the lazy single-sample (LSS) mechanism is at least half of that of VCG with lazy monopoly reserve (LMR).*

Proof: Fix $\mathbf{v}_{-i}$. We argue that the expected revenue from agent i in LSS is at least half of that in LMR. LSS and LMR are deterministic and ex post 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 t_i be the critical quantile of VCG. In LSS, i's critical quantile is $\min(t_i, q)$ for $q \sim U[0, 1]$. In LMR, i's critical quantile $\min(t_i, q^*)$, where q^* is the quantile of the monopoly reserve.

Now consider the induced revenue curve in LSS from agent i with value $v_i \sim F$ as a function of q in the case where $t_i \leq q^*$ and where $t_i > q^*$ (Figure 6.2). Again we show, via the geometric interpretation, that in each case LSS is a 2-approximation. Notice that if $q \leq t_i$ then LSS's revenue from i is $R(q)$, otherwise it is $R(t_i)$. The LMR revenue from i is $R(\max(t_i, q^*))$. By concavity of $R(\cdot)$ and geometry the theorem follows. ■

6.3.3 Single-sample versus Optimal

We have shown that lazy random reserve pricing is almost as good as optimal lazy reserve pricing. We now connect optimal lazy 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 monopoly reserve pricing in matroid settings. Also, Theorem 5.19 showed that monopoly reserve pricing is optimal. We conclude the following corollary. (Of course, k-unit auctions are a special case of matroid settings.)

Corollary 6.7 *For any i.i.d., regular distribution and any matroid setting, the single-sample mechanism gives a 2-approximation to the revenue of the optimal mechanism.*

For downward-closed settings we have slightly more work to do. Theorem 5.16 showed that for monotone-hazard-rate distributions VCG with (eager) monopoly reserves is a 2-approximation to the optimal mechanism. For downward-closed settings, eager and lazy

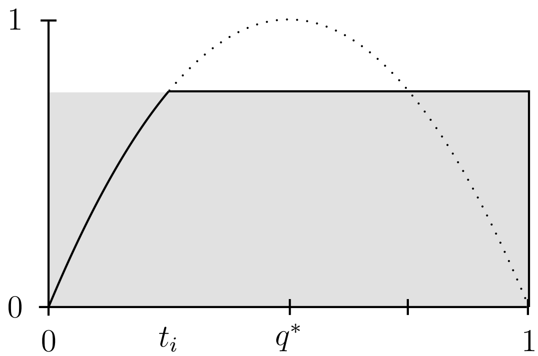

(a) Case 1: $t_i \leq q^*$, LMR revenue (shaded)

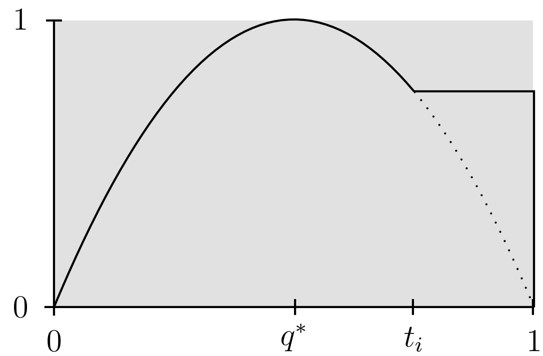

(b) Case 2: $t_i \geq q^*$, LMR revenue (shaded)

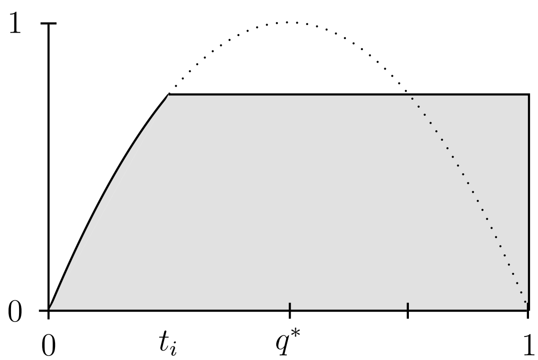

(c) Case 1: $t_i \leq q^*$, LSS revenue (shaded)

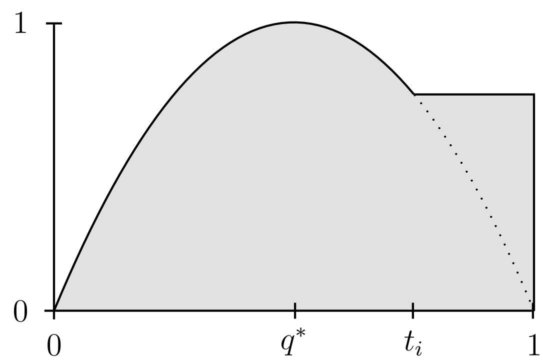

(d) Case 2: $t_i \geq q^*$, LSS revenue (shaded)

Figure 6.2: In each case, the revenue curve (dotted), the induced revenue curve of LSS (solid), and the LSS and LMR revenues (shaded).

reserve pricing are not identical. Instead we will use Theorem 5.13, which proves that the revenue of VCG with lazy monopoly reserves is an e-approximation to the optimal surplus. Clearly, then its revenue is an e-approximation to the optimal revenue. We can then conclude the following.

Corollary 6.8 *For any i.i.d., monotone-hazard-rate distribution and any downward-closed setting, the lazy single-sample mechanism gives a $2e$-approximation to the revenue of the optimal mechanism.*

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

6.4 Average-case Mechanisms

We now turn to mechanism that are completely prior-free. 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 6.9 A mechanism $\mathcal{M}$ is an *average-case prior-free β-approximation* if

$$\forall \mathbf{F}, \quad \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathcal{M}(\mathbf{v})] \geq \frac{\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathrm{Mye}_{\mathbf{F}}(\mathbf{v})]}{\beta}$$

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

The central idea behind the design of prior-free 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:

1. Randomly select an agent i^*.

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

3. Reject agent 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 6.7). For $n \geq 2$ this mechanism guarantees a 4-approximation. The same approach can be applied to matroid and downward-closed settings as well; however, we will not discuss them in more detail here.

6.4.1 The Digital Good Setting

An important single-dimensional agent setting is that of *digital goods*. 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:

- 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).

There are also randomized variants of each of these auctions, namely the *random pairing auction* and the *random circuit auction*.

Theorem 6.10 *For i.i.d. and regular distributions, any digital goods auction wherein each agent is offered the price of another random or arbitrary (but not value dependent) agent is a 2-approximation to the revenue of the optimal auction.*

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. The point of this exercise is to notice that it is relatively easy, within a mechanism, to get samples from the distribution.

6.4.2 General Settings

We now adapt the result for digital goods to general settings. The main idea here is to replace the lazy single-sample reserve with a lazy circuit or pairing mechanism. Notice that in downward-closed settings we can view the lazy reserve price (used with VCG) as a digital good auction. VCG outputs a feasible outcome. Since all subsets of feasible outcomes are feasible, the setting is essentially one of digital goods.

Two deterministic ex post IC 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, t'_i and t''_i, respectively. Consider the composite mechanism $\mathcal{M}$ in which i's critical value is $t_i = \max(t'_i, t''_i)$. 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 and by definition of the critical values it is IC. Notice that VCG with lazy reserves is the composition, in this manner, of VCG with the mechanism that simply posts the reserve price to each agent.

Instead of composing VCG with the reserve pricing, we can compose it with either the pairing or circuit auctions.

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

Both of the theorems below follow from analyses similar to that of the single-sample mechanism.

Theorem 6.11 *For i.i.d. and regular distributions in matroid settings, the paring and circuit mechanisms give 2-approximations to the optimal mechanism.*

Theorem 6.12 *For i.i.d. and monotone-hazard-rate distributions in downward-closed settings, the pairing and circuit mechanisms give* $2e$*-approximations to the optimal mechanism.*[1]

It is possible to get good average-case prior-free mechanisms for irregular distributions in downward-closed settings; however, these mechanisms use techniques developed for worst-case prior-free mechanism so we defer discussion of this general case to the next section.

6.5 Worst-case Mechanisms

We now drop the assumption that the valuations are drawn from any distribution at all. The first hurdle in performing such an analysis is that without a distribution there is no natural benchmark to which our mechanism's performance can be compared. However, we could, for any given benchmark, $\mathcal{G}$, consider the ability of a mechanism to approximate it on any valuation profile.

Definition 6.13 A mechanism $\mathcal{M}$ is an *worst-case prior-free* β*-approximation* to benchmark $\mathcal{G}$ if

$$\forall \mathbf{v}, \quad \mathcal{M}(\mathbf{v}) \geq \frac{\mathcal{G}(\mathbf{v})}{\beta}.$$

It is important to approach the worst-case prior-free setting from the context of the Bayesian and average-case approaches already discussed. For instance, ideally a mechanism that was a good worst-case approximation would also be a good average-case approximation. If not, our worst-case prior-free analysis would not have good economic foundation.

Definition 6.14 The *Bayesian optimal benchmark* is

$$\mathcal{G}(\mathbf{v}) = \sup_{\mathbf{F} \in \mathrm{IID}} \mathrm{Mye}_{\mathbf{F}}(\mathbf{v}).$$

Notice that an auction that is a worst-case prior-free β-approximation to the Bayesian optimal benchmark $\mathcal{G}$ is also an average-case prior-free β-approximation to the Bayesian optimal auction for any i.i.d. distribution. In fact, the worst-case prior-free approximation simultaneously approximates the performance of all Bayesian optimal mechanisms.

[1] This approximation factor can be improved to four.

Theorem 6.15 *If $\mathcal{M}$ is a worst-case prior-free β-approximation to the Bayesian optimal benchmark then $\mathcal{M}$ is an average-case prior-free β-approximation.*

Unfortunately it is impossible to obtain a constant approximation to this benchmark in the special case where $n = 1$ (for the objective of profit maximization). Notice that a mechanism in such a setting must just make the single agent an uninformed take-it-or-leave-it offer. No such offer can be better than a $(1 + \ln h)$-approximation when the value of the agent is known to be within $[1, h]$; when nothing is known approximation is only worse. In other words, when most of the benchmark revenue comes from a single agent, it is impossible to approximate the benchmark. We leave explicit proof of this impossibility as an exercise.

This impossibility is viewed as a technicality. To avoid it, we slightly weaken the benchmark to exclude from it the case where any agent contributes more than half of the benchmark revenue.

Definition 6.16 The profit benchmark $\mathcal{G}^{(2)}$ is $\mathcal{G}^{(2)}(\mathbf{v}) = \mathcal{G}(\mathbf{v}_{-(1)}, v_{(2)})$. I.e., it is the benchmark on the valuation profile with the second highest value duplicated in place of the highest value.

This benchmark allows the possibility that there is a single high-valued agent, but does not allow this agent to pay more than $v_{(2)}$. In many relevant cases, i.e., the ones where the it is not optimal to sell only to the agent with the highest value, $\mathcal{G}(\mathbf{v}) = \mathcal{G}^{(2)}(\mathbf{v})$. Furthermore, as we will see, it is possible to approximate this prior-free benchmark.

6.5.1 The Digital Good Setting

We start with the special case of digital good auctions where there is no feasibility constraint. Our goal is to approximate $\mathcal{G}(\mathbf{v}) = \sup_{\mathbf{F}\in\mathrm{IID}} \mathrm{Mye}_{\mathbf{F}}(\mathbf{v})$ (actually, $\mathcal{G}^{(2)}(\mathbf{v})$) so the first question to address is what is $\mathcal{G}(\mathbf{v})$ for the digital good setting?

The optimal auction for digital good settings with i.i.d. distribution F is to just post the monopoly price as a take-it-or-leave-it offer to each agent. Therefore, the supremum over optimal mechanisms is exactly the optimal revenue from a single posted price. Clearly the optimal such price will be one of the agent values. We conclude that $\mathcal{G}(\mathbf{v}) = \max_i iv_{(i)}$ and $\mathcal{G}^{(2)}(\mathbf{v}) = \max_{i\geq 2} iv_{(i)}$. Intuitively the latter is the revenue of the optimal single posted price with at least two winners.

We now consider approximating this benchmark. The main idea that enables approximation of the benchmark is that when figuring out a price to offer agent i we can use statistics from the values of all other agents $\mathbf{v}_{-i}$.

Deterministic Auctions

The above discussion motivates the following mechanism.

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

It is possible to show that DOP is an average-case prior-free constant approximation. However, we are focusing on worst-case prior-free auctions here and DOP is not a good worst-case approximation. For example, consider the valuation profile with ten high-valued agents, with value ten, and 90 low-valued agents, with value one. What does DOP do on such a valuation profile? The offer to a high-valued agent is $t_h = 1$, as $\mathbf{v}_{-h}$ consists of 90 low-valued agents and 9 high-valued agents. The revenue from the high price is 90; while the revenue from the low price is 99. The offer to a low-valued agent is $t_1 = 10$, as $\mathbf{v}_{-1}$ consists of 89 low-valued agents and 10 high-valued agents. The revenue from the high price is 100; while the revenue from the low price is 99. Clearly with these offers all high-valued agents will win and pay one, while all low-valued agents will lose. The total revenue is ten, a far cry from the benchmark revenue of $\mathcal{G}^{(2)}(\mathbf{v}) = \mathcal{G}(\mathbf{v}) = 100$. In fact, this deficiency of DOP is one that is fundamental to all *anonymous* (a.k.a., symmetric) deterministic auctions.

Theorem 6.17 *No anonymous, deterministic digital good auction is better than a worst-case prior-free n-approximation.*

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

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

1. Observe that for any auction that is a good approximation, it must be that for all m, $t(m, 0) = h$. Otherwise, on the all h's input, the auction only achieves profit n while the optimal profit is $\mathcal{G}^{(2)} = hn$. Thus, the auction would be at most an h-approximation which is not constant.

2. Likewise, observe that for any auction that is a good approximation, it must be that for all m, $t(0, m) = 1$. Otherwise, on the all 1's input, the auction achieves no profit and is clearly not an approximation of $\mathcal{G}^{(2)} = n$.

3. For the final argument, consider taking m sufficiently large and looking at $t(k, m-k)$. As we have argued for $k = 0$, $t(k, m-k) = 1$. Consider increasing k until $t(k, m-k) = h$. This must occur since $t(k, m-k) = h$ when $k = m$. Let $k^* = \min\{k : t(k, m-k) = h\}$ be this transition point. Now consider an $n = m + 1$ agent valuation profile with $n_h(\mathbf{v}) = k^*$ and $n_1(\mathbf{v}) = m - k^* + 1$. Consider separately the offer prices to high- and low-valued agents:

    - For low-valued agents: $t(n_h(\mathbf{v}_{-1}), n_1(\mathbf{v}_{-1})) = t(k^*, m - k^*) = h$. Thus, all low-valued agents are rejected and contribute nothing to the auction profit.