4
Bayesian Approximation

One of the most intriguing conclusions from the preceding chapter is that for i.i.d. regular single-item environments the second-price auction with a reservation price is revenue optimal. This result is compelling as the solution it proposes is quite simple, therefore, making it easy to prescribe. Furthermore, reserve-price-based auctions are often employed in practice so this theory of optimal auctions is also descriptive. Unfortunately, i.i.d. regular single-item environments are hardly representative of the scenarios in which we would like to design good mechanisms. Furthermore, if any of the assumptions are relaxed, reserve-price-based mechanisms are not optimal.

Another point of contention is that auctions, even simple ones like the second-price auction, can be a slow and inconvenient way to allocate resources. In many contexts posted pricings are preferred to auctions. As we have seen, posted pricings are not optimal unless there is only a single consumer. In addition to being preferred for their speed and simplicity, posted pricings also offer robustness to out-of-model phenomena such as collusion. Therefore, approximation results for posted pricings imply that good collusion resistant mechanisms exist.

In this chapter we address these deficiencies by showing that while posted pricings and reserve-price-based mechanisms are not generally optimal, they are approximately optimal in a wide range of environments. Furthermore, these approximately optimal mechanisms are more robust, less dependent on the details of the distribution, and sometimes provide more conceptual understanding than their optimal counterparts. The approximation factor obtained by most of these approximation mechanisms is two. Meaning, for the worst distributional assumptions, the mechanism's expected performance is within a factor two of the optimal mechanism. Of course, in any particular environment these mechanisms may perform better than this worst-case guarantee.

A number of properties of the environment will be crucial for enabling good approximation mechanisms. As in Chapter 3 these are: independence of the distribution of preferences for the agents, distributional regularity as implied by the concavity of the price-posting revenue curve, and downward closure of the designer's feasibility constraint. In addition, two new structural restrictions on the environment will be introduced.

A *matroid set system* is one that is downward closed and satisfies an additional "augmentation property." An important characterization of the matroid property is that the surplus maximizing allocation (subject to feasibility) is given by the *greedy-by-value* algorithm: sort the agents by value, then consider each agent in-turn and serve the agent if doing so is feasible given the set of agents already being served. The optimality of greedy-by-value implies that the order of the agents' values is important for finding the surplus maximizing outcome, but the relative magnitudes of their values are not.

The *monotone hazard rate* condition is a refinement of the regularity property of a distribution of values. Intuitively, the monotone hazard rate condition restricts how heavy the tail of the distribution is, i.e., how much probability mass is on very high values. An important consequence of the monotone hazard rate assumption is that the optimal revenue and optimal social surplus are within a factor of $e \approx 2.718$ of each other. This will enable mechanism that optimize social surplus to give good approximations to revenue.

Mathematical Note Subsequently we will consider using monopoly reserve prices for distributions where these prices are not unique. For these distributions we should always assume the worst tie-breaking rule as it is always possible to perturb the distribution slightly to make that worst monopoly price unique. For example, recall that a regular distribution can be equivalently specified by its distribution function or its revenue curve. For instance the *equal revenue distribution* has constant revenue curve, $R^{\mathrm{EQR}}(q) = 1$, and therefore any price on $[1, \infty)$ is optimal. A sufficient perturbation to make the price of one the unique monopoly price is given by revenue curve $R(q) = 1 - \epsilon(1 - q)$ which is uniquely maximized at monopoly quantile $\hat{q}^\star = 1$ with corresponding monopoly price $\hat{v}^\star = V(1) = R(1) = 1$.

In the previous two chapters, with the characterization of Bayes-Nash equilibrium (Theorem 2.2) and the characterization of profit-optimal mechanisms (Corollary 3.21), we assumed that the values of the agents were drawn from continuous distributions. In this chapter, especially when describing examples that show that the assumptions of a theorem are necessary, it will sometimes more expedient to work with discrete distributions. A discrete distribution is specified by a set of values and probabilities for these values.

There are two ways to relate these discrete examples to the continuous environments we have heretofore been considering. First, we could rederive Theorem 2.2 and Corollary 3.21 (and their variants) for discrete distributions (see Exercise 2.2 and Exercise 3.6, respectively). Importantly, via such a rederivation, it is apparent that discrete and continuous environments are intuitively similar. Second, we could consider a continuous perturbation of the discrete distribution which will exhibit the same phenomena with respect to optimization and approximation. For example, one such perturbation is, for a sufficiently small ϵ, to replace any value v from the discrete distribution with a uniform value from $[v, v + \epsilon]$.

4.1 Monopoly Reserve Pricing

We start our discussion of simple mechanisms that are approximately optimal by showing that a natural generalization of the second-price auction with monopoly reserve continues to be approximately optimal for regular but asymmetric distributions. Recall that monopoly prices are a property of virtual value functions which are a property of the distributions from which agents' values are drawn (Definition 3.7). When the agents' values are drawn from distinct distributions their monopoly prices are generally distinct. The following definition generalizes the second-price auction with a single reserve price to one with *discriminatory*, i.e., agent-specific, reserve prices.

Definition 4.1 The *second-price auction with (discriminatory) reserves* $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$ is:

(i) reject each agent i with $v_i < \hat{v}_i$,
(ii) allocate the item to the highest valued agent remaining (or none if none exists), and
(iii) charge the winner her critical price.

With non-identical distributions the optimal single-item auction indeed needs the exact marginal revenue functions to determine the optimal allocation (see Example 4.1). This contrasts to the i.i.d. regular case where all we needed was a single number, the monopoly price for the distribution, and reserve pricing with this number is optimal. Figure 4.1 compares allocations of the (asymmetric) optimal auction with those of the second-price auction with (asymmetric) monopoly reserves.

Example 4.1 Consider a two-agent single-item auction where agent 1 (Alice) and agent 2 (Bob) have values distributed uniformly on $[0, 2]$ and $[0, 3]$, respectively. The virtual value functions are $\phi_1(v_1) = 2v_1 - 2$ and $\phi_2(v_2) = 2v_2 - 3$. Alice's monopoly price one; Bob's monopoly price is $3/2$. Alice has a higher virtual value than Bob when $v_1 > v_2 - 1/2$. The optimal auction is asymmetric. It serves an agent only if one is above their respective monopoly price. If both are above their respective monopoly reserves, it serves the highest valued agent with a penalty of $1/2$ against Bob (cf. Example 3.11, page 67). In contrast the monopoly-reserves auction is the same but with no penalty for Bob. See Figure 4.1.

In the remainder of this section we show that if the agents' values are drawn from regular distributions then the (single item) monopoly-reserves auction is a two approximation to the optimal revenue. We will then show that, except for the consideration of more general feasibility constraints, this result is tight. The approximation bound of two is tight: we show by example that there is a non-identical regular distribution where the ratio of the optimal to monopoly-reserves revenue is two. The regularity assumption is tight: for irregular distributions the approximation ratio of monopoly reserves can be as bad as linear (i.e., it grows with the number of agents). Thus, we conclude that this two-approximation result for regular distributions in single-item environments is essentially the right answer. Later in the chapter we will consider the extent to which this result generalizes beyond single-item environments.

4.1.1 Approximation for Regular Distributions

The main result of this section shows that, though distinct, the monopoly-reserves auction and the revenue-optimal auction have similar revenues.

Theorem 4.2 *For single-item environments and agents with values drawn independently from (non-identical) regular distributions, the second-price auction with (asymmetric) monopoly reserve prices obtains at least half the revenue of the (asymmetric) optimal auction.*

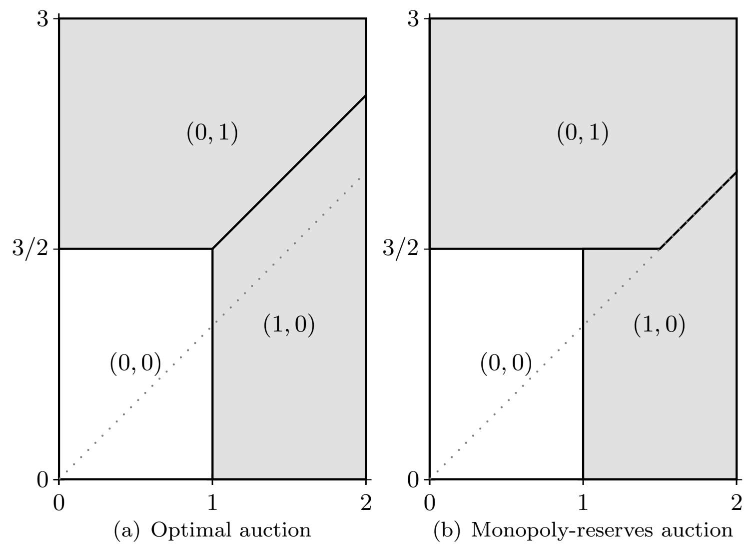

Figure 4.1 In Example 4.1 Agent 1 has value $v_1 \sim U[0, 2]$; agent 2 has value $v_2 \sim U[0, 3]$. In the space of valuation profiles $\boldsymbol{v} \in [0, 2] \times [0, 3]$, with agent 1's value on the horizontal axis and agent 2's value on the vertical axis, the allocation $\boldsymbol{x} = (x_1, x_2)$ for the (asymmetric) optimal auction and (asymmetric) monopoly-reserves auction are depicted.

The proof of Theorem 4.2 is enabled by the following three properties of regular distributions and virtual value functions. First, Corollary 3.27 shows that for a regular distribution, a monotone allocation rule, and virtual value given by the marginal revenue curve, the expected revenue is equal to the expected virtual surplus. The second and third properties are given by the two lemmas below.

Lemma 4.3 *For any virtual value function, the virtual values corresponding to values that exceed the monopoly price are non-negative.*

Proof The lemma follows immediately from the definition of virtual value functions which requires their monotonicity (Definition 3.6) . □

Lemma 4.4 *For any distribution, the value of an agent is at least her virtual value for revenue.*

Proof We prove the lemma for regular distributions (as is necessary for Theorem 4.2) and leave the general proof to Exercise 4.3. For regular distributions, where the virtual values for revenue are given by the formula $\phi(v) = v - \frac{1-F(v)}{f(v)}$, the lemma follows as both $1 - F(v)$ and $f(v)$ are non-negative. □

Our goal will be to show that the expected revenue of the monopoly-reserves auction is approximately an upper bound on the expected virtual surplus of the optimal auction (which is equal to its revenue). Consider running both auctions on the same random input. Notice that conditioned on the event that both auctions serve the same agent, both auctions obtain the same (conditional) expected virtual surplus. Notice also that conditioned on the event that the auctions serve distinct agents, the monopoly-reserves auction has higher expected payments than the optimal auction. It is not correct to bound revenue by combining conditional virtual values with conditional payments as the amortized analysis that defines virtual values is only correct under unconditional expectations. Therefore, for the second case we will instead relate the payment of monopoly reserves to the virtual value of the winner in the optimal auction (for which it gives an upper bound).

Proof of Theorem 4.2 Let REF denote the optimal auction and its expected revenue and APX denote the second-price auction with monopoly reserves and its expected revenue. Clearly, REF $\geq$ APX; our goal is to give an approximate inequality in the opposite direction by showing that $2\,\mathrm{APX} \geq \mathrm{REF}$. Let I be the winner of the optimal auction and J be the winner of the monopoly reserves auction. I and J are random variables. Notice that neither auctions sell the item if and only if all virtual values are negative; in this situation define $I = J = 0$. With these definitions and Corollary 3.27, $\mathrm{REF} = \mathbf{E}[\phi_I(v_I)]$ and $\mathrm{APX} = \mathbf{E}[\phi_J(v_J)]$.

We start by simply writing out the expected revenue of the optimal auction as its expected virtual surplus conditioned on $I = J$ and $I \neq J$.

$$\mathrm{REF} = \underbrace{\mathbf{E}[\phi_I(v_I) \mid I = J]\,\mathbf{Pr}[I = J]}_{\mathrm{REF}_=} + \underbrace{\mathbf{E}[\phi_I(v_I) \mid I \neq J]\,\mathbf{Pr}[I \neq J]}_{\mathrm{REF}_{\neq}}.$$

We will prove the theorem by showing that both the terms on the right-hand side are bounded from above by APX. Thus, $\mathrm{REF} \leq 2\,\mathrm{APX}$. For the first term:

$$\begin{aligned}
\mathrm{REF}_= &= \mathbf{E}[\phi_I(v_I) \mid I = J]\,\mathbf{Pr}[I = J] \\
&= \mathbf{E}[\phi_J(v_J) \mid I = J]\,\mathbf{Pr}[I = J] \\
&\leq \mathbf{E}[\phi_J(v_J) \mid I = J]\,\mathbf{Pr}[I = J] + \mathbf{E}[\phi_J(v_J) \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&= \mathrm{APX}.
\end{aligned}$$

The inequality in the above calculation follows from Lemma 4.3 as even when $I \neq J$ the virtual value of J must be non-negative. Therefore, the term added is non-negative. For the second term:

$$\begin{aligned}
\mathrm{REF}_{\neq} &= \mathbf{E}[\phi_I(v_I) \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[v_I \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[p_J(\boldsymbol{v}) \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[p_J(\boldsymbol{v}) \mid I \neq J]\,\mathbf{Pr}[I \neq J] + \mathbf{E}[p_J(\boldsymbol{v}) \mid I = J]\,\mathbf{Pr}[I = J] \\
&= \mathrm{APX}\,.
\end{aligned}$$

The first inequality in the above calculation follow from values upper bounding virtual values (Lemma 4.4). The second inequality follows because, among agents who meet their reserve, J is the highest valued agent and I is a lower valued agent. Therefore, as APX is a second-price auction, the winner J's payment is at least the loser I's value. The third inequality follows because payments are non-negative so the term added is non-negative. $\square$

Theorem 4.2 shows that when agent values are non-identically distributed at least half of the revenue of the optimal asymmetric auction which is parameterized by complicated virtual value functions can be obtained by a simple auction which is parameterized by natural statistical quantities, namely, each distribution's monopoly price. The theorem holds for a broad class of distributions that satisfy the regularity property. While for specific distributions the approximation bound may be better than two, we will see subsequently, by example, that if the only assumption on the distribution is regularity then the approximation factor of two is tight.

Definition 4.2 The *equal-revenue distribution* has distribution function $F^{\mathrm{EQR}}(z) = 1 - {}^1\!/\!_z$ and density function $f^{\mathrm{EQR}}(z) = {}^1\!/\!_{z^2}$ on support $[1, \infty)$.

The equal-revenue distribution is so called because the revenue obtained from posting any price is the same. Consider posting price $\hat{v} > 1$. The expected revenue from such a price is $\hat{v} \cdot (1 - F^{\mathrm{EQR}}(\hat{v})) = 1$. As the price-posting revenue curve is the constant function $P^{\mathrm{EQR}}(\hat{q}) = 1$, the distribution is on the boundary between regularity and irregularity. As it is the boundary between regularity and irregularity, it often provides an extremal example for results that hold for regular distributions.

Lemma 4.5 *There is an (non-identical) regular two-agent single-item* *environment where the optimal auction obtains twice the revenue of the second-price auction with (discriminatory) monopoly reserves.*

Proof For any $\epsilon > 0$ we will give a distribution and show that there is an auction with expected revenue strictly greater than $2 - \epsilon$ but the revenue of the monopoly reserves auction is precisely one.

Consider the asymmetric two-agent single-item environment where agent 1 (Alice) has value (deterministically) one and agent 2 (Bob) has value distributed according to the equal-revenue distribution. The monopoly price for the equal-revenue distribution is ill-defined because every price is optimal, but a slight perturbation of the distribution has a unique monopoly price of $\hat{v}_2^\star = 1$ (see Mathematical Note on page 101). Thus the monopoly prices are $\hat{\boldsymbol{v}}^\star = (1, 1)$ and the expected revenue of the second-price auction with monopoly reserves is one.

Of course, for this distribution it is easy to see how we can do much better. Offer Bob a high price h. If he rejects this price then offer Alice a price of 1. Notice that by the definition of the equal-revenue distribution, Bob's expected payment is one, but still Bob rejects the offer with probability $1 - 1/h$ and the item can be sold to Alice. The expected revenue of the mechanism is $h \cdot 1/h + 1 \cdot (1 - 1/h) = 2 - 1/h$. Choosing $h > 1/\epsilon$ gives the claimed result. $\square$

While the monopoly-reserves auction (parameterized by n monopoly prices) is significantly less complex than the optimal auction (parameterized by n virtual value functions), it is not often used in practice. In practice, even in asymmetric environments, auctions are often parameterized by a single *anonymous* reserve price. For regular, non-identical distributions anonymous reserve pricing continues to give a good approximation to the optimal auction. This and related results are discussed in Section 4.4.

4.1.2 Inapproximability Irregular Distributions

The second-price auction with monopoly reserve prices only guarantees a two approximation for regular distributions. The proof of Theorem 4.2 relied on regularity crucially when it invoked Corollary 3.27 to calculate revenue in terms of virtual surplus for all monotone allocation rules. Recall that for irregular distributions, revenue is only equal to virtual surplus for allocation rules that are constant where the virtual value functions are constant. For irregular distributions there are two challenges for that the monopoly-reserves auction must confront. First, even

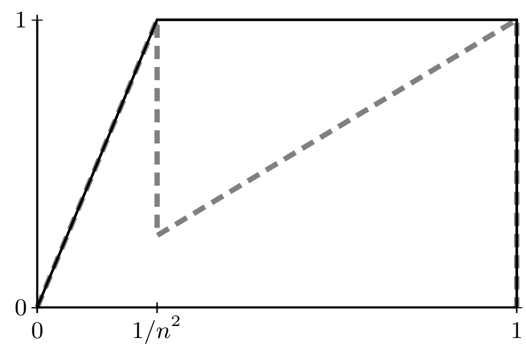

Figure 4.2 The revenue curve (thin, solid, black) and price-posting revenue curve (gray, thick, dashed) for the discrete two-point equal revenue distribution from the proof of Proposition 4.6 with $h = 2$. As usual for revenue curves, the horizontal axis is quantile.

if the distributions are identical, the optimal auction is not the second-price auction with monopoly reserves; it irons (see Section 3.3.3). Second, the distributions may not be identical. We show here that even for i.i.d. irregular distributions this trivial bound cannot be improved (Proposition 4.6), and that this lower bound is tight as the monopoly-reserves auction for (non-identical) irregular distributions is, trivially, an n approximation (Proposition 4.7).

Of course, irregular distributions that are "nearly regular" do not exhibit the above worst case behavior. For example, Exercise 4.6 formalizes a notion of near regularity under which reasonable approximation bounds can be proven.

Proposition 4.6 *For (irregular) i.i.d. n-agent single-item environments, the second-price auction with monopoly reserve is at best an n approximation.*

Proof Consider the discrete equal-revenue distribution on $\{1, h\}$, i.e., with $f(h) = 1/h$ and $f(1) = 1 - 1/h$, slightly perturbed so that the monopoly price is one (see Mathematical Note on page 101). With a monopoly reserve of $\hat{v}^\star = 1$ and all values at least one, the reserve is irrelevant for the second-price auction.

Consider the expected revenues of the second-price auction $\text{APX}(h)$ and the optimal auction $\text{REF}(h)$ as a function of h. We show the following limit result which implies the proposition.

$$\mathrm{APX} = \lim_{h\to\infty} \mathrm{APX}(h) = 1, \text{and} \tag{4.1}$$

$$\mathrm{REF} = \lim_{h\to\infty} \mathrm{REF}(h) = n. \tag{4.2}$$

An agent is high-valued with probability $1/h$ and low valued with probability $(1 - 1/h)$. The probability that there are exactly k high valued agents is:

$$\mathbf{Pr}[\text{exactly } k \text{ are high valued}] = \binom{n}{k} \cdot h^{-k} \cdot (1 - 1/h)^{n-k}.$$

For constant n and k and in the limit as h goes to infinity, the first term is constant and the last term is one. The middle term goes to zero at a rate of h^{-k}. Thus,

$$\lim_{h\to\infty} h^k \cdot \mathbf{Pr}[\text{exactly } k \text{ are high valued}] = \binom{n}{k}, \text{ and} \tag{4.3}$$

$$\lim_{h\to\infty} h^k \cdot \mathbf{Pr}[\text{at least } k \text{ are high valued}] = \binom{n}{k}. \tag{4.4}$$

For the discrete equal-revenue distribution, $\phi(1) = 0$ and $\phi(h) = h$ (see Figure 4.2 and Exercise 3.6). Now we can calculate $\mathrm{REF} = \lim_{h\to\infty} \mathrm{REF}(h)$ as $\phi(1)$ times the probability that there are no high-valued agents plus $\phi(h)$ times the probability that there are one or more high-valued agents. $\mathrm{REF} = 0 + \binom{n}{1} = n$.

We can similarly calculate $\mathrm{APX} = \lim_{h\to\infty} \mathrm{APX}(h)$ as one times the probability that there are one or fewer high-valued agents plus h times the probability that there are two or more high-valued agents. By equation (4.3) with $k = 0$ and 1, the first term is one; by equation (4.4) with $k = 2$, the second term is zero. Thus, $\mathrm{APX} = 1$. □

Proposition 4.7 *For (non-identical, irregular) n-agent single-item environments, the second-price auction with monopoly reserve is at worst an n approximation.*

Proof Let REF and APX and denote the monopoly-reserve auction and the optimal auction and their revenue, respectively, in an n-agent, single-item environment.

As usual for approximation bounds when the optimal mechanism REF is complex, we will formulate an upper bound that is simple. Denote by UB the optimal auction and its revenue for the n-agent, n-unit environment (a.k.a. a digital good). Clearly, $\mathrm{UB} \geq \mathrm{REF}$ as this auction could discard all but one unit and then simulate the outcome REF (the optimal single-unit auction). UB is also very simple. As there are n units and n

agents there is no competition between the agents and the optimization problem decomposes into n independent monopoly pricing problems. Denote by $\boldsymbol{R}^\star = (R_0^\star, \ldots, R_n^\star)$ the profile of monopoly revenues. The revenue of the optimal n-unit auction is:

$$\mathrm{UB} = \sum\nolimits_i R_i^\star.$$

We now get a lower bound on the monopoly-reserves revenue APX. Consider the mechanism LB that chooses, before asking for agent reports, the agent $i^\star$ with the highest monopoly revenue and offers this agent her monopoly price $\hat{v}_{i^\star}^\star$. LB obtains revenue

$$\mathrm{LB} = \max\nolimits_i R_i^\star.$$

Moreover, $\mathrm{APX} \geq \mathrm{LB}$ as if $i^\star$ would accept price her monopoly price $\hat{v}_{i^\star}^\star$ then some agent in APX must accept a price of at least $\hat{v}_{i^\star}^\star$ (either agent $i^\star$ or an agent beating out agent $i^\star$).

Finally, we make the simple observation that $n{\cdot}\mathrm{LB} \geq \mathrm{UB}$ which proves the proposition. □

4.2 Oblivious Posted Pricings and the Prophet Inequality

Two problematic aspects of employing auctions to allocate resources is that (a) they require multiple rounds of communication (i.e., they are slow) and (b) they require all agents to be present at the time of the auction. Often both of these requirements are prohibitive. In routing in computer networks a packet needs to be routed, or not, quickly and, if the network is like the Internet, without state in the routers. Therefore, auctions are unrealistic for congestion control. In a supermarket where you go to buy lettuce, we should not hope to have all the lettuce buyers in the store at once. Finally, in selling goods on the Internet, eBay has found empirically that posted pricing via the "buy it now" option is more appropriate than a slow (days or weeks) ascending auction.

Posted pricings give very robust revenue guarantees. For instance, their revenue guarantees are impervious to many kinds of collusive behavior on the part of the agents. Moreover, the prices (to be posted) can also be used as reserve prices for the first- and second-price auctions and this only improves on the revenue from price posting.

In a posted pricing, distinct prices can be posted to the agents with first-come-first-served and while-supplies-last semantics. In this section we show that *oblivious posted pricing*, where agents arrive and consider their respective prices in any arbitrary order, gives a two approximation to the optimal auction. In the next section, we show that *sequential posted pricing*, where the mechanism chooses the order in which the agents are permitted to consider their respective posted prices, gives an improved approximation of $e/e-1 \approx 1.58$. Both results hold for objectives of revenue and social surplus and for any independent distribution on agent values (i.e., regularity is not assumed).

There are several challenges to the design and analysis of oblivious posted pricings. First, for any particular n-agent scenario, an oblivious posted pricing potentially requires optimization of n distinct prices. In high dimensions (i.e., large n) this optimization problem is computationally challenging. Moreover, it is not immediately clear that the resulting optimal prices would perform well in comparison to the optimal auction. To justify usage of posted pricings over auctions, we must be able to easily find good prices and these prices should give revenue that compares favorably to that of the optimal auction. The approach of this section is to solve both problems at once by identifying a class of easy-to-find posted pricings that perform well.

4.2.1 The Prophet Inequality

The oblivious posted pricing theorem we present is an application of a *prophet inequality* theorem from optimal stopping theory. Consider the following scenario. A gambler faces a series of n games, one on each of n days. Game i has prize v_i distributed independently according to distribution F_i. The order of the games and distribution of the prize values is fully known in advance to the gambler. On day i the gambler *realizes* the prize $v_i \sim F_i$ of game i and must decide whether to keep this prize and *stop* or to return the prize and *continue* playing. In other words, the gambler is only allowed to keep one prize and must decide whether or not to keep a given prize immediately on realizing the prize and before any future prizes are realized.

The gambler's optimal strategy can be calculated by *backwards induction*. On day n the gambler should stop with whatever prize is realized. This results in expected value $\mathbf{E}[v_n]$. On day $n-1$ the gambler should stop if the prize has greater value than $\hat{v}_{n-1} = \mathbf{E}[v_n]$, the expected value of the prize from the last day. On day $n-2$ the gambler should stop with if the prize has greater value than $\hat{v}_{n-2}$, the expected value of the strategy for the last two days. Proceeding in this manner the gambler can calculate a threshold $\hat{v}_i$ for each day where the optimal strategy is to stop with prize i if and only if $v_i \geq \hat{v}_i$.

This optimal strategy suffers from many of the drawbacks of optimal strategies. It is complicated: it takes n numbers to describe it. It is sensitive to small changes in the game, e.g., changing of the order of the games or making small changes to distribution i strictly above $\hat{v}_i$. It does not allow for much intuitive understanding of the properties of good strategies. Finally, it does not generalize well to give solutions to other similar kinds of games, e.g., that of our oblivious posted pricing problem.

Approximation gives a crisper picture. A *uniform threshold strategy* is given by a single threshold $\hat{v}$ and requires the gambler to accept the first prize i with $v_i \geq \hat{v}$. Threshold strategies are clearly suboptimal as even on day n if prize $v_n < \hat{v}$ the gambler will not stop and will, therefore, receive no prize. We refer to the prize selection procedure when multiple prizes are above the threshold as the *tie-breaking rule*. The tie-breaking rule implicit in the specification of the gambler's game is lexicographical, i.e., by "smallest i."

Theorem 4.8 *For any product distribution on prize values $\boldsymbol{F} = F_1 \times \cdots \times F_n$, there exists a uniform threshold strategy such that the expected prize of the gambler is at least half the expected value of the maximum prize; moreover, the bound is invariant with respect to the tie-breaking rule; moreover, for continuous distributions with non-negative support one such threshold strategy is the one where the probability that the gambler receives no prize is exactly $1/2$.*

Theorem 4.8 is a *prophet inequality*: it suggest that even though the gambler does not know the realizations of the prizes in advance, she can still do half as well as a prophet who does. While this result implies that the optimal (backwards induction) strategy satisfies the same performance guarantee, this guarantee was not at all clear from the original formulation of the optimal strategy.

Unlike the optimal (backwards induction) strategy this prophet inequality provides substantial conclusions. Most obviously, it is a very simple strategy. The result is clearly driven by trading off the probability of not stopping and receiving no prize with the probability of stopping early with a suboptimal prize. Notice that the order of the games makes no difference in the determination of the threshold, and if the distribution above or below the threshold changes, neither the bound nor suggested strategy is affected. Moreover, the invariance of the theorem to the tie-breaking rule suggests the bound can be applied to other related scenarios. The profit inequality is quite robust.

Proof of Theorem 4.8 Let REF denote prophet and her expected prize, i.e., the expected maximum prize, $\mathbf{E}[\max_i v_i]$, and APX denote a gambler with threshold strategy $\hat{v}$ and her expected prize. Define $\hat{q}_i = 1-F_i(\hat{v}) = \mathbf{Pr}[v_i \geq \hat{v}]$ as the probability that prize i is above the threshold $\hat{v}$ and $\chi = \prod_i(1-\hat{q}_i)$ as the probability that the gambler rejects all prizes. The proof follows in three steps. In terms of the threshold $\hat{v}$ and failure probability χ, we get an upper bound on the expected prophet's payoff. Likewise, we get a lower bound on expected gambler's payoff. Finally, we choose $\hat{v}$ so that $\chi = 1/2$ to obtain the bound. If there is no $\hat{v}$ with $\chi = 1/2$, which is possible if the distributions $\boldsymbol{F}$ are not continuous, we give a slightly more sophisticated method for choosing $\hat{v}$.

In the analysis below, the notation "$(v_i - \hat{v})^+$" is shorthand for "$\max(v_i - \hat{v}, 0)$." The prophet is allowed not to pick any prize, e.g., if all prizes have negative value, to denote this outcome we add a prize indexed 0 with value deterministically $v_0 = 0$; all summations are over prizes $i \in \{0, \ldots, n\}$.

(i) An upper bound on $\mathrm{REF} = \mathbf{E}[\max_i v_i]$:

The prophet's expected payoff is

$$\begin{aligned}\mathrm{REF} = \mathbf{E}[\max_i v_i] &= \hat{v} + \mathbf{E}[\max_i (v_i - \hat{v})] \\ &\leq \hat{v} + \mathbf{E}[\max_i\,(v_i - \hat{v})^+] \\ &\leq \hat{v} + \sum\nolimits_i \mathbf{E}[(v_i - \hat{v})^+]\,. \end{aligned} \tag{4.5}$$

The last inequality follows because $(v_i - \hat{v})^+$ is non-negative.

(ii) A lower bound on $\mathrm{APX} = \mathbf{E}[\text{prize of gambler with threshold } \hat{v}]$:

We will split the gambler's payoff into two parts, the contribution from the first $\hat{v}$ units of the prize and the contribution, when prize i is selected, from the remaining $v_i - \hat{v}$ units of the prize. The first part is $\mathrm{APX}_1 = (1-\chi)\cdot\hat{v}$. To get a lower bound on the second part we consider only the contribution from the no-tie case. For any i, let $\mathcal{E}_i$ be the event that all other prizes j are below the threshold $\hat{v}$ (but v_i is unconstrained). The bound is:

$$\begin{aligned}\mathrm{APX}_2 &\geq \sum\nolimits_i \mathbf{E}[(v_i - \hat{v})^+ \mid \mathcal{E}_i]\,\mathbf{Pr}[\mathcal{E}_i] \\ &\geq \chi \cdot \sum\nolimits_i \mathbf{E}[(v_i - \hat{v})^+]\,.\end{aligned}$$

The second line follows because $\mathbf{Pr}[\mathcal{E}_i] = \prod_{j\neq i}(1-\hat{q}_j) \geq \prod_j(1-\hat{q}_j) = \chi$ and because the conditioned variable $(v_i - \hat{v})^+$ is independent from the conditioning event $\mathcal{E}_i$. Therefore, the gambler's payoff is at least:

$$\text{APX} \geq (1-\chi)\cdot\hat{v} + \chi \cdot \sum\nolimits_i \mathbf{E}[(v_i - \hat{v})^+]\,. \tag{4.6}$$

(iii) Plug in $\hat{v}$ with $\chi = 1/2$:

From the upper and lower bounds of equations (4.5) and (4.6), if there is a non-negative $\hat{v}$ such that $\chi = 1/2$ then, for this $\hat{v}$, APX $\geq$ REF $/2$.

For discontinuous distributions, e.g., ones with point masses, χ as a function of $\hat{v}$, denoted $\chi(\hat{v})$, may be discontinuous. Therefore, there may be no $\hat{v}$ with $\chi(\hat{v}) = 1/2$. For distributions that have negative values in their supports the $\hat{v}$ with $\chi(\hat{v}) = 1/2$ may be negative. For these cases there is another method for finding a suitable threshold $\hat{v}$. Observe that the two common terms of equations (4.5) and (4.6), namely $\hat{v}$ and $\sum_i \mathbf{E}[(v_i - \hat{v})^+]$ are continuous functions of $\hat{v}$. The former is strictly increasing from $\hat{v} = 0$, the latter strictly decreases to zero; therefore they must cross at some non-negative $\hat{v}^\dagger$. For $\hat{v}^\dagger$ satisfying $\hat{v}^\dagger = \sum_i \mathbf{E}[(v_i - \hat{v}^\dagger)^+]$, regardless of the corresponding $\chi \in [0,1]$, the right-hand side of equation (4.5) is exactly twice that of equation (4.6). For this $\hat{v}^\dagger$ the two-approximation bound holds. □

The prophet inequality is tight in the sense that a better approximation bound cannot generally by obtained by a uniform threshold strategy (Exercise 4.9).

As alluded to above, the invariance to the tie-breaking rule implies that the prophet inequality gives approximation bounds in scenarios similar to the gambler's game. In an oblivious posted pricing agents arrive in a worst-case order and the first agent who desires to buy the item at her offered price does so. We now use the prophet inequality to show that there is are *oblivious posted pricings* that guarantee half the optimal surplus and half the optimal auction revenue, respectively.

4.2.2 Oblivious Posted Pricing

Consider attempting to allocate a resource to maximize the social surplus. We know from Corollary 1.4 that the second-price auction obtains the optimal surplus of $\max_i v_i$. Suppose we wish to instead us a simpler posted pricing mechanism. A uniform posted price corresponds to a uniform threshold in value space. In worst case arrival order, the agent with the lowest value above the posted price is the one who buys. This corresponds to a game like the gambler's with tie-breaking by smallest value v_i. The invariance of the prophet inequality to the tie-breaking rule allows the conclusion that posting an uniform (a.k.a. anonymous) price gives a two-approximation to the optimal social surplus.

Proposition 4.9 *In single-item environments there is an anonymous pricing whose expected social surplus under any order of agent arrival is at least half of that of the optimal social surplus.*

Not consider the objective of revenue. The revenue-optimal single-item auction select the winner with the highest (positive) virtual value (for revenue). To draw a connection between the auction problem and the gambler's problem, we note that the gambler's problem in prize space is similar to the auctioneer's problem in virtual-value space (with virtual value functions given by the marginal revenue curves of the agents' distributions). The gambler aims to maximize expected prize while the auctioneer aims to maximize expected virtual value. A uniform threshold in the gambler's prize space corresponds to a *uniform virtual price* in virtual-value space. Note, however, in value space uniform virtual prices correspond to non-uniform (a.k.a., discriminatory) prices.

Definition 4.3 A *virtual price* $\hat{\phi}$ corresponds to *uniform virtual pricing* $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$ satisfying $\phi_i(\hat{v}_i) = \hat{\phi}$ for all i.

Now compare uniform virtual pricing to the gambler's threshold strategy in the stopping game. The difference is the tie-breaking rule. For uniform virtual pricing, we obtain the worst revenue when the agents arrive in order of increasing price (in value space). Thus, the uniform virtual pricing revenue implicitly breaks ties by smallest posted price $\hat{v}_i$. The gambler's threshold strategy breaks ties by the ordering assumption on the games (i.e., lexicographically by smallest i). Recall, though, that irrespective of the tie-breaking rule the bound of the prophet inequality holds.

Theorem 4.10 *In single-item environments there is a uniform virtual pricing (for virtual values equal to marginal revenues) whose expected revenue under any order of agent arrival is at least half of that of the optimal auction.*

Proof A uniform virtual price $\hat{\phi}$ corresponds to non-uniform prices (in value space) $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$. The outcome of such a posted pricing, for the worst-case arrival order of agents, is as follows. When there is only one agent i with value v_i that exceeds her offered price $\hat{v}_i$, the revenue is precisely $\hat{v}_i$. When there are multiple agents S whose values exceed their offered prices, the one with the lowest price arrives first and pays her offered price of $\min_{i \in S} \hat{v}_i$. In other words, with respect to the gambler's game, the tie-breaking rule is by smallest $\hat{v}_i$.

To derive a bound on the revenue of is uniform virtual pricing with the worst-case arrival order we will relate its revenue to its virtual surplus. For the aforementioned outcome of a uniform virtual pricing (with virtual values as the marginal revenue) satisfies the conditions of Theorem 3.18. In particular, the induced allocation rule for each agent is constant wherever the marginal revenue is constant. Therefore, the expected revenue of a uniform virtual pricing is equal to its expected virtual surplus.

By the prophet inequality (Theorem 4.8), there is a uniform virtual price that obtains a virtual surplus of at least half the maximum virtual value (i.e., the optimal virtual surplus for single-item environments). Thus, the revenue of the corresponding price posting is at least half the optimal revenue. $\square$

In Chapter 1 we saw that that an anonymous posted pricing can be a $e/e-1 \approx 1.58$ approximation to the optimal mechanism for social surplus for i.i.d. distributions (Theorem 1.5). This approximation factor also holds for revenue and i.i.d., regular distributions. In the next section we will give a more general result that shows that if the mechanism is allowed to order the agents (i.e., in the best-case order instead of the worst-case order as above) then this better $e/e-1$ bound can be had even for asymmetric distributions. In this context of best-case versus worst-case order, the i.i.d. special case is precisely the one where symmetry renders the ordering of agents irrelevant.

4.3 Sequential Posted Pricings and Correlation Gap

In this section we consider sequential posted pricings, i.e., where the mechanism posts prices to the agents in an order that it specifies. See Section 4.2 for additional motivation for posted pricings.

One of the main challenges in designing and analyzing simple approximation mechanisms is that the optimal mechanism is complex and, therefore, difficult to analyze. For single-item auctions, this complexity arises from virtual values which come from arbitrary monotone functions. The main approach for confronting this complexity is to derive a simple upper bound on the optimal auction and then exploit the structure suggested by this bound to construct an simple approximation mechanism.

4.3.1 The Ex Ante Relaxation

One method for obtaining a simple upper bound for an optimization problem is to relax some of the constraints in the problem. For example, ex post feasibility for a single-item auction requires that, in the outcome selected by the auction, at most a single agent is served. In other words, the feasibility constraint binds ex post. For Bayesian mechanism design problems, we can relax the feasibility constraint to bind ex ante. The corresponding ex ante constraint for a single-item environment is that the expected (over randomization in the mechanism and the agent types) number of agents served is at most one.

Definition 4.4 The *ex ante relaxation* of mechanism design problem is the optimization problem with the ex post feasibility constraint replaced with a constraint that holds in expectation over randomization of the mechanism and the agents' types. The solution to the ex ante relaxation is the *optimal ex ante mechanism.*

Proposition 4.11 *The optimal ex ante mechanism's performance upper bounds the optimal (ex post) mechanism's performance.*

To see what the optimal ex ante mechanism is, consider any mechanism and denote by $\hat{\boldsymbol{q}} = (\hat{q}_1, \ldots, \hat{q}_n)$ the ex ante probabilities that each of the agents is served by this mechanism. By linearity of expectation the expected number of agents served is $\sum_i \hat{q}_i$. For a single-item environment the ex ante feasibility constraint then requires that $\sum_i \hat{q}_i \leq 1$. Notice that as far as the ex ante constraint is concerned, the agents only impose externalities on each other via their ex ante allocation probability. If we fix attention to mechanisms for which agent i is allocated with ex ante probability $\hat{q}_i$ then the remaining allocation probability for the other agents is fixed to at most $1 - \hat{q}_i$. Any method of serving agent i with probability $\hat{q}_i$ can be combined with any other method for serving an expected $1 - \hat{q}_i$ number of the remaining agents. Thus, the relaxed optimization problem with an ex ante feasibility constraint decomposes across the agents.

Considering an agent i, one way to serve the agent with ex ante probability $\hat{q}_i$ is to use the ex ante optimal lottery pricing (Definition 3.12).

The expected payment of the agent is given by her revenue curve as $R_i(\hat{q}_i)$. Thus, for ex ante allocation probabilities $\hat{\boldsymbol{q}}$ the optimal revenue is $\sum_i R_i(\hat{q}_i)$. Recall that for regular distributions, this optimal pricing is simply to post the price $V_i(\hat{q}_i)$ which has probability $\hat{q}_i$ of being accepted by the agent. Therefore, for regular distributions the optimal ex ante mechanism is a posted pricing.

The optimal ex ante mechanism design problem is identical to the classical microeconomic problem of optimizing the amount of a unit supply of a good (e.g., grain) to fractionally allocate across each of several markets. Each market i has a concave revenue curve as a function of the faction of the supply allocated to it. Both of these optimization problem are given by the following convex program:

$$\begin{aligned} \max_{\hat{\boldsymbol{q}}} & \sum\nolimits_i R(\hat{q}_i) \\ \text{s.t.} & \sum\nolimits_i \hat{q}_i \leq 1. \end{aligned} \tag{4.7}$$

As described previously, the marginal revenue interpretation provides a simple method for solving this program. The optimal solution equates marginal revenues, i.e., $R_i'(\hat{q}_i) = R_j'(\hat{q}_j)$ for i and j with $\hat{q}_i$ and $\hat{q}_j$ strictly larger than zero. We conclude with the following proposition.

Proposition 4.12 *The optimal ex ante mechanism is a uniform virtual pricing (with virtual values defined as marginal revenues).*

Because, at least for regular distributions, the optimal ex ante mechanism is a price posting, it provides a convenient upper bound for determining the extent to which price posting (with the ex post constraint) approximates the optimal (ex post) auction. In particular, if we post the exact same prices then the difference between the ex ante and ex post posted pricing is in how violations of the ex post feasibility constraint are resolved. In the former, violations are ignored, in the latter they must be addressed. In the terminology of the previous section, we must address how ties, i.e., multiple agents desiring to buy at their respective prices, are to be resolved to respect the ex post feasibility constraint. Unlike the previous section where the oblivious ordering assumption required breaking ties in worst-case order, in this section we break ties in the mechanisms favor.

Consider the special-case where the distribution is regular and that the optimal ex ante revenue of $R_i(\hat{q}_i) = \hat{q}_i\hat{v}_i$ from agent i is obtained by posting price $\hat{v}_i = V_i(\hat{q}_i)$. The best order to break ties is in favor of higher prices, i.e., by larger $\hat{v}_i$. For general (possibly irregular distributions) this corresponds to ordering the agents by $R_i(\hat{q}_i)/\hat{q}_i$, i.e., the agent's bang-per-buck. The goal of this section is to prove an approximation bound on this sequential price posting.

4.3.2 The Correlation Gap

The sequential posted pricing theorem we present is an application of a *correlation gap* theorem from stochastic optimization. Consider a non-negative real-valued set function g over subsets S of an n element ground set $N = \{1, \ldots, n\}$ and a distribution over subsets given by $\mathcal{D}$. Let $\hat{q}_i$ be the ex ante[1] probability that element i is in the random set $S \sim \mathcal{D}$ and let $\mathcal{D}^I$ be the distribution over subsets induced by independently adding each element i to the set with probability equal to its ex ante probability $\hat{q}_i$. The *correlation gap* is then the ratio of the expected value of the set function for the (correlated) distribution $\mathcal{D}$, i.e., $\mathbf{E}_{S\sim\mathcal{D}}[g(S)]$, to the expected value of the set function for the independent distribution $\mathcal{D}^I$, i.e., $\mathbf{E}_{S\sim\mathcal{D}^I}[g(S)]$ A typical analysis of correlation gap will consider specific families of set functions g in worst case over distributions $\mathcal{D}$.

We show below that for any values $\hat{\boldsymbol{v}}$ the *maximum-weight-element* set function $g^{\mathrm{MWE}}(S) = \max_{i\in S} \hat{v}_i$ has a correlation gap of $e/e-1$.

Lemma 4.13 *The correlation gap for any maximum-weight-element set function and any distribution over sets is $e/e-1$.*

Proof This proof proceeds in three steps. First, we argue that it is without loss to consider distributions $\mathcal{D}$ over singleton sets. Second, we argue that it is without loss to consider set functions where the weights are uniform, i.e., the one-or-more set function. Third, we show that for distributions over singleton sets, the one-or-more set function has a correlation gap of $e/e-1$.

(i) We have a set function $g^{\mathrm{MWE}}(S) = \max_{i\in S} \hat{v}_i$. Add a dummy element 0 with weight $\hat{v}_0 = 0$; if $S = \emptyset$ then changing it to $\{0\}$ affects neither the correlated value nor the independent value. Moreover, the correlated value $\mathbf{E}_{S\sim\mathcal{D}}\left[g^{\mathrm{MWE}}(S)\right]$ is unaffected by changing the set to only ever include its highest weight element. This change to the distribution only (weakly) decreases the ex ante probabilities $\hat{\boldsymbol{q}} = (\hat{q}_1, \ldots, \hat{q}_n)$

[1] In probability theory, this probability is also known as the marginal probability of $i \in S$; however to avoid confusion with usage of the term "marginal" in economics, we will refer to it via its economic interpretation as an ex ante probability as if S was the feasible set output by a mechanism.

and the independent value $\mathbf{E}_{S\sim\mathcal{D}^I}\left[g^{\mathrm{MWE}}(S)\right]$ is monotone increasing in the ex ante probabilities. Therefore, this transformation only makes the correlation gap larger. We conclude that it is sufficient to bound the correlation gap for distributions $\mathcal{D}$ over singleton sets for which the ex ante probabilities sum to one, i.e., $\sum_i \hat{q}_i = 1$.

(ii) With set distribution $\mathcal{D}$ over singletons and a maximum-weight-element set function $g^{\mathrm{MWE}}(S) = \max_{i\in S} \hat{v}_i$, the correlated value simplifies to $\mathbf{E}_{S\sim\mathcal{D}}\left[g^{\mathrm{MWE}}(S)\right] = \sum_i \hat{q}_i\hat{v}_i$. Scaling the weights $\hat{\boldsymbol{v}} = (\hat{v}_1,\ldots,\hat{v}_n)$ by the same factor has no effect on the correlation gap; therefore, it is without loss to normalize so that the correlated value is $\sum_i \hat{q}_i\hat{v}_i = 1$. We now argue that among all such normalized weights $\hat{\boldsymbol{v}}$, the ones that give the largest correlation gap are the uniform weights $\hat{v}_i = 1$ for all i. This special case of the maximum-weight-element set function is the one-or-more set function, $g^{\mathrm{OOM}}(S) = 1$ if $|S| \geq 1$ and otherwise $g^{\mathrm{OOM}}(S) = 0$.

Sort the elements by $\hat{v}_i$ and let $c_i = \prod_{j<i}(1-\hat{q}_j)$ denote the probability that no element with higher weight than i is in S and, therefore, i's contribution to the independent value is $c_i\hat{q}_i\hat{v}_i$. Let $\delta_i = \hat{q}_i\cdot(\hat{v}_i - 1)$ be the additional contribution in excess of one to the correlated value of i with value $\hat{v}_i$. Importantly, by our normalization assumption that $\sum_i \hat{q}_i\hat{v}_i = 1$, the sum of these excess contributions is zero, i.e., $\sum_i \delta_i = 0$. The expected independent value for the maximum-weight-element set function is

$$\sum\nolimits_i c_i\hat{q}_i\hat{v}_i = \sum\nolimits_i c_i\cdot(\hat{q}_i + \delta_i) \geq \sum\nolimits_i c_i\hat{q}_i. \tag{4.8}$$

where the inequality follows from monotonicity of c_i and the fact that $\sum_i \delta_i = 0$. The right-hand side of (4.8) is the expected independent value of the one-or-more set function. The correlated value is one for both (normalized) general weights and uniform weights, so uniform weights give no lower correlation gap.

(iii) The correlation gap of the one-or-more set function g^{OOM} on any distribution $\mathcal{D}$ over singletons can be bounded as follows. First, the expected correlated value is one. Second, the expected independent value is, for $S\sim\mathcal{D}^I$,

$$\begin{aligned}\mathbf{E}\left[g^{\mathrm{OOM}}(S)\right] &= \mathbf{Pr}[|S|\geq 1] = 1 - \mathbf{Pr}[|S| = 0] = 1 - \prod\nolimits_i(1-\hat{q}_i)\\ &\geq 1 - (1 - {}^1\!/\!_n)^n \geq 1 - {}^1\!/\!_e,\end{aligned}$$

where the first inequality follows because $\sum_i \hat{q}_i = 1$ and because the product of a set of positive numbers with a fixed sum is maximized when the numbers are equal. The last inequality follows as $(1 - 1/n)^n$ is monotonically increasing in n and it is $1/e$ in the limit as n goes to infinity.[2] $\square$

4.3.3 Sequential Posted Pricings

The correlation gap is central to the theory of approximation for sequential posted pricings. Contrast the revenue of the optimal ex ante mechanism (a price posting) with the revenue from sequentially posting the same prices. The optimal ex ante mechanism has total ex ante service probability $\sum_i \hat{q}_i \leq 1$ (by definition). If we could coordinate the randomization (by adding correlation to the randomization of agents' types and the mechanism) then we could obtain this optimal revenue and satisfy ex post feasibility. In a sequential posted pricing, of course, no such coordination is permitted. Instead, ex post feasibility is satisfied by serving the agent that arrives first in the specified sequence.

Given any $\hat{\boldsymbol{q}}$ with $\sum_i \hat{q}_i \leq 1$, consider the correlated distribution $\mathcal{D}$ that selects the singleton set $\{i\}$ with probability $\hat{q}_i$ and the empty set $\emptyset$ with probability $1 - \sum_i \hat{q}_i$. The induced ex ante probabilities of this correlated distribution are exactly $\hat{q}_i$ for each agent i. Assume for now that the distribution is regular and that the revenue of $R_i(\hat{q}_i) = \hat{q}_i \hat{v}_i$ is obtained by posting price $\hat{v}_i = V_i(\hat{q}_i)$. For the maximum-weight-element set function, i.e., $g^{\mathrm{MWE}}(S) = \max_{i \in S} \hat{v}_i$. For $S \sim \mathcal{D}$ the expected value of this set function is precisely the optimal ex ante revenue $\sum_i \hat{v}_i \hat{q}_i$.

On the other hand, consider sequentially posting prices $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$ to agents ordered by largest $\hat{v}_i$. Let S denote the set of agents whose values are at least their prices, i.e., $S = \{i : v_i \geq \hat{v}_i\}$. Each agent i is in S independently with probability $\hat{q}_i$. Importantly, S may have cardinality larger than one, but when it does, the ordering of agents by price implies that the agent $i \in S$ with the highest price wins. The revenue of the sequential posted pricing is given by the expected value of the maximum-weight-element set function $g^{\mathrm{MWE}}(S)$ on $S \sim \mathcal{D}^I$.

For regular distributions, the translation from the solution to the optimal ex ante mechanism which is given by $\hat{\boldsymbol{q}}$ to a sequential pricing is direct. As described above, the prices $\hat{v}_i = V_i(\hat{q}_i)$ are posted to agents in decreasing order of $\hat{v}_i$. For irregular distributions the $\hat{q}_i$ optimal lottery

[2] The last part of this analysis is identical to the proof of Theorem 1.5. Again, $(1 - 1/n)^n \leq 1/e$ is a standard observation that can be had by taking the natural logarithm and then applying L'Hopital's rule for evaluating the limit.

for agent i is not necessary a posted pricing. It may be, via Theorem 3.28, a lottery over two prices. These lottery pricings arise when $\hat{q}_i$ is in an interval where the revenue curve has been ironed and is therefore locally linear. The marginal revenue (i.e., virtual value) is constant on this interval. If we break ties in the optimization of program (4.7) lexicographically, then for the optimal ex ante probabilities $\hat{\boldsymbol{q}}$ at most one is contained strictly within an ironed interval. Recall that the marginal revenues of any agents who have non-zero ex ante allocation probability are equal. At this marginal revenue, the lexicographical tie breaking rule requires that we increase the allocation probability to the early agents before later agents. We stop when we run out of ex ante allocation probability and at this stopping point the ex ante allocation probabilities can be within at most one agents ironed interval.

By the above discussion, the suggested sequential pricing potentially has one agent receiving a lottery over two prices. The expected revenue of this pricing satisfies the approximation bound guaranteed by the correlation gap theorem. Of course, it cannot be the case that both the pricings in the support of the randomized pricing have revenue below the expected revenue of the lottery pricing. Therefore, the pricing with the higher revenue gives the desired approximation. Notice that the lexicographical ordering and derandomization steps may result in prices (in value space) that are discriminatory even in the case that the environment is symmetric (i.e., for i.i.d. distributions).

Theorem 4.14 *For any single-item environment, there is sequential posted pricing (ordered by price) with uniform virtual prices that obtains a revenue that is an* $e/e-1 \approx 1.58$ *approximation to the optimal auction revenue (and the optimal ex ante mechanism revenue).*

Proof By Proposition 4.11 the optimal ex ante revenue upper bounds the optimal auction revenue. The upper bound on the approximation ratio then follows directly from the correspondence between the revenues of the optimal ex ante mechanism and the sequential posted pricing revenue and the correlated and independent values for the maximum weight element set system (Lemma 4.13). The prices correspond to a uniform virtual pricing by the characterization of the optimal ex ante mechanism (Proposition 4.12). $\square$

The construction and analysis of Theorem 4.14 can similarly be applied to the objective of social surplus (see Exercise 4.10) to obtain an

$e/e-1$ by a sequential posted pricing that generalizes Theorem 1.5 to non-identical distributions.

4.4 Anonymous Reserves and Pricings

Thus far we have shown that simple posted pricings and reserve-price-based auctions approximate the optimal auction. Unfortunately, these prices are generally discriminatory and, thus, may be impractical for many scenarios, especially ones where agents could reasonably expect some degree of fairness of the auction protocol. We therefore consider the extent to which an *anonymous posted price* or an auction with an *anonymous reserve price*, i.e., the same for each agent, can approximate the revenue of the optimal, perhaps discriminatory, auction.

For instance, in the eBay auction the buyers are not identical. Some buyers have higher *ratings* and these ratings are public knowledge. The value distributions for agents with different ratings may generally be distinct and, therefore, the eBay auction may be suboptimal. Surely though, if the eBay auction was very far from optimal, eBay would have switched to a better auction. The theorem below gives some justification for eBay sticking with the second-price auction with anonymous reserve.

Our approach to approximation for (first- or second-price) auctions with anonymous reserve will be to show that anonymous price posting gives a good approximation and then to argue via the following proposition, that the auction revenue pointwise dominates the pricing revenue. While there is not a succinct close-form expression for the best anonymous reserve price for the second-price auction; the best anonymous posted price is precisely the monopoly price for the distribution of the maximum value. Notice that with distribution functions $F_1, \ldots, F_n$, the distribution of the maximum value has distribution function $F_{\max}(z) = \prod_i F_i(z)$. From this formula, the monopoly price can be directly calculated.

Proposition 4.15 *In any single-item environment, the revenues from the first- and second-price auctions with an anonymous reserve price is at least the revenue from the anonymous posted pricing with the same price.*

Proof Recall that a posted pricing of $\hat{v}$ obtains revenue $\hat{v}$ if and only if there is an agent with value at least $\hat{v}$. For the auction, the utility an agent receives for bidding strictly below $\hat{v}$ is zero, while individual

	regular auction	regular pricing	irregular
identical	1	$\approx e/e-1$	2
non-identical	$[2,4]$	$[2,4]$	n

Figure 4.3 Approximation bounds are given for the second-price auction with anonymous reserve and for anonymous posted pricing. If a number is given, then the bound is tight in worst case, if a range is given then the bound is not known to be tight. For irregular distributions, the auction and pricing bounds are the same. For i.i.d. regular distributions, the approximation ratio of anonymous pricing is upper bounded by $e/e-1$ for all n; for small n the bound can be improved, e.g., for $n = 1$ pricing is optimal, for $n = 2$ it is a $4/3$ approximation. A nearly matching lower bound is the subject of Exercise 4.12.

rationality implies that an agent with value $v \geq \hat{v}$ will have a non-negative utility from bidding on $[\hat{v}, v]$. Thus, the auction sells at a price of at least $\hat{v}$ if and only if there is an agent with value at least $\hat{v}$. □

4.4.1 Identical Distributions

We start with results for anonymous posted pricing and identical distributions; these bounds are summarized by the first row of Figure 4.3. For i.i.d. regular distributions the second-price auction with an anonymous reserve is optimal (Corollary 3.12). For anonymous posted pricing, Theorem 4.14 implies a $e/e-1 \approx 1.58$ approximation for regular distributions and Theorem 4.10 implies a two approximation for irregular distributions. Notice that while Theorem 4.14 holds for irregular distributions, for identical irregular distributions the prices for which the result holds may not be anonymous (due to the derandomization step).

Corollary 4.16 *For i.i.d. regular single-item environments, anonymous posted pricing is an $e/e-1$ approximation to the optimal auction; this bound is nearly tight.*

Proof For i.i.d. distributions, the optimization problem of program (4.7) is symmetric and convex and, therefore, always admits a symmetric optimal solution. For regular distributions, this symmetric optimal solution corresponds to an anonymous posted pricing. Theorem 4.14 shows that this anonymous posted pricing is a $e/e-1$ approximation. For tightness, see Exercise 4.12. □

Corollary 4.17 *For i.i.d. (irregular) single-item environments, both anonymous posted pricing and the second-price auction with anonymous* *reserve are two approximations to the optimal auction revenue; these bounds are tight.*

Proof For any (possibly irregular) distribution, Theorem 4.10 shows that posting a uniform virtual price gives a two approximation to the revenue of the optimal auction. For i.i.d. distributions where the virtual value functions are identical, uniform virtual prices are anonymous. The price-posting result follows. By Proposition 4.15, using this anonymous price as a reserve price in the second-price auction only improves the revenue.

To see that this bound of two is tight, we give an i.i.d. irregular distribution for which the approximation ratio of anonymous reserve pricing for n agents is $2 - 1/n$. Consider the discrete distribution and $h \gg n$ where

$$v = \begin{cases} h \text{ (high valued)} & \text{w.p. } 1/h\text{, and} \\ n \text{ (low valued)} & \text{otherwise.} \end{cases}$$

We then analyze the optimal auction revenue, REF, and the second-price auction with any reserve, APX, for n agents and in the limit as h goes to ∞. We show that REF $= 2n - 1$ and APX $= n$; the result follows. For any given value of h, the probability that there are k high-valued agents and $n - k$ low valued agents is the same as in the proof of Proposition 4.6; the analysis below makes use of equations (4.3) and (4.4) from its proof.

We start by analyzing REF. The virtual values are $\phi(h) = h$ and, as h goes to ∞, $\phi(n) = n - 1$. The optimal auction has virtual surplus $n - 1$ if there are no high-valued agents and virtual surplus h if there is one or more high-valued agents. The former case happens with probability that goes to one and so the expected virtual surplus is $n-1$; and in the limit, h times the probability of the latter case goes to n. Thus, REF $= 2n-1$.

We now analyze APX. We show that both a reserve of n and a reserve of h give the same revenue of n in the limit. For the first case: a reserve of n is never binding. The second-price auction has revenue h if there are two or more high-valued agents and a revenue of n if there are one or fewer. In the limit (as h goes to infinity) the contribution to the expected revenue of the first term is zero and that of the second term is n. For the second case: a reserve of h gives revenue of h when there is one or more high-valued agent, and otherwise zero. As above, the product of h and this probability is n in the limit. Thus, APX $= n$. □

4.4.2 Non-identical Distributions

We now turn to asymmetric distributions. For asymmetric distributions, the challenge with anonymous pricing comes from the asymmetry in the environment. For non-identical regular distributions, an anonymous posted pricing gives a constant approximation (implying the same for anonymous reserve pricing). For non-identical irregular distributions, anonymous posted and reserve pricing are n approximations. We begin with lower and upper bounds for regular distributions.

Lemma 4.18 *Anonymous reserve or posted pricing is at best a two approximation to the optimal revenue.*

Proof This lower bound is exhibited by an $n = 2$ agent example where agent 1's value is a point-mass at one and agent 2's value is drawn from the equal revenue distribution (Definition 4.2) on $[1, \infty)$, i.e., $F_2(z) = 1 - 1/z$. Recall that, for the equal revenue distribution, posting any price $\hat{v} \geq 1$ gives an expected revenue of one. For this asymmetric setting the revenue of the second-price auction with any anonymous reserve is exactly one. On the other hand, an auction could first offer the item to agent 2 at a very high price (for expected revenue of one), and if (with very high probability) agent 2 declines, then it could offer the item to agent 1 at a price of one. The expected revenue of this mechanism in the limit is two. □

Theorem 4.19 *For single-item environments and agents with values drawn independently from regular distributions, anonymous reserve and posted pricings give a four approximation to the revenue of the optimal auction. One such anonymous price is the monopoly price for the distribution of the maximum value.*

Proof This proof combines elements from the proof of the prophet inequality (Section 4.2.1, page 111) theorem with the upper bound on the optimal auction given by the ex ante relaxation (Section 4.3.1, page 117). Let REF $= \sum_i \hat{v}_i \hat{q}_i$ denote the optimal ex ante mechanism which posts prices $\hat{v}_i = V_i(\hat{q}_i)$ and, with out loss of generality, satisfies $\sum_i \hat{q}_i = 1$. Let APX denote the revenue from posting an anonymous price $\hat{v}$. A key part of the proof is to use regularity (i.e., convexity of the price-posting revenue curve) to derive a lower bound on the probability that an agent i with $\hat{v}_i$ (from the optimal ex ante mechanism, above) has value at least the anonymous price $\hat{v}$. The full proof is left to Exercise 4.13. □

We now give a tight inapproximation bound for anonymous reserves and pricings with irregular distributions. Recall the proof of Proposition 4.7 which implies that, for (non-identical) irregular distributions, posting an anonymous price that corresponds to the monopoly reserve price of the agent with the highest monopoly revenue gives an n approximation to the optimal auction. This is, in fact, the best bound guaranteed by the second-price auction with an anonymous reserve or an anonymous posted pricing.

Theorem 4.20 *For (non-identical, irregular) n-agent single-item environments the second-price auction with anonymous reserve and anonymous posted pricing are n approximations to the optimal auction revenue; these bounds are tight.*

Proof The upper bound can be seen by adapting the proof of Proposition 4.7 as per the above discussion. The lower bound can be seen by analyzing the optimal revenue and the revenue of the second-price auction with any anonymous reserve on the following discrete distribution in the limit as parameter h approaches infinity. Agent i's value is drawn as:

$$v_i = \begin{cases} h^i & \text{w.p. } h^{-i}\text{, and} \\ 0 & \text{otherwise.} \end{cases}$$

The details of this analysis are left to Exercise 4.15. □

4.5 Multi-unit Environments

The simplest environment we could consider generalizing approximation results to are multi-unit environments. In a multi-unit environment, there are multiple units of a single item for sale and each agent desires a single unit. Denote by k the number of units. For k-unit environments the surplus maximization mechanism is simply the $(k+1)$st-price auction where the k agents with the highest bids win and are required to pay the $(k+1)$st bid. Except for the anonymous reserve pricing result for non-identical regular distributions, all of the single-item results extend to multi-unit environments.

Consider extending the results for monopoly reserve pricing to multi-unit environments. For regular (non-identical) k-unit environments, the $(k+1)$st-price auction with monopoly reserves continues to be a two approximation to the revenue optimal auction. We defer the statement and proof this result to Section 4.6 where it is a special case of Theorem 4.28. For irregular distributions the tight approximation bound for single-unit environments of Proposition 4.6 and Proposition 4.7 generalize to k-unit environments where the approximation ratio of monopoly reserve pricing is ${}^{n}\!/_{k}$ (see Exercise 4.16).

It is possible to generalize and improve the prophet inequality to show that a gambler who is able to select k prizes can, with a uniform threshold, obtain a $(1 + \sqrt{{}^{8}\!/_{k} \ln k})$ approximation to the prophet (i.e., the expected maximum value of k prizes) for sufficiently large k. From this generalized prophet inequality, the same bound holds for oblivious posted pricing.

Proposition 4.21 *For k-unit environments with sufficiently large k, there is an oblivious posted pricing that is a $(1+\sqrt{{}^{8}\!/_{k} \ln k})$ approximation to the optimal auction.*

Sequential posted pricing bounds generalize to multi-unit environments and the bound obtained improves with k and asymptotically approach one, i.e., optimal. The proof of this generalization follows from considering the correlation gap of the k-maximum-weight-elements set function, reducing its correlation gap to that of the k-capped-cardinality set function $g(S) = \min(k, |S|)$ (the one-or-more set function is the 1-capped-cardinality), and showing that this set function's correlation gap in the limit as n approaches infinity is $(1 - ({}^{k}\!/_{e})^{k} \cdot {}^{1}\!/_{k!})^{-1}$ which, by Stirling's approximation[3] is $(1 - {}^{1}\!/_{\sqrt{2\pi k}})^{-1}$ (see Exercise 4.17).

Proposition 4.22 *For k-unit environments, there is a sequential posted pricing that is a $(1 - {}^{1}\!/_{\sqrt{2\pi k}})^{-1}$ approximation to the optimal auction.*

An anonymous reserve price continues to be revenue optimal for i.i.d. regular multi-unit environments. For i.i.d. regular multi-unit environments the correlation-gap-based sequential posted pricing result (Proposition 4.22, above) implies the same bound is attained by an anonymous pricing because for i.i.d. regular distributions, a uniform virtual pricing is an anonymous pricing (in value space). For i.i.d. irregular multi-unit environments the prophet-inequality-based oblivious posted pricing result (Proposition 4.21, above) implies the same bound by an anonymous pricing (and consequently for the $(k+1)$st price auction with an anony-

[3] Stirling's approximation is $k! = ({}^{k}\!/_{e})^{k}\sqrt{2\pi k}$. This approximation is obtained by approximating the natural logarithm as $\ln(k!) = \ln(1) + \ldots + \ln(k)$ by an integral instead of a sum.

mous reserve), because for i.i.d. distributions the uniform virtual pricing identified corresponds to an anonymous pricing (in value space).

The one result that does not generalize from single-item environments to multi-unit environments is the anonymous posted and reserve pricing for non-identical distributions. In fact, this lower bound holds more generally for any set system where where it is possible to serve k agents (see Lemma 4.23, below). For irregular, non-identical distributions the n-approximation bound of Theorem 4.20 for single-item environments generalizes and is tight.

Lemma 4.23 *For any (non-identical) regular environment where it is feasible to simultaneously serve k agents, anonymous pricing and anonymous reserve pricing are at best an $\mathcal{H}_k \approx \ln k$ approximation to the optimal mechanism revenue, where $\mathcal{H}_k$ is the kth* harmonic number $\mathcal{H}_k = \sum_{i=1}^{k} 1/i$.

Proof Fix a set of k agents that are feasible to simultaneously serve and reindex them without loss of generality to be $\{1, \ldots, k\}$. The value distribution that gives this bound is the one where F_i is a pointmass at $1/i$ for agents $i \in \{1, \ldots, k\}$ and a pointmass at zero for agents $i > k$. For such a distribution, competition does not increase the price above the reserve, therefore anonymous reserve pricing is identical to anonymous posted pricing. For any $i \in \{1, \ldots, k\}$, anonymous pricing of $1/i$ to all agents obtains revenue $i \cdot 1/i = 1$ as there are i agents with values that exceed $1/i$. On the other hand, the optimal auction posts a discriminatory price to the top k agents of $1/i$ for agent i; its revenue is the kth harmonic number $\sum_{i=1}^{k} 1/i = \mathcal{H}_k$. The kth harmonic number can be approximated by the integral $\int_1^k 1/i \, \mathrm{d}i$ and satisfies $\ln k - 1 \leq \mathcal{H}_k \leq \ln k$. □

To summarize the generalization of the single-item results to multi-unit environments: all approximation and inapproximation results generalize (and some improve) except for the anonymous pricing result for non-identical, regular distributions.

4.6 Ordinal Environments and Matroids

In Chapter 3 we saw that the second-price auction with the monopoly reserve was optimal for i.i.d. regular single-item environments. In the first section of this chapter we showed that the second-price auction with monopoly reserves is a two approximation for (non-identical) regular single-item environments. We now investigate to what extent the constraint on the environment to single-item feasibility can be relaxed while still preserving these approximation results. In this section we give equivalent algorithmic and combinatorial answers to this question. The algorithmic answer is "when the greedy-by-value algorithm works;" the combinatorial answer is "when the set system satisfies a augmentation property (i.e., matroids)."

Definition 4.5 The *greedy-by-value algorithm* is

(i) Sort the agents in decreasing order of value (and discard all agents with negative value).
(ii) $\boldsymbol{x} \leftarrow \mathbf{0}$ (the null assignment).
(iii) For each agent i (in sorted order),
  if $(1, \boldsymbol{x}_{-i})$ is feasible, $x_i \leftarrow 1$.
  (I.e., serve i if i can be served alongside previously served agents.)
(iv) Output $\boldsymbol{x}$.

Notice that the greedy-by-value algorithm is optimal for single-item environments. To optimize surplus in a single-item environment we wish to serve the agent with the highest value (when it is non-negative, and none otherwise). The greedy-by-value algorithm does just that. Notice also that the optimality of the greedy-by-value algorithm for all profiles of values implies that, for the purpose of selecting the optimal outcome, the relative magnitudes of the agents' values do not matter, only the order of the of the values (and zero) matters.

Definition 4.6 An environment is *ordinal* if for all valuation profiles, the greedy-by-value algorithm optimizes social surplus.

Recall the argument for i.i.d. regular single-item environments that showed that the optimal auction is the second-price auction with the monopoly reserve price (Corollary 3.12). An agent, Alice, had to satisfy two properties to win. She must have the highest virtual value and her virtual value must be non-negative. Having a non-negative virtual value is equivalent having a value of at least the monopoly price. Having the highest virtual value, by regularity and symmetry, is equivalent to having the highest value. Thus, Alice wins when she has the highest value and is at least the monopoly price. This auction is precisely the second-price auction with the monopoly reserve price. For general environments, the non-negativity of virtual value again suggests any agents who do not have values at least the monopoly reserve price should be rejected. For an ordinal environment with values drawn i.i.d. from a regular distribution, maximization of virtual surplus for the remaining agents gives the same outcome as maximizing the surplus of the remaining agents as symmetry and strictly increasing virtual value functions imply that the relative order values is identical to that of virtual values. We conclude with the following proposition.

Proposition 4.24 *For i.i.d. regular ordinal environments, surplus maximization with the monopoly reserve price optimizes expected revenue.*

We will see in the remainder of this section that ordinality is a sufficient condition on the feasibility constraint of the environment to permit the extension of several of the single-item results from the preceding sections. In particular, for regular (non-identical) distributions, surplus maximization with (discriminatory) monopoly reserves continues to be a two approximation. For general distributions a sequential posted pricing continues to be an $e/e-1$ approximation. Neither anonymous posted prices or reserve prices generalize (as they do not generalize even for the special case of multi-unit environments, see Section 4.5).

Definition 4.7 The surplus maximization mechanism with reserves $\hat{\boldsymbol{v}}$ is:

(i) filter out agents who do not meet their reserve price, $\boldsymbol{v}^\dagger \leftarrow \{\text{agents with } v_i \geq \hat{v}_i\}$
(ii) simulate the surplus maximization mechanism on the remaining agents, and

$$(\boldsymbol{x}, \boldsymbol{p}^\dagger) \leftarrow \mathrm{SM}(\boldsymbol{v}^\dagger)$$

(iii) set prices $\boldsymbol{p}$ from critical values as:

$$p_i \leftarrow \begin{cases} \max(\hat{v}_i, p_i^\dagger) & \text{if } x_i = 1, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

where SM is the surplus maximization mechanism with no reserves.

4.6.1 Matroid Set Systems

As ordinal environments enable good approximation mechanisms, it is important to be able to understand and identify environments that are ordinal. For general feasibility environments (Definition 3.1) subsets of agents that can be simultaneously served are given by a set system.

We will see shortly, that set systems that correspond to ordinal environments, i.e., where the greedy-by-value algorithm optimizes social surplus, are matroid set systems. Checking ordinality of the environment then is equivalent to checking whether the matroid conditions hold.

Definition 4.8 A set system is $(N, \mathcal{I})$ where N is the *ground set* of elements and $\mathcal{I}$ is a set of feasible subsets of N.[4] A set system is a *matroid* if it satisfies:

- *downward closure*: subsets of feasible sets are feasible.
- *augmentation*: given two feasible sets, there is always an element from the larger whose union with the smaller is feasible.

$$\forall I, J \in \mathcal{I},\ |J| < |I| \Rightarrow \exists i \in I \setminus J,\ \{i\} \cup J \in \mathcal{I}.$$

The augmentation property trivially implies that all maximal feasible sets of a matroid have the same cardinality. These maximal feasible sets are referred to as *bases* of the matroid; the cardinality of the bases is the *rank* of the matroid. To get some more intuition for the role of the augmentation property, the following lemma shows that if the set system is not a matroid then the greedy-by-value algorithm is not always optimal.

Lemma 4.25 *The greedy-by-value algorithm selects the feasible set with largest surplus for all valuation profiles only if feasible sets are a matroid.*

Proof The lemma follows from showing for any non-matroid set system that there is a valuation profile $\boldsymbol{v}$ that gives a counterexample. First, we show that downward closure is necessary and then, for downward-closed set systems, that the augmentation property is necessary.

If the set system is not downward closed there are subsets $J \subset I$ with $I \in \mathcal{I}$ and $J \notin \mathcal{I}$. Consider the valuation profile $\boldsymbol{v}$ with

$$v_i = \begin{cases} 2 & \text{if } i \in J, \\ 1 & \text{if } i \in I \setminus J, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

The optimal outcome is to select set I which is feasible and contains all the elements with positive value. The greedy-by-value algorithm will

[4] For matroid set systems the feasible sets are often referred to as *independent sets*. To avoid confusion with independent distributions and to promote the connection between the set system and a designer's feasibility constraint, we will prefer the former term.

start adding elements $i \in J$. As J is not feasible, it must fail to add at least one of these elements. This element is permanently discarded and, therefore, the set selected by greedy is not equal to I and, therefore, not optimal.

Now, assume that the set system is downward-closed but does not satisfy the augmentation property. In particular there exists sets $J, I \in \mathcal{I}$ with $|J| < |I|$ but there is no $i \in I \setminus J$ that can be added to J, i.e., such that $J \cup \{i\} \in \mathcal{I}$. Consider the valuation profile $\boldsymbol{v}$ with (for a ground set N of size n)

$$v_i = \begin{cases} n+1 & \text{if } i \in J, \\ n & \text{if } i \in I \setminus J, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

The greedy-by-value algorithm first attempts to and succeeds at adding all the elements of J. As there are no elements in $I \setminus J$ that are feasible when added to J, the algorithm terminates selecting exactly the set J. Because I has at least one more element than J, the value of I exceeds the value of J, and the optimality of the algorithm is contradicted. $\square$

The following matroids will be of interest.

- In a *k-uniform matroid* all subsets of cardinality at most k are feasible. The 1-uniform matroid corresponds to a single-item auction; the k-uniform matroid corresponds to a k-unit auction.
- In a *transversal matroid* the ground set is the set of vertices of part A of the bipartite graph $G = (A, B, E)$ (where vertices A are adjacent to vertices B via edges E) and feasible sets are the subsets of A that can be simultaneously matched. E.g., if A is people, B is houses, and an edge from $a \in A$ to $b \in B$ suggests that b is acceptable to a; then the feasible sets are subsets of people that can simultaneously be assigned acceptable houses with no two people assigned the same house. Notice that k-uniform matroids are the special case where $|B| = k$ and all houses are acceptable to each person. Therefore, transversal matroids represent a generalization of k-unit auctions to a market environment where not all units are acceptable to every agent.
- In a *graphical matroid* the ground set is the set of edges E in graph $G = (V, E)$ and feasible sets are acyclic subgraphs (i.e., a *forest*). Maximal feasible sets in a connected graph are spanning trees. The greedy-by-value algorithm for graphical matroids is known as *Kruskal's algorithm.*

The matroid properties characterize the set systems for which the greedy-by-value algorithm optimizes social surplus. Typically the most succinct method for arguing that matroid/ordinal environments have good properties is by using the fact that the greedy-by-value algorithm is optimal. Typically the most succinct method for arguing that an environment is matroid/ordinal is by showing that it satisfies the augmentation property (and is downward closed).

Theorem 4.26 *The greedy-by-value algorithm selects the feasible set with largest surplus for all valuation profiles if and only if feasible sets are a matroid.*

Proof The "only if" direction was shown above by Lemma 4.25. The "if" direction is as follows. Let r be the *rank* of the matroid. Let $I = \{i_1, \ldots, i_r\}$ be the set of agents selected in the surplus maximizing assignment, and let $J = \{j_1, \ldots, j_r\}$ be the set of agents selected by greedy-by-value. The surplus from serving a subset S of the agents is $\sum_{i \in S} v_i$.

Assume for a contradiction that the surplus of set I is strictly more than the surplus of set J, i.e., greedy-by-value is not optimal. Index the agents of I and J in decreasing order of value. With respect to this ordering, there must exist a first index k such that $v_{i_k} > v_{j_k}$. Let $I_k = \{i_1, \ldots, i_k\}$ and let $J_{k-1} = \{j_1, \ldots, j_{k-1}\}$. Applying the augmentation property to sets I_k and J_{k-1} we see that there must exist some agent $i \in I_k \setminus J_{k-1}$ such that $J_{k-1} \cup \{i\}$ is feasible. Of course, by the ordering of I_k, $v_i \geq v_{i_k} > v_{j_k}$ which means that agent i was considered by greedy-by-value before it selected j_k. By downward closure and feasibility of $J_{k-1} \cup \{i\}$, when agent i was considered by greedy-by-value it was feasible. By definition of the algorithm, agent i should have been added; this is a contradiction. □

To verify that an environment is ordinal/matroid the most direct approach is to verify the augmentation property. As an example we show that constrained matching markets (a.k.a., the transversal matroid) are indeed a matroid.

Lemma 4.27 *For matching agents $N = \{1, \ldots, n\}$ to items $K = \{1, \ldots, k\}$ via bipartite graph $G = (N, K, E)$ where an agent $i \in N$ can be matched to an item $j \in K$ if edge $(i, j) \in E$, the subsets of agents N that correspond to matchings in G are the feasible sets of a matroid on ground set N.*

Proof Consider any two subsets $N^\dagger$ and $N^\ddagger$ of N that are feasible, i.e., that correspond to matching in G, with $|N^\dagger| < |N^\ddagger|$. We argue that there exists an $i \in N^\ddagger \setminus N^\dagger$ such that $N^\dagger \cup \{i\}$ is feasible.

A matching M corresponds to a subset of edges E such each vertex (either an agent in N or an item in K) in the induced subgraph (N, K, M) has degree (i.e., number of adjacent edges in M) at most one. Denote the matching that witnesses the feasibility of $N^\dagger$ by $M^\dagger$, and likewise, $M^\ddagger$ for $N^\ddagger$. Consider the induced subgraph $(N, K, M^\dagger \cup M^\ddagger)$. The vertices in this subgraph have degree at most two. A graph of degree at most two is a collection of paths and cycles.

There must be a path that starts at a vertex corresponding to an agent $i \in N^\ddagger \setminus N^\dagger$ and ends with a vertex corresponding to an item $j \in K$. This is because paths that start with agents $i \in N^\ddagger \setminus N^\dagger$ can only end at items or at agents $i \in N^\dagger \setminus N^\ddagger$. By the assumption $|N^\dagger| < |N^\ddagger|$, there are more agents in $N^\ddagger \setminus N^\dagger$ than $N^\dagger \setminus N^\ddagger$ and so a path ending in an item must exist.

This path that ends at an item must alternate between edges in $M^\ddagger$ and $M^\dagger$. This path has an odd number of edges as it starts with an agent and ends with an item. As it starts with an agent matched by $M^\ddagger$. It has one more edge from $M^\ddagger$ than $M^\dagger$. In matching theory and with respect to matching $M^\dagger$ this path is an *augmenting path* as swapping the edges between the matchings results in a new matching for $M^\dagger$ with one more matched edge, and consequently one more agent is matched. This additional matched agent is i. The existence of this new matching implies that $N^\dagger \cup \{i\}$ is feasible. Thus, the matroid augmentation property is satisfied. □

4.6.2 Monopoly Reserve Pricing

In matroid environments that are inherently asymmetric, the i.i.d. assumption is unnatural and therefore restrictive. As in single-item environments, the surplus maximization mechanism with (discriminatory) monopoly reserves continues to be a good approximation even when the agents' values are non-identically distributed.

Theorem 4.28 *In regular, matroid environments the revenue of the surplus maximization mechanism with monopoly reserves is a two approximation to the optimal mechanism revenue.*

There are two very useful facts about the surplus maximization mechanism in ordinal environments that enable the proof of Theorem 4.28.

The first shows that the critical value (which determine an agent's payment) for an agent is the value of the agent's "best replacement." The second shows that the surplus maximization mechanism is pointwise revenue monotone, i.e., if the values of any subset of agents increases the revenue of the mechanism does not decrease. These properties are summarized by Lemma 4.29 and Theorem 4.30, below. We will prove Lemma 4.29 and leave the formal proofs of Theorem 4.28 and Theorem 4.30 for Exercise 4.19 and Exercise 4.20, respectively.

Definition 4.9 If $I \cup \{i\} \in \mathcal{I}$ is surplus maximizing set containing i then the *best replacement* for i is $j = \text{argmax}_{\{k\,:\,I\cup\{k\}\in\mathcal{I}\}}\, v_k$.

Definition 4.10 A mechanism is *revenue monotone* if for all valuation profiles $\boldsymbol{v} \geq \boldsymbol{v}^\dagger$ (i.e., for all i, $v_i \geq v_i^\dagger$), the revenue of the mechanism on $\boldsymbol{v}$ is at least its revenue on $\boldsymbol{v}^\dagger$.

Lemma 4.29 *In matroid environments, the surplus maximization mechanism on valuation profile $\boldsymbol{v}$ has the critical values $\hat{\boldsymbol{v}}$ satisfying, for each agent i, $\hat{v}_i = v_j$ where j is the* best replacement *for i.*

Proof The greedy-by-value algorithm is ordinal, therefore we can assume without loss of generality that the cumulative values of all subsets of agents are distinct. To see this, add a $U[0, \epsilon]$ random perturbation to each agent value, the event where two subsets sum to the same value has measure zero, and as $\epsilon \to 0$ the critical values for the perturbation approach the critical values for the original valuation profile, i.e., from equation (4.9) below.

To proceed with the proof, consider two alternative calculations of the critical value for player i. The first is from the proof of Lemma 3.1 where $\text{OPT}(0, \boldsymbol{v}_{-i})$ and $\text{OPT}_{-i}(\infty, \boldsymbol{v}_{-i})$ are optimal surplus from agents other than i with i is not served and served, respectively.

$$\hat{v}_i = \text{OPT}(0, \boldsymbol{v}_{-i}) - \text{OPT}_{-i}(\infty, \boldsymbol{v}_{-i}). \tag{4.9}$$

The second is from the greedy algorithm. Sort all agents except i by value, then consider placing agent i at any position in this ordering. Clearly, i is served when placed first. Let j be the first agent after which i would not be served. Then,

$$\hat{v}_i = v_j. \tag{4.10}$$

Now we compare these the two formulations of critical values given by equations (4.9) and (4.10). Consider i ordered immediately before and immediately after j and suppose that i is served in former order and not served in the later order. In the latter order, it must be that j is served as this is the only possible difference between the outcomes of the greedy algorithm for these two orderings up to the point that both i and j have been considered. Therefore, agent j must be served in the calculation of $\mathrm{OPT}(0, \boldsymbol{v}_{-i})$. Let $J \cup \{j\}$ be the agents served in $\mathrm{OPT}(0, \boldsymbol{v}_{-i})$ and let $I \cup \{i\}$ be the agents served in $\mathrm{OPT}(\infty, \boldsymbol{v}_{-i})$. We can deduce from equations (4.9) and (4.10) that,

$$\begin{aligned} v_j &= \hat{v}_i \\ &= \mathrm{OPT}(0, \boldsymbol{v}_{-i}) - \mathrm{OPT}_{-i}(\infty, \boldsymbol{v}_{-i}) \\ &= v_j + v(J) - v(I), \end{aligned}$$

where $v(S)$ denotes $\sum_{k \in S} v_k$. We conclude that $v(I) = v(J)$ which, by the assumption that the cumulative values of distinct subsets are distinct, implies that $I = J$. Meaning: j is a replacement for i; furthermore, by optimality of $J \cup \{j\}$ for $\mathrm{OPT}(0, \boldsymbol{v}_{-i})$, j must be the best, i.e., highest valued, replacement. □

Theorem 4.30 *In matroid environments, the surplus maximization mechanism is revenue monotone.*

4.6.3 Oblivious and Adaptive Posted Pricings

Recall that an oblivious posted pricing predetermines prices to offer each agent and its revenue must be guaranteed in worst case over the order that the agents arrive. It is conjectured that oblivious posted pricing is a constant approximation for any matroid environment. In contrast, an *adaptive posted pricing* is one that, for any arrival order of the agents, calculates the price to offer each agent when she arrives. The calculated price can be a function of the agents identity, the agents that have previously arrived and the agents that are currently being served by the mechanism. The proof of the following theorem is based on a *matroid prophet inequality* (that we will not cover in this text).

Theorem 4.31 *For (non-identical, irregular) matroid environments, there is an adaptive posted pricing that is a two approximation to the optimal mechanism revenue.*

4.6.4 Sequential Posted Pricings

The $e/e-1$ approximation for single-item sequential posted pricing and its proof via correlation gap extends to matroid environments. To present this extension, we first extend the definition of the optimal ex ante mechanism to matroids. We then relate the sequential posted pricing question to the optimal ex ante mechanism via the correlation gap. Finally, we conclude with a necessary extra step for adapting the pricing to irregular distributions.

Consider a matroid set system $(N, \mathcal{I})$. Previously we defined the rank of a matroid as the maximum cardinality of any feasible set. We can similarly define the rank of a not-necessarily-feasible subset S of the ground set N as the maximum cardinality of any feasible subset of it. In other words, it is the rank of the induced matroid on $(S, \mathcal{I})$. Let $\mathrm{rank}(S)$ denote this matroid rank function.

A profile of ex ante probabilities $\hat{\boldsymbol{q}} = (\hat{q}_1, \ldots, \hat{q}_n)$ is *ex ante feasible*, if there exists a distribution $\mathcal{D}$ over feasible sets $\mathcal{I}$ of the matroid that induces these ex ante probabilities. This definition is cumbersome; however, it is simplified by the following characterization. For any distribution $\mathcal{D}$ over feasible sets and any not-necessarily-feasible set S it must be that the expected number of agents served by $\mathcal{D}$ is at most the rank of that set. I.e., for all $S \subset N$,

$$\sum\nolimits_{i \in S} \hat{q}_i \leq \mathrm{rank}(S). \tag{4.11}$$

This inequality follows as the left-hand side is the expected number of agents in S that are served and the right hand side is the maximum number of agents in S that can be simultaneously served. It is impossible for this expected number to be higher than this maximum possible. In fact, this necessary condition is also sufficient.

Proposition 4.32 *For a matroid set system $(N, \mathcal{I})$, a profile of ex ante probabilities $\hat{\boldsymbol{q}}$ is ex ante feasible (i.e., there is a distribution $\mathcal{D}$ over feasible sets $\mathcal{I}$ that induces ex ante probabilities $\hat{\boldsymbol{q}}$) if and only if $\sum_{i \in S} \hat{q}_i \leq \mathrm{rank}(S)$ holds for all subsets S of N.*

From the above characterization of ex ante feasibility, we can write the optimal ex ante pricing program as follows.

$$\begin{aligned} &\max_{\hat{\boldsymbol{q}}} \sum\nolimits_i R(\hat{q}_i) && \qquad (4.12) \\ &\text{s.t.} \sum\nolimits_{i \in S} \hat{q}_i \leq \mathrm{rank}(S), && \forall S \subset N. \end{aligned}$$

If the objective were given by linear weights instead of concave revenue curves, this program would be optimized easily by the greedy-by-value algorithm (with values equal to weights).[5] With convex revenue curves,

[5] Readers familiar with convex optimization will note that the matroid rank

the marginal revenue approach enables this program to be optimized via a simple greedy-by-value based algorithm.[6]

Suppose for now that the distribution over agent values is regular. The revenue curve for an agent with inverse demand curve $V(\cdot)$ is consequently given by $R(\hat{q}) = \hat{q} \cdot \hat{v}$ for $\hat{v} = V(\hat{q})$ since, for a regular distribution, the $\hat{q}$ optimal ex ante pricing posts price $\hat{v}$. The optimal ex ante revenue from program (4.12) is thus $\sum_i \hat{q}_i \hat{v}_i$.

The ex ante optimal revenue can be interpreted as the correlated value of a set function as follows. Consider the *matroid weighted rank* function $\mathrm{rank}_{\hat{\boldsymbol{v}}}(\cdot)$ for weights $\hat{\boldsymbol{v}}$ defined for a feasible set $S \in \mathcal{I}$ as $\sum_{i \in S} \hat{v}_i$ and in general for not-necessarily-feasible set $S \subset N$ as that maximum over feasible subsets of S of the weighted rank of that subset. As $\hat{\boldsymbol{q}}$ is ex ante feasible, there exists a correlated distribution $\mathcal{D}$ over feasible sets which induces ex ante probabilities $\hat{\boldsymbol{q}}$. The correlated value of this distribution for the matroid weighted rank set function is exactly the optimal ex ante revenue.

Now consider the sequential posted pricing that orders the agents by decreasing price $\hat{v}_i$. When an agent i arrives in this order, if it is feasible to serve the agent along with the set of agents who have been previously served, then offer her price $\hat{v}_i$; otherwise, offer her a price of infinity (i.e., reject her). Consider the outcome of this process for valuation profile $\boldsymbol{v}$ where the set of agents willing to buy at their respective price is $S = \{i : v_i \geq \hat{v}_i\}$ (which may not be feasible). The revenue from this sequential posted pricing is given by the matroid weighted rank function as $\mathrm{rank}_{\hat{\boldsymbol{v}}}(S)$.

We conclude that the approximation factor of sequential posted pricing with respect to the optimal ex ante revenue (which upper bounds the optimal revenue for ex post feasibility) is given by the correlation gap of the matroid weighted rank set function. Thus, it remains to analyze the correlation gap of the matroid weighted rank set function. An approach, which we will discuss here to analyze the correlation gap of the matroid weighted rank set functions, is to observe that the matroid

function is submodular and therefore the constraint imposed by ex ante feasibility is that of a polymatroid.

[6] Discretize quantile space $[0, 1]$ into Q evenly sized pieces. Consider the Q-wise union of the matroid set system (the class of matroid set systems is closed under union). Calculate marginal revenues of each discretized quantile of each agent. Run the greedy-by-marginal-revenue algorithm. Calculate $\hat{q}_i$ as the total quantile of agent i that is served by algorithm, i.e., $1/Q$ times the number of i's discretized pieces that are served.

weighted rank function is *submodular* and that the correlation gap of any submodular function is $^e/_{e-1}$.

For ground set N, consider a real valued set function $g : 2^N \to \mathbb{R}$. Intuitively, *submodularity* corresponds to diminishing returns. Adding an element i to a large set increases the value of the set function less than it would for adding it to a smaller subset.

Definition 4.11 A set function g is *submodular* if for $S^\dagger \subset S^\ddagger$ and $i \notin S^\ddagger$,

$$g(S^\dagger \cup \{i\}) - g(S^\dagger) \geq g(S^\ddagger \cup \{i\}) - g(S^\ddagger).$$

Importantly, the matroid rank and weighted-rank functions are submodular (Definition 4.11). Therefore, the matroid structure imposes diminishing returns.

Theorem 4.33 *The matroid rank function is submodular; for any real valued weights, the matroid weighted-rank function is submodular.*

Proof We prove the special case of uniform weights (equivalently: that the matroid rank function is submodular; for the general case, see Exercise 4.21). Consider $S^\dagger \subset S^\ddagger$ and $i \notin S^\ddagger$ and the weights $\boldsymbol{v}_{-i}$ as

$$v_j = \begin{cases} 4 & \text{if } j \in S^\dagger, \\ 2 & \text{if } j \in S^\ddagger \setminus S^\dagger, \\ 0 & \text{otherwise.} \end{cases}$$

Consider the case that $v_i = 1$ and $v_i = 3$. If i is added by greedy-by-value when $v_i = 1$ then i is certainly added by greedy-by-value when $v_i = 3$: moving i earlier in the greedy ordering only makes it more plausible that it is feasible to add i at the time i is considered. Therefore, difference in rank of $S^\dagger$ with and without i is at least the difference in rank of $S^\ddagger$ with and without i. Hence, the defining equation (Definition 4.11) for submodularity holds. $\square$

We omit the proof of the following theorem and instead refer readers to the simpler proof that the maximum value element set function has correlation gap $^e/_{e-1}$ (see Lemma 4.13, Section 4.3).

Theorem 4.34 *The correlation gap for a submodular set function and any distribution over sets is* $^e/_{e-1}$.

For regular distributions and by the above discussion, the ex ante service probabilities from the ex ante program (4.12) corresponds to a sequential posted pricing that has approximation factor bounded by the correlation gap. The same bound can be obtained for irregular distributions as well (see Section 4.3 and Exercise 4.22).

Theorem 4.35 *For matroid environments, there is a sequential posted pricing with revenue that is a $e/e-1$ approximation to the optimal auction revenue.*

4.6.5 Anonymous Reserves

While Proposition 4.24 showed that anonymous reserves are optimal for i.i.d. regular matroid environments, this is the extent to which anonymous reserves give good approximation for matroid environments. Of course, all lower bounds for multi-unit environments extend to matroids (where the k-unit auction result generalizes to rank k matroids). In addition there two new lower bounds. For i.i.d. regular matroid environments, anonymous posted pricing does not give a constant approximation. For (irregular) i.i.d. matroid environments, neither anonymous reserve nor posted pricing gives a constant approximation (Exercise 4.23).

4.6.6 Beyond Ordinal Environments

Generalizing reserve and posted pricing approximation beyond ordinal environments is difficult because in general environments (even downward-closed ones) the optimal mechanism may choose to serve one agent over a set of other agents, or vice versa. For example, this would happen when the first agents virtual value exceeds the sum of the other agents' virtual values. Recall that the matroid property discussed previously guarantees that tradeoffs between serving agents is always done one for one (e.g., via Lemma 4.29). There are two, in fact opposite, effects we should be worried about when proceeding to general environments. First, in a general downward-closed environment one agent could potentially block many agents with each with comparable payments. Second, many agents with minimal payments could potentially block a few agents who would have made significant payments.

We illustrate the first effect with an impossibility result for posted pricing mechanisms.

Lemma 4.36 *For (i.i.d., regular) downward-closed environments the approximation ratio of posted pricing (oblivious or sequential) is at best $\Omega(\log n/ \log\log n)$.*

Proof Fix an integer h, set $n = h^{h+1}$, and partition the n agents into h^h parts of size h each. Consider the one-part-only feasibility constraint that forbids simultaneously serving agents in distinct parts, but allows and number of agents in the same part to be served. The agents' values are i.i.d. from the equal revenue distribution on $[1, h]$, i.e., with $F(z) = 1 - 1/z$ and a pointmass of $1/h$ at value h. Call an agent high-valued if her value is h and, otherwise, low-valued. We show that the approximation factor is at least $h/2 \cdot e-1/e$ and conclude that the approximation factor is $\Omega(h) = \Omega(\log n/\log\log n)$.[7]

To get a lower bound on the optimal revenue, REF, consider the mechanism that serves a part only if all agents in the part are high valued, charges each of the agents in the part h, and obtains a total revenue of h^2. As there are h^h parts and each part has probability h^{-h} of being all high valued, the probability that one or more of these parts is all high valued is given by the correlation gap of the one-or-more set function as $e-1/e$ (Lemma 4.13). Thus, the optimal revenue is at least $\mathrm{REF} \geq h^2 \cdot e-1/e$.

To get an upper bound on the revenue of any posted pricing, notice that once one agent accepts a price, only agents in that same part as this agent can be simultaneously served. Since the distribution is equal revenue, the revenue from serving these remaining agents totals exactly $h - 1$ (one from each of $h - 1$ agents). The best revenue we can get from the first agent in the part is h. Thus, any posted pricing mechanism's revenue is upper bounded by $2h - 1$, and so $\mathrm{APX} \leq 2h$. $\square$

Before we illustrate the second effect (many low-paying agents blocking a few high-paying agents), notice that the tradeoffs of optimizing virtual values (for revenue) can be much different from the tradeoffs of optimizing values (for social surplus). Therefore, the outcome from surplus maximization could be much different from that of virtual surplus maximization.

Example 4.37 The expected value the equal revenue distribution on $[1, h]$ is $\ln h - 1$ (for the unbounded equal revenue distribution it is infinite). This can be calculated from the formula $\mathbf{E}[v] = \int_0^\infty (1 - F^{\mathrm{EQR}}(z))\,\mathrm{d}z$ with $F^{\mathrm{EQR}}(z) = 1 - 1/z$. On the other hand, the monopoly revenue for the equal revenue distribution is one. Therefore, the optimal

[7] To see the asymptotic behavior of the approximation ratio in terms of n, notice that by definition $\log n = (h+1)\log h$, so (a) rearranging $h = \log n/\log h - 1$ and (b) taking the logarithm $\log\log n > \log(h+1) + \log\log h$. From (b), $\log\log n = \Theta(\log h)$ and plugging this into (a) $h = \Theta(\log n/\log\log n)$.

social surplus and optimal revenue for a regular single-agent environment can be arbitrarily separated.

Because of the difference between social surplus and potential revenue (i.e., virtual surplus) can be large, there may be a set of agents with high social surplus that collectively block another set of agents from whom a large revenue could be obtained. In the surplus maximization mechanism with reserves, the payment an agent makes is either her reserve price or the externality she imposes on the other agents. In the scenario under consideration it may be that none of the agents in the first set is individually responsible for other agents being rejected, consequently none impose any externality. Therefore, the revenue they contribute need not exceed the revenue that could have been obtained by serving the second set. We illustrate this phenomenon with an impossibility result for surplus maximization with monopoly reserves in regular downward-closed environments.

Lemma 4.38 *For (regular) downward-closed environments the approximation factor of the second-price auction with monopoly reserves is* $\Omega(\log n)$.

Proof Consider a one-versus-many set system on $n+1$ agents where it is feasible to serve agent 1 (Alice) or any subset of the remaining agents $2,\ldots,n+1$ (the Bobs). This set system is downward closed.

A sketch of the argument is as follows. The Bobs' values are distributed i.i.d. from an equal revenue distribution. If we decide to sell to the Bobs the best we can get is a revenue of n total (one from each). Of course, the social surplus of the Bobs is much bigger than the revenue that selling to them would generate (see Example 4.37, above). We then set Alice's value deterministically to a large value that is $\Theta(n \log n)$ but with high probability below the social surplus of the Bobs. The optimal auction could always sell to Alice at her high value; thus, REF is $\Theta(n \log n)$. Unfortunately, the monopoly reserves for the Bobs are one and, therefore, not binding. Surplus maximization with monopoly reserves will with high probability not serve Alice, and therefore derive most of its revenue from the Bobs. The maximum expected revenue obtainable from the Bobs is n; thus, APX $= \Theta(n)$. See Exercise 4.24 for the details. $\square$

In the next section we show; for a large class of important distributions that, intuitively, do not have tails that are too heavy; that virtual values and values are close. Consequently, maximizing surplus is similar enough to maximizing virtual surplus that monopoly reserve pricing gives a good approximation to the optimal mechanism.

4.7 Monotone-hazard-rate Distributions

An important property of electronic devices, such as light bulbs or computer chips, is how long they will operate before failing. If we model the lifetime of such a device as a random variable then the failure rate, a.k.a., *hazard rate*, for the distribution at a certain point in time is the conditional probability (actually: density) that the device will fail in the next instant given that it has survived thus far. Device failure is naturally modeled by a distribution with a monotone (non-decreasing) hazard rate, i.e., the longer the device has been running the more likely it is to fail in the next instant. The uniform, normal, and exponential distributions all have monotone hazard rate. The equal-revenue distribution (Definition 4.2) does not.

Definition 4.12 The *hazard rate* of distribution F (with density f) is $h(z) = \frac{f(z)}{1-F(z)}$. The distribution has *monotone hazard rate (MHR)* if $h(z)$ is monotone non-decreasing.

Intuitively distributions with monotone hazard rate are not *heavy tailed*. In fact, the exponential distribution, with $F^{\mathrm{EXP}}(z) = 1 - e^{-z}$, $f^{\mathrm{EXP}}(z) = e^{-z}$, and $h^{\mathrm{EXP}}(z) = 1$ is the boundary between monotone hazard rate and non; its hazard rate is constant. Hazard rates are clearly important for revenue-optimal auctions as the definition of virtual valuations (for revenue), expressed in terms of the hazard rate, is

$$\phi(v) = v - {}^{1}\!/_{h(v)}. \tag{4.13}$$

It is immediately clear from equation (4.13) that monotone hazard rate implies regularity (i.e., monotonicity of virtual value; Definition 3.4).

An important property of monotone hazard rate distributions that will enable approximation by the surplus maximization mechanism with monopoly reserves is that the optimal revenue is within a factor of $e \approx 2.718$ of the optimal surplus. We illustrate this bound with the exponential distribution (Example 4.39), prove it for the case of a single-agent environments, and defer general downward-closed environments to Exercise 4.25. Contrast these results to Example 4.37, above, which shows that for non-monotone-hazard-rate distributions, the ratio of surplus to revenue can be unbounded.

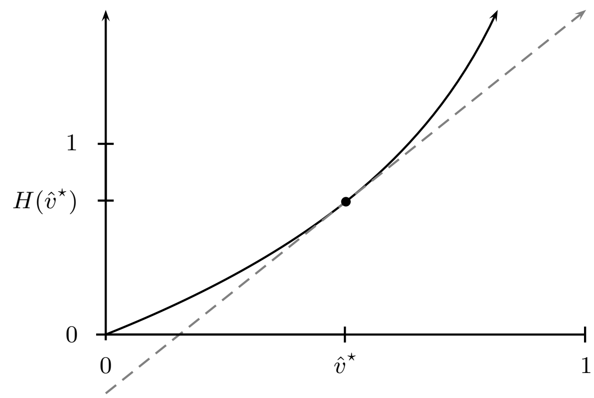

Figure 4.4 The cumulative hazard rate function (solid, black) for the uniform distribution is $H(v) = -\ln(1 - v)$ and it is lower bounded by its tangent (dashed, gray) at $\hat{v}^\star = 1/2$.

Example 4.39 The expected value the exponential distribution (with rate one) is one. This can be calculated from the formula $\mathbf{E}[v] = \int_0^\infty (1 - F^{\mathrm{EXP}}(z))\,\mathrm{d}z$ with $F^{\mathrm{EXP}}(z) = 1 - e^{-z}$. Since the exponential distribution has hazard rate $h^{\mathrm{EXP}}(z) = 1$, the virtual valuation formula for the exponential distribution is $\phi^{\mathrm{EXP}}(v) = v - 1$. The monopoly price is one. The probability that the agent accepts the monopoly price is $1 - F^{\mathrm{EXP}}(1) = 1/e$ so its expected revenue is $1/e$. The ratio of the expected surplus to expected revenue is e.

Theorem 4.40 *For any downward-closed, monotone-hazard-rate environment, the optimal expected revenue is an $e \approx 2.718$ approximation to the optimal expected surplus.*

Lemma 4.41 *For any monotone-hazard-rate distribution its expected value is at most e times more than the expected monopoly revenue.*

Proof Let REF $= \mathbf{E}[v]$ be the expected value and APX $= \hat{v}^\star \cdot (1 - F(\hat{v}^\star))$ be the expected monopoly revenue. Let $H(v) = \int_0^v h(z)\,\mathrm{d}z$ be the *cumulative hazard rate* of the distribution F. We can write

$$1 - F(v) = e^{-H(v)}, \tag{4.14}$$

an identity that can be easily verified by differentiating the natural logarithm of both sides of the equation.[8] Recall of course that the expectation

[8] We have $\frac{\mathrm{d}}{\mathrm{d}v} \ln(1 - F(v)) = \frac{-f(v)}{1-F(v)}$ and $\frac{\mathrm{d}}{\mathrm{d}v} \ln\left(e^{-H(v)}\right) = -h(v)$.

of $v \sim \boldsymbol{F}$ is $\int_0^\infty (1-F(z))\,\mathrm{d}z$. To get an upper bound on this expectation we need to upper bound $e^{-H(v)}$ or equivalently lower bound $H(v)$.

The main difficulty is that the lower bound must be tight for the exponential distribution where optimal expected value is exactly e times more than the expected monopoly revenue. Notice that for the exponential distribution the hazard rate is constant; therefore, the cumulative hazard rate is linear. This observation suggests that perhaps we can get a good lower bound on the cumulative hazard rate with a linear function.

Let $\hat{v}^\star = \phi^{-1}(0)$ be the monopoly price. Since $H(v)$ is a convex function (it is the integral of a monotone function), we can get a lower bound $H(v)$ by the line tangent to it at $\hat{v}^\star$. See Figure 4.4. I.e.,

$$\begin{aligned} H(v) &\geq H(\hat{v}^\star) + h(\hat{v}^\star)(v - \hat{v}^\star) \\ &= H(\hat{v}^\star) + \tfrac{v-\hat{v}^\star}{\hat{v}^\star}. \end{aligned} \tag{4.15}$$

The second part follows because $\hat{v}^\star = 1/h(\hat{v}^\star)$ by the choice of monopoly price $\hat{v}^\star$ and equation (4.13). Now we use this bound to calculate a bound on the expectation.

$$\begin{aligned} \mathrm{REF} &= \int_0^\infty (1-F(z))\,\mathrm{d}z = \int_0^\infty e^{-H(z)}\,\mathrm{d}z \\ &\leq \int_0^\infty e^{-H(\hat{v}^\star) - z/\hat{v}^\star + 1}\,\mathrm{d}z = e \cdot e^{-H(\hat{v}^\star)} \cdot \int_0^\infty e^{-z/\hat{v}^\star}\,\mathrm{d}z \\ &= e \cdot e^{-H(\hat{v}^\star)} \cdot \hat{v}^\star = e \cdot (1 - F(\hat{v}^\star)) \cdot \hat{v}^\star = e \cdot \mathrm{APX}\,. \end{aligned}$$

The first and last lines follow from equation (4.14); the inequality follows from equation (4.15). □

Shortly we will show that the surplus maximization mechanism with monopoly reserve prices is a two approximation to the optimal mechanism for monotone-hazard-rate downward-closed environments. This result is derived from the intuition that revenue and surplus are close. For revenue and surplus to be close, it must be that virtual values and values are close. Notice that the monotone-hazard-rate condition, via equation (4.13), implies that for higher values (which are more important for optimization) virtual value is even closer to value than for lower values (see Figure 4.5). The following lemma reformulates this intuition.

Lemma 4.42 *For any monotone-hazard-rate distribution F and $v \geq \hat{v}^\star$, $\phi(v) + \hat{v}^\star \geq v$.*

Proof Since $\hat{v}^\star = \phi^{-1}(0)$ it solves $\hat{v}^\star = 1/h(\hat{v}^\star)$. By MHR, $v \geq \hat{v}^\star$ implies

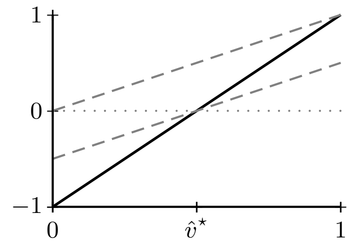

Figure 4.5 The virtual value for the uniform distribution is depicted. For $v \geq \hat{v}^\star$ the virtual value $\phi(v)$ (solid, black) is sandwiched between the value v (dashed, gray) and value less the monopoly price $v - \hat{v}^\star$ (dashed, gray).

$h(v) \geq h(\hat{v}^\star)$. Therefore,

$$\phi(v) + \hat{v}^\star = v - 1/h(v) + 1/h(\hat{v}^\star) \geq v. \qquad \square$$

Theorem 4.43 *For any monotone-hazard-rate downward-closed environment, the revenue of the surplus maximization mechanism with monopoly reserves is a two approximation to the optimal mechanism revenue.*

Proof Let APX denote the surplus maximization mechanism with monopoly reserves (and its expected revenue) and let REF denote the revenue-optimal mechanism (and its expected revenue). We start with two bounds on APX and then add them.

$$\begin{aligned} \mathrm{APX} &= \mathbf{E}[\text{APX's virtual surplus}]\,, \text{ and} \\ \mathrm{APX} &\geq \mathbf{E}[\text{APX's winners' reserve prices}]\,. \end{aligned}$$

Sum these two equations and let $\boldsymbol{x}(\boldsymbol{v})$ denote the allocation rule of APX,

$$\begin{aligned} 2 \cdot \mathrm{APX} &\geq \mathbf{E}[\text{APX's winners' virtual values + reserve prices}] \\ &= \mathbf{E}\Big[\sum\nolimits_i (\phi_i(v_i) + \hat{v}_i^\star) \cdot x_i(\boldsymbol{v})\Big] \\ &\geq \mathbf{E}\Big[\sum\nolimits_i v_i \cdot x_i(\boldsymbol{v})\Big] = \mathbf{E}[\text{APX's surplus}] \\ &\geq \mathbf{E}[\text{REF's surplus}] \geq \mathbf{E}[\text{REF's revenue}] = \mathrm{REF}\,. \end{aligned}$$

The second inequality follows from Lemma 4.42. By downward closure, neither REF nor APX sells to agents with negative virtual values. Of course, APX maximizes the surplus subject to not selling to agents with negative virtual values. Hence, the third inequality. The final inequality follows because the revenue of any mechanism is never more than its surplus. $\square$

We have seen in this section that, for monotone-hazard-rate distributions in downward closed environments, the optimal social surplus and optimal revenue are close. We then used this fact to show that a the monopoly-reserves auction is a good approximation to the optimal auction. Because surplus and revenue are close, the optimal surplus can be used as an upper bound on the optimal revenue. Finally, we showed that the monopoly-reserves auction has a revenue that approximates the optimal surplus. This approach of comparing revenue to surplus is somewhat brute-force, and there is thus a sense that these approximation bounds could be considered trivial.

Exercises

4.1 In Chapter 1 we saw that a lottery (Definition 1.2) was an n approximation to the optimal social surplus. At the time we claimed that this approximation guarantee was the best possible by a mechanism without transfers. Prove this claim.

4.2 Consider a two-agent single-item auction where agent 1 and agent 2 have values distributed uniformly on $[0, 2]$ and $[0, 3]$, respectively. Calculate and compare the expected revenue of the (asymmetric) revenue-optimal auction and the second-price auction with (asymmetric) monopoly reserves. In other words, calculate the expected revenues for the allocation rules of Example 3.11 which are depicted in Figure 4.1.

4.3 Finish the proof of Lemma 4.4 by showing that for any irregular distribution, the value of an agent is at least her virtual value for revenue. Hint: start by observing that with respect to the price-posting revenue curve $P(q) = q \cdot V(q)$, $V(q)$ is the slope of the line from the origin to the point $(q, P(q))$ on the curve, and that the lemma for the regular case implies that lines from the origin cross the curve only once.

4.4 Define a distribution to be *prepeak monotone* if its revenue curve is monotone non-decreasing on $[0, \hat{q}^\star]$, i.e., at values above the monopoly price. Notice that prepeak monotonicity is a weaker condition than regularity. First, it requires nothing of the distribution below the monopoly price. Second, above the monopoly price the price-posting revenue curve does not need to be concave. Reprove Theorem 4.2 with a weaker assumption that the agents' distributions are prepeak monotone.

4.5 Calculate the expected revenue of the optimal auction in an n-agent k-unit environment with values drawn i.i.d. from the equal revenue distribution (Definition 4.2; distribution function $F^{\mathrm{EQR}}(z) = 1 - 1/z$). Express your answer in terms of n and k.

4.6 Show that the revenue from the single-item monopoly-reserves auction smoothly degrades as the distribution becomes more irregular. To show this you will need to formally define near regularity. One reasonable definition is as follows. A distribution F is *α-nearly regular* if there is a regular distribution $F^{\dagger}$ such that price-posting revenue curves of these distributions satisfy $P(q) \geq P^{\dagger}(q) \geq 1/\alpha P(q)$ for all q.

(a) Explain why the definition above is a good definition for near regularity.

(b) Prove an approximation bound the second-price auction with monopoly reserves in α-nearly regular environments.

4.7 Generalize the prophet inequality theorem to the case where both the prophet and the gambler face an ex ante constraint $\hat{q}$ on the probability that they accept any prize.

4.8 Show that another method for choosing the threshold in the prophet inequality is to set $\hat{v} = 1/2 \cdot \mathbf{E}[\max_i v_i]$. Hint: for this choice of $\hat{v}$, prove that $\hat{v} \leq \sum_i \mathbf{E}[(v_i - \hat{v})^+]$.

4.9 Show that the prophet inequality is tight in two senses.

(a) Show that there is a distribution over prizes such that the expected prize of the optimal backwards induction strategy is half of the prophet's.

(b) Show that there is a distribution over prizes such that the expected prize of any uniform threshold strategy is at most half of the optimal backwards induction strategy.

4.10 Adapt the statement and proof of Theorem 4.14 to the objective of social surplus. Be explicit about the prices and ordering of agents in the sequential posted pricing of your construction.

4.11 For two agents with values drawn from the uniform distribution, calculate and compare the price postings from:

(a) the prophet inequality based oblivious posted pricing,

(b) the correlation gap based sequential posted pricing, and

(c) the optimal anonymous price posting.

4.12 For i.i.d. regular single-item environments, give a lower bound lower bound for the approximation ratio of anonymous pricing that

that nearly matches the upper bound. Hint: consider the regular distribution with revenue curve $R(q) = (1 - 1/n)q + 1//n$.

4.13 Prove Theorem 4.19 by adapting the analysis of the prophet inequality (Theorem 4.8) to show, for any (non-identical) regular single-item environment, that there exists an anonymous price (i.e., the same for each agent) such that price-posting obtains four approximation to the optimal ex ante mechanism revenue.

4.14 Show that there exists an i.i.d. distribution and a matroid for which the surplus maximization mechanism with an anonymous reserve is no better than an $\Omega(\log n/\log\log n)$ approximation to the optimal mechanism revenue.

4.15 Show that for (non-identical, irregular) n-agent single-item environments the second-price auction with anonymous reserve and anonymous posted pricing are at best n approximations to the optimal auction revenue (i.e., prove the lower bound of Theorem 4.20). To do so, analyze the revenue of the optimal auction and the second-price auction with any anonymous reserve when the agents values distributed as:

$$v_i = \begin{cases} h^i & \text{w.p. } h^{-i}\text{, and} \\ 0 & \text{otherwise.} \end{cases}$$

and parameter h approaches infinity. Hint: the analysis of Proposition 4.6 is similar.

(a) Show that the optimal auction has an expected revenue of n in the limit of h.

(b) Show that posting anonymous price h^i (for $i \in \{1, \ldots, n\}$) has an expected revenue of one in the limit of h.

(c) Show that for the second-price auction and anonymous reserve price h^i (for $i \in \{1, \ldots, n\}$) has an expected revenue of one in the limit of h. Hint: notice that conditioned on their being exactly one agent with a positive value, anonymous reserve pricing and anonymous posted pricing give the same revenue.

(d) Combine the above three steps to prove the theorem.

4.16 Generalize Proposition 4.7 and Proposition 4.6 to show that for n-agent k-unit irregular environments the $(k+1)$st-price auction with monopoly reserves is a n/k approximation and give a matching lower bound, respectively.

4.17 Prove Proposition 4.22, i.e., for k-unit environments that there is a sequential posted pricing that is a $(1 - 1/\sqrt{2\pi k})^{-1}$ approximation to the optimal auction, by completing the following steps.

(a) Reduce the correlation gap of the k-maximum-weight-elements set function, i.e., for weights $\hat{\boldsymbol{v}} = (\hat{v}_1, \ldots, \hat{v}_n)$ the value of $g^{\mathrm{kMWE}}(S)$ for subset S is the sum of the k largest weight elements of S, and arbitrary correlated distributions to correlated distributions over sets of cardinality exactly k.

(b) Reduce the correlation gap of the k-maximum-weight-elements set function on correlated distributions over sets of cardinality k to the correlation gap of the k-capped-cardinality set function $g^{\mathrm{kCC}}(S) = \min(k, |S|)$ (over the same class of distributions).

(c) Show that the correlation gap of the k-capped-cardinality set function on correlated distributions over sets of cardinality k is $(1 - k/e)^k \cdot 1/k!)^{-1}$.

(d) Apply the correlation gap to obtain a bound on the approximation ratio of the revenue of a uniform virtual pricing for (non-identical, irregular) k-unit environments with respect to the optimal auction revenue. Explain exactly how to find an appropriate pricing.

4.18 Recall that a feasible set of a matroid is maximal if there is no element that can be added to it such that the union is feasible. It is easy to see that the augmentation property implies that all maximal feasible sets of a matroid have the same cardinality. Rederive this result directly from the fact that greedy-by-value is optimal.

4.19 Show that in regular, matroid environments the surplus maximization mechanism with monopoly reserves gives a two approximation to the optimal mechanism revenue, i.e., prove Theorem 4.28. Hint: This result can be proved using Lemma 4.29 and Theorem 4.30 and a similar argument to the proof of Theorem 4.2.

4.20 A mechanism $\mathcal{M}$ is *revenue monotone* if for all pairs of valuation profiles $\boldsymbol{v}$ and $\boldsymbol{v}^\dagger$ such that for all i, $v_i \geq v_i^\dagger$, the revenue of $\mathcal{M}$ on $\boldsymbol{v}$ is at least its revenue on $\boldsymbol{v}^\dagger$. It is easy to see that the second-price auction is revenue monotone.

(a) For single-dimensional linear agents, give a downward-closed environment for which the surplus maximization mechanism (Mechanism 3.3) is not revenue monotone.

(b) Prove that the surplus maximization mechanism is revenue monotone in matroid environments.

4.21 Prove, directly from the fact that greedy-by-value is optimal for matroid set systems, that the matroid rank function is submodular. I.e., complete the proof of Theorem 4.33.

4.22 Consider sequential posted pricings for irregular matroid environments.

(a) Show that there is a sequential posted pricing that is an $e/e-1$ approximation to the revenue optimal auction.

(b) Give an algorithm for finding such a sequential posted pricing. Assume you are given the ex ante service probabilities $\hat{\boldsymbol{q}}$ that optimizes program (4.12). Assume you are given oracle access to the single-agent optimal ex ante pricing problems for each agent, i.e., for any agent i and service probability $\hat{q}_i$ the oracle will tell you the revenue-optimal lottery pricing that this agent with ex ante probability $\hat{q}_i$. Finally, assume you have blackbox access to a procedure that for any sequential posted pricing $\hat{\boldsymbol{v}}$ will tell you the sequential posted pricing's expected revenue (assuming prices are offered to agents in decreasing order). Your algorithm should run in linear time in the number n of agents, i.e., it should have at most a linear number of basic computational steps and calls to any of the above oracles.

4.23 Show the following inapproximability results for anonymous reserve and posted pricing in i.i.d. matroid environments.

(a) For i.i.d. regular matroid environments, anonymous posted pricing does not give a constant approximation.

(b) For (irregular) i.i.d. matroid environments, neither anonymous reserve nor posted pricing gives a constant approximation.

4.24 Complete the proof of Lemma 4.38 by showing that there is a family of regular downward-closed environments that demonstrates that the surplus maximization mechanism with monopoly reserves is an $\Omega(\log n)$ approximation to the optimal revenue. Hint: to set the value of Alice such that with high probability the social surplus of the Bobs exceeds Alice's value you can truncate the equal revenue distribution to a finite value h and then employ a standard Chernoff-Hoeffding *concentration bound* that shows that the sum of i.i.d. random variables on $[0, h]$ is concentrated around its expectation. For a sum S of i.i.d. random variables on $[0, h]$:

$$\mathbf{Pr}[|S - \mathbf{E}[S]| \geq \delta] \leq 2e^{-2\delta^2/nh^2}.$$

4.25 Consider the following *surplus maximization mechanism with lazy* *monopoly reserves* where, intuitively, we run the surplus maximization mechanism SM and then reject any winner i whose value is below her monopoly price $\hat{v}_i^\star$:

(a) $(\boldsymbol{x}^\dagger, \boldsymbol{p}^\dagger) \leftarrow \mathrm{SM}(\boldsymbol{v})$,

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

(c) $p_i = \max(\hat{v}_i^\star, p_i^\dagger)$.

Prove that the revenue of this mechanism is an e approximation to the optimal social surplus in any downward-closed, monotone-hazard-rate environment. Conclude Theorem 4.40 as a corollary.

Chapter Notes

For non-identical, regular, single-item environments, the proof that the second-price auction with monopoly reserves is a two approximation is from Chawla et al. (2007). The generalization of monopoly reserve pricing to general environments is from Hartline and Roughgarden (2009). They showed that it is a two approximation for regular matroid environments and for monotone-hazard-rate downward-closed environments. For single-item environments, the second-price auction with an anonymous reserve was shown to be between and two and four approximation by Hartline and Roughgarden (2009).

The prophet inequality theorem was proven by Samuel-Cahn (1984) and the connection between prophet inequalities and mechanism design was first made by Taghi-Hajiaghayi et al. (2007). Chawla et al. (2010) studied approximation of the optimal mechanism via oblivious and sequential posted pricings. They showed, via the prophet inequality, that a uniform virtual pricing is a two approximation for single-item environments. For k-unit environments, Taghi-Hajiaghayi et al. (2007) give a generalized prophet inequality with an upper bound of $(1+\sqrt{8/k \ln k})$ for sufficiently large k; an analogous approximation bound for uniform virtual pricing holds. Beyond single- and multi-unit environments, Chawla et al. (2010) showed that oblivious posted pricings give a three approximation for graphical matroid environments and upper bounded the approximation factor for general matroids of rank k as logarithmic in k. As of this writing, it is unknown whether there is an oblivious posted pricing give constant approximations for general matroids. On the other hand, Kleinberg and Weinberg (2012) show that there is an adaptive posted pricing that obtains a two approximation for any arrival order of the agents. This adaptive posted pricing determines the price to offer an agent when it arrives and this price can be based on the set of agents who have previously arrived and potentially been served.[9] See Alaei (2011) for a general framework for adaptive posted pricing.

The usage of the optimal ex ante mechanism as an upper bound on the optimal mechanism is from Chawla et al. (2007) and Alaei (2011). The approximation factor of sequential posted pricings were first studied by Chawla et al. (2010) they proved the $e/e-1$ approximation for single-item environments, a two approximation for matroid environments, and constant approximations for several other environments. The connection to correlation gap and the $e/e-1$ approximation for matroid environments was observed by Yan (2011) by way of the correlation gap theorem of Agrawal et al. (2010) for submodular set functions. Yan also gave the improved analysis for multi-unit auctions which shows that as the number k of available units increases the approximation factor from sequential posted pricing converges to one.

The non-game-theoretic analysis of the optimality of the greedy-by-value algorithm under matroid feasibility was initiated by Joseph Kruskal (1956) and there are books written solely on the structural properties of matroids, see e.g., Oxley (2006) or Welsh (2010). Mechanisms based on the greedy-by-value algorithm were first studied by Lehmann et al. (2002) who showed that even when these algorithms are not optimal, mechanisms derived from them are incentive compatible (cf. Chapter 8). The first comprehensive study of the revenue of the surplus maximizing mechanism in matroid environments was given by Talwar (2003); for instance, he proved critical values for matroid environments are given by the best replacement. The revenue monotonicity for matroid environments and non-monotonicity for non-matroids is discussed by Ausubel and Milgrom (2006), Day and Milgrom (2007), and Dughmi et al. (2009).

The amenability to approximation of environments with value distributions satisfying the monotone hazard rate as been observed several times, e.g., by Hartline et al. (2008), Hartline and Roughgarden (2009), and Bhattacharya et al. (2010). The structural comparison that shows that the optimal revenue is an $e \approx 2.718$ approximation to the optimal social surplus for for downward-closed, monotone-hazard-rate environments was given by Dhangwatnotai et al. (2010).

[9] Note that both the sequential posted pricings and oblivious posted pricings considered in this chapter fix the prices that each agent will receive before the mechanism is run.