Chapter 7

Multi-dimensional Approximation

Throughout the majority of this text we have assumed that the agents' private preferences are given by a single value for receiving an abstract service, i.e., that agents' types are single dimensional. We now turn to multi-dimensional environments where the agents' preferences are given by a multi-dimensional type. E.g., a home buyer may have distinct values for different houses on the market; an Internet user may have distinct values for various qualities of service; an advertiser on an Internet search engine may value traffic for search phrase "mortgage" higher than that for "loan", etc.

One of the most important example environments for multi-dimensional mechanism design is that of *combinatorial auctions*. In combinatorial auctions each agent has a valuation function that is defined across all bundles. I.e., if agent i receives bundle $S \subset \{1, ..., m\}$ then she has value $v_i(S)$. A combinatorial auction assigns to agent i bundle S_i and payment p_i. For such an outcome, agent i's utility is given by $v_i(S_i) - p_i$, i.e., it is *quasi-linear*.

For the objective of social surplus, the single-dimensional-agent surplus maximization mechanism (Mechanism 3.1) generalizes and is optimal. In this generalization, agents report their multi-dimensional preferences, the mechanism chooses the outcome that maximizes social surplus for the reported preferences, and it charges each agent the externality imposed on the remaining agents. The proof of the following theorem follows in a similar fashion to that of Theorem 3.7 and Corollary 3.8.

Theorem 7.1. *For agents with (generally multi-dimensional) quasi-linear preferences, the surplus maximization mechanism is dominant strategy incentive compatible and maximizes the social surplus.*

Even though the surplus maximization mechanism is optimal (for social surplus), it is sometimes infeasible to run. For instance, in many environments posted-pricing mechanisms are used in place of auction-like mechanisms. We will show that posted-pricing mechanisms can approximate the optimal social surplus in some relevant environments, though not for general combinatorial auctions.

For the objective of profit, there are no general descriptions of optimal mechanisms for environments where agents have multi-dimensional preferences. Essentially, mechanisms for multi-dimensional environments are complex and optimizing over them does not yield concise or intuitive descriptions, nor does it yield practical mechanisms. In this section we will explore approximation for the objective of profit maximization. In particular, we will show that both surplus maximization with reserve prices and posted-pricing mechanisms can approximate the optimal mechanism. Furthermore, the prices in these mechanisms that perform well can be easily calculated and interpreted.

We will use as a running example in this chapter the environment of *matching markets*. In a matching market there are n agents and m items (e.g., houses). Each agent i has a value v_{ij} for house j. The agents are unit-demand, i.e., each wants at most one house, and the houses are unit-supply, i.e., each can be sold to at most one agents. Agent values are drawn independently at random, e.g., with $v_{ij} \sim F_{ij}$.

7.1 Item Pricing

We start with the special case of the matching markets where there is only one agent, i.e., $n = 1$. In this environment an important optimization problem is to identify revenue-optimal pricings. I.e., a pricing $\mathbf{p} = (p_1, \ldots, p_m)$ such that when the agent buys the item that generates the highest positive utility, i.e., the j that maximizes $v_j - p_j$, the revenue of the seller is maximized.

Unfortunately, there is no concise economic understanding of optimal pricings and their revenue. Therefore, in pursuit of goal approximately optimal pricings, the first hurdle is in finding concise understanding of an upper bound on the revenue of an optimal pricing. Then, if a pricing approximates this upper bound, it also approximates the optimal pricing.

The main idea in obtaining an upper bound is from the thought experiment where we imagine that instead of one agent with unit-demand preferences over the m items that we have m (single-dimensional) agents who each want their specific item, but with the constraint that at most one can be served. In this latter environment the optimal selling mechanism would be the optimal single-item auction derived in Chapters 3. Notice that while, in the pricing problem, the seller can only post a price on each item, in the auction problem, competition between agents can drive the price up. Therefore, intuition suggests that the revenue in the (single-dimensional) auction environment may be an upper bound on the revenue in the (multi-dimensional) pricing environment. This is indeed the case.

Theorem 7.2. *For any product distribution $\mathbf{F} = F_1 \times \cdots \times F_m$, the expected revenue of the optimal single-agent, m-item pricing when the agent's values for the items are drawn from $\mathbf{F}$ is at most that of the optimal single-item, m-agent auction when the agents' values for the item are drawn from $\mathbf{F}$.*

Proof. Any item pricing $\mathbf{p}$ can be converted into a single-item auction $\mathcal{A}_{\mathbf{p}}$ such that the expected revenue from the item pricing is at most that of the auction. For convenience define $v_0 = p_0 = 0$. The auction $\mathcal{A}_{\mathbf{p}}$ assigns the item to the agent j that maximizes $v_j - p_j$. For any fixed values of the other agents, $\mathbf{v}_{-j}$, this allocation rule is monotone in agent j's value and therefore ex post incentive compatible. It is also deterministic, so by Corollary 2.18

there is a critical value τ_j for agent j which is the infimum of values for which the agent wins the auction; the agent pays exactly this critical value on winning. Of course $\tau_j \geq p_j$.

Now notice that the allocation rule of the auction $\mathcal{A}_{\mathbf{p}}$ is identical to the allocation rule of the pricing $\mathbf{p}$. For the pricing the agent chooses the item that maximizes $v_j - p_j$; for the auction the winner is selected to maximize $v_j - p_j$. Furthermore, the revenue for the pricing is exactly the p_j that corresponds to this j whereas in the auction it is τ_j which, as discussed, is at least p_j. Therefore, the auction $\mathcal{A}_{\mathbf{p}}$ obtains at least revenue of the pricing $\mathbf{p}$.

Therefore, the optimal auction obtains at least the revenue of the optimal pricing. □

With the upper bound from optimal single-item auctions in hand, our goal of approximating the optimal pricing can be refined to approximating this optimal single-item auction revenue. In fact, the desired approximation result is an immediate consequence of Theorem 4.10 for single-item auctions, i.e., that for any $\mathbf{F}$ a sequential posted pricing with constant ironed virtual prices is a 2-approximation to the optimal single-item auction revenue. Of course, the revenue of our single-agent, m-item environment is no worst than that of a single-item, m-agent sequential posted pricing (because the sequential posted pricing revenue is, by definition, from the worst possible ordering of the agents).

Corollary 7.3. *For any independent, unit-demand, single-agent environment, a pricing with uniform ironed virtual prices is a 2-approximation to the optimal pricing revenue.*

For single-agent environments item pricings are equivalent to deterministic mechanisms. This equivalence follows from a multi-dimensional variant of Corollary 2.18 which is generally known as the *taxation principle*. Therefore, an approximation to the optimal pricing revenue is equivalently an approximation of the optimal deterministic mechanism. (We defer discussion of approximation of randomized mechanisms to Section 7.3.)

7.2 Reduction: Unit-demand to Single-dimensional Preferences

It should be noted that the construction in the preceding section can be viewed as a reduction from multi-dimensional unit-demand preferences to single-dimensional preferences. We can conclude that from the perspective of approximation, the multi-dimensional unit-demand preferences are similar enough to single-dimensional preferences that a good approach to unit-demand environments is to simulate the outcome of the corresponding single-dimensional environment. We now make that connection and the reduction precise. (Crucial to this connection is the independence of the agents' values.)

Formally, consider the following *general unit-demand environment*. There are n agents and m services each agent i has value v_{ij} for service j. An outcome is an assignment of agents to services (perhaps with some agents left unassigned). We will denote this assignment by the indicator $\mathbf{x}$ with $x_{ij} = 1$ if i receives service j and 0 otherwise. There is an arbitrary feasibility constraint over such assignments which we denote, as before, with a cost function $c(\cdot)$ which is zero or infinity for feasibility problems. We assume, without loss of generality, the implicit feasibility constraint that each agent can only receive one service, i.e., $\mathbf{x}$ such that $x_{ij} = x_{i'j} = 1$ for $i \neq i'$ have $c(\mathbf{x}) = \infty$.

A unit-demand environment is thus specified by the distribution $\mathbf{F}$ indexed by agent-service pairs and the cost function $c(\cdot)$ over outcomes $\mathbf{x}$, also indexed by agent-item pairs. In all of the results described herein, the agents will be independently distributed; in most of the results the items will also be independently distributed.

7.2.1 Single-dimensional Analogy

As in the pricing environment we can define the single-dimensional analog to any general unit-demand environment. In this analog, each unit-demand agent is replaced with a single-dimensional representative for eached desired service. Notice that in the single-dimensional analog the implicit feasibility constraint that a unit-demand agent can receive at most one service is translated to the constraint that only one of its representatives can be served at once.

Definition 7.4. *The* representative environment *for the n agent, m service unit-demand environment given by $\mathbf{F}$ and $c(\cdot)$ is the single-dimensional environment given by $\mathbf{F}$ and $c(\cdot)$ with nm single-dimensional agents indexed by coordinates ij.*

7.2.2 Upper bound

The restriction that only one representative of each unit-demand agent can be served at once induces competition between representatives. Intuitively this competition should result in an increased revenue in the optimal mechanism for the representative environment over the original unit-demand environment. Were this the whole story, the optimal revenue in the representative environment would be an upper bound on the optimal revenue in the original environment. In fact it is almost the whole story. The optimal mechanism for the representative environment (which is deterministic) is an upper bound on the optimal deterministic mechanism for the original (unit-demand) environment.

Detailed discussion of randomized mechanisms for multi-dimensional environments are deferred to Section 7.3, where we will see that, while a randomized mechanism for the unit-demand environment can obtain more revenue than the optimal mechanism for the representative environment, it is only by a constant factor more, e.g., a factor of two for single-agent environments. Therefore, a constant times the revenue of the optimal mechanism for the representative environment is an upper bound on the optimal (randomized) unit-demand mechanism. Such a bound is sufficient for obtaining constant approximations via the reduction described here.

Theorem 7.5. *For any independent, unit-demand environment, the optimal deterministic mechanism's revenue is at most that of the optimal mechanism for the single-dimensional representative environment.*

Proof. The proof of this theorem is similar to that of Theorem 7.2. See Exercise 7.2. ∎

	Item 1	Item 2
Agent 1	$v_{11} = 4$	$v_{12} = 5$
Agent 2	$v_{21} = 10$	$v_{22} = 8$

(a) agent values

	Item 1	Item 2
Agent 1	$p_{11} = \mathbf{2}$	$p_{12} = 4$
Agent 2	$(p_{21} = 5)$	$p_{22} = \mathbf{4}$

(b) agent specific item prices

Figure 7.1: The tables above depict agent values and posted prices in a two-agent two-item matching environment. When agent 1 arrives before agent 2, then agent 1 buys 1, agent 2 buys 2, and the revenue is 6 (purchace prices depicted in boldface). If the agents arrive in the opposite order a higher revenue is obtained.

7.2.3 Reduction

The goal of this section it to reduce the problem of designing a mechanism that approximates the optimal unit-demand mechanism to a single-dimensional-agent approximation problem. Following the techniques developed in Chapter 4 it may then be possible to instantiate the reduction by solving the single-dimensional-agent approximation problem.

For the unit-demand single-agent item-pricing example of Section 7.1, Corollary 7.3, which states that item pricing can approximate the Bayesian optimal auction in the single-agent unit-demand environment, follows from Theorem 4.10, which states that sequential posted pricings, i.e., where the agents arrive in any order, approximate the optimal multi-agent single-item auction. To see why this is, compare the tie-breaking rules in these two environments. In the unit-demand pricing problem the item is allocated that maximizes $v_j - p_j$. In the sequential posted pricing problem ties are broken in worst-case order, i.e., to maximize $-p_i$. Clearly, the expected revenue from multi-dimensional pricing is no worse than that of the single-dimensional pricing.

Extend the definition of sequential posted pricings to unit-demand environments with multiple agents (i.e., to generalize item prices). A sequential posted pricing is given by prices $\mathbf{p}$ with p_{ij} the price offered to agent i for service j. After the valuations are realized, the agents arrive in sequence and take their utility maximizing service that is still feasible, given the actions of preceding agents in the sequence. The revenue of such a process clearly depends on the sequence and we pessimisticly assume the worst-case. See Figure 7.1 for an example.

Definition 7.6. *A* sequential posted pricing *is an pricing of services (specialized) for each agent with the semantics that agents arrive in any order and take their favorite service that remains feasible. The revenue of such a pricing is given by the worst-case ordering.*

Consider the sequential posted pricing problem in both the original unit-demand environment and the representative single-dimensional environment. Suppose you had the choice of being the seller in one of these two environments, given the same distribution and costs, which environment would you choose? I.e., which environment gives a higher expected revenue? Whereas when considering auction problems, you would prefer the representative environment because of the increased competition, for sequential posted pricings there is no benefit from competition. In fact, the seller in the representative environment is at a disadvantage because the agents are in a worst case order and there are more possible orderings of the agents in the nm-agent representative environment than the n-agent original environment.

Theorem 7.7. *The expected revenue of a sequential posted pricing for unit-demand environments is at least the expected revenue of the same pricing in the representative single-dimensional environment.*

Proof. Compare sequential posted pricings for unit-demand environments (i.e., with n unit-demand agents) with sequential posted pricings for their representative environments (i.e., with nm single-dimensional agents). The difference between these two environments with respect to sequential posted pricings is that in the representative environment the nm agents can arrive in any order whereas in the original environment the an agent arrives and considers the prices on services ordered by utility. Thus, the set of orders in which the nm prices are considered in the representative environment contains the set of orders in the original environment. For worst-case sequences, then, the representative environment is worse. ☐

Combining this lower bound with the upper bound from Theorem 7.5 we have our reduction: approximation of the optimal mechanism by multi-dimensional sequential posted pricing reduces to that of single-dimensional sequential posted pricing.

Corollary 7.8. *If a sequential posted pricing is approximately optimal in the representative (single-dimensional) environment it is approximately optimal in the original (unit-demand) environment.*

7.2.4 Instantiation

It remains to instantiate the reduction from sequential posted pricing approximation in unit-demand environments to single-dimensional environments. I.e., we need to show that there are good sequential posted pricing mechanisms for single-dimensional environments. Here we will give such an instantiation for independent, regular, matching markets, i.e., where the services are items, and each item has only one unit of supply.

The representative environment for matching markets is one where there are nm agents and agent ij with value $v_{ij} \sim F_{ij}$ desires item j. For any original agent i and all j at most one representative ij can win. For any item j and all i at most one representative ij can win. The virtual surplus maximization mechanism, denoted VSM, is optimal for this single-dimensional environment.

Let q_{ij}^{VSM} be the probability that VSM serves representative ij. Let $p_{ij}^{\mathrm{VSM}} = F_{ij}^{-1}(1-q_{ij}^{\mathrm{VSM}})$ be the corresponding price at which, if posted to representative ij, would be accepted with probability q_{ij}^{VSM}. Now consider the pricing $p_{ij} = F_{ij}^{-1}(1 - q_{ij})$ for $q_{ij} = q_{ij}^{\mathrm{VSM}}/2$. These probabilities and prices can be calculated, for instance, by simulating the optimal mechanism.

Definition 7.9. *For representative matching market environments, the* simulation prices, $\mathbf{p}$, *satisfy* $p_{ij} = F_{ij}^{-1}(1 - \frac{1}{2}\mathbf{Pr}[$*the optimal mechanism serves* $ij])$ *for all* i *and* j.

We claim that sequential posted pricing with the simulation prices give an 8-approximation to the optimal mechanism's revenue. The theorem is proven in two steps, the first gives an upper bound on the revenue of the optimal mechanism in terms of the above prices and probability, the second gives a lower bound on the sequential pricing revenue in terms of the same. As will be evident from the proof, this bound is not tight; improving the bound is left for Exercise 7.3.

Theorem 7.10. *For regular distributions in the representative matching market environment, the sequential posted pricing with the simulation prices* $\mathbf{p}$ *is an 8-approximation to the revenue of the optimal mechanism.*

Lemma 7.11. *For regular distributions in the representative matching market environment, the expected revenue of the optimal mechanism,* VSM*, is at most* $\sum_{ij} p_{ij}^{\text{VSM}} q_{ij}^{\text{VSM}}$.

Proof. The proof of this lemma follows from a standard approach. Consider an "unconstrained" mechanism that allocates to each representative ij with probability at most q_{ij}^{VSM} but is not constrained by the original feasibility constraints, i.e., that only one representative ij of each agent i is served and that each item j is only allocated to at most one representative ij. In such an unconstrained environment the representatives do not interact at all. Furthermore, by regularity and the fact that the original p_{ij}^{VSM} are at least the monopoly price, the optimal unconstrained mechanism simply posts price p_{ij}^{VSM} to each representative ij. Its expected revenue is $\sum_{ij} p_{ij}^{\text{VSM}} q_{ij}^{\text{VSM}}$. Finally, VSM, the optimal mechanism for the constrained environment, is a valid solution to the unconstrained environment, therefore the optimal unconstrained mechanism revenue gives an upper bound on its revenue. □

Lemma 7.12. *For regular distributions in the representative matching market environment, the expected revenue from the sequential posted pricing of the simulation prices is at least* $\frac{1}{8}\sum_{ij} p_{ij}^{\text{VSM}} q_{ij}^{\text{VSM}}$.

Proof. If the sequential posted pricing is able to make an offer to agent ij then its expected revenue is $q_{ij}p_{ij} \geq q_{ij}^{\text{VSM}} p_{ij}^{\text{VSM}}/2$. This inequality follows because the $q_{ij} = q_{ij}^{\text{VSM}}/2$ and $p_{ij} \geq p_{ij}^{\text{VSM}}$ (since prices only increase with a lower selling probability). We now show that the probability that the sequential posted pricing is able to make the offer to representative ij is at least $1/4$. As a consequence the expected revenue from representative ij is $q_{ij}^{\text{VSM}} p_{ij}^{\text{VSM}}/8$; and summing over all representatives ij gives the lemma.

To show that the probability that it is feasible to offer service to representative ij is at least $1/4$, consider the worst-case ordering for this probability, i.e., the ordering where representative ij is last. Representative ij can be served if for all $j' \neq j$ representatives $j'i$ are not served, and for all $i' \neq i$ representatives $i'j$ are not served. The first event certainly happens if $v_{i'j} < p_{i'j}$ for all $i' \neq i$ and the second if $v_{ij'} < p_{ij'}$ for all $j' \neq j$. We now show that each of these events happens with probability at least $1/2$; since the events are independent the probability that both occur is at least $1/4$.

Consider the event that all $v_{i'j} < p_{i'j}$ for all $i' \neq i$ (the probability of the other event can be analized with the same approach). With respect to this event the possibility that $v_{i'j} \geq p_{i'j}$ is a bad event that happens with probability $q_{i'j}$. The probability of any of these bad events occuring can be bounded using the union bound by $\sum_{i'} q_{i'j}$. Of course, $\sum_{i'} q_{i'j}^{\mathrm{VSM}} \leq 1$ since the optimal mechanism allocates to one of these $i'j$ representatives with probability at most one (by the feasibility constraint) so $\sum_{i'} q_{i'j} \leq 1/2$. Therefore, probability that none of the bad events happen is at least $1/2$. $\square$

This instantiation of the reduction above covers matching markets with regular distributions. Similar instantiations can be applied to generalizations that include irregular distributions and environments with feasibility constraints induced by matroids. Sequential posted pricings do not give good approximations in general downward-closed environments.

7.3 Lottery Pricing and Randomized Mechanisms

Thus far in this chapter we have showed that there are pricing mechanisms that approximate the optimal deterministic mechanism in multi-dimensional unit-demand environments. These results are a little unsatisfying because we would really like a mechanism that approximates the optimal, potentially randomized, mechanism. Even in the simple single-agent environments described previously in this chapter, the optimal mechanism may not be a deterministic pricing of items. Instead, it might price randomized outcomes, a.k.a., lotteries.

This distinction raises a sharp contrast with (Bayesian) single-dimensional environments where there is always an optimal mechanism that is deterministic. For instance, with a lexicographical tie-breaking rule, the ironed virtual surplus maximization mechanism has a deterministic allocation rule.

Consider the single-agent unit-demand problem of designing a mechanism to maximize the revenue of the seller. Deterministic mechanisms are equivalent to item pricings whereas randomized mechanisms are equivalent to *lottery pricings*. A lottery is a probability distribution over outcomes. For instance, for the $m = 2$ item case, a lottery could assign either item 1 or item 2 with probability $1/2$ each. Lotteries do not have to be uniform, i.e., they can be biased in favor of some items, and they do not have to be complete, i.e., there may be some probability of assigning no item. A lottery pricing is then a set of lotteries and prices for each. For such a lottery pricing, the agent then chooses the lottery and price that give her highest utility for her given valuations for the items.

The following example shows that lottery pricings can give higher revenue than item pricings. There are two items (and one agent). The agent's value for each item is distributed independently and uniformly from the interval $[5, 6]$. The optimal item pricing for this environment to set a uniform price of 5.097 for each item. I.e., the agent is offered the option to buy item 1 at price 5.097 or to buy item 2 at price 5.097. The agent then buys the item that she values most as long as her value for that item is at most 5.097. Such an allocation rule is depicted in Figure 7.2(a) with $p = 5.097$. Now consider adding the additional option of buying at price 5.057 a lottery that realizes to item 1 or item 2 each

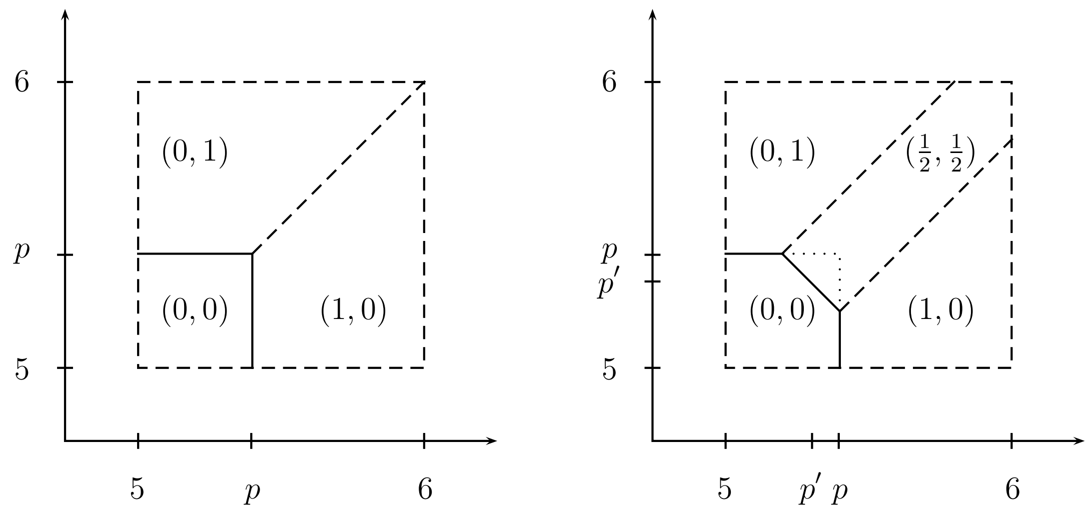

(a) Item pricing.

(b) Lottery pricing.

Figure 7.2: Depicted are the allocation regions for item pricing (p, p) and lottery pricing $\{((0,1), p), ((1,0), p), (\frac{1}{2}, \frac{1}{2}), p')\}$. The pricing and lotteries divide the valuation space into regions based on the preferred outcome of the agent. The diagonal line that gives the lower left boundary of the region where the lottery is preferred is the solution to the equation $v_1 + v_2 = p'$.

with probability 1/2. Now if the agent is nearly indifferent between the two items then she will buy the lottery and pay the lower price. Without the lottery option if the agent had average value bigger than 5.057 but no individual value over 5.097, the agent would buy nothing. Therefore, by adding this lottery option revenue is lost for some valuations of the agent and gained for others. One can calculate these losses and gains to conclude that the lottery pricing increases the expected revenue. Figure 7.2(b) with $p' = 5.057$ depicts the allocation rule that additionally offers the lottery option.

We would like to have a theory for approximating the optimal (possibly randomized) mechanism in multi-dimensional environments. Again, a crucial step in this endeavor is in identifying an analytically tractable upper bound on the optimal mechanism. Recall for the single-agent environment that the single-dimensional representative environment gave such an upper bound on optimal deterministic mechanisms (Theorem 7.2). The intuition for this bound was that the increased competition over the representative environment allowed the optimal mechanism for it to obtain more revenue than that of the original unit-demand environment. This intuition turns out to not be entirely correct when randomized mechanisms are allowed. In particular, there are examples where the optimal lottery pricing obtains more revenue than the optimal single-item auction for the representative environment.

To get some intuition for the failure of single-item auctions to provide an upper bound for lottery pricings consider the $m = 2$ item environment and the fair lottery which assigns item 1 or 2 each with probability 1/2. What is the value that the agent has for this lottery? It is the average of the agent's values. Averages of independent random variables concentrate around their expectations, therefore, the agent's value for this lottery has less randomness (in particular, a lower standard deviation) than her value for either of the individual items. A lottery pricing mechanism can take advantage of this sort of concentration. (In fact, later in this chapter we will apply this same intuition to environments with additive valuations, i.e., the value for a bundle of items is the sum of independently distributed values for each item in the bundle.)

We now show that the advantage that a lottery pricing has over a single-item auction in the representative environment is at most a factor of two.

Theorem 7.13. *For any product distribution $\mathbf{F} = F_1 \times \cdots \times F_m$, the expected revenue of the optimal single-agent, lottery pricing when the agent's values for the items are drawn from $\mathbf{F}$ is at most twice that of the optimal single-item, m-agent auction with the agents' values for the item drawn from $\mathbf{F}$.*

Proof. Our goal is to take a lottery pricing $\mathcal{L}$ and construct an auction $\mathcal{M}_{\mathcal{L}}$ for the representative environment such that the sum of its revenue with the revenue of the second-price auction is at least the revenue of the original lottery pricing. Both the constructed auction $\mathcal{M}_{\mathcal{L}}$ and the second-price auction revenues can be bounded from above by the optimal auction revenue; therefore, twice the optimal auction revenue is at least the revenue of the lottery pricing.

We construct $\mathcal{M}_{\mathcal{L}}$ as follows. Consider a valuation profile $\mathbf{v}$ for which the multi-dimensional agent selects lottery $l = ((q_1, \ldots, q_m), p)$. This agent receives utility $\sum_j v_j q_j - p$. As usual we denote by (j) the index of the jth highest value, i.e., $v_{(1)} \geq \ldots \geq v_{(m)}$. For such a valuation profile $\mathcal{M}_{\mathcal{L}}$ serves representative (1) with probability $q_{(1)}$ and charges this representative $p - \sum_{j \neq (1)} q_j v_j$ (always). This representative's utility is therefore perfectly aligned with our original multi-dimensional agent. Since the original agent preferred this lottery over all others, so does the representative; i.e., $\mathcal{M}_{\mathcal{L}}$ is incentive compatible.

By definition revenue of $\mathcal{M}_{\mathcal{L}}$ in the above environment is $p - \sum_{j \neq (1)} q_j v_j$. Consider the second part of this formula. This is the rebate we need to give representative (1) in order to incentivize the representative to prefer this lottery over all others. By definition the total probability to which the original agent is served by this lottery is at most one. Therefore, $\sum_{j \neq (1)} q_j v_j \leq v_{(2)}$, the second highest value. The revenue of $\mathcal{M}_{\mathcal{L}}$ is at least $p - v_{(2)}$. Recall that the revenue of the second-price auction is exactly $v_{(2)}$. Taking expectations over all valuation profiles, the expected revenue of $\mathcal{M}_{\mathcal{L}}$ is at least that of the original lottery pricing less that of the second price auction. Rearranging gives the desired inequality. □

A similar theorem can be proven in environments with multiple agents such as that of matching markets. The proof uses matroid properties and the basic intuition from the single-agent case, above. We omit the full proof from this text.

Theorem 7.14. *For independent, matching market environments, the optimal (randomized) mechanism's revenue is at most five times that of the optimal mechanism for the single-dimensional representative environment.*

7.4 Beyond Independent Unit-demand Environments

Up to this point the chapter has focused on independent unit-demand environments, i.e., ones where there is some set of services available, each agent desires at most one service, and the agent's value for each service are independent random variables. Unfortunately, not much is known about general distributions of preferences. In particular, environments where an agent desires more than one service or environments where an agent's value for distinct services are correlated. There are two notable exceptions, one from each of these classes.

The first exception is for "common base value" distributions, i.e., ones where the agents value for a service j is $v_0 + v_j$ and for $0 \leq j \leq m$, v_j are independent. The value v_0 is referred to as the base value because it offsets the values for each service. To this environment most of the preceding theorems can be extended, albeit with worse approximation factors. Unfortunately, the proofs of these extensions are brute-force and do not yield much additional understanding of the structure of good mechanisms in the common base value model.

The second exception is for additive preferences, i.e., where the agent's value for a bundle of services is the sum of the agent's value for each service. Again the agent's value for each individual service is independently distributed. In this environment, again for simple reasons, the optimal mechanism can be approximated. Sums of independent random variables tend to concentrate around their expectation. Therefore, it is possible to offer an agent a posted price for the grand bundle of items that is close to but below the this expectation and nearly the full surplus can be extracted.

Beyond these two cases, not much is known about approximately optimal mechanisms for general preferences. A major challenge in this research area is in identifying reasonable, analytically tractable upper bounds on the optimal multi-dimensional mechanism.

7.5 Optimal Lottery-pricing via Linear Programming

While there is little economic understanding of optimal mechanisms when agents' preferences are multi-dimensional that does not necessarily mean that the optimization problem is intractable. For the distributions on preference discussed heretofore, e.g., when the value that an agent has for various items is distributed independently, then an exponentially large type space can be described succinctly. In particular, each single-dimensional distribution need only be described. For such a distribution, a mechanism that was brute-force, i.e., its calculation explicitly considers every type in the type space, would be intractable.

On the other hand, if the type space is small enough that the distribution can be given explicitly, i.e., each type is given with its associated probability, then mechanisms that are brute-force in the type space may be reasonable. For the m-item, single-agent environment, for instance, the optimal lottery pricing in such a situation can be easily calculated.

Suppose the type space is given explicitly as follows. The agent has type $t \in \{1, \ldots, N\}$; denote by $\mathbf{v}$ the $N \times m$ matrix of values; let v_{tj} be the agent's value for item j when her type is t; and let π_t be the probability her type is t. We can write the optimization problem now as a linear program. The linear program will associate with each type t a lottery given by a price p_t and the probabilities for receiving each of the m items $(x_{t1}, \ldots, x_{tm})$. Notice that the agent's utility with type t for the lottery predesignated for type t' is $\sum_t v_{tj} x_{t'j} - p_{t'}$. The linear program will maximize expected payments (weighted by the distribution) subject to incentive constraints, individual rationality constraints, and probabilities summing to at most one (feasibility).

Maximize:

$$\sum_t \pi_t p_t \qquad \text{(expected revenue)}$$

Subject to:

$$\begin{aligned}
&\forall t, t' && \sum_t v_{tj} x_{tj} - p_t \geq \sum_t v_{tj} x_{t'j} - p_{t'} && \text{(incentive compatibility)}\\
&\forall t && \sum_t v_{tj} x_{tj} - p_t \geq 0 && \text{(individual rationality)}\\
&\forall t && \sum_j x_{tj} \leq 1 && \text{(feasibility)}
\end{aligned}$$

It is easy to see that when the type space and distribution are given explicitly that this program can be easily solved for the optimal set of lotteries to offer.

Lottery pricings correspond to fractional solutions of the linear program above; when the variables x_{tj} are integer these are simply (determinstic) item pricings. Calculating optimal item pricings in correlated environments, i.e., solving this mixed-integer program where x_{tj}s are constrained to be integral, is extremely challenging. For this problem, obtaining any approximation factor that is asymptotically better than linear in the number of items is computationally intractable under reasonable assumptions. Of course, a linear factor approximation is trivial. Formally:

Theorem 7.15. *Under complexity-theoretic assumptions, the problem of computing prices that $o(m)$-approximate the revenue of the optimal item pricing is computationally intractable.*

The computational intractability of a problem, and this perspective is discussed more in Chapter 8, suggests that there is inherent inability to make important structural observations.

Exercises

7.1 Consider the design of prior-free incentive-compatible mechanisms with revenue that approximates the (optimal) social-surplus benchmark, i.e., $\mathrm{OPT}(\mathbf{v})$, when all values are known to be in a bounded interval $[1, h]$. For general (multi-dimensional) combinatorial auctions, i.e., there are m items and each agent i has a value $v_i(S') \in [1, h]$ for each subset $S' \subseteq S = \{1, \ldots, m\}$ of the m items, give a prior-free $\Theta(\log h)$-approximation mechanism.

7.2 Prove Theorem 7.5: For any independent, unit-demand environment, the optimal deterministic mechanism's revenue is at most that of the optimal mechanism for the single-dimensional representative environment.

7.3 Recall that Theorem 7.10 shows that for the representative matching market environment, a sequential posted pricing gives an 8-approximation to the optimal (single-dimensional) mechanism. This bound can be improved.

(a) Give an improved bound.

(b) Assume that the agents are identically distributed (but not necessarily the items) and give an improved bound.

(c) Assume that both the agents and the items are identically distributed and given an improved bound.

7.4 Consider the design of prior-independent mechanisms for (multi-dimensional) unit-demand agents. Suppose there are n agents and $m = n$ houses and agent i's value for house j is drawn independently from a regular distribution F_j. (I.e., the agents are i.i.d., but the houses are distinct.) Give a prior-independent mechanism that approximates the Bayesian optimal mechanism. What is your mechanism's approximation factor?

Chapter Notes

There is a long history of study of multi-dimensional pricing and mechanism design in economics. Wilson's text *Nonlinear Pricing* is a good reference for this area (Wilson, 1997).

Algorithmic questions related to item-pricing for unit-demand agents were initiated by Aggarwal et al. (2004) and Guruswami et al. (2005) in an environment where the agent's values are correlated. The hardness of $o(m)$-approximation for such an m-item environment, i.e., Theorem 7.15, is due to Briest (2008). On the other hand Briest et al. (2010) show that optimal lottery pricings can be calculated via a linear program that is polynomially big in the support of the (correlated) distribution of the agent's valuations.

Approximation for item-pricings when the agent's values are independent were first studied by Chawla et al. (2007) where a 3-approximation was given. The 2-approximation via prophet inequalities that is presented in this chapter is due to Chawla et al. (2010a). Cai and Daskalakis (2011) show that it is computationally tractable to construct a pricing that approximates the revenue of the optimal pricing to within any multiplicative factor. The example presented herein that shows that a lottery pricing can give more revenue than the optimal item pricing was given by Thanassoulis (2004). Lottery pricings and the theorem that shows that the optimal lottery pricing is at most a factor of two more than the optimal mechanism's revenue in the single-dimensional representative environment is due to Chawla et al. (2010b).

The study of sequential posted pricing mechanisms in multi-dimensional environments that is discussed in this chapter is given by Chawla et al. (2010a); these sequential posted pricings are constant approximations to the optimal deterministic mechanisms. Alaei (2011) gives a refined analysis and approach. Extensions of these results to bound the revenue of the sequential posted pricing in terms of the optimal (randomized) mechanism's revenue are from Chawla et al. (2010b). Neither the bound of two (for single-agent lottery pricing) or five (for matching markets) is known to be tight.

Extensions from product distributions to the common base value environment are given in Chawla et al. (2010b). Briest et al. (2010) study general environments with correlated values and show that when more than $m = 4$ services are available then the ratio between the optimal lottery pricing (i.e., randomized mechanism) and the optimal item pricing (i.e., deterministic mechanism) is unbounded. This contrasts starkly with environment with independent values where Theorem 7.13 shows that the ratio is at most two. Finally, the independent additive values case, where pricing the grand bundle gives an asymptotically optimal revenue, was studied by Armstrong (1996).