

9

Approximation for Multi-dimensional and Non-linear Preferences

Approximation for multi-dimensional and non-linear agents is more urgent than for single-dimensional linear agents. For single-dimensional linear agents, Chapter 3 described optimal mechanisms. Chapter 4 then considered whether simple or practical mechanisms were approximately optimal. In contrast Chapter 8, which presented optimal mechanisms for agents with multi-dimensional or non-linear preferences, only solved families of environments where the optimization problem can be sufficiently simplified to be analytically tractable. In this chapter we show that these results can be generalized to a much broader family of environments with approximation.

Approximation can expose the implicit structure of multi-dimensional and non-linear mechanism design problems. As a first example, for the objective of revenue maximization, we will see that unit-demand agents who desire one of several alternatives are similar enough to single-dimensional representatives who compete with each other to purchase their corresponding alternative on behalf of the original agent (cf. the representative environment of Definition 8.6.3 on page 297). Single-dimensional agents are, of course, well understood by the developments of Chapter 3. As a second example, for revenue maximization in service constrained environments, we will see that even in environments with agents that are highly non-revenue-linear, marginal revenue is an important quantity that guides the design of (approximately) optimal mechanisms.

Approximation can also show that simple and practical mechanisms are pretty good relative to complex and impractical optimal mechanisms. For example, posted pricing mechanisms like those considered in Chapter 4 are approximately optimal in many multi-dimensional environments. Of course, the surplus maximization mechanism with reserve

prices always does as well as a posted pricing (that obtains its performance guarantee under any arrival order of the agents, i.e., obliviously).

Prior-independent approximation is also possible for agents with multi-dimensional and non-linear preferences (cf. Chapter 5). For example, for single-dimensional but non-linear agents the first-price auction is (prior-independent and) often a good approximation to the optimal mechanisms for both the objectives of welfare and revenue. For multi-dimensional matching environments, supply limiting mechanisms, e.g., where the surplus maximization mechanism is used with an additional constraint that only half the agents are served, are prior-independent and give good approximation to the optimal revenue. Both of these results can be viewed as generalizations of the Bulow-Klemperer Theorem (Theorem 5.2.1 on page 160). Generalizations of the (prior-free) random sampling auction (Definition 7.4.3 on page 240) to multi-dimensional and non-linear preferences are also known. These results, however, will not be further discussed in this text.

9.1 Single-agent Approximation

Our first order of business will be to address the complexity of the (unconstrained) single-agent problem for unit-demand preferences. The optimal mechanism for a single agent with unit-demand preferences is generally a lottery pricing that is given by menu of probability distributions over the alternatives with corresponding prices. This optimization problem can be described by a linear program (Section 8.8.7), but does not generally have closed-form solutions. Lottery pricings are complex and, in many circumstances, unnatural; while deterministic pricings of the alternatives are simple and prevalent in practice. For m -alternative environments such a deterministic pricing offers a menu with m outcomes

Chapter 9: Topics Covered.

- item and bundle pricing,
- multi-dimensional virtual values (revisited),
- the marginal revenue mechanism (revisited),
- the representative environment for unit-demand agents (revisited),
- multi-dimensional posted-pricing mechanisms, and
- (multi-dimensional) matching markets.

that correspond to the pricing of each alternative (as well as the trivial outcome of obtaining none of the alternatives at a price of zero). When the alternatives correspond to items that are for sale, such a pricing is an *item pricing*.

An important special case is when the agent's values for each alternative are drawn independently. In this case there is a natural item pricing that is a two approximation to the optimal item pricing. The importance of this result is that the natural item pricing is given by a closed-form solution and can be understood in terms of the single-dimensional theory of virtual values developed in Chapter 3. Moreover, the optimal item pricing approximates the revenue of optimal lottery pricing. Of course, these two results can be combined to show that a simple item pricing approximates the optimal lottery pricing.

An important model in between the cases where the agent's values for distinct alternatives are independently distributed or arbitrarily correlated is the case where there are k items, $m = 2^k - 1$ non-trivial bundles of items, the agent's value for each item is independently distributed, and the agent's value for a bundle is the sum of her values for the individual items in the bundle.¹ In this model either posting individual prices for each item or posting a single price for the grand bundle (i.e., all items together) is a constant approximation to the optimal revenue.

Some form of independence of the agent's values for the alternatives is required for the existence of simple approximation mechanisms. When the agent's values are drawn from general correlated distributions, the above results generally cease to hold. There are not simple, intuitive item pricings that give good approximations to the optimal item pricing. Moreover and even for $m = 2$ alternatives, neither does the optimal item pricing approximate the optimal lottery pricing nor, more generally, does any finite menu of lotteries.

In the subsections below these results are formally developed.

9.1.1 Item Pricing

Consider mechanisms that post a price for each alternative and permit the agent to purchase her favorite alternative at its price. When the alternatives correspond to items, this form of mechanism is called an item pricing. Any deterministic mechanism for a single unit-demand agent is a deterministic pricing of alternatives.

¹ The term *bundle* simply refers to any subset of the items.

Identifying good item pricings is difficult. Unlike the optimal lottery pricing which is given by the solution to a linear program, the optimal item pricing problem is not convex. Specifically, convex combinations of item pricings are not item pricings. One consequence of this non-convexity is that the optimal item pricing in a symmetric environment may not be symmetric (see Exercise 9.1). Another consequence of this non-convexity is that there is no simple description of optimal item pricings. Therefore, the first issue we encounter in an attempt to find an approximation for this single-agent problem is that we really do not understand the optimal mechanism very well. We will take the usual approach to such a situation and identify an upper bound that is analytically tractable.

Consider the following thought experiment. We have a seller facing a single unit-demand agent who wishes to buy one of m items. Consider the *representative environment* where this unit-demand agent is replaced with m single-dimensional representatives with the unit-demand preference reinterpreted as a feasibility constraint that at most one representative can be served (recall Definition 8.6.3). Would a seller prefer to be in the original environment or in the representative environment; i.e., in which environment can the seller obtain a higher revenue? Intuition suggests that the seller should prefer the representative environment as competition between representatives should produce a higher revenue. The following theorem shows that this intuition is correct for deterministic mechanisms.

Theorem 9.1.1. *For a unit-demand agent with independent values, the revenue of the optimal auction for the (single-dimensional) representative environment is at least the revenue for the optimal item pricing in the original (unit-demand) environment.*

Proof. This proof follows a common approach for this chapter. The approach is to construct from any item pricing for the original environment an auction for the representative environment and to argue that the auction obtains at least the revenue of the item pricing.

Consider the item pricing $\hat{t} = (\hat{t}_1, \dots, \hat{t}_m)$ for the original environment. Notice that the allocation rule of the item pricing is to serve the agent the alternative j^* that maximizes $t_{j^*} - \hat{t}_{j^*}$ if non-negative and, otherwise, nothing. Fixing all coordinates of the type except for j , the allocation of alternative j is monotone in \hat{t}_j . Therefore, Theorem 2.5.1 implies that the same allocation rule, for the representative environment, corresponds to a dominant strategy mechanism. Importantly, in this mechanism, the

representative who wins is the one who corresponds to the alternative that the original unit-demand agent selects. The revenue in the original environment is simply \hat{t}_{j^*} ; the revenue in the representative environment is the minimum bid that representative j^* must make to win. This minimum bid is at least \hat{t}_{j^*} . Therefore, the auction revenue exceeds the item-pricing revenue. \square

While revenue of the optimal lottery pricing may exceed the revenue of the optimal item pricing, it is nonetheless interesting to consider whether it is easy to approximate the optimal item pricing (e.g., via an item pricing that is easy to find and understand). Using the optimal representative revenue as an upper bound on the optimal item pricing revenue, we can easily attain such a bound. In fact, the revenue of an item pricing is lower bounded by the revenue we would obtain if the agent buys the cheapest priced alternative instead of her utility maximizing alternative, i.e., when her values for several alternatives are above their corresponding prices, if instead of choosing to maximize $t_j - \hat{t}_j$ she chose to minimize \hat{t}_j . For single-dimensional agents, Section 4.2.1 discussed the factor by which oblivious posted pricings (i.e., tie-breaking by minimum price) approximate optimal auctions. Theorem 4.2.3, specifically, employed the prophet inequality (Theorem 4.2.1) to show that uniform virtual pricing is a two approximation to the optimal auction. The following is a corollary of this theorem and the above upper bound (Theorem 9.1.1). Recall that a uniform virtual pricing $\hat{\phi}$ gives prices $(\hat{t}_1, \dots, \hat{t}_m)$ satisfying $\phi_j(\hat{t}_j) = \hat{\phi}$ for all j where ϕ_j is the single-dimensional virtual value function for distribution F_j of alternative j .

Corollary 9.1.2. *For a unit-demand agent with independent values, the revenue of a uniform virtual item pricing is a two approximation to the optimal representative revenue (and the the optimal item-pricing revenue).*

Generally – for asymmetric type distributions – the uniform virtual item pricing of Corollary 9.1.2 is asymmetric. From the discussion of Section 4.4, however, we saw that that anonymous pricing is often a good approximation to the optimal auction in the representative environment. Specifically, the analyses of anonymous pricings given in Corollary 4.4.2, Corollary 4.4.3, and Theorem 4.4.5 can also be combined with the unit-demand upper bound (Theorem 9.1.1) to give bounds on the revenue of uniform item pricing, i.e., posting the same price for each item, to unit-demand agents. These bounds approximately extend the optimality

of uniform item pricing for a unit-demand agents with type from the uniform distribution on $[0, 1]^m$ and, more generally, from item-symmetric ratio-monotone distributions (Section 8.8; Definition 8.8.2).

Corollary 9.1.3. *For a unit-demand agent with independent values, a uniform item pricing is often a constant approximation to the optimal item-pricing revenue (and the optimal representative revenue). Specifically, its approximation is at most*

- (i) $e/e-1 \approx 1.582$ for regular and identical distributions,
- (ii) two for identical (but irregular) distributions, and
- (iii) $e \approx 2.718$ for regular (but non-identical) distributions.

9.1.2 Revenue Upperbounds via Amortization

Consider the single-agent unit-demand problem of designing a mechanism to maximize the revenue of the seller. Deterministic mechanisms are equivalent to the item pricings that were discussed in the previous section whereas randomized mechanisms are equivalent to *lottery pricings*. A lottery is a probability distribution over outcomes. For instance, for the $m = 2$ alternative case, a lottery could assign either alternative 1 or alternative 2 with probability $1/2$ each. Lotteries do not have to be uniform, i.e., they can be biased in favor of some alternatives, and they do not have to be complete, i.e., there may be some probability of assigning no alternative. A lottery pricing is then a set of lotteries and prices for each. For such a lottery pricing, the agent chooses the lottery and price that give her highest utility given her valuation of the alternatives (see Section 8.8.7 on page 331).

The following example shows that lottery pricings can give higher revenue than item pricings. There are two alternatives (and one agent). The agent's value for each alternative is distributed independently and uniformly from the interval $[5, 6]$. The optimal item pricing for this environment to set a uniform price of about 5.097 for each alternative. I.e., the agent is offered the option to buy alternative 1 at price 5.097 or to buy alternative 2 at price 5.097. For this item pricing, the agent buys the alternative that she values most as long as her value for that alternative is at most 5.097. Such a two-dimensional allocation rule is depicted in Figure 9.1(a). Now consider adding the additional option of buying at a lower price of 5.057 a lottery that realizes to alternative 1 or alternative 2 each with probability $1/2$. Notice that if the agent was previously buying one of the items but is nearly indifferent between the




Figure 9.1. Depicted are the allocation regions for item pricing (\hat{v}, \hat{v}) and lottery pricing $\{((0, 1), \hat{v}), ((1, 0), \hat{v}), (1/2, 1/2), \hat{v}^\dagger)\}$. 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 $t_1 + t_2 = \hat{v}^\dagger$.

two alternatives then she will prefer the lottery at the lower price. On the other hand, 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 types of the agent and gained for others. The losses and gains can be compared to conclude that adding the lottery increases the expected revenue of the pricing. Figure 9.1(b) depicts the allocation rule that additionally offers the lottery option. (A similar example is described in detail in Section 8.8.3 beginning on page 311.)

This non-optimality of deterministic mechanisms in multi-dimensional environments contrasts with single-dimensional environments where there is always an optimal mechanism that is deterministic. For instance, with a lexicographical tie-breaking rule, the virtual surplus maximization mechanism of Chapter 3 chooses a winner deterministically. Nonetheless, the gap between the item pricing revenue and the lottery pricing revenue in the example above is small. In the next few sections we will prove that the gap between simple deterministic mechanisms remains

small when the agent's values for the different alternatives are independently distributed.

Proceeding to develop a theory for approximating the optimal (possibly randomized) mechanism in multi-dimensional environments, a crucial step is in identifying an analytically tractable upper bound on the revenue of the optimal mechanism. We just saw that the single-dimensional representative environment gave such an upper bound on optimal deterministic mechanisms (Theorem 9.1.1). The intuition for this bound was that the increased competition of 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.

The approach of multi-dimensional virtual values described in Section 8.8 beginning on page 305 facilitates the identification of an analytically tractable upper bound. Per Definition 8.8.4, a virtual value function must satisfy three properties.

- (i) Amortization of revenue: the expected virtual surplus of any mechanism must exceed its expected revenue;
- (ii) incentive compatibility: a point-wise virtual surplus maximizer is incentive compatible; and
- (iii) tightness: the expected virtual surplus of the virtual surplus maximizer is equal to its expected revenue.

While identifying a virtual value function, which by definition satisfies all three of these properties, is difficult; identifying an amortization of revenue, i.e., a function that satisfies only the first property, is easy. Specifically, per Definition 8.8.6 there is a canonical amortization of revenue corresponding to any decomposition of the type space into paths. Such an amortization is sufficient for identifying an upper bound on the optimal revenue.

For product distributions the following definition of multi-dimensional amortizations of revenue that are constructed from the single-dimensional ironed and non-ironed virtual values will enable the bounding of the optimal multi-dimensional mechanism's revenue, like Theorem 9.1.1, in terms of the optimal mechanism for the single-dimensional representative environment.

The subsequent analysis of this section will apply generally to agents

with linear utility for subsets of k items. Linear utility for a type t , allocation x , and payments p is expressed as $t \cdot x - p$ where $t \cdot x$ is the dot product $\sum_j \{t\}_j \{x\}_j$. In Section 9.1.3 this analysis will be applied to unit-demand agents (which have been under sole consideration up to this point; unit-demand agents are additive with the additional constraint that at most one item is allocated). In Section 9.1.4 this analysis will be applied to additive agents (where a bundle of items can be allocated to the agent). In Section 9.1.5 we show that the assumption that the distribution is independent across the items is crucial. Generally with correlated items no simple (e.g., deterministic) mechanism approximates the optimal mechanism.

Definition 9.1.1. For product distribution $F = F_1 \times \dots \times F_k$, the multi-dimensional extensions of the single-dimensional non-ironed and ironed virtual value functions are the vector fields ϕ^{MD} and $\bar{\phi}^{MD}$, respectively, defined as follows:

- (i) $\{\phi^{MD}(t)\}_{j^*} = \phi_{j^*}^{SD}(\{t\}_{j^*})$, $\{\bar{\phi}^{MD}(t)\}_{j^*} = \bar{\phi}_{j^*}^{SD}(\{t\}_{j^*})$, and
- (ii) $\{\phi^{MD}(t)\}_j = \{\bar{\phi}^{MD}(t)\}_j = \{t\}_j$,

where item $j^* \in \operatorname{argmax}_j \{t\}_j$ is the favorite item, item $j \neq j^*$ ranges over all other items, and where $\phi_j^{SD}(v)$ and $\bar{\phi}_j^{SD}(v)$ are the single-dimensional non-ironed and ironed virtual value functions, specifically, $\phi_j^{SD}(v) = v - 1 - F_j(v)/f_j(v)$ and $\bar{\phi}_j^{SD}$ is derived from ϕ_j^{SD} by ironing (see Section 3.3.5).

Theorem 9.1.4. For an agent with linear utility and values drawn from a product distribution, the multi-dimensional extensions of the single-dimensional non-ironed and ironed virtual value functions (ϕ^{MD} and $\bar{\phi}^{MD}$ in Definition 9.3.2) are amortizations of revenue.

The high-level approach to proving this theorem is given below with key lemmas stated and proved later in this section.

Proof. Consider the special case of the theorem where the value for item 1 is drawn according to distribution F_1 and all other values are constant. On this special case the type space is a path (more precisely, a line).

- (i) On the path special case, Lemma 9.1.5 (which follows directly from the analysis in Section 8.8.3, specifically Theorem 8.8.3) implies that the vector field defined by the single-dimensional marginal price-posting revenue on each coordinate is an amortization of revenue. This vector field is identical to ϕ^{MD} (for the path special case).

- (ii) On the path special case, Lemma 9.1.7 implies that the vector field defined by the single-dimensional marginal revenue on each coordinate is an amortization of revenue. This vector field is identical to $\bar{\phi}^{\text{MD}}$ (for the path special case).
- (iii) Lemma 9.1.8 implies that for any threshold \hat{v} and single-dimensional distribution (e.g., F_1), the single dimensional non-ironed virtual value functions, for both the distribution and the distribution conditioned on values at least \hat{v} , are identical on values at least \hat{v} .
- (iv) Lemma 9.1.9 extends Lemma 9.1.8 to ironed virtual value functions where the ironed virtual value of the conditional distribution is at most that of the unconditioned distribution.
- (v) Lemma 9.1.10 shows that a vector field for a type space and distribution is an amortization of revenue if there exists a partitioning of type space such that it is an amortization of revenue on each part.

These points combine as follows. Partition type space into the paths (actually lines) where the value for the favorite item is constant. For these paths the distribution of the value for the favorite item is the conditioned on the value being at least the value of any other item. By (i) and conditioned on each such path, the single-dimensional non-ironed virtual value for each coordinate is an amortization of revenue. By (iii) the conditional single-dimensional non-ironed virtual value for the favorite item is the same as the unconditional non-ironed virtual value for the favorite item. Thus, the vector field ϕ^{MD} defined for the full space is an amortization of revenue on each path. Finally, by (v), ϕ^{MD} is an amortization of revenue for the full type space.

Similarly, (ii), (iv), and (v) imply that $\bar{\phi}^{\text{MD}}$ is an amortization of revenue for the full type space. \square

The following notation will facilitate stating and proving the lemmas required in the proof of Theorem 9.1.4. For our original problem type t is drawn from the product distribution $F = F_1 \times \cdots \times F_m$. A path problem is specified in quantile space by $\tau : [0, 1] \rightarrow \mathcal{T}$ (see Definition 8.8.1). The path induces distributions for each item F_1, \dots, F_m . For example $\{\tau(q)\}_j$ with quantile $q \sim U[0, 1]$ is distributed according to F_j , i.e., $F_j(\{\tau(q)\}_j) = 1 - q$. Single-dimensional price-posting revenue curves for the induced distribution for each item are, e.g., $P_j(q) = q \{\tau(q)\}_j$, i.e., posted price $\{\tau(q)\}_j$ for item j bought with probability q . The single-dimensional non-ironed virtual value for item j is $\phi_j^{\text{SD}}(v) = v - 1 - F_j(v)/f_j(v)$ and, for q satisfying $v = \{\tau(q)\}_j$, it equals the marginal price-posting revenue $P_j'(q) = \{\tau(q) + q\tau'(q)\}_j$. (Recall that $\tau(q)$ is

non-increasing in quantile and, thus, $\tau'(q)$ is vector with non-positive coordinates.) The revenue curve $R_j(q)$ is the smallest concave function that upper bounds the price-posting revenue curve $P_j(q)$. The ironed virtual value is $\bar{\phi}_j^{\text{SD}}(v)$ equal to $R'_j(q)$ for $v = \{\tau(q)\}_j$. With this notation and discussion, we now prove the lemmas that comprise each step of the proof of Theorem 9.1.4.

The correctness of ϕ^{MD} , the multi-dimensional extension of the single-dimensional non-ironed virtual value functions, as an amortization of revenue for the special case of paths is a fairly directly consequence of Theorem 8.8.3 which shows that for paths that are monotonically non-increasing in each coordinate of type space, the vector field of the marginal price-posting revenues is an amortization of revenue.

Lemma 9.1.5. *For an agent with value for item 1 drawn from distribution F_1 and values for other items specified deterministically (e.g., by distributions that are entirely a pointmass), the multi-dimensional extension of the single-dimensional non-ironed virtual value functions (ϕ^{MD} in Definition 9.3.2) is an amortization of revenue.*

Proof. For paths, Theorem 8.8.3 characterizes the expected revenue of any mechanism (x, p) as:

$$\mathbf{E}[p(\tau(q))] = \mathbf{E}[x(\tau(q)) \cdot P'(q)] - u(\tau(1))$$

where $P'(q) = (P'_1(q), \dots, P'_m(q))$ is the vector of marginal price-posting revenue curves corresponding to quantile q . Since utility is non-negative for individually rational mechanisms, we obtain the bound:

$$\mathbf{E}[p(\tau(q))] \leq \mathbf{E}[x(\tau(q)) \cdot P'(q)]$$

Observe for pointmass distributions $j \neq 1$, that the marginal price-posting revenue is identically equal to the value of pointmass (formulaically we have $P'_j(q) = \{\tau(q) + q\tau'(q)\}_j$ where $\{\tau'(q)\}_j = 0$); thus, $P'_j(q) = \phi_j^{\text{SD}}(\{\tau(q)\}_j) = \{\tau(q)\}_j$. For alternative 1, by the equivalence of marginal price-posting revenues and non-ironed virtual values, we have $P'_1(q) = \phi_1^{\text{SD}}(\{\tau(q)\}_1)$. Thus, the vector field ϕ^{MD} is an amortization of revenue for the path special case of the lemma. \square

The one way to show the correctness of $\bar{\phi}^{\text{MD}}$, the multi-dimensional extension of the single-dimensional ironed virtual value functions, as an amortization of revenue for the special case of paths, is to show that under any incentive compatible mechanism the expected amortized surplus according to $\bar{\phi}^{\text{MD}}$ is at least that of ϕ^{MD} . To do so we need to

better understand the constraint imposed by incentive compatibility on a multi-dimensional allocation rule.

Lemma 9.1.6. *Each coordinate j of the allocation rule x of an incentive compatible mechanism is monotonically non-decreasing in that coordinate, i.e., $\partial\{x(t)\}_j/\partial\{t\}_j \geq 0$ for all j .*

Proof. By Theorem 2.5.1 the utility u is convex and its gradient is the allocation rule. Considering coordinate j and fixing other coordinates, the utility $u(t)$ is convex in $\{t\}_j$ and, thus, the allocation probability $\{x(t)\}_j = \partial u(t)/\partial\{t\}_j$ of item j is monotonically non-decreasing in $\{t\}_j$. \square

Lemma 9.1.7. *For an agent with type for item 1 drawn from distribution F_1 and other items specified deterministically (e.g., by distributions that are entirely a pointmass), the multi-dimensional extension of the single-dimensional ironed virtual value functions ($\bar{\phi}^{\text{MD}}$ in Definition 9.3.2) is an amortization of revenue.*

Proof. Denote by $y(q)$ the multidimensional allocation rule in quantile space, i.e., $y(q) = x(\tau(q))$ and write $y'(q) = \frac{dy(q)}{dq}$ as the vector of derivatives of the allocation probabilities of each alternative with respect to quantile. Write the difference in amortized surplus as:

$$\begin{aligned} & \mathbf{E}_q \left[[\bar{\phi}^{\text{MD}}(\tau(q)) - \phi^{\text{MD}}(\tau(q))] \cdot x(\tau(q)) \right] \\ &= \mathbf{E}_q [R'(q) - P'(q)] \cdot y(q) \\ &= \left[(R(q) - P(q)) \cdot y(q) \right]_0^1 - \mathbf{E}_q [(R(q) - P(q)) \cdot y'(q)] \\ &\geq 0. \end{aligned}$$

The second line transforms to quantile space. The third line follows from integration by parts. The fourth line follows from the equality of R and P at the endpoints $q = 0$ and $q = 1$, the inequality $R(q) \geq P(q)$ on the interior $q \in (0, 1)$, and the non-positivity of y' (from Lemma 9.1.6). \square

We now consider any single-dimensional distribution and compare the mapping from types to virtual values for the distribution and the distribution conditioned on values exceeding a given threshold. Specifically, we show that the non-ironed virtual values of the two distributions are the same on values above the threshold, and the ironed virtual values for the unconditioned distribution are at least the ironed virtual values of the conditioned distribution. Both of these observations are relatively straight forward.

Lemma 9.1.8. *For any single-dimensional distribution F and any threshold \hat{v} , the non-ironed virtual values of the conditional distribution (on values at least the threshold) equal those of the original distribution.*

Proof. Denote the single-dimensional mapping from quantiles to values by $\nu(q) = F^{-1}(1 - q)$. Non-ironed virtual values for the original distribution are $\phi^{\text{SD}}(\nu(q)) = P'(q) = \nu(q) + q\nu'(q)$. Let \hat{q} satisfy $\nu(\hat{q}) = \hat{v}$ where conditioning on values exceeding \hat{v} corresponds to conditioning on quantiles lying below \hat{q} . Denote the conditioned distribution by F^\dagger . Thus, the marginal price-posting revenue for the conditional distribution is $\frac{d}{dq}P^\dagger(q) = \nu^\dagger(q) + q\frac{d\nu^\dagger(q)}{dq}$. The mapping from values to quantiles for the conditional distribution scales as $\nu^\dagger(q) = \nu(q\hat{q})$. Notice that $\frac{d\nu^\dagger(q)}{dq} = \frac{d\nu(q\hat{q})}{dq} = q\hat{q}\nu'(q\hat{q})$. Thus, $\frac{d}{dq}P^\dagger(q) = \nu(q\hat{q}) + q\hat{q}\nu'(q\hat{q}) = P'(q\hat{q})$. \square

Lemma 9.1.9. *For any single-dimensional distribution F and any threshold \hat{v} , the ironed virtual values of the conditional distribution (on values at least the threshold) are at most those of the original distribution.*

Proof. Let \hat{q} be the quantile corresponding to \hat{v} . Compare the price-posting revenue curve P on the full quantile interval $[0, 1]$ with the price-posting revenue curve on the truncated interval $[0, \hat{q}]$. The smallest concave function that upper bounds P on $[0, 1]$ (namely, R) is at least the smallest concave function that upper bounds P on $[0, \hat{q}]$. Specifically, if there is an ironed interval that contains \hat{q} , the ironed curve for the truncated interval will be lower than the ironed curve for the full interval. This lower curve will have a smaller derivative (corresponding to a smaller ironed virtual value). On the other hand, for ironed intervals that do not contain \hat{q} , the concave upper bound on the truncated interval and the full interval are the same. Thus, the marginal revenue at any quantile on the truncated interval is no more than the marginal revenue on the full interval.

By the arguments of Lemma 9.1.8 this relation between the marginal revenues for the truncated interval $[0, \hat{q}]$ and the full interval $[0, 1]$ correspond to the same relation for the ironed virtual value functions. \square

The final ingredient of the proof of Theorem 9.1.4 is to show that a vector field is an amortization of revenue on the full type space if it is an amortization of revenue on each part of a partition of type space. This statement is formalized and proved as follows.

Lemma 9.1.10. *Given type space \mathcal{T} , distribution over types F , partitioning of types into subspaces $\mathcal{T} = S^{(1)} \cup S^{(2)} \cup \dots$, and conditional distributions $F^{(\ell)}$ of types in subspace $S^{(\ell)}$ for each ℓ ; a vector field ϕ is an amortization of revenue for \mathcal{T} and F if it is an amortization of revenue for $S^{(\ell)}$ and $F^{(\ell)}$ for each ℓ .*

Proof. Fix any incentive compatible mechanism (x, p) for the whole type space \mathcal{T} . This mechanism is incentive compatible for each of its subspaces as each type $t \in \mathcal{T}$ not preferring the outcome of each type $s \in \mathcal{T}$ to her own implies that $t \in \mathcal{T}^{(\ell)}$ does not prefer the outcome of $s \in \mathcal{T}^{(\ell)}$ to her own. The revenue of this mechanism is the convex combination of the revenue of the mechanism on each subspace. The amortization of $\phi^{(\ell)}$ on type space $\mathcal{T}^{(\ell)}$ with distribution $F^{(\ell)}$ implies that ϕ is an amortization of revenue for type space \mathcal{T} and distribution F . Specifically:

$$\begin{aligned} \mathbf{E}[p(t)] &= \sum_{\ell} \mathbf{E}[p(t)|t \in S^{(\ell)}] \mathbf{Pr}[t \in S^{(\ell)}] \\ &\leq \sum_{\ell} \mathbf{E}[\phi(t) \cdot x(t)|t \in S^{(\ell)}] \mathbf{Pr}[t \in S^{(\ell)}] \\ &= \mathbf{E}[\phi(t) \cdot x(t)]. \quad \square \end{aligned}$$

9.1.3 Item Pricing versus Lottery Pricing

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 and that the advantage lottery pricing has over item pricing is at most a factor of four. These two results are corollaries of the following theorem.

Theorem 9.1.11. *For a unit-demand agent with independent values, the sum of revenues of the second price auction with and without lazy monopoly reserves in the single-dimensional representative environment upper-bounds the revenue for the optimal lottery pricing for the original unit-demand environment.*

Proof. Denote by REF the optimal amortized surplus of the multi-dimensional extension of the single-dimensional ironed virtual value functions ($\bar{\phi}^{\text{MD}}$ from Definition 9.3.2) with a unit-demand constraint. Let FAVE denote the optimal amortized surplus according to $\bar{\phi}$ from only selling the favorite item. I.e., FAVE is the optimal amortized surplus according to vector field $\bar{\phi}^{\text{FAVE}}$ that equals $\bar{\phi}^{\text{MD}}$ on the coordinate of the favorite item and is zero on other the coordinates of non-favorite items. Notice that FAVE is identically the revenue from the second-price auction with lazy monopoly reserves: the virtual surplus of the second price

auction with lazy monopoly reserves is the non-negative virtual value of the representative with the highest value. Let NONFAVE denote the optimal amortized surplus according to $\bar{\phi}^{\text{MD}}$ without selling the favorite item (but possibly selling any of the other items). I.e., NONFAVE is the optimal amortized surplus according to $\bar{\phi}^{\text{NONFAVE}}$ that equals $\bar{\phi}^{\text{MD}}$ on the coordinate of the non-favorite items and is zero on the coordinates of the favorite item. Notice that NONFAVE is the revenue of the second price auction in the representative environment. Of course the optimal amortized surplus REF is at most the optimal amortized surplus from selling the favorite item FAVE plus the optimal amortized surplus from selling the non-favorite items NONFAVE. Thus, the theorem holds. \square

Corollary 9.1.12. *For a unit-demand agent with independent values, the revenue of the optimal auction for the single-dimensional representative environment is a two approximation to the revenue for the optimal lottery pricing for the original unit-demand environment.*

Proof. The proof of this corollary follows from Theorem 9.1.11 and the fact the the optimal revenue for the representative environment upper bounds the second-price revenue with and without lazy reserves. \square

Corollary 9.1.13. *For a unit-demand agent with independent values, a uniform virtual pricing is a four approximation to the optimal lottery pricing.*

Proof. The proof of this corollary follows from Corollary 9.1.12 and Corollary 9.1.2, the latter of which states that a uniform virtual pricing is a two approximation to the optimal mechanism for the representative environment. \square

As with the previous comparison to optimal item pricing, uniform item pricings also give constant approximations to the optimal lottery pricing under various distributional assumptions.

Corollary 9.1.14. *For a unit-demand agent with independent values, a uniform item pricing is often a constant approximation to the optimal lottery revenue. Specifically, its approximation is at most*

- (i) $2e/e-1 \approx 3.164$ for regular and identical distributions,
- (ii) four for identical (but irregular) distributions, and
- (iii) $2e \approx 5.437$ for regular (but non-identical) distributions.

The following theorem improves on the first part of Corollary 9.1.14 and approximately extends the optimality of uniform pricing for ratio

monotone distributions (Theorem 8.8.7 on page 318) to environments where the agent's values for each alternative is independently, identically, and regularly distributed. The main idea in the proof of this theorem is that the amortization of revenue given by the multi-dimensional extension of the favorite alternative projection (Definition 8.8.8) can be used to show that the sum revenues of the second-price auction and the optimal uniform pricing upper bound the optimal lottery revenue.

Theorem 9.1.15. *For a unit-demand agent with independently, identically, and regularly distributed values, a uniform item pricing is a $2e-1/e-1 \approx 2.58$ approximation to the optimal lottery pricing revenue. For $k = 2$ alternatives, the bound improves to $7/3 \approx 2.33$.*

Proof. See Exercise 9.2. □

9.1.4 Additive Values

Consider selling $m = 2^k - 1$ alternatives that correspond to non-trivial bundles of k items to an agent who has additive values for the items. In other words, the agent's value for a bundle of items is the sum of the agent's values for the individual items in the bundle. Thus, agent's type is given by a k -tuple which expresses the value that she assigns to each item. Like in the previous sections, the agent's values for each item will be independently but not identically distributed.

Item pricing is a natural and prevalent selling mechanism for such an environment. An item pricing specifies a price for each item and the agent then chooses to buy all the items for which her value exceeds the price. This mechanism is used predominantly in retail, both online and and in brick and mortar stores.

Bundling all of the items together is also a common selling mechanism. As an example, online movie and music streaming services often charge a flat monthly rate that allows any movie or song in the online catalog to be viewed or listened to. Intuitively, bundle pricing can be better than item pricing for additive agents. The agent's value for a bundle is the sum of her value for each item in the bundle and, given that the agent's values for the items are independent, her value for the bundle is more concentrated around its expectation than the individual items are. Bundle pricing is often able to take advantage of this concentration to extract most of the surplus.

In this section we show that one of these two common approaches is always pretty good.

Theorem 9.1.16. *For an additive agent with values for each item drawn independently, either selling the items separately at the optimal price for each item or bundling the items together and optimally setting the bundle price is a six approximation to the revenue of the optimal mechanism.*

As usual, the proof of this theorem comes from comparing an upper bound on the revenue of the optimal mechanism to the revenue of an approximation mechanism. One such upper bound is given by the amortized surplus of the multi-dimensional extension of the single-dimensional ironed virtual value functions (Definition 9.3.2; Theorem 9.1.4). Recall that this amortization of revenue is defined by vector field $\bar{\phi}^{\text{MD}}$ with $\{\bar{\phi}^{\text{MD}}(t)\}_{j^*} = \bar{\phi}_j^{\text{SD}}(\{t\}_{j^*})$ and $\{\bar{\phi}^{\text{MD}}(t)\}_j = \{t\}_j$ for $j^* = \text{argmax}_j \{t\}_j$ and $j \neq j^*$.

It will be helpful to separate the upper bound of the optimal revenue from the identified amortization into three parts as each part will be analyzed using different methods. One part is the contribution to the amortization from the favorite item. The other two parts partition non-favorite-item values into high values and low values by decomposition threshold \hat{v} . The threshold \hat{v} is chosen to be low enough that the bound it imposes on values below it implies that the sum of those values concentrates, and high enough that the contribution to the amortization from non-favorite-item values above it is small.

Lemma 9.1.17. *For an agent with additive values drawn from a product distribution and decomposition threshold \hat{v} , the revenue of the optimal mechanism REF is upper bounded by the sum of three terms:*

- CORE: the expected surplus from values that are at most \hat{v} ,
- TAIL: the expected surplus from non-favorite-items with values at least \hat{v} , and
- FAVE: the expected amortized surplus from non-negative favorite-item amortized values.

Proof. This lemma follows from mapping the types to the amortized values given by the multi-dimensional extension of the single-dimensional ironed virtual value functions (Definition 9.3.2) and considering the outcome that optimizes amortized surplus. This outcome allocates to the agent each item j for which the j th coordinate of the amortized value is non-negative. CORE and TAIL decompose the amortized surplus for the non-favorite items while FAVE accounts for the favorite items. (There may be incidental double counting of amortized value from favorite

items in FAVE with the amortized values from values less than \hat{v} in CORE.) \square

Our approximation mechanism, APX, has expected revenue that is the maximum of the expected revenue from optimally pricing individual items, ITEM, and the expected revenue from optimally pricing the grand bundle, BNDL; i.e., $\text{APX} = \max(\text{ITEM}, \text{BNDL})$. We will now bound the individual terms of the bound on $\text{REF} \leq \text{CORE} + \text{TAIL} + \text{FAVE}$ by the individual components of APX, namely ITEM and BNDL, for decomposition threshold chosen to be $\hat{v} = \text{ITEM}$.

Lemma 9.1.18. *The revenue from selling only the favorite item is no more than the revenue from item pricing: $\text{FAVE} \leq \text{ITEM}$.*

Proof. FAVE sells the agent her favorite item at that item's monopoly price. ITEM sells the agent each item at the item's monopoly price. The latter has higher amortized surplus and thus exceeds the former. \square

Lemma 9.1.19. *For decomposition threshold $\hat{v} = \text{ITEM}$, the surplus from high-valued non-favorite items is at most the revenue from item pricing: $\text{TAIL} \leq \text{ITEM}$.*

Proof. This proof is based on identifying item pricings within the calculation of TAIL and upper bounding their revenue by the revenue of the optimal item pricing ITEM. Consider the contribution of item j to TAIL. Item j with value $v \geq \hat{v}$ contributes to TAIL when it is not the favorite item. Thus, its contribution to TAIL is

$$\begin{aligned} \text{TAIL}_j(v) &= v \cdot \Pr_{t \sim F} \left[\exists j^\dagger \neq j, \{t\}_{j^\dagger} > v \right] \\ &\leq v \cdot \mathbf{E}_{t \sim F} \left[\left| \left\{ j^\dagger : \{t\}_{j^\dagger} > v \right\} \right| \right] \\ &= \text{“the revenue from uniformly pricing with price } v\text{”} \\ &\leq \text{ITEM}. \end{aligned}$$

The total contribution from item j for all values $v \geq \hat{v} = \text{ITEM}$ is thus

$$\begin{aligned} \text{TAIL}_j &\leq \text{ITEM} \cdot \Pr_{v \sim F_j} [v \geq \text{ITEM}] \\ &= \text{“the revenue from posting price ITEM for item } j\text{”} \end{aligned}$$

Summing over all items

$$\begin{aligned} \text{TAIL} &= \sum_j \text{TAIL}_j \\ &\leq \text{“the revenue from uniformly pricing with price ITEM”} \\ &\leq \text{ITEM}. \quad \square \end{aligned}$$

The final step in the proof of Theorem 9.1.16 is bounding CORE in terms of the revenues of item pricing ITEM and bundle pricing BNDL. This bound is based on showing that CORE, a sum of the agent’s values, is concentrated. When an item j contributes to CORE it is bounded from above by \hat{v} . To show that CORE concentrates, we show that the contribution of each item to the variance of CORE is low. Specifically, the variance of a distribution is bounded by twice the product of its range and its monopoly revenue.

Lemma 9.1.20. *The variance of a distribution F on support $[0, h]$ with monopoly revenue R^* is at most $2 h R^*$.*

Proof. The variance $\mathbf{E}[(v - \mathbf{E}[v])^2]$ of a random variable is upper bounded by its second moment $\mathbf{E}[v^2]$. A distribution F on range $[0, h]$ with monopoly revenue R^* is stochastically dominated by the equal revenue distribution on $[R^*, h]$ which has distribution function $F^\dagger(z) = 1 - R^*/z$, density function $f^\dagger(z) = R^*/z^2$, and a pointmass of R^*/h at h . Thus, the second moment of F is upper bounded by the second moment of F^\dagger , which is:

$$\mathbf{E}_{v \sim F^\dagger}[v^2] = h^2 R^*/h + \int_{R^*}^h z^2 f^\dagger(z) dz = hR^* + \int_{R^*}^h R^* dz \leq 2 h R^*. \quad \square$$

Lemma 9.1.21. *For decomposition threshold $\hat{v} = \text{ITEM}$, the surplus from low-valued items is at most the revenue from item pricing plus three times the revenue from bundle pricing: $\text{CORE} \leq \text{ITEM} + 3 \text{BNDL}$.*

Proof. First, the variance of CORE is at most 2ITEM^2 . For each item j , consider truncating distribution F_j at \hat{v} by moving probability mass from all $v > \hat{v}$ to zero; call this truncated distribution F_j^\dagger . By definition $\text{CORE} = \mathbf{E}_{t^\dagger \sim F^\dagger}[\sum_j \{t^\dagger\}_j]$. This operation bounds the resulting distribution by \hat{v} and only lowers its monopoly revenue. Therefore, by Lemma 9.1.20, the variance of the truncated distribution is at most twice \hat{v} times the monopoly revenue of the original distribution, namely $2 \hat{v} R_j^*$. As CORE is the expectation of the sum of values contributed from each item, the variance of this sum is bounded by the sum of the bounds on the variances of each of these values, i.e., $2 \hat{v} \sum_j R_j^*$. The decomposition

threshold \hat{v} is assumed in the lemma to be the optimal item pricing ITEM and, by definition, $\text{ITEM} = \sum_j R_j^*$. Thus, the variance of the sum of values that contribute to CORE is at most 2ITEM^2 .

Second, a posted price of $\hat{v}^\dagger = \text{CORE} - \text{ITEM}$ for the grand bundle is accepted with probability at least $1/3$. By Cantelli's inequality (Mathematical Note on page 356) we bound the probability that the agent accepts the bundle price by

$$\begin{aligned} \Pr_{t \sim F} \left[\sum_j \{t\}_j \geq \text{CORE} - \text{ITEM} \right] &\geq \Pr_{t \sim F^\dagger} \left[\sum_j \{t\}_j \geq \text{CORE} - \text{ITEM} \right] \\ &\geq 1 - \frac{2 \text{ITEM}^2}{2 \text{ITEM}^2 + \text{ITEM}^2} = \frac{1}{3}. \end{aligned}$$

Concluding, $\text{BNDL} \geq 1/3 (\text{CORE} - \text{ITEM})$. Rearranging, we have the statement of the lemma: $3 \text{BNDL} + \text{ITEM} \geq \text{CORE}$. \square

The proof of Theorem 9.1.16 simply puts the preceding lemmas together.

Proof of Theorem 9.1.16. Combining Lemma 9.1.17, Lemma 9.1.18, Lemma 9.1.19, and Lemma 9.1.21 with $\hat{v} = \text{ITEM}$ and with approximation mechanism $\text{APX} = \max(\text{ITEM}, \text{BNDL})$, we have

$$\begin{aligned} \text{REF} &\leq \text{FAVE} + \text{TAIL} + \text{CORE} \\ &\leq \text{ITEM} + \text{ITEM} + 3 \text{BNDL} + \text{ITEM} \\ &\leq 6 \text{APX}. \end{aligned} \quad \square$$

9.1.5 Inapproximability for Correlated Distributions

Previously in this section we showed that, for an agent with values for items drawn from an independent distribution and either unit-demand or additive preferences, simple mechanisms like item pricings are a constant approximation to the revenue-optimal mechanism (which generally

Mathematical Note. Cantelli's inequality, also known as the *one-sided Chebyshev inequality* gives a tail bound on a random variable Z with mean μ and variance σ^2 as

$$\Pr[X \geq \mu - \lambda] \geq 1 - \frac{\sigma^2}{\sigma^2 + \lambda^2}$$

for non-negative λ .

prices lotteries, i.e., randomized allocations). In this section we show that, generally for correlated distributions, no simple mechanism is a good approximation to the optimal mechanism.

We will discuss first the case of an additive agent where type t and outcome and payment (x, p) gives utility $t \cdot x - p$ where $t \cdot x = \sum_j \{t\}_j \{x\}_j$ is the *dot product* between t and x , both treated as vectors. At the end of the section we extend the result to unit-demand agents.

The main idea behind this construction is finding a sequence of allocations $x \in [0, 1]^k$ such that (a) the dot product of the allocation with itself is significantly greater than the dot product of an allocation with any previous allocation in the sequence and (b) the minimum difference between these terms, defined as the “gap” below, is non-decreasing. These properties are helpful because they will enable the construction of a corresponding sequence of types and prices such that (a) the type that corresponds to an outcome does not prefer earlier outcomes in the sequence and (b) later outcomes in the sequence are higher priced.

Definition 9.1.2. *Given a sequence of allocations $\mathbf{x} = (x_1, x_1, \dots)$ in $[0, 1]^k$, a sequence $\delta = (\delta_1, \delta_2, \dots)$ in $[0, k]$ is a gap sequence for \mathbf{x} if it satisfies*

$$\delta_i \leq x_i \cdot x_i - x_i \cdot x_{i^\dagger}$$

for each i and $i^\dagger < i$. The gap sequence’s total gap is $\delta = \sum_i \delta_i$.

Lemma 9.1.22. *For a sequence of allocations \mathbf{x} and its corresponding gap sequence δ , an agent with type $t = v x_i$ prefers outcome (x_i, p_i) with $p_i = v \delta_i$ to the outcome of any previous allocation at any non-negative price.*

Proof. Type t ’s utility for outcome (x_i, p_i) is $v(x_i \cdot x_i - \delta_i)$. Type t ’s utility for any previous outcome $(x_{i^\dagger}, p_{i^\dagger})$ in the sequence is upper bounded by her surplus for the outcome $v x_i \cdot x_{i^\dagger}$. By the definition of gap δ_i , $x_i \cdot x_i - \delta_i \geq x_i \cdot x_{i^\dagger}$; thus, she prefers outcome (x_i, p_i) . \square

Lemma 9.1.23. *If a sequence of allocations has non-decreasing gap, then for a non-decreasing sequence of values (v_1, v_2, \dots) , prices of the form $p_i = v_i \delta_i$ are non-decreasing.*

Proof. Immediate. \square

Example 9.1.1. *Consider the all-bundles sequence $\mathbf{x} = (x_1, \dots, x_m)$ with all $m = 2^k - 1$ non-trivial bundles of the k items ordered from smallest to largest bundle size, i.e., by $\sum_j \{x_i\}_j$.*

To identify a gap sequence δ for \mathbf{x} , notice that the dot product $x \cdot x^\dagger$ for vectors in $\{0, 1\}^k$, representing bundles of the k items, simply counts the number of items that are contained in both bundles. Thus, we can choose $\delta_i = 1$: A bundle x_i with size ℓ has $x_i \cdot x_i = \ell$ and $x_i \cdot x_{i^\dagger} \leq \ell - 1$ for all $i^\dagger < i$. For the second part, notice that the bundle x_{i^\dagger} must have only have the same or fewer items number of items in it. In both cases, the number of items in common with x_i is at most $\ell - 1$. (This construction is tight, there is an $i^\dagger < i$ with $x_i \cdot x_{i^\dagger} = \ell - 1$.)

The total gap of the sequence is $\delta = \sum_{i=1}^m \delta_i = m = 2^k - 1$.

Our goal now is to use a sequence of allocations with high total gap to define a distribution over types where both item pricing and grand bundle pricing have small revenue but a more complex pricing has high revenue. To do so, we will be defining a type t_i and a price p_i for every allocation x_i in the sequence and choosing a suitable distribution over types F given by probability mass function $f_i = \mathbf{Pr}_{t \sim F}[t = t_i]$. The allocations and prices give a menu from which each type would select the outcome that maximizes its utility.

The construction will satisfy three properties: the expected revenue for item pricing is low; the expected revenue from outcomes in the sequence, if each type chooses its corresponding outcome, is high; and the i th type t_i prefers the i th outcome (x_i, p_i) to all lower priced outcomes. The revenue from the optimal mechanism is at least the revenue from the constructed menu which is at least the revenue that would be obtained of each type selected its corresponding outcome.

Theorem 9.1.24. *For a single-agent k -item environment and any sequence of allocations \mathbf{x} in $[0, 1]^k$ with non-decreasing gaps δ (Definition 9.1.2), there is a type distribution for an additive agent such that the revenue from pricing k individual items or bundling the k items together is at most $2k$ while the optimal revenue is the total gap $\delta = \sum_i \delta_i$.*

Proof. The construction of the type distribution is the following:

$$\begin{aligned} v_i &= 2^i, \\ t_i &= v_i x_i, \\ f_i &= 2^{-i}. \end{aligned}$$

Consider the menu of outcomes $\{(x_1, p_1), (x_2, p_2), \dots\}$ with p_i defined as:

$$p_i = v_i \delta_i.$$

By Lemma 9.1.22, the i th type does not prefer the outcome of any of the previous types $i^\dagger < i$. By Lemma 9.1.23, the prices of outcomes of subsequent types are non-decreasing. Consequently, each type i contributes at least $p_i = v_i \delta_i$ to the revenue. The expected revenue is at least $\sum_i p_i f_i = \sum_i v_i \delta_i f_i = \sum_i \delta_i = \delta$.

Now consider the revenue from posting a price for item j . The value of the agent for item j is $v_i \{x_i\}_j \leq v_i$ with probability f_i . The probability that the agent buys item j if it is priced at $v_i = 2^i$ is at most $\sum_{i^\dagger \geq i} f_{i^\dagger} \leq 2^{-i+1}$. Thus, the expected revenue from the item is, for any price, at most 2. Summing over all k items, the total revenue of any item pricing is bounded by $2k$.

The argument for bounding the revenue from posting a price for the grand bundle is similar. Type t_i 's value for the grand bundle is $v_i \sum_j \{x_i\}_j \leq v_i k$. The probability that the agent has value at most $v_i k$ for the grand bundle is at most $\sum_{i^\dagger > i} f_{i^\dagger} \leq 2^{-i+1}$. The revenue from any pricing of the grand bundle is at most $2k$. \square

The following corollary, which shows that simple pricing can be exponentially bad, is immediate from applying Theorem 9.1.24 to the all-bundles sequence of Example 9.1.1 which has total gap $\delta = 2^k - 1$.

Corollary 9.1.25. *For single-agent k -item additive-values environments, there is a correlated distribution on the agent's values for the items such that neither pricing individual items or pricing the grand bundle is better than a $2^{k-1}/2k$ approximation to the optimal pricing of the $2^k - 1$ non-trivial bundles.*

The all-bundles allocation sequence gives a separation between the revenues of simple pricings, of either individual items or the grand bundle, and the pricing of the $m = 2^k - 1$ non-trivial bundles (all deterministic allocations). The optimal mechanism, however, may price randomized allocations, i.e., lotteries. The following theorem shows that there is an infinite separation in revenue between simple mechanisms with a finite number of outcomes and complex mechanisms with an infinite number of (randomized) outcomes even in the special case where there are only $k = 2$ items. The proof is based on constructing an infinite sequence of

allocations which has infinite gap. For this sequence, unlike the one described above, the gaps are decreasing; however, they decrease at a slower rate than the values are increasing in the construction of Theorem 9.1.24 and, thus, prices are decreasing as is sufficient for the theorem's proof.

Theorem 9.1.26. *For single-agent 2-item additive-value environments, there is a correlated distribution on the agent's values for the items such the expected revenue of the optimal mechanism is infinite, while any mechanism that prices only a finite number of allocations has only finite revenue.*

Thus far this section has considered agents with additive valuations. Specifically the agent's value for an allocation is linear and given by $t \cdot x$ for allocations x that satisfy the supply constraint $\{x\}_j \leq 1$ for each item j . The main results of this section; specifically Theorem 9.1.24, Corollary 9.1.25, and Theorem 9.1.26; extend to unit-demand agents for which, relative to an additive agent, the supply constraint is tightened so that at most one of the k items is allocated, i.e., $\sum_j \{x\}_j \leq 1$. Given a mechanism for additive agents it can be converted to a mechanism for unit-demand agents by dividing the allocation probabilities and prices by k , the number of items. This change reduces the revenue of the mechanism by exactly a factor of k . As the values of the agent are defined as a scaled allocation, i.e., $t_i = v_i x_i$, in the construction of Theorem 9.1.24; reducing the allocation probabilities by a factor of k also reduces the agent's values for individual items by a factor of k which reduces the revenue of any item pricing by a factor of k . Thus, the approximation factor of Theorem 9.1.24 is unchanged in its modification to unit-demand agents.

Theorem 9.1.27. *For any sequence of allocations \mathbf{x} in $[0, 1]^k$ and non-decreasing gaps δ with total gap $\delta = \sum_i \delta_i$ (Definition 9.1.2), there is a type distribution for a unit-demand agent such that the revenue from item pricing is at most 2 while the optimal revenue is δ/k .*

As both Corollary 9.1.25 and Theorem 9.1.26 are proved by exhibiting sequences of allocations with large total gap and applying Theorem 9.1.24; both of these results extend to unit-demand agents.

Theorem 9.1.28. *For single-agent 2-item unit-demand environments, there is a correlated distribution on the agent's values for the items such the expected revenue of the optimal mechanism is infinite, while any*

mechanism that prices only a finite number of allocations has only finite revenue.

9.2 Multi-agent Approximation in Service Constrained Environments

This section considers the approximation of optimal mechanisms for agents with multi-dimensional and non-linear preferences in service constrained environments (Definition 8.3.1). One example of such an environment is that of selling a single item to one of several agents with public budget preferences (Example 8.2.1). Another example is selling a single item that has multiple configurations (i.e., alternatives) to one of several unit-demand agents (for example, a digital movie download that can be with or without subtitles in various languages).

Recall from Section 8.4 beginning on page 265 that, when the single-agent problems are revenue linear, there is a relatively simple mechanism that is optimal, namely the marginal revenue mechanism (Definition 8.4.3). Optimization of marginal revenue is a guiding principle of microeconomics and, though it is not always optimal to do so, in this section we will see that it is often approximately optimal.

The marginal revenue framework from Section 8.4 relates the single agent mechanism design problems with interim and ex ante allocation constraints. Recall that optimal single-agent revenue given an interim allocation constraint \hat{y} is denoted by $\mathbf{Rev}[\hat{y}]$ and is the maximum revenue obtained by mechanisms with allocation rule y that is no stronger than \hat{y} , i.e., for all \hat{q} the cumulative allocation rules satisfy $Y(\hat{q}) \leq \hat{Y}(\hat{q})$. The marginal revenue for allocation constraint \hat{y} is $\mathbf{MargRev}[\hat{y}] = \mathbf{E}_q[R'(q) \hat{y}(q)]$ where $R(\hat{q})$ is the \hat{q} ex ante optimal revenue. Equivalently, $R(\hat{q}) = \mathbf{Rev}[\hat{y}^{\hat{q}}]$ where the allocation constraint $\hat{y}^{\hat{q}}(\cdot)$ is defined as the reverse step function from 1 to 0 at \hat{q} .

Definition 9.2.1. *With n agents with revenue curves $\mathbf{R} = (R_1, \dots, R_n)$, optimal surplus given by $\mathbf{OPT}(\cdot)$, and quantiles $\mathbf{q} = (q_1, \dots, q_n)$ uniformly drawn from $[0, 1]^n$; the optimal marginal revenue is $\mathbf{E}_q[\mathbf{OPT}(\mathbf{R}'(\mathbf{q}))]$.*

The next two sections give two methods for showing that the optimal marginal revenue is a good approximation to the optimal revenue. The first method uses properties of the feasibility constraint imposed by the service-constrained environment and the second method is based on

showing that the single-agent interim problems are approximately linear. The third section shows that approximations to the ex ante optimal mechanism can also be used in the marginal revenue framework. An example result from this section is that unit-demand service constrained environments with fairly permissive distributional assumptions can be approximately solved by projecting the agents' multi-dimensional types to a single dimension given by the agents' values for their favorite alternatives. The resulting mechanism is very simple. The final section describes implementation of the marginal revenue mechanism, which is not generally as simple as it has been in previous discussions.

9.2.1 Marginal Revenue and the Ex Ante Relaxation

This section gives bounds on the approximation factor of the marginal revenue mechanism when the service constrained environment exhibits nice structure. Specifically, we show that the optimal mechanism for the ex ante relaxation (which is a marginal revenue mechanism by definition) does not violate ex post feasibility too much and that there is a natural way to address its violations of feasibility and construct from it a mechanism that is ex post feasible and does not lose too much of its revenue.

A framework for addressing this question was developed in Section 4.3. If the optimal mechanism for the ex ante relaxation serves the agents with ex ante probabilities \hat{q} , its revenue is $\sum_i R_i(\hat{q}_i)$. The sequential posted pricing approach resolves the ex post feasibility constraint with the first-come-first-served principle. To optimize revenue with sequential posted pricing, the agents are ordered by “bang per buck”, i.e., $R_i(\hat{q}_i)/\hat{q}_i$. When an agent i is considered, if it is feasible to serve the agent given the previously served agents then the agent is offered the menu of the \hat{q}_i optimal mechanism; otherwise, the agent is not served. In many natural environments for mechanism design these sequential posted pricing mechanisms have good revenue. For single-item environments, i.e., where at most one agent can be served, Theorem 4.3.4 shows that sequential posted pricing is an $e/e-1$ approximation (and the same bound extends to matroid environments, see Theorem 4.6.12).

A mechanism “is a marginal revenue mechanism” if to each agent, fixing the behavior of other agents, the agent is offered the menu of an ex ante optimal mechanism for some ex ante constraint. Both the ex ante relaxation and the sequential posted pricing, described above, are ex ante mechanisms. The ex ante relaxation, which is a relaxation of the

ex post feasibility constraint of the original mechanism design problem, upper bounds the optimal revenue. The sequential posted pricing is ex post feasible and, as a marginal revenue mechanism, lower bounds the revenue of the (optimal) marginal revenue mechanism. The following corollary of Theorem 4.6.12 follows from this discussion.

Corollary 9.2.1. *For matroid service-constrained environments and agents with (multi-dimensional, non-linear preferences and) independently distributed types, the optimal marginal revenue is an $\frac{e}{e-1}$ approximation to the optimal mechanism.*

Observe that the sequential posted pricing identified in the discussion above (and the proof of Theorem 4.6.12) gives a simple mechanism that achieves the same guarantee of Corollary 9.2.1. Of course the optimal marginal revenue mechanism will sometimes improve on this guarantee of the corollary.

9.2.2 Marginal Revenue and Approximate Revenue Linearity

When the single-agent interim-optimal mechanism design problem is revenue linear, i.e., $\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}]$, the optimal marginal revenue is the optimal revenue. This optimality smoothly degrades with the revenue non-linearity of the single-agent problems. The optimal marginal revenue is close to the optimal revenue when the interim optimal revenues are close to linear.

Proposition 9.2.2. *For downward-closed service-constrained environments and agents with independent types and multi-dimensional, non-linear preferences that satisfy $\mathbf{MargRev}[\hat{y}] \geq 1/\beta \mathbf{Rev}[\hat{y}]$ for all interim allocation constraints \hat{y} , the optimal marginal revenue is a β approximation to the optimal revenue.*

Proof. Consider the profile of interim allocation rules \mathbf{y} of the optimal mechanism. By the assumption of the proposition, the marginal revenue for this profile $\sum_i \mathbf{E}_{q_i} [R'_i(q_i) y_i(q_i)]$ is within a β fraction of its (optimal) revenue $\sum_i \mathbf{Rev}[y_i]$. The optimal marginal revenue, by definition, can only do better. \square

To instantiate Proposition 9.2.2, it suffices to (a) identify a linear upper bound on the interim optimal revenue for all interim constraints \hat{y} and (b) show that an ex ante mechanism approximates this upper bound

for all ex ante constraint \hat{q} . This approach is carried out below for agents with unit-demand preferences drawn from a product distribution.

For part (a), consider an allocation constraint \hat{y} for the representative environment (see Definition 8.6.3 and Section 9.1) where the multi-dimensional agent is replaced with single-dimensional representatives. The optimal revenue for the representative environment is easy to describe. Consider the distribution of the maximum virtual value and serve the agent with the maximum virtual value with the probability specified by $\hat{y}(\cdot)$ for the quantile corresponding to this virtual value (with respect to the distribution of the maximum virtual value). By the analysis of virtual values in single-dimensional environments of Chapter 3, specifically Definition 3.3.3 (virtual values) and Theorem 3.4.5 (virtual surplus maximization is optimal), the expected virtual surplus is equal to the optimal revenue for the representative environment. Furthermore, virtual surplus is linear. It is relatively straightforward to generalize Theorem 9.1.1, which shows that twice the representative revenue upper bounds the optimal unit-demand revenue for unconstrained single-agent problems, to single-agent environments with interim constraints.

Theorem 9.2.3. *For a unit-demand agent with independent values and any interim allocation constraint \hat{y} , the revenue of the optimal auction for the representative environment is at least half the revenue of the optimal lottery pricing in the original unit-demand environment; moreover, its revenue is linear in \hat{y} .*

Proof. See Exercise 9.4. □

For part (b), we must show that there is an \hat{q} ex ante mechanism that approximates the upper bound given by twice the representative revenue with the same ex ante constraint \hat{q} . This representative revenue can be calculated as the expected virtual surplus from serving the top \hat{q} measure of draws from distribution of the maximum virtual value. Just as we obtained a four approximation to the unconstrained upper bound via the prophet inequality (Theorem 4.2.1), there is a straightforward adaptation of the prophet inequality to the case where both the prophet and the gambler have an ex ante constraint \hat{q} . In this generalization the $\hat{q} = 1$ case, i.e., the original prophet inequality, gives the worst approximation bound of two. As a consequence, via the same argument as Corollary 9.1.2 but with an ex ante allocation constraint, a uniform virtual price gives a good approximation to the optimal representative revenue.

Theorem 9.2.4. *For a unit-demand agent with independent values and ex ante allocation constraint \hat{q} , the revenue of a uniform virtual item pricing is a two approximation to the optimal representative revenue.*

Proof. See Exercise 9.5. \square

We conclude that a unit-demand agent with independent values is approximately revenue linear and therefore the optimal marginal revenue for a collection of such agents is approximately optimal. Notice that for matroid environments, the $\frac{e}{e-1}$ bound via the feasibility constraint is better than the bound we get via the unit-demand single-agent problem; however, approximate linearity gives a bound more generally for downward-closed service-constrained environments.

Corollary 9.2.5. *For unit-demand agents with independent values and downward-closed service-constrained environments, the optimal marginal revenue marginal is a four approximation to the optimal revenue.*

9.2.3 The Marginal Revenue Mechanism with Ex Ante Approximations

The marginal revenue framework constructs multi-agent mechanisms from single-agent mechanisms that solve the ex ante mechanism design problems. These multi-agent mechanisms are, thus, only as simple as the ex ante mechanisms are simple. As we saw in Section 8.8, however, the optimal unconstrained mechanism for a single unit-demand agent, and thus the ex ante optimal mechanisms, can be quite complex. This section shows how the marginal revenue framework can be adapted to construct approximately optimal multi-agent mechanisms from approximately optimal single-agent ex ante mechanisms.

The proof of the approximate optimality of the optimal marginal revenue for unit-demand agents, above, was based on Theorem 9.2.4 which showed that a simple uniform virtual pricing for the ex ante problem approximates the linear upper bound given by twice the representative revenue. (The optimal ex ante mechanism is no worse, thus the bound is established.) This proof approach, however, suggests that the same guarantee for the optimal marginal revenue constructed from the optimal single-agent ex ante mechanisms also holds for optimal marginal revenue constructed from this family of single-agent ex ante approximation mechanisms, namely uniform virtual pricings.

This approach can be formalized as follows. Denote by $\tilde{P}(\hat{q})$ the rev-

enue of a family of ex ante approximation mechanisms as a function of the ex ante constraint \hat{q} . Denote by $\tilde{R}(\hat{q})$ the smallest concave function that upper bounds $\tilde{P}(\hat{q})$. Refer to $\tilde{R}(\cdot)$ as the *pseudo-revenue curve*. Notice that uniform virtual pricings induce an ordering on types (cf. Definition 8.4.1) and therefore the *marginal pseudo-revenue mechanism* following Definition 8.4.3, the marginal revenue mechanism for orderable agents, is well defined. Specifically:

- (i) Map agent types \mathbf{t} to quantiles \mathbf{q} . The quantile of a type t_i is the infimum quantile q_i for which the type is served by agent i 's q_i ex ante mechanism.
- (ii) Calculate the marginal pseudo-revenues for each agent $\tilde{\mathbf{R}}(\mathbf{q})$.
- (iii) Serve the set of agents that maximizes the surplus of marginal pseudo revenue, $\mathbf{x} = \text{OPT}(\tilde{\mathbf{R}}(\mathbf{q}))$.
- (iv) For each agent i , offer the menu corresponding to the \hat{q}_i ex ante mechanism where \hat{q}_i is set from \mathbf{q}_{-i} as the supremum quantile where i is served by $\text{OPT}(\tilde{\mathbf{R}}(\hat{q}_i, \mathbf{q}_{-i}))$.

This mechanism, by the same argument as Corollary 9.2.5, is a four approximation to the optimal mechanism.

Corollary 9.2.6. *For unit-demand agents with independent values in a downward-closed service-constrained environment, the marginal pseudo-revenue mechanism defined by uniform virtual pricing is a four approximation to the optimal mechanism.*

Recall that for the single-dimensional agents of Chapter 3, when the agents' types are identically distributed, i.e., $F_i = F$ for all i , then the marginal revenue mechanism is simply the second price auction with the monopoly reserve price. The marginal pseudo-revenue mechanism similarly simplifies when the agents' types are identically distributed. The agent with the overall highest positive single-dimensional virtual value (for any alternative) wins and buys her utility maximizing (i.e., value minus price) alternative under the uniform virtual prices corresponding to the overall highest virtual value of the other agents or a virtual reserve price, whichever is higher.

The marginal pseudo-revenue framework can also be applied to uniform pricings which are approximately optimal via Corollary 9.1.14. A uniform pricing always sells any agent her favorite alternative or nothing. Thus, mechanisms based on uniform pricing can be seen as projecting an agent's multi-dimensional type down to a single-dimension given by the agent's value for her favorite alternative. The proof of the following

theorem follows from an ex ante constrained version of Corollary 9.1.14 and Theorem 9.1.15.

Theorem 9.2.7. *For unit-demand agents with independent (but non-identical) type distributions in a downward-closed service-constrained environment, the optimal mechanism for the favorite-alternative projection is often a constant approximation to the optimal mechanism. Specifically, its approximation is at most*

- (i) $2e-1/e-1 \approx 2.58$ if each agent's type distribution is identical across alternatives, and
- (ii) $2e \approx 5.437$ if each agent's type distribution is regular (but non-identical) across alternatives.

The optimal mechanism for the favorite-alternative projection is a single-dimensional mechanism and, consequentially, the optimal mechanism based on uniform pricings and further approximations follow from the developments of Chapter 3-Chapter 6. Here are a collection of simple observations for single-item environments with type distributions that are identical across agents, regular across alternatives (but not identical), and regular for the distribution of any agent's value for her favorite alternative. The second-price auction with reserve for the winner's favorite alternative is a $2e$ approximation to the optimal mechanism. Without a reserve, the second price auction is a $2en/n-1$ approximation. (See Exercise 9.7.) Furthermore, relaxing the assumption that the agents' types are identically distributed, the second price auction with a reserve remains a $2e^2$ approximation to the optimal mechanism. Posted pricing mechanisms are also approximately optimal.

9.2.4 Implementation of the Marginal Revenue Mechanism

The marginal revenue mechanism for revenue-linear agents and the marginal pseudo-revenue mechanism described in the previous section have simple implementations because the ex ante mechanisms induce an ordering on types that can be used to define quantiles. Implementing the marginal revenue mechanism without this orderability can be complex. See Section 3.4 and Section 8.4 for additional discussion of the marginal revenue framework that is pertinent to the developments of this section.

Given the agents' revenue curves \mathbf{R} the allocation rules of the marginal revenue mechanism can be identified as follows. For each agent i and

quantile q_i , allocation rule $y_i(q_i)$ is the probability that i is served by $\text{OPT}(\mathbf{R}(\mathbf{q}))$ when the other agent quantiles \mathbf{q}_{-i} are drawn uniformly from $[0, 1]^{n-1}$. The interim mechanism agent i faces, with allocation rule y_i thus constructed, is given by the convex combination of the \hat{q}_i ex ante mechanisms with \hat{q}_i drawn with cumulative distribution function given by $G_i(z) = 1 - y_i(z)$. Its revenue, as desired, is $\mathbf{E}_{\hat{q}_i \sim G_i}[R(\hat{q}_i)] = \mathbf{E}_q[R'(q) y(q)]$.

Given a mechanism that maps profiles of quantiles to an ex post feasible allocation rule – for the marginal revenue optimization this mechanism is $\text{OPT}(\mathbf{R}(\cdot))$ – and a profile of interim mechanisms that have allocation rules that are no stronger than the ones induced by the ex post mechanism, there is a generic construction for combining these into an ex post feasible mechanism wherein each agent faces her interim mechanism. The formal details of this construction are given by Definition 8.5.6 and Theorem 8.5.13 in Section 8.5.4.

In the remainder of this section we develop a simpler approach to construct the marginal revenue mechanism when the single-agent ex ante mechanisms exhibit additional structure. A running example that these methods apply to is that of the public budget agent discussed in Section 8.7. The ex ante optimal mechanisms for a public budget agent satisfy the following monotonicity property (to be formally proven at the end of this section).

Definition 9.2.2. *An agent has monotone ex ante mechanisms if the allocation rules (in type space) of the optimal ex ante mechanisms, denoted $x^{\hat{q}}(t)$, are monotonic functions of \hat{q} for all fixed types $t \in \mathcal{T}$, i.e., $x^{\hat{q}}(t) \leq x^{\hat{q}^\dagger}(t)$ for $\hat{q} < \hat{q}^\dagger$.*

There are two challenges with generalizing the marginal revenue mechanism for orderable agents (where the ex ante mechanisms define a partial ordering on types). The first challenge is how to map types to quantiles. The second challenge is selecting a consistent outcome for the mechanism when all agents face one of their ex ante mechanisms. When the agents are orderable, both steps are easy. Types are mapped to quantiles according to the ordering and type distribution. When the ex ante mechanisms are deterministic, as they are for orderable agents, there is only one service possibility for a given type and ex ante mechanism; i.e., for any agent, type t , and ex ante constraint \hat{q} , the allocation probability is $x^{\hat{q}}(t) \in \{0, 1\}$.² Critically, with only monotone ex ante mechanisms, the

² Technically, this statement is restricted to ex ante constraints where the revenue

allocation is not generally a deterministic function of the ex ante constraint and the type. For example, allocation probabilities are not zero or one in the ex ante mechanisms for the public budget agent where the budget constraint is binding (see Lemma 9.2.11 and Figure 9.2, below). The following definition of the marginal revenue mechanism resolves both of the above issues.

Definition 9.2.3. *The marginal revenue mechanism for agents with monotone ex ante mechanisms works as follows:*

- (i) Define quantile distribution $Q_i^{t_i}$ with cumulative distribution function $Q_i^{t_i}(z) = x_i^z(t_i)$ for each agent i with ex ante mechanisms given by x_i^q .
- (ii) Map the profile of agents' types \mathbf{t} to a profile of quantiles \mathbf{q} by sampling $q_i \sim Q_i^{t_i}$.
- (iii) Calculate the profile of marginal revenues for the profile of quantiles $\mathbf{R}'(\mathbf{q})$.
- (iv) Calculate a feasible allocation to optimize the surplus of marginal revenue, i.e., $\mathbf{x} = \text{OPT}(\mathbf{R}'(\mathbf{q})) \in \{0, 1\}^n$. For each agent i , calculate the supremum quantile \hat{q}_i she could possess for which she would be allocated in the above calculation of \mathbf{x} .
- (v) For each agent i , an outcome distribution is given by her type and the \hat{q}_i ex ante mechanism. Sample from this outcome distribution conditioned on whether or not she is served in \mathbf{x} , i.e., $x_i \in \{0, 1\}$.³

There are two key properties that are sufficient for the mechanism of Definition 9.2.3 to implement the optimal marginal revenue. First, the quantiles of the agents that are calculated within the mechanism should be independently and uniformly distributed on $[0, 1]$. Second, conditioned on agent i 's critical quantile \hat{q}_i , the outcome for agent i should faithfully implement the \hat{q}_i ex ante mechanism. These two properties are proved in the two lemmas below. The subsequent theorem then shows that the defined mechanism is incentive compatible and implements the optimal marginal revenue.

Lemma 9.2.8. *For an agent with $t \sim F$, ex ante mechanisms with allocation rules x^q satisfying $\mathbf{E}_{t \sim F}[x^q(t)] = \hat{q}$ for $\hat{q} \in [0, 1]$, and quantile*

curve is not ironed, i.e., $R''(\hat{q}) \neq 0$. Recall, however, that quantiles \hat{q} where $R''(\hat{q}) = 0$ cannot be critical quantiles in the marginal revenue mechanism.

³ Recall, in the general service constrained environments of Section 8.3, some outcomes for an agent are designated as service outcomes and others are designated a no-service outcomes.

distributions $Q^t(z) = x^z(t)$; for $t \sim F$ and $q \sim Q^t$ the distribution of q is uniform on $[0, 1]$.

Proof. For a fixed type, by definition of the cumulative distribution function $\Pr_{q \sim Q^t}[q \leq \hat{q}] = Q^t(\hat{q}) = x^{\hat{q}}(t)$. Taking expectation over $t \sim F$ of this probability we obtain

$$\Pr_{t \sim F; q \sim Q^t}[q \leq \hat{q}] = \mathbf{E}_{t \sim F}[x^{\hat{q}}(t)] = \hat{q}.$$

Thus, q is uniformly distributed. \square

Lemma 9.2.9. *The marginal revenue mechanism of Definition 9.2.3, for agent i and conditioned on \hat{q}_i , faithfully implements the \hat{q}_i ex ante mechanism.*

Proof. Consider agent i with type t_i and fix \hat{q}_i as determined by the random quantiles \mathbf{q}_{-i} of other agents. Viewing the \hat{q}_i ex ante mechanism as a menu of distributions over outcomes, type t_i chooses her favorite distribution over outcomes from this menu. Some of the outcomes in the support of the chosen distribution are service outcomes and some are no-service outcomes. The probability of a service outcome is, by definition, exactly $x_i^{\hat{q}_i}(t_i)$. One way to draw an outcome from the appropriate distribution is draw an outcome conditioned on service ($x_i = 1$) with probability $x_i^{\hat{q}_i}(t_i)$ or an outcome conditioned on no-service with probability $1 - x_i^{\hat{q}_i}(t_i)$. The mechanism of the definition draws $q_i \sim Q_i^{t_i}$ and serves i when $q_i \leq \hat{q}_i$. It suffices to verify that this probability of service is exactly $x_i^{\hat{q}_i}(t_i)$. By the definition of the distribution function $\Pr[q_i \leq \hat{q}_i] = Q_i^{t_i}(\hat{q}_i) = x_i^{\hat{q}_i}(t_i)$. The lemma follows. \square

Theorem 9.2.10. *The marginal revenue mechanism for agents with monotone ex ante mechanisms has revenue equal to the optimal marginal revenue and is dominant strategy incentive compatible.*

Proof. Consider agent i with type t_i and fixed \hat{q}_i as determined by the random quantiles \mathbf{q}_{-i} of other agents (which come from the other agents' reports). The \hat{q}_i ex ante mechanism is incentive compatible and, by Lemma 9.2.9, faithfully implemented by the mechanism of the definition. Therefore, the mechanism of the definition is dominant strategy incentive compatible.

The optimal marginal revenue is $\mathbf{E}_{\mathbf{q}}[\text{OPT}(\mathbf{R}'(\mathbf{q}))]$. To see that the revenue of the defined mechanism is the optimal marginal revenue, note that in the defined mechanism agent i faces the \hat{q}_i ex ante mechanism where \hat{q}_i is the maximum quantile at which agent i is served when the

other agent quantiles \mathbf{q}_{-i} , by Lemma 9.2.8, are drawn from the uniform distribution on $[0, 1]^{n-1}$. The same is true of the optimal marginal revenue. As the defined mechanism faithfully implements the ex ante mechanism for each \hat{q}_i (Lemma 9.2.9), the contributions from agent i to the optimal marginal revenue and the revenue of the defined mechanism are both $\mathbf{E}_{\hat{q}_i}[R_i(\hat{q}_i)]$. The distribution over outcomes of the mechanism and the optimal marginal revenue is also the same; thus, any cost incurred by the designer for the outcome produced is the same. \square

Now that we have seen how to construct the marginal revenue mechanism for agents with monotone ex ante mechanisms, we will show that single-dimensional agents with public budgets have monotone ex ante mechanisms. Useful in the proof of the theorem is the lemma, below, which characterizes the ex ante mechanisms for public budget agents.

Lemma 9.2.11. *The \hat{q} ex ante optimal mechanism for a single-dimensional agent with public budget B and type drawn from public-budget-regular distribution F (Definition 8.7.1) is either:*

- (i) *Budget binds: Post the price B for allocation probability $B/t^\dagger \leq 1$ with \hat{t}^\dagger set to satisfy $\hat{q} = B/t^\dagger(1 - F(\hat{t}^\dagger))$. Types $t \geq \hat{t}^\dagger$ select the lottery.*
- (ii) *Allocation probability binds: Post price $\hat{t} = F^{-1}(1 - \hat{q})$ for allocation probability one.*
- (iii) *Neither bind: Post the monopoly price of distribution F for allocation probability one.*

Proof. If the budget is not binding then the Lagrangian relaxation of the budget constraint has Lagrangian parameter $\lambda = 0$ and the optimal mechanism is the same as the optimal mechanism for a single-dimensional linear agent. Thus, parts (ii) and (iii) follow from the results of Chapter 3. The remainder of the proof focuses on the part (i) where the budget binds.

When the budget binds, Theorem 8.7.5 on page 304 shows that the optimal mechanism is given by optimizing the Lagrangian revenue curve that, for the public budget regular distributions, is characterized by Proposition 8.7.4 as having a single ironed interval for the strongest quantiles, i.e., quantiles in $[0, \hat{q}^\dagger]$. The interim allocation constraint $\hat{y}^{\hat{q}}$ that corresponds to an ex ante allocation constraint \hat{q} is the reverse step function from 1 to 0 at \hat{q} . Either the ironed interval $[0, \hat{q}^\dagger]$ spans \hat{q} or it does not. (In fact, it will span \hat{q} in the case that the budget binds, though we do not need to directly prove this fact.) In either case, the

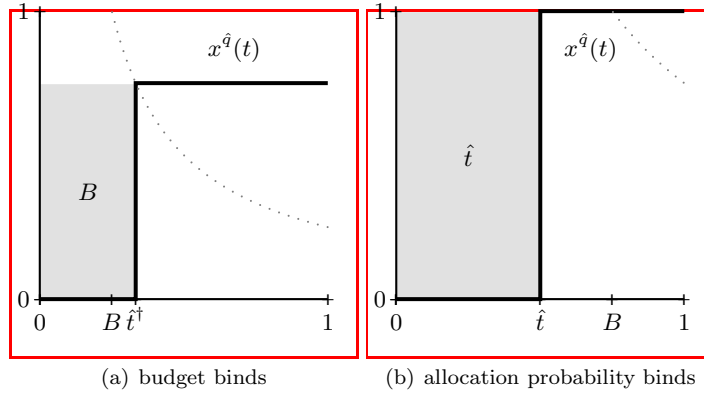


Figure 9.2. The allocation rule of the ex ante optimal mechanism for an agent with a public budget is depicted. Subfigure (a) shows the case where the budget binds; the shaded area is the payment of the highest type and is equal to the budget $p^q(1) = B$. The allocation rules of the ex ante mechanism where the budget binds step at \hat{t}^\dagger from 0 to B/\hat{t}^\dagger ; the pointset $\{(\hat{t}^\dagger, B/\hat{t}^\dagger) : \hat{t}^\dagger \in [B, \infty)\}$ is depicted with a dotted line. Subfigure (b) shows the case where the allocation probability constraint binds. The payment of the high types is equal to $\hat{t} = F^{-1}(1 - \hat{q}) \leq B$. Depicted is the special case where the type distribution is uniform on $[0, 1]$ and the ex ante constraint is $\hat{q} = 1/2$. Subfigure (a) has budget $B = 1/4$; subfigure (b) has budget $B = 3/4$. For the uniform distribution the area under the curve in type space, i.e., $\int_0^1 x^q(z) dz$, is equal to the area under the curve in quantile space, a.k.a., the ex ante allocation probability \hat{q} . Thus, in subfigure (a), $\hat{t}^\dagger = B/B + \hat{q}$.

allocation rule must be a reverse step function (in quantile space) and a step function (in type space). The allocation rule in type space – where the payment of the highest type is the area above the curve and equals the budget B – must step at some \hat{t}^\dagger to from 0 to B/\hat{t}^\dagger (see Figure 9.2). The step function of this form that has ex ante allocation probability \hat{q} sets \hat{t}^\dagger to solve $B/\hat{t}^\dagger(1 - F(\hat{t}^\dagger)) = \hat{q}$. \square

Theorem 9.2.12. *A public budget agent with type drawn from a public budget regular distribution (Definition 8.7.1) has monotone ex ante mechanisms.*

Proof. By Lemma 9.2.11, for any ex ante constraint \hat{q} where the budget binds, the allocation rule in type space steps at some \hat{t}^\dagger from 0 to B/\hat{t}^\dagger . For this form of mechanism, as the ex ante probability \hat{q} is increased, \hat{t}^\dagger decreases and the probability of service for the higher types B/\hat{t}^\dagger increases (see Figure 9.2). The change in allocation probabilities are as follows:

Types that were served before are now served with higher probability and some types that were not served before are now served. Thus, the agent has monotone ex ante mechanisms.

For ex ante constraints where the budget does not bind (specifically, where instead the allocation constraint binds), the analysis of linear agents from Chapter 3 implies the theorem. \square

9.3 Multi-agent Approximation with Multi-dimensional Externalities

A fundamental problem in multi-dimensional mechanism design is that of identifying good mechanisms for selling multiple items to multiple agents. Of course, if one agent obtains any of the items, these items cannot be allocated to other agents. The externality that such an agent imposes on the others is multi-dimensional. Rarely do we see such goods sold by auctions, even for rare goods, instead posted prices tend to be preferred. This section shows that posting prices is approximately optimal quite broadly. These results extend those from Chapter 4 for single-dimensional agents to multi-dimensional agents with multi-dimensional externalities.

Consider the canonical multi-dimensional matching environment (previously discussed in Section 8.6 beginning on page 296). In this environment the agents are unit demand, each desires at most one of the m items; and the items are in unit supply, each can be sold to at most one of the n agents. The agents' multi-dimensional types are independently distributed and, moreover, each agent's values for the items are independently distributed. We will show that there is an easy to identify posted pricing that is a constant approximation to the optimal mechanism. When the agents are symmetric, the posted pricing is anonymous (as is the optimal mechanism), i.e., each agent is offered the same menu of items and prices; when the agents are asymmetric the posted pricing is generally discriminatory (as is the optimal mechanism).

The proof approach generalizes the methods of Section 9.1. We obtain upper bounds on the optimal mechanism for unit-demand agents from the single-dimensional representative environment. We obtain lower bounds on the revenue of a posted pricing mechanism by the revenue of the same posted pricing in the representative environment. The remaining question of relating these upper and lower bounds is a question of the effectiveness of posted prices as approximation mechanisms for

the single-dimensional representative environment. The main results of this section are immediate corollaries. These results are formally stated below with their proofs deferred to the subsequent discussion.

These first three theorems establish the upper and lower bounds.

Theorem 9.3.1. *For unit-demand agents with independent values, the revenue of the optimal mechanism for the representative environment upper bounds the revenue for the optimal deterministic dominant-strategy incentive-compatible mechanism for the original unit-demand environment.*

Proof. See Exercise 9.8. □

Theorem 9.3.2. *For unit-demand agents with independent values, the revenue of the optimal mechanism for the ex ante relaxation of the single-dimensional representative environment is a three approximation to the revenue of the optimal mechanism for the original unit-demand environment.*

Theorem 9.3.3. *For unit-demand agents with independent values, the revenue from oblivious posted pricing in the unit-demand environment is at least that of oblivious posted pricing in the representative environment.*

The approach of this section can be viewed as a reduction from unit-demand environments to the single-dimensional representative environment. Specifically, the problem of approximating the optimal mechanism by a posted pricing in an unit-demand environment reduces to the problem of approximating the optimal mechanism by a posted pricing in the single-dimensional representative environment. The following corollary, which combines the previous theorems, makes this reduction precise.

Corollary 9.3.4. *For any unit-demand environment, if an oblivious posted pricing is a β approximation to the optimal ex ante relaxation of the representative environment then the same oblivious posted pricing is a β approximation to the optimal deterministic mechanism and a 3β approximation to the optimal mechanism.*

The following theorem instantiates the reduction of Corollary 9.3.4 for unit-demand matching environments.

Theorem 9.3.5. *For single-dimensional bipartite matching environments (where agents are edges and feasible outcomes are matchings), for any independent distribution of values, there is an oblivious posted*

pricing that is a nine approximation to the optimal revenue of the ex ante relaxation (and the optimal revenue).

The following corollary states the unit-demand approximation result that follows from the reduction.

Corollary 9.3.6. *For unit-demand unit-supply matching environments with independently distributed values, there is an oblivious posted pricing that is a nine approximation to the optimal deterministic mechanism and a 27 approximation to the optimal mechanism.*

The remaining agenda for this section is to prove the theorems above. These proofs reinforce many of the topics already covered in this text. The proof of the upper bound for deterministic mechanisms (Theorem 9.3.1) is similar to the single-agent proof of Theorem 9.1.1 and will be left for Exercise 9.8. The proof of the upper bound for potentially randomized mechanisms (Theorem 9.3.2) follows a similar analysis to that of a single unit-demand agent that was given in Section 9.1.3. The proof of the approximate optimality of oblivious posted pricings (Theorem 9.3.5) for the single-dimensional representative environment borrows analysis methods from Chapter 4.

9.3.1 Multi-service Service-constrained Environments, Revisited

The multi-dimensional matching environment is a special case of the more general multi-service service-constrained environments described in Section 8.6. By Definition 8.6.1 (restated below), the matching environment is given by feasible outcomes of the matching polytope $\mathcal{X} = \{\mathbf{x} \in [0, 1]^{N \times M} : \forall j, \sum_i x_{ij} \leq 1 \wedge \forall i, \sum_j x_{ij} \leq 1\}$.

Definition 8.6.1. *In a multi-service service-constrained environment there are n agents N and m services M . The subset of agent-service pairs that can be simultaneously assigned is given by $\mathcal{X} \subset \{0, 1\}^{N \times M}$.*

A unit-demand multi-service service-constrained environment is one where the feasible outcomes, without loss, restrict to serving each agent at most one of the services, i.e., for $\mathbf{x} \in \mathcal{X}$ and all agents i , $\sum_j x_{ij} \leq 1$.

9.3.2 Oblivious Posted Pricing

Consider oblivious posted pricings in unit-demand multi-service service-constrained environments. An oblivious posted pricing is given by prices

	Item 1	Item 2		Item 1	Item 2
Agent 1	$t_{11} = 4$	$t_{12} = 5$	Agent 1	$\hat{t}_{11} = \mathbf{2}$	$\hat{t}_{12} = 4$
Agent 2	$t_{21} = 10$	$t_{22} = 8$	Agent 2	$(\hat{t}_{21} = 5)$	$\hat{t}_{22} = \mathbf{4}$
(a) agent types			(b) discriminatory posted pricing		

Figure 9.3. 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 item 1, agent 2 buys item 2, and the revenue is 6 (purchase prices depicted in boldface). If the agents arrive in the opposite order a higher revenue is obtained.

\hat{t} with \hat{t}_{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 order of the agents and we pessimistically assume the worst-case. See Figure 9.3 for an example.

Definition 9.3.1. An oblivious posted pricing is a pricing of services (discriminatory) for each agent with the semantics that agents arrive in any order and take their favorite service that remains feasible. For a distribution over agent types, the revenue of an oblivious posted pricing is given by the worst ordering for the realized types taken in expectation of the distribution.

Consider the oblivious 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 oblivious 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 9.3.7. The expected revenue of an oblivious posted pricing for unit-demand environments is at least the expected revenue of the same pricing in the single-dimensional representative environment.

Proof. Compare oblivious posted pricings for unit-demand environments (i.e., with n unit-demand agents) with oblivious 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 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. The worst-case sequences, then, for the representative environment are worse than those of the original unit-demand environment. \square

9.3.3 Posted Pricing for the Representative Environment

For the unit-demand matching market the single-dimensional representative environment is given by a bipartite graph where representatives correspond to edges and an outcome is feasible if it is a matching. In this single-dimensional representative environment there is an oblivious posted pricing that is a good approximation to the optimal mechanism (and the ex ante relaxation). As this is a single-dimensional approximation problem, we will adopt the notation and terminology of Chapter 4. However, for convenience we will index the nm agents, which correspond to edges in the bipartite matching, by their ij coordinates. Representative ij 's value is v_{ij} drawn independently from distribution F_{ij} , her quantile is q_{ij} , her revenue curve is R_{ij} , etc.

The optimal revenue from the ex ante relaxation is given by the program:

$$\begin{aligned} \max_{\hat{q}} \quad & \sum_{ij} R_{ij}(\hat{q}_{ij}) & (9.3.1) \\ \text{s.t.} \quad & \sum_i \hat{q}_{ij} \leq 1 & \forall j, \\ & \sum_j \hat{q}_{ij} \leq 1 & \forall i. \end{aligned}$$

For the purpose of preliminary discussion, assume the value distributions are regular. The solution \hat{q} to the ex ante program (9.3.1) corresponds to a price posting \hat{v} with $\hat{v}_{ij} = 1/\hat{q}_{ij} R_{ij}(\hat{q}_{ij})$. Posting these prices does not generally lead to a good revenue when the agents arrive in a worst-case order. To make sure a good revenue is attained when the agents

arrive arbitrarily, consider instead the solution \hat{q}^\dagger with $\hat{q}_{ij}^\dagger = \hat{q}_{ij}/2$. This solution is ex ante feasible, does not degrade the revenue by too much (at most a factor of two), and guarantees, even if agent ij arrives last in the order, that the probability that it is ex post feasible to serve the agent is at least a quarter. This calculation follows for agent ij because the probability that item j is sold already is at most one half and the probability that a representative of agent i has been already served is at most one half. If neither of these events have happened then it is feasible to serve agent ij . Since each agent can be served with probability at least a quarter and the prices posted obtain at least half the ex ante optimal revenue, the expected revenue of this oblivious posted pricing is at least an eighth of the expected revenue of the ex ante relaxation.

When the value distributions are irregular, the revenue curve of an agent is not equal to the price-posting revenue curve. Therefore, we may not be able to obtain both the revenue and ex ante allocation probability of the desired point on a given revenue curve by posting a deterministic price. The lemma below demonstrates that the construction above can be modified to find a posted price that approximates the desired point on the revenue curve. Specifically, it gives nearly the same revenue and its ex ante probability of allocation is no higher.

Lemma 9.3.8. *For any single-dimensional agent, ex ante probability \hat{q} , and factor $\beta > 1$, there is a posted price \hat{v}^\dagger that corresponds to $\hat{q}^\dagger \leq \hat{q}/\beta$ for which the price-posting revenue at \hat{q}^\dagger is a $\beta + 1$ approximation to the original revenue at \hat{q} , i.e., $P(\hat{q}^\dagger) \geq 1/(\beta + 1) R(\hat{q})$.*

Proof. We would like to offer a posted price that corresponds to quantile \hat{q}/β . For regular distributions (Definition 3.3.1 on page 64) where price-posting and optimal revenue are equal, this price gives a β approximation as $P(\hat{q}/\beta) = R(\hat{q}/\beta) \geq 1/\beta R(\hat{q})$ (by concavity of $R(\cdot)$ and because $R(0) \geq 0$). The challenge then is irregular distributions.

This proof will exploit the geometry of revenue curves; specifically, that effective prices $1/\hat{q} R(\hat{q})$ and $1/\hat{q} P(\hat{q})$ are non-increasing in the ex ante probability of allocation \hat{q} . Specifically, when selling with a higher probability the average price cannot increase. For the price posting revenue curve, the effective price is the actual price, $1/\hat{q} P(\hat{q}) = \hat{v}$. The proof proceeds by a case analysis on \hat{q}^\ddagger , the highest quantile below \hat{q} for which the price-posting revenue and optimal revenue are equal, i.e., $P(\hat{q}^\ddagger) = R(\hat{q}^\ddagger) = \hat{v}^\ddagger \hat{q}^\ddagger$.

If $\hat{q}^\ddagger \geq \hat{q}/\beta$ then, as follows, the price \hat{v}^\ddagger corresponding to ex ante probability $\hat{q}^\dagger = \hat{q}/\beta$ satisfies the conditions of the lemma. By the geom-




Figure 9.4. A geometric depiction of the quantities in Lemma 9.3.8. The revenue curve R (thin, black, solid line) is the convex hull of the price-posting revenue curve P (thick, grey, dashed line). The quantiles \hat{q} and \hat{q}/β are within an ironed interval with lower bound \hat{q}^\dagger . The dashed lines that connect to points on the price-posting revenue curve have slope equal to the price posted, i.e., $P(\hat{q}) = \hat{v}\hat{q}$ by definition. Importantly for the proof, the price-posting revenue of \hat{q}/β is at least a β fraction of the price-posting revenue for \hat{q} , i.e., $P(\hat{q}/\beta) \geq 1/\beta P(\hat{q})$, and the price-posting revenue of \hat{q}^\dagger is at least the difference in the optimal revenue and price-posting revenue of \hat{q} , i.e., $P(\hat{q}^\dagger) \geq R(\hat{q}) - P(\hat{q})$.

entry of revenue curves, the price \hat{v}^\dagger is at least \hat{v}^\dagger which is at least the effective price of optimal revenue at \hat{q} , namely $1/\hat{q} R(\hat{q})$; consequently, the revenues satisfy $P(\hat{q}/\beta) = \hat{v}^\dagger \hat{q}/\beta \geq \hat{v}^\dagger \hat{q}/\beta \geq 1/\hat{q} R(\hat{q}) \hat{q}/\beta = 1/\beta R(\hat{q})$ as desired.

If $\hat{q}^\dagger < \hat{q}/\beta$, as depicted in Figure 9.4, then the ex ante probability \hat{q}^\dagger in the statement of the lemma is either \hat{q}^\dagger or \hat{q}/β whichever has larger price-posting revenue. Partition the optimal revenue $R(\hat{q})$ into two pieces, that of the price-posting revenue $P(\hat{q})$ and the difference $R(\hat{q}) - P(\hat{q})$. The price-posting revenue from \hat{q}/β is at least a β fraction of $P(\hat{q}) = \hat{v}\hat{q}$ (as its corresponding price is only higher than \hat{v} and its probability of service is exactly a β factor lower). The slope of the revenue curve at \hat{q} , i.e., the marginal revenue $R'(\hat{q})$, is always at most \hat{v} . By geometry, then, the price-posting revenue of \hat{q}^\dagger is at least the remainder $R(\hat{q}) - P(\hat{q})$. Thus,

we have:

$$\begin{aligned} P(\hat{q}^\dagger) &\geq R(\hat{q}) - P(\hat{q}), \\ P(\hat{q}/\beta) &\geq 1/\beta P(\hat{q}). \end{aligned}$$

Adding the first line and β times the second line, we have:

$$P(\hat{q}^\dagger) + \beta P(\hat{q}/\beta) \geq R(\hat{q}).$$

Bounding both $P(\hat{q}^\dagger)$ and $P(\hat{q}/\beta)$ by their maximum, we have:

$$(1 + \beta) \max(P(\hat{q}^\dagger), P(\hat{q}/\beta)) \geq R(\hat{q}).$$

The lemma follows. \square

We now give the proof of Theorem 9.3.5, restated below. This proof applies the basic construction described previously, but with slightly different constants, and incorporates Lemma 9.3.8 to address the possible irregularity of the value distributions.

Theorem 9.3.5. *For single-dimensional bipartite matching environments (where agents are edges and feasible outcomes are matchings), for any independent distribution of values, there is an oblivious posted pricing that is a nine approximation to the optimal revenue of the ex ante relaxation (and the optimal revenue).*

Proof. Let \hat{q} be the revenue-optimal profile of probabilities that is ex ante feasible; i.e., \hat{q} optimizes program (9.3.1). Invoking Lemma 9.3.8 with $\beta = 3$ we obtain a profile of quantiles \hat{q}^\dagger with at most one third the allocation probability, i.e., $\hat{q}_{ij}^\dagger \leq \hat{q}_{ij}/3$, and at least one fourth the revenue, i.e., $P(\hat{q}_{ij}^\dagger) \geq 1/4 R(\hat{q}_{ij})$ for all ij .

By definition, \hat{q}^\dagger satisfies $\sum_{ij} P_{ij}(\hat{q}_{ij}^\dagger) \geq 1/4 \sum_{ij} R_{ij}(\hat{q}_{ij})$ and $\sum_i \hat{q}_{ij} \leq 1/3$ for all j and $\sum_j \hat{q}_{ij} \leq 1/3$ for all i . Moreover, if ij is feasible at the time ij arrives then the corresponding posted pricing of \hat{v}^\dagger attains revenue $P_{ij}(\hat{q}_{ij}^\dagger)$. To get a lower bound on the revenue from representative ij , imagine ij is the last representative to arrive. By the union bound, the probability of the event that another representative $i^\dagger j$ for $i^\dagger \neq i$ was previously served is at most $1/3$, likewise for the event that another representative ij^\dagger for $j^\dagger \neq j$ was previously served. Independence of these two events implies that the probability neither happens is $(2/3)^2 = 4/9$. Therefore, the revenue we can expect from representative ij under any ordering is at least $4/9 \cdot 1/4 \cdot R_{ij}(\hat{q}_{ij})$. Summing over all representatives gives the desired nine approximation. \square

Though the section focused on approximation by oblivious posted pricings, bounds can also be obtained for posted pricings with less pessimistic assumptions on the order of the agents. For example, Exercise 9.11 proves that when agents arrive in a random order and the distributions are independent and regular then the prices from the ex ante relaxation give a two approximation to the optimal mechanism.

9.3.4 Upper bound via the Representative Environment

A key step in proving that a simple and practical approximation mechanism performs nearly as well as the complex optimal mechanism is identifying an upper bound on the performance of the optimal mechanism. For deterministic mechanisms this upper bound, stated in Theorem 9.3.1 and with proof left for Exercise 9.8, is the optimal revenue from the single-dimensional representative environment. For general (randomized) mechanisms, Theorem 9.3.2, as proved below, shows that the three times optimal revenue for the ex ante relaxation of the representative environment is an upper bound.

To simplify the proof of this theorem, as we did for approximation for single-dimensional agents, we will represent the competition between agents by ex ante constraints. Consider a unit-demand agent i . In the optimal mechanism the agent receives each service j with some ex ante probability \hat{q}_{ij} . The unit-demand constraint requires $\sum_j \hat{q}_{ij} \leq 1$. Denote a multi-dimensional ex ante constraint by $\hat{q}_i = (\hat{q}_{i1}, \dots, \hat{q}_{im})$. Define the optimal revenue that can be obtained from this agent when selling with at most these ex ante probabilities by $R_i(\hat{q}_i)$. The following proposition shows that the optimal revenue for unit-demand agents is bounded by the optimal revenue from the ex ante relaxation.

Proposition 9.3.9. *For a unit-demand multi-service service-constrained environment (Definition 8.6.1), the optimal revenue is at most the optimal ex ante revenue $\sup_{\hat{q} \in \hat{\mathcal{Q}}} \sum_i R_i(\hat{q}_i)$, where $\hat{\mathcal{Q}}$ is the space of ex ante feasible probability profiles and $R_i(\hat{q}_i)$ is the unit-demand optimal revenue with ex ante allocation probabilities $\hat{q}_i = (\hat{q}_{i1}, \dots, \hat{q}_{im})$.*

Proof. Denote by \hat{q} the profile of ex ante probabilities induced by the optimal mechanism, i.e., where \hat{q}_{ij} is the ex ante probability that agent i is allocated service j . By definition this profile is ex ante feasible. The optimal way to serve each agent i with at most the prescribed probabilities gives revenue $R_i(\hat{q}_{i1}, \dots, \hat{q}_{im})$. The sum of these ex ante revenues for this ex ante feasible profile \hat{q} is at most the the sum of the

ex ante revenues for the ex ante feasible profile $\hat{\mathbf{q}}^*$ that optimizes ex ante revenue. \square

Proposition 9.3.9 reduces the multi-agent analysis to a single agent analysis. To make the notation more clear for the remainder of this section we drop the agent subscript i and consider the single-agent problem using the same notation as Section 9.1. It suffices to bound the optimal single-agent unit-demand revenue for any given multi-dimensional ex ante constraint $\hat{q} = (\hat{q}_1, \dots, \hat{q}_m)$ by the optimal revenue of the ex ante relaxation of the single-dimensional representative environment with the same ex ante constraint.

The single-agent unit-demand revenue is given by a mechanism that satisfies the ex ante constraint given by \hat{q} as well as an ex post constraint required by the unit-demand assumption on the agent's preference. Specifically, a unit-demand agent always receives at most one item ex post. To give an upper bound on the unit demand revenue in terms of the representative revenue we identify an amortization of revenue and relate the optimal expected surplus of the amortization to the optimal representative revenue, cf. Theorem 9.1.4 and Theorem 9.1.11.

Theorem 9.3.10. *For a unit-demand agent with multi-dimensional ex ante allocation constraint $\hat{q} = (\hat{q}_1, \dots, \hat{q}_m)$, the optimal revenue from the ex ante relaxation of the representative environment is a three approximation to the optimal revenue.*

The proof of Theorem 9.3.2 (that the optimal mechanism for the ex ante relaxation of the representative environment is a three approximation to the optimal mechanism for an original multi-agent unit-demand environment) follows from Proposition 9.3.9 and Theorem 9.3.10. We prove Theorem 9.3.10 below. The proof will be based on a generalization of the amortization of revenue from Section 9.1.2.

Definition 9.3.2. *For product distribution $F = F_1 \times \dots \times F_m$, the multi-dimensional extension of the single-dimensional ironed virtual value functions for prices \hat{t} is the vector field $\bar{\phi}^{MD}$ defined as follows:*

- (i) $\{\bar{\phi}^{MD}(t)\}_{j^*} = \bar{\phi}_{j^*}^{SD}(\{t\}_{j^*})$, and
- (ii) $\{\bar{\phi}^{MD}(t)\}_j = \{t\}_j$,

where item $j^* \in \operatorname{argmax}_j \{t - \hat{t}\}_j$ is the favorite item at prices \hat{t} , item $j \neq j^*$ ranges over all other items, and where $\bar{\phi}_{j^*}^{SD}(v)$ is the single-dimensional ironed virtual value function for distribution F_{j^*} .

Theorem 9.3.11. *The multi-dimensional extension of the single-dimensional ironed virtual value functions for prices \hat{t} is an amortization of revenue, i.e., for any agent, the expected revenue of any incentive compatible mechanism (x, p) is at most the expected amortized surplus, i.e., $\mathbf{E}_{t \sim F}[p(t)] \leq \mathbf{E}_{t \sim F}[\bar{\phi}^{MD}(t) \cdot x(t)]$.*

Proof. See Exercise 9.12. □

The proof of Theorem 9.3.10 follows from Theorem 9.3.11 and by decomposing the amortization into the sum of three terms and then upper bounding the surplus with respect to each term by the revenue of a mechanism for the representative environment. This decomposition is stated and proved in Lemma 9.3.12, below.

Lemma 9.3.12. *For a unit-demand agent with ex ante allocation constraint $\hat{q} = (\hat{q}_1, \dots, \hat{q}_m)$ that correspond to prices $\hat{t} = (\hat{t}_1, \dots, \hat{t}_m)$, the optimal revenue is at most the sum of the revenues of (a) the ex ante relaxation of the representative environment, (b) the auction that serves the representative with the highest positive value of $\{t\}_j - \{\hat{t}\}_j$, and (c) the auction that serves the representative with the highest positive value of $\{t\}_j - \{\hat{t}\}_j$ with lazy monopoly reserves.*

The difficulty of proving this lemma, in contrast to the proof of the analogous result without the ex ante constraint (Theorem 9.1.11), is that we need to compare the optimal mechanism to mechanisms that also satisfies the ex ante constraint. Specifically, in the proof of Theorem 9.1.11, it is important that the revenue of the second-highest valuation is attainable by the second-price auction (in the representative environment). With an ex ante constraint, an auction for the representative environment is not always able to sell to the agent with the highest value and, thus, the price charged may not always be boundable by the second-highest value.

Proof. Upper bound the amortization $\bar{\phi}^{SD}$ with the sum of three amortization:

$$\begin{aligned} \{\bar{\phi}^{FAVE}(t)\}_j &= \bar{\phi}_j^{SD}(\{t\}_j), & j = j^* \wedge \{t - \hat{t}\}_j \geq 0 \\ \{\bar{\phi}^{SP}(t)\}_j &= \{t - \hat{t}\}_j, & j \neq j^* \wedge \{t - \hat{t}\}_j \geq 0 \\ \bar{\phi}^{PP}(t) &= \hat{t}, \end{aligned}$$

where $j^* \in \operatorname{argmax}_j \{t - \hat{t}\}_j$ and all coordinates not explicitly specified

are set to zero. We now compare the amortized surplus from optimizing with respect to each of these amortizations to the revenue of the mechanisms listed in the theorem statement.

The amortization $\bar{\phi}^{\text{PP}}$ pointwise equals the prices \hat{t} . Given the ex ante constraint \hat{q} , the optimal amortized surplus is $\hat{t} \cdot \hat{q}$. This amortized surplus is exactly the revenue of the ex ante mechanism that posts prices \hat{t} . Mechanism (a), the optimal ex ante relaxation, achieves only higher revenue (e.g., if any of the single-dimensional ironed virtual values are negative at the prices \hat{t} then the ex ante optimal mechanism can improve on the revenue by posting an even higher price. Thus, the revenue of auction (a) upper bounds the optimal surplus from $\bar{\phi}^{\text{PP}}$.

The amortization $\bar{\phi}^{\text{SP}}$ is non-zero only on non-favorite item at prices \hat{t} and only when the value for the favorite item is at least its price. The optimal amortized surplus is precisely the revenue of the second-price auction on representative with values $v_j = \{t - \hat{t}\}_j$. The auction (b) for the representative environment that serves the representative with the highest positive value of $\{t\}_j - \{\hat{t}\}_j$, namely j^* , attains this revenue plus $\{\hat{t}\}_{j^*}$ which is only larger. Thus, the revenue of auction (b) upper bounds the optimal surplus from $\bar{\phi}^{\text{SP}}$.

The amortization $\bar{\phi}^{\text{FAVE}}$ is non-zero only on the favorite item j^* at prices \hat{t} and only when the value for this item is at least its price. The amortized surplus maximizing outcome is the same as the auction (c) for the representative environment that serves the representative with the highest positive value of $\{t\}_j - \{\hat{t}\}_j$ with lazy monopoly reserves. Moreover, an auction's revenue in the representative environment equals the its amortized surplus. Thus, the revenue of auction (c) upper bounds the optimal surplus from $\bar{\phi}^{\text{FAVE}}$. \square

Exercises

- 9.1 Even for a single unit-demand agent with values for the alternatives drawn from a symmetric distribution, there may be no symmetric optimal item pricing. Identify a distribution F for which there is no symmetric optimal item pricing for an agent values for $m = 2$ alternatives are drawn i.i.d. from F .
- 9.2 Prove Theorem 9.1.15: *For a unit-demand agent with independently, identically, and regularly distributed values, a uniform item pricing is a $2e-1/e-1 \approx 2.58$ approximation to the optimal lottery*

pricing revenue. For $k = 2$ alternatives, the bound improves to $7/3 \approx 2.33$.

- 9.3 Consider selling k items to an agent with additive values on support $[0, h]$ and show that the revenue of pricing the grand bundle approaches the optimal revenue as k approaches infinity.
- 9.4 Prove Theorem 9.2.3: *For a unit-demand agent with independent values and any interim allocation constraint \hat{y} , the revenue of the optimal auction for the representative environment is at least half the revenue of the optimal lottery pricing in the original unit-demand environment; moreover, its revenue is linear in \hat{y} .*
- 9.5 Prove Theorem 9.2.4: *For a unit-demand agent with independent values and ex ante allocation constraint \hat{q} , the revenue of a uniform virtual item pricing is a two approximation to the optimal representative revenue.*
- 9.6 Prove part (a) of Theorem 9.2.7: *For unit-demand agents in a downward-closed service-constrained environment with distributions that are independent and non-identical across agents but independent and identical across alternatives, the optimal mechanism for the favorite-alternative projection is a $2e-1/e-1 \approx 2.58$ approximation to the optimal mechanism.*
- 9.7 Consider unit-demand agents in single-item service constrained environments with type distributions that are identical across agents, regular across alternatives (but not identical), and regular for the distribution of any agent's value for her favorite alternative. Prove that the second-price auction with reserve for the winner's favorite alternative is a $2e$ approximation to the optimal mechanism and the second price auction (without a reserve) is a $2en/n-1$ approximation to the optimal mechanism.
- 9.8 Prove Theorem 9.3.1: *For unit-demand agents with independent values, the revenue of the optimal mechanism for the representative environment upper bounds the revenue for the optimal deterministic dominant-strategy incentive-compatible mechanism for the original unit-demand environment.*
- 9.9 Given an improved bound for Theorem 9.3.5 under the assumption that the value distributions are regular. Specifically, show that for single-dimensional bipartite matching environments (where agents are edges and feasible outcomes are matchings), for any independent and regular distribution of values, there is an oblivious posted pricing that is a constant approximation to the optimal revenue. The constant in your proof should be strictly less than nine.

- 9.10 Give an improved bound for Corollary 9.3.6 under the assumption that the order of the agents is random not oblivious (i.e., not worse case). Specifically, show that for unit-demand unit-supply matching environments with independently distributed values, there is a posted pricing that, for agents who arrive in a uniformly random order, is a constant approximation to the optimal revenue. The constant in your proof should be strictly less than 27.
- 9.11 Consider a single-dimensional bipartite matching environments where agents correspond to edges in a bipartite graph and feasible outcomes are matchings. For any independent and regular distribution of values, show that there is a posted pricing that is a two approximation to the optimal revenue when the agents arrive in a uniformly random order.
- 9.12 Prove Theorem 9.3.11: *The multi-dimensional extension of the single-dimensional ironed virtual value functions for prices \hat{t} is an amortization of revenue, i.e., for any agent, the expected revenue of any incentive compatible mechanism (x, p) is at most the expected amortized surplus, i.e., $\mathbf{E}_{t \sim F}[p(t)] \leq \mathbf{E}_{t \sim F}[\bar{\phi}^{MD}(t) \cdot x(t)]$.*

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).

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 two approximation via prophet inequalities that is presented in this chapter is due to Chawla et al. (2010b). 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 originally from Chawla et al. (2010a); however, the proof given here employs the framework for amortized analysis of revenue from Haghpanah and Hartline (2015) with a proof approach from Cai et al. (2016).

The study of approximation for additive agents was initiated by Armstrong (1999) who showed, for example, that pricing the grand bundle is asymptotically (in the number of items) when the agent's values for the items are independent and identically distributed. The idea that bundling can give higher revenue than pricing the items individually is related to the idea, more generally in mechanism design, of the linking of independent decisions (Jackson and Sonnenschein, 2007). The study of simple approximation mechanisms for additive agents and a small number of items was systematically considered by Hart and Nisan (2012). When the agent's values are independently distributed, Babaioff et al. (2014) showed that the better of pricing individual items and pricing the grand bundle is a constant approximation to the optimal, perhaps randomized, mechanism. This result has been extended in a number of directions including (a) subadditive valuations, (b) multiple agents, and (c) partial correlation between item values. The proof given here of the six approximation is from Cai et al. (2016).

The inapproximability of the optimal mechanism for additive and unit-demand agents when the agent's values are arbitrarily correlated was proven by Briest et al. (2010) for unit-demand agents and $m = 3$ alternatives. Hart and Nisan (2013) generalize the construction to $m = 2$ alternatives (of course for $m = 1$ alternative, the optimal mechanism is an item pricing). The construction presented in this chapter is from Hart and Nisan (2013).

The marginal revenue framework and its approximate optimality for general service constrained environments was developed by Alaei et al. (2013).

The study of sequential posted pricing mechanisms in multi-dimensional environments that is discussed in this chapter is given by Chawla et al. (2010b); these sequential posted pricings are constant approximations to the optimal deterministic mechanisms. Alaei (2011) gives a refined analysis and approach with improved approximation bounds. 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. (2010a) and Chawla and Miller (2016). Neither the bound of four (for single-agent lottery pricing) or 27 (for matching markets) is known to be tight.