

Mechanism Design and Approximation

Jason D. Hartline

Manuscript date: January 19, 2016

Copyright © 2011–2016
<http://jasonhartline.com/MDnA/>

Contents

1	Mechanism Design and Approximation	<i>page</i> 1
1.1	An Example: Congestion Control and Routing	2
1.1.1	Non-monetary payments	9
1.1.2	Posted Pricing	11
1.1.3	General Routing Mechanisms	14
1.2	Mechanism Design	15
1.3	Approximation	17
1.3.1	Philosophy of Approximation	18
1.3.2	Approximation Factors	20
2	Equilibrium	24
2.1	Complete Information Games	24
2.2	Incomplete Information Games	26
2.3	Bayes-Nash Equilibrium	27
2.4	Single-dimensional Games	29
2.5	Characterization of Bayes-Nash Equilibrium	31
2.6	Characterization of Dominant Strategy Equilibrium	36
2.7	Revenue Equivalence	37
2.8	Solving for Bayes-Nash Equilibrium	38
2.9	Uniqueness of Equilibria	42
2.10	The Revelation Principle	46
	Exercises	48
3	Optimal Mechanisms	53
3.1	Single-dimensional Environments	54
3.2	Social Surplus	56
3.3	Profit	59
3.3.1	High-level Approach: Amortized Analysis	61

3.3.2	The Virtual Surplus Maximization Mechanism	64
3.3.3	Single-item Environments	66
3.3.4	Quantile Space, Price-posting Revenue, and Derivation of Virtual Values	71
3.3.5	Virtual Surplus Maximization Subject to Monotonicity	75
3.4	Multi- to Single-agent Reduction	79
3.4.1	Revenue Curves	80
3.4.2	Optimal and Marginal Revenue	81
3.4.3	Downward Closure and Pricing	82
3.4.4	Revenue Linearity	83
3.4.5	Optimal Ex Ante Pricings, Revisited	85
3.4.6	Optimal Interim Pricings, Revisited	87
3.5	Social Surplus with a Balanced Budget	89
3.5.1	Lagrangian Relaxation	90
3.5.2	Monotone Lagrangian Virtual Values	92
3.5.3	Non-monotone Lagrangian Virtual Values and Partial Ironing	95
	Exercises	96
4	Bayesian Approximation	100
4.1	Monopoly Reserve Pricing	101
4.1.1	Approximation for Regular Distributions	104
4.1.2	Inapproximability Irregular Distributions	107
4.2	Oblivious Posted Pricings and the Prophet Inequality	110
4.2.1	The Prophet Inequality	111
4.2.2	Oblivious Posted Pricing	114
4.3	Sequential Posted Pricings and Correlation Gap	116
4.3.1	The Ex Ante Relaxation	117
4.3.2	The Correlation Gap	119
4.3.3	Sequential Posted Pricings	121
4.4	Anonymous Reserves and Pricings	123
4.4.1	Identical Distributions	124
4.4.2	Non-identical Distributions	126
4.5	Multi-unit Environments	127
4.6	Ordinal Environments and Matroids	129
4.6.1	Matroid Set Systems	131
4.6.2	Monopoly Reserve Pricing	135
4.6.3	Oblivious and Adaptive Posted Pricings	137

	4.6.4	Sequential Posted Pricings	137
	4.6.5	Anonymous Reserves	141
	4.6.6	Beyond Ordinal Environments	141
	4.7	Monotone-hazard-rate Distributions	144
		Exercises	148
5		Prior-independent Approximation	155
	5.1	Motivation	156
	5.2	“Resource” Augmentation	158
	5.2.1	Single-item Environments	158
	5.2.2	Multi-unit and Matroid Environments	159
	5.3	Single-sample Mechanisms	160
	5.3.1	The Geometric Interpretation	161
	5.3.2	Monopoly versus Single-sample Reserves	163
	5.3.3	Optimal versus Lazy Single-sample-reserve Mechanism	164
	5.4	Prior-independent Mechanisms	165
	5.4.1	Digital Good Environments	166
	5.4.2	General Environments	167
		Exercises	169
6		Bayes-Nash Approximation	172
	6.1	Social Surplus of Winner-pays-bid Mechanisms	173
	6.1.1	The Geometry of Best Response	175
	6.1.2	Utility Approximates Value	176
	6.1.3	Revenue Covering Approximation	177
	6.1.4	Social Surplus in Bayes-Nash Equilibrium	180
	6.2	Beyond Winner-pays-bid Mechanisms	181
	6.3	Simultaneous Composition	183
	6.4	Reserve Prices	185
	6.4.1	Surplus Approximates Value	187
	6.4.2	Revenue Covering Approximation	187
	6.4.3	Social Surplus in Bayes-Nash Equilibrium	188
	6.5	Analysis of Revenue	189
	6.6	Revenue Covering Optimization	191
	6.6.1	Non-bossiness, Approximation, and Greedy Algorithms	191
	6.6.2	Single-minded Combinatorial Auctions	194
		Exercises	198
7		Prior-free Mechanisms	201
	7.1	The Framework for Prior-free Mechanism Design	203

7.2	The Digital-good Environment	206
7.2.1	The Envy-free Benchmark	207
7.2.2	Deterministic Auctions	210
7.2.3	The Random Sampling Auction	212
7.2.4	Decision Problems for Mechanism Design	217
7.3	The Envy-free Benchmark	221
7.3.1	Envy-free Pricing	222
7.3.2	Envy-free Optimal Revenue	224
7.3.3	Envy freedom versus Incentive Compatibility	226
7.3.4	Permutation Environments	229
7.3.5	The Envy-free Benchmark	230
7.4	Multi-unit Environments	232
7.4.1	Reduction to Digital Goods	233
7.4.2	Combination of Benchmarks and Auctions	236
7.4.3	The Random Sampling Auction	237
7.5	Matroid Permutation and Position Environments	238
7.6	Downward-closed Permutation Environments	243
	Exercises	247
8	Multi-dimensional and Non-linear Preferences	251
8.1	Optimal Single-agent Mechanisms	252
8.1.1	Public Budget Preferences	254
8.1.2	Unit-demand Preferences	258
8.2	Service Constrained Environments	261
8.3	The Ex Ante Reduction	262
8.3.1	Example: Uniform Unit-demand Preferences	263
8.3.2	Orderability	266
8.3.3	The Marginal Revenue Mechanism	270
8.4	The Interim Reduction	271
8.4.1	Symmetric Single-item Environments	272
8.4.2	Interim Feasibility	276
8.4.3	Interim Feasibility by Stochastic Weighted Optimization	283
8.4.4	Combining Ex post Feasibility and Bayesian Incentive Compatibility	288
8.5	Multi-dimensional Externalities	292
8.6	Public Budget Preferences	295
8.7	Unit-demand Preferences	301
8.7.1	Warmup: The Uniform Distribution	302

	<i>Contents</i>	vii
8.7.2	Multi-dimensional Characterization of Incentive Compatibility	304
8.7.3	Optimal Mechanisms for Paths	306
8.7.4	Uniform Pricing for Ratio-monotone Dis- tributions	311
8.7.5	Multi-dimensional Virtual Values	313
8.7.6	Reverse Solving for Virtual Values	319
	Exercises	325
	Appendix Mathematical Reference	329
	<i>References</i>	333

Author's Note

This text is suitable for advanced undergraduate or graduate courses; it has been developed at Northwestern U. as the primary text for such a course since 2008.

This text provides a look at select topics in economic mechanism design through the lens of approximation. It reviews the classical economic theory of mechanism design wherein a Bayesian designer looks to find the mechanism with optimal performance in expectation over the distribution from which the preferences of the participants are drawn. It then adds to this theory practical constraints such as simplicity, tractability, and robustness. The central question addressed is whether these practical mechanisms are good approximations of the optimal ones. The resulting theory of approximation in mechanism design is based on results that come mostly from the theoretical computer science literature. The results presented are the ones that are most directly compatible with the classical (Bayesian) economic theory and are not representative of the entirety of the literature.

– Jason D. Hartline

1

Mechanism Design and Approximation

Our world is an interconnected collection of economic and computational systems. Within such a system, individuals optimize their actions to achieve their own, perhaps selfish, goals; and the system combines these actions with its basic laws to produce an outcome. Some of these systems perform well, e.g., the national residency matching program which assigns medical students to residency programs in hospitals, e.g., auctions for online advertising on Internet search engines; and some of these systems perform poorly, e.g., financial markets during the 2008 meltdown, e.g., gridlocked transportation networks. The success and failure of these systems depends on the basic laws governing the system. Financial regulation can prevent disastrous market meltdowns, congestion protocols can prevent gridlock in transportation networks, and market and auction design can lead to mechanisms for allocating and exchanging goods or services that yield higher profits or increased value to society.

The two sources for economic considerations are the preferences of individuals and the performance of the system. For instance, bidders in an auction would like to maximize their gains from buying; whereas, the performance of the system could (i.e., from the perspective of the seller) be measured in terms of the revenue it generates. Likewise, the two sources for computational considerations are the individuals who must optimize their strategies, and the system which must enforce its governing rules. For instance, bidders in the auction must figure out how to bid, and the auctioneer must calculate the winner and payments from the bids received. While these calculations may seem easy when auctioning a painting, they both become quite challenging when, e.g., the Federal Communications Commission (FCC) auctions cell phone spectrum for which individual lots have a high degree of complementarities.

These economic and computational systems are complex. The space

of individual strategies is complex and the space of possible rules for the system is complex. Optimizing among strategies or system rules in complex environments should lead to complex strategies and system rules, yet the individuals' strategies or system rules that are successful in practice are often remarkably simple. This simplicity may be a consequence of individuals and designers preference for ease of understanding and optimization (i.e., tractability) or robustness to variations in the scenario, especially when these desiderata do not significantly sacrifice performance.

This text focuses on a combined computational and economic theory for the study and design of mechanisms. A central theme will be the tradeoff between optimality and other desirable properties such as simplicity, robustness, computational tractability, and practicality. This tradeoff will be quantified by a theory of approximation which measures the loss of performance of a simple, robust, and practical approximation mechanism in comparison to the complicated and delicate optimal mechanism. The theory provided does not necessarily suggest mechanisms that should be deployed in practice, instead, it pinpoints salient features of good mechanisms that should be a starting point for the practitioner.

In this chapter we will explore mechanism design for routing and congestion control in computer networks as an example. Our study of this example will motivate a number of questions that will be addressed in subsequent chapters of the text. We will conclude the chapter with a formal discussion of approximation and the philosophy that underpins its relevance to the theory of mechanism design.

1.1 An Example: Congestion Control and Routing in Computer Networks

We will discuss novel mechanisms for congestion control and routing in computer networks to give a preliminary illustration of the interplay between strategic incentives, approximation, and computation in mechanism design. In this discussion, we will introduce basic questions that will be answered in the subsequent chapters of this text.

Consider a hypothetical computer network where network users reside at computers and these computers are connected together through a network of routers. Any pair of routers in this network may be connected by a network link and if such a network link exists then each router

can route a message directly through the other router. We will assume that the network is completely connected, i.e., there is a path of network links between all pairs of users. The network links have limited capacity; meaning, at most a fixed number of messages can be sent across the link in any given interval of time. Given this limited capacity the network links are a resource that may be over demanded. To enable the sending of messages between users in the network we will need mechanisms for *congestion control*, i.e., determining which messages to route when a network link is over-demanded, and *routing*, i.e., determining which path in the network each message should take.

There are many complex aspects of this congestion control problem: dynamic demands, complex networks, and strategic user behavior. Let us ignore the first two issues at first and focus on the latter: strategic user behavior. Consider a static version of this routing problem over a single network link with unit capacity: each user wishes to send a message across the link, but the link only has capacity for one message. How shall the routing protocol select which message to route?

That the resource that the users (henceforth: agents) are vying for is a network link is not important; we will therefore cast the problem as a more general single-item resource allocation problem. An implicit assumption in this problem is that it is better to allocate the item to some agents over others. For instance, we can model the agents as having value that each gains for receiving the item and it would be better if the item went to an agent that valued it highly.

Definition 1.1. The *single-item allocation* problem is given by

- a single indivisible *item* available,
- n strategic *agents* competing for the item, and
- each agent i has a *value* v_i for receiving the item.

The objective is to maximize the *social surplus*, i.e., the value of the agent who receives the item.

The social surplus is maximized if the item is allocated to the agent with the highest value, denoted $v_{(1)}$. If the values of the agent are publicly known, this would be a simple allocation protocol to implement. Of course, e.g., in our routing application, it is rather unlikely that values are publicly known. A more likely situation is that each agent's value is known privately to that agent and unknown to all other parties. A mechanism that wants to make use of this private information must

then solicit it. Consider the following mechanism as a first attempt at a single-item allocation mechanism:

- (i) Ask the agents to report their values (\Rightarrow agent i reports b_i),
- (ii) select the agent i^* with highest report ($\Rightarrow i^* = \operatorname{argmax}_i b_i$), and
- (iii) Allocate the item to agent i^* .

Suppose you were one of the agents and your value was \$10 for the item; how would you bid? What should we expect to happen if we ran this mechanism? It should be pretty clear that there is no reason your bid should be at all related to your value; in fact, you should always bid the highest number you can think of. The winner is the agent who thinks of and reports the highest number. The unpredictability of the outcome of the mechanism will make it hard to reason about its performance. There are two natural ways to try to address this unpredictability. First, we can accept that the bids are meaningless, ignore them (or not even solicit them), and pick the winner randomly. Second, we could attempt to penalize the agents for bidding a high amount, for instance, with a monetary payment.

Definition 1.2. The *lottery mechanism* is:

- (i) select a uniformly random agent, and
- (ii) allocate the item to this agent.

The *social surplus* of a mechanism is total value it generates. In this routing example the social surplus is the value of the message routed. It is easy to calculate the expected surplus of the lottery. It is $1/n \sum_i v_i$. This surplus is a bit disappointing in contrast to the surplus available in the case where the values of the messages were publicly known, i.e., $v_{(1)} = \max_i v_i$. In fact, by setting $v_1 = 1$; $v_i = \epsilon$ (for $i \neq 1$); and letting ϵ go to zero we can observe that the surplus of the lottery approaches $v_{(1)}/n$; therefore, its worst-case is at best an n approximation to the optimal surplus $v_{(1)}$. Of course, the lottery always obtains at least an n th of $v_{(1)}$; therefore, its worst-case approximation factor is exactly n . It is fairly easy to observe, though we will not discuss the details here, that this approximation factor is the best possible by any mechanism without payments.

Theorem 1.1. *The surplus of the lottery mechanism is an n approximation to the highest agent value.*

If instead it is possible to charge payments, such payments, if made proportionally to the agents' bids, could discourage low-valued agents from making high bids. This sort of dynamic allocation and pricing mechanism is referred to as an *auction*.

Definition 1.3. A *Single-item auction* is a solution to the single-item allocation problem that solicits bids, picks a winner, and determines payments.

A natural allocation and pricing rule that is used, e.g., in government procurement auctions, is the *first-price auction*.

Definition 1.4. The *first-price auction* is:

- (i) ask agents to report their values (\Rightarrow agent i reports b_i),
- (ii) select the agent i^* with highest report ($\Rightarrow i^* = \operatorname{argmax}_i b_i$),
- (iii) allocate the item to agent i^* , and
- (iv) charge this winning agent her bid, b_{i^*} .

To get some appreciation for the strategic elements of the first price auction, note that an agent who wins wants to pay as little as possible, thus bidding a low amount is desirable. Of course, if the agent bids too low, then she probably will not win. Strategically, this agent must figure out how to balance this tradeoff. Of course, since agents may not report their true values, the agent with the highest bid may not be the agent with the highest-valued message. See Figure 1.1.

We will be able to analyze the first-price auction and we will do so in Chapter 2. However, for two reasons, there is little hope of generalizing it beyond the single-network-link special case (i.e., to large asymmetric computer networks) which is our eventual goal. First, calculating equilibrium strategies in general asymmetric environments is not easy; consequently, there would be little reason to believe that agents would play by the equilibrium. Second, it would be a challenge to show that the equilibrium is any good. Therefore, we turn to auctions that are strategically simpler.

The ascending-price auction is a stylized version of the auction popularized by Hollywood movies; art, antiques, and estate-sale auction houses such as Sotheby's and Christie's; and Internet auction houses such as eBay.

Definition 1.5. The *ascending-price auction* is:¹

¹ The ascending-price auction is also referred to as the *English auction* and it contrasts to the Dutch (descending-price) auction.

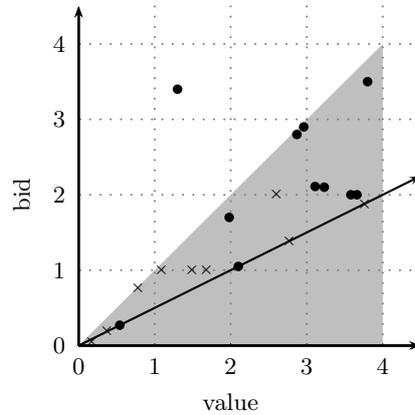


Figure 1.1. An in-class experiment: 21 student were endowed with values uniformly drawn from the interval $[0, 4]$ (denoted as $v_i \sim U[0, 4]$), they were told their own values and that the distribution of values was $U[0, 4]$, they were asked to submit bids for a two-agent one-item first-price auction. The bids of the students were collected and randomly paired for each auction; the winner was paid the difference between his value and his bid in dollars (real money). Winning bids are shown as “•” and losing bids are shown as “×”. The grey area denotes strategies that are not dominated. The black line $b = v/2$ denotes the equilibrium strategy in theory. In economic experiments, just like our in class experiment, bidders tend to overbid the equilibrium strategy. A few students knew the equilibrium strategy in advance of the in-class experiment.

- (i) gradually raise an offered price up from zero,
- (ii) allow agents to drop out when they no longer wish to win at the offered price,
- (iii) stop at the price where the second-to-last agent drops out, and
- (iv) allocate the item to the remaining agent and charges her the stopping price.

Strategically this auction is much simpler than the first-price auction. What should an agent with value v do? A good strategy would be “drop when the price exceeds v .” Indeed, regardless of the actions of the other agents, this is a good strategy for the agent to follow, i.e., it is a *dominant strategy*. It is reasonable to assume that an agent with an obvious dominant strategy will follow it.

Since we know how agents are behaving, we can now make conclusions as to what happens in the auction. The second-highest-valued agent will drop out when the ascending prices reaches her value, $v_{(2)}$. The highest-valued agent will win the item at this price. We can conclude that this

auction maximizes the *social surplus*, i.e., the sum of the utilities of all parties. Notice that the utility of losers are zero, the utility of the winner is $v_{(1)} - v_{(2)}$, and the utility of the seller (e.g., the router in the congestion control application) is $v_{(2)}$, the payment received from the winner. The total is simply $v_{(1)}$, as the payment occurs once positively (for the seller) and once negatively (for the winner) and these terms cancel. Of course $v_{(1)}$ is the optimal surplus possible; we could not give the item to anyone else and get more value out of it.

Theorem 1.2. *The ascending-price auction maximizes the social surplus in dominant strategy equilibrium.*

What is striking about this result is that it shows that there is essentially no loss in surplus imposed by the assumption that the agents' values are privately known only to themselves. Of course, we also saw that the same was not true of routing mechanisms that cannot require the winner to make a payment in the form of a monetary *transfer* from the winner to the seller. Recall, the lottery mechanism could be as bad as an n approximation. A conclusion we should take from this exercise is that transfers are very important for surplus maximization when agents have private values.

Unfortunately, despite the good properties of the ascending-price auction there are two drawbacks that will prevent our using it for routing and congestion control in computer networks. First, mechanisms for sending messages in computer networks must be very fast. Ascending auctions are slow and, thus, impractical. Second, the ascending-price auction does not generalize to give a routing mechanisms in networks beyond the single-network-link special case. Challenges arise because ascending prices would not generally find the social surplus maximizing set of messages to route. A solution to these problems comes from Nobel laureate William Vickrey who observed that if we simulate the ascending-price auction with sealed bids we arrive at the same outcome in equilibrium without the need to bother with an ascending price.

Definition 1.6. The *second-price auction* is:²

- (i) accept sealed bids,
- (ii) allocate the item to the agent with the highest bid, and
- (iii) charge this winning agent the second-highest bid.

² The second-price auction is also referred to as the Vickrey auction.

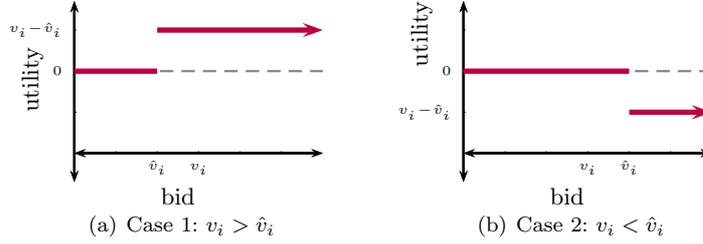


Figure 1.2. Utility as a function of bid in the second-price auction.

In order to predict agent behavior in the second-price auction, notice that its outcome can be viewed as a simulation of the ascending-price auction. Via this viewpoint, there is a one-to-one correspondence between bidding b in the second-price auction and dropping out at price b in the ascending-price auction. Since the dominant strategy in the ascending-price auction is for an agent to drop out when the price exceeds her value; it is similarly a dominant strategy for the agent to bid her true value in the second-price auction. While this intuitive argument can be made formal, instead we will argue directly that *truthful* bidding is a dominant strategy in the second-price auction.

Theorem 1.3. *Truthful bidding is a dominant strategy in the second-price auction.*

Proof. We show that truthful bidding is a dominant strategy for agent i . Fix the bids of all other agents and let $\hat{v}_i = \max_{j \neq i} v_j$. Notice that given this \hat{v}_i there are only two possible outcomes for agent i . If she bids $b_i > \hat{v}_i$ then she wins, pays \hat{v}_i (which is the second-highest bid), and has utility $u_i = v_i - \hat{v}_i$. On the other hand, if she bids $b_i < \hat{v}_i$ then she loses, pays nothing, and has utility $u_i = 0$. This analysis allows us to plot the utility of agent i as a function of her bid in two relevant cases, the case that $v_i < \hat{v}_i$ and the case that $v_i > \hat{v}_i$. See Figure 1.2.

Agent i would like to maximize her utility. In Case 1, this is achieved by any bid greater than \hat{v}_i . In Case 2, it is achieved by any bid less than \hat{v}_i . Notice that in either case bidding $b_i = v_i$ is a good choice. Since the same bid is a good choice regardless of which case we are in, the same bid is good for any \hat{v}_i . Thus, bidding truthfully, i.e., $b_i = v_i$, is a dominant strategy. \square

Notice that, in the proof of the theorem, \hat{v}_i is the infimum of bids that the bidder can make and still win, and the price charge to such a

winning bidder is exactly \hat{v}_i . We henceforth refer to \hat{v}_i as agent i 's *critical value*. It should be clear that the proof above can be easily generalized, in particular, to any auction where each agent faces such a critical value that is a function only of the other agents' reports. This observation will allow the second-price auction to be generalized beyond single-item environments.

Corollary 1.4. *The second-price auction maximizes the social surplus in dominant strategy equilibrium.*

Proof. By the definition of the second-price auction, the agent with the highest bid wins. By Theorem 1.3 is a dominant strategy equilibrium for agents to bid their true values. Thus, in equilibrium the agent with the highest bid is identically the agent with the highest value. The social surplus is maximized. \square

In the remainder of this section we explore a number of orthogonal issues related to practical implementations of routing and congestion control. Each of these vignettes will conclude with motivating questions that will be addressed in the subsequent chapters. First, we address the issue of payments. The routing protocol in today's Internet, for instance, does not allow the possibility of monetary payments. How does the routing problem change if we also disallow monetary payments? The second issue we address is speed. While the second-price auction is faster than the ascending-price auction, still the process of soliciting bids, tallying results, and assigning payments may be too cumbersome for a routing mechanism. A simpler posted-pricing mechanism would be faster, but how can we guarantee good performance with a posted pricing? Finally, the single-link case is far from providing a solution to the question of routing and congestion control in general networks. How can we extend the second-price auction to more general environments?

1.1.1 Non-monetary payments

Most Internet mechanisms, including its congestion control mechanisms, do not currently permit monetary transfers. There are historical, social, and infrastructural reasons for this. The Internet was initially developed as a research platform and its users were largely altruistic. Since its development, the social norm is for Internet resources and services to be free and unbiased. Indeed, the "net neutrality" debates of the early 2000's were largely on whether to allow differentiated service in routers

based on the identity of the source or destination of messages (and based on contracts that presumably would involve payments). Finally, micropayments in the Internet would require financial infrastructure which is currently unavailable at reasonable monetary and computational overhead.

One solution that has been considered, and implemented (but not widely adopted) for similar resource allocation tasks (e.g., filtering unsolicited electronic mail, a.k.a., spam) is *computational payments* such as “proofs of work.” With such a system an agent could “prove” that her message was high-valued by having her computer perform a large, verifiable, but otherwise, worthless computational task. Importantly, unlike monetary payments, computational payments would not represent utility transferred from the winner to the router. Instead, computational payments are utility lost to society.

The *residual surplus* of a mechanism with computational payments is the total value generated less any payments made. The residual surplus for a single-item auction is thus the value of the winner less her payment. For the second-price auction, the residual surplus is $v_{(1)} - v_{(2)}$. For the lottery, the residual surplus is $\frac{1}{n} \sum_i v_i$, which is the same as the surplus as there are no payments.

While the second-price auction maximizes surplus (among all mechanisms) regardless of the values of the agents, for the objective of residual surplus it is clear that neither the second-price auction nor the lottery mechanism is best regardless of agent values. Consider the bad input for the lottery, where $v_1 = 1$ and $v_i = \epsilon$ (for $i \neq 1$). If we let ϵ go to zero, the second-price auction has residual surplus $v_{(1)} = 1$ (which is certainly optimal) and the lottery has expected surplus $1/n$ (which is far from optimal). On the other hand, if we consider the all-ones input, i.e., $v_i = 1$ for all i , then the residual surplus of the second-price auction is $v_{(1)} - v_{(2)} = 0$ (which is far from optimal), whereas the lottery surplus is $v_{(1)} = 1$ (which is clearly optimal). Of course, on the input with $v_1 = v_2 = 1$ and $v_i = \epsilon$ (for $i \geq 3$) both the lottery and the second-price auction have residual surplus far from what we could achieve if the values were publicly known or monetary transfers were allowed.

The underlying fact in the above discussion that separates the objectives of surplus and residual surplus is that for surplus maximization there is a single mechanism that is optimal for any profile of agent values, namely the second-price auction; whereas there is no such mechanism for residual surplus. Since there is no absolute optimal mechanism we must trade-off performance across possible profiles of agent values. There are

two ways to do this. The first approach is to assume a distribution over value profiles and then optimize residual surplus in expectation over this distribution. Thus, we might trade off low residual surplus on a rare input for high residual surplus on a likely input. This approach results in a different “optimal mechanism” for different distributions. The second approach begins with the solution to the first approach and asks for a single mechanism that best approximates the optimal mechanism in worst-case over distributions. This second approach may be especially useful for applications of mechanism design to computer networks because it is not possible to change the routing protocol to accommodate changing traffic workloads.

Question 1.1. In what settings does the second-price auction maximize residual surplus? In what settings does the lottery maximize residual surplus?

Question 1.2. For any given distribution over agent values, what mechanism optimizes residual surplus for the distribution?

Question 1.3. If the optimal mechanism for a distribution is complicated or unnatural, is there a simple or natural mechanism that approximates it?

Question 1.4. In worst-case over distributions of agent values, what single mechanism best approximates the optimal mechanism for the distribution?

1.1.2 Posted Pricing

Consider again the original single-item allocation problem to maximize surplus (with monetary payments). Unfortunately, even a single-round, sealed-bid auction such as the second-price auction may be too complicated and slow for congestion control and routing applications. An even simpler approach would be to just post a take-it-or-leave-it price. Consider the following mechanism.

Definition 1.7. For a given price \hat{v} , the *uniform-pricing* mechanism serves the first agent willing to pay \hat{v} (breaking ties in arrival order randomly).

For instance, if we assumed all agents arrive at once and $\hat{v} = 0$ this uniform pricing mechanism is identical to the aforementioned lottery.

Recall that the lottery mechanism is very bad when there are many low-valued agents and a few high-valued agents. The bad example had one agent with value one, and the remaining $n - 1$ agents with value ϵ . This uniform-pricing mechanism, however, is more flexible. For instance, for this example we could set $\hat{v} = 2\epsilon$, only the high-valued agent will want to buy, and the surplus would be one. Such a posted-pricing mechanism is very practical and, therefore, especially appropriate for our application to Internet routing.

Of course, the price \hat{v} needs to be chosen well. Fortunately in the routing example where billions of messages are sent every day, it is reasonable to assume that there is some distributional knowledge of the demand. Imagine that the value of each agent i is drawn independently and identically from distribution F . The *cumulative distribution function* for random variable v drawn from distribution F specifies the probability that it is at most z , denoted $F(z) = \Pr_{v \sim F}[v < z]$. For example the uniform distribution on interval $[0, 1]$ is denoted $U[0, 1]$ and its cumulative distribution function is $F(z) = z$.

There is a very natural way to choose \hat{v} : mimic the outcome of the second-price auction as much as possible. Notice that with n identically distributed agents, the *ex ante* (meaning: before the values are drawn) probability that any particular agent wins is $1/n$. To mimic the outcome of the second-price auction on any particular agent we could set a price \hat{v} so that the probability that the agent's value is above \hat{v} is exactly $1/n$, this price can be found by inverting the cumulative distribution function as $\hat{v} = F^{-1}(1 - 1/n)$. For the uniform distribution, the solution to this inverse is $\hat{v} = 1 - 1/n$. Unlike the second-price auction, posting a uniform price of \hat{v} may result in no winners (if all agent values are below \hat{v}) or an agent other than that with the highest value may win (if there are more than one agents with value above \hat{v}).

Theorem 1.5. *For values drawn independently and identically from any distribution F , the uniform pricing of $\hat{v} = F^{-1}(1 - 1/n)$ is an $e/e-1 \approx 1.58$ approximation to the optimal social surplus.*

Proof. The main idea of this proof is to compare three mechanisms. Let REF denote the second-price auction and its surplus (our reference mechanism). Let APX denote the uniform pricing and its surplus (our approximation mechanism). The second-price auction, REF, optimizes surplus, subject to the *ex post* (meaning: after the mechanism is run) *supply constraint* that at most one agent wins, and chooses to sell to each agent with *ex ante* probability $1/n$. Consider for comparison a third

mechanism UB that maximizes surplus subject to the constraint that each agent is served with ex ante probability at most $1/n$, but has no supply constraint, i.e., UB can serve multiple agents if it so chooses.

The first step in the proof is the simple observation that UB upper bounds REF, i.e., $UB \geq REF$. This is clear as both mechanisms serve each agent with ex ante probability $1/n$, but REF has an ex post supply constraint whereas UB does not. UB could simulate REF and get the exact same surplus, or it could do something even better. Conclude,

$$UB \geq REF. \quad (1.1)$$

In fact, UB will do something better than REF. First, observe that UB's optimization is independent between agents. Second, observe that the socially optimal way to serve an agent with ex ante probability $1/n$ is to offer her price $\hat{v} = F^{-1}(1 - 1/n)$. We now wish to calculate UB's *expected surplus*. Let $\mathbf{E}[v \mid v \geq \hat{v}]$ denote the expected value of an agent given that her value v is above the price \hat{v} . If we sell to an agent and all we know is that her value is above the price, this quantity is the expected surplus generated. By the choice of price \hat{v} , the probability that an agent has a value v that exceeds the price \hat{v} is $\Pr[v \geq \hat{v}] = 1/n$, and when an agent's value is below the price her surplus is zero. Thus, her (total) expected surplus in UB is exactly $\mathbf{E}[v \mid v \geq \hat{v}] \cdot \Pr[v \geq \hat{v}]$. By linearity of expectation, UB's (total) expected surplus is just the sum over the n agents of the surplus of each agent's surplus. Therefore,

$$\begin{aligned} UB &= n \cdot \mathbf{E}[v \mid v \geq \hat{v}] \cdot \Pr[v \geq \hat{v}] \\ &= \mathbf{E}[v \mid v \geq \hat{v}]. \end{aligned} \quad (1.2)$$

Finally, we get a lower bound on APX's surplus that we can relate to REF via its upper bound UB. Recall that the price in the uniform-pricing mechanism is selected so that the probability that any given agent has value exceeding the price is exactly $1/n$. The probability that there are no agents who are above the price is equal to the probability that all agents are below the price, which is equal to the product of the probabilities that each agent is below the threshold, i.e., $(1 - 1/n)^n \leq 1/e$.³ Therefore, the probability that the item is sold by uniform pricing is at least $1 - 1/e$. If the item is sold, it is sold to an arbitrary agent with value conditioned to be at least \hat{v} , and the expected value of any such agent

³ The *natural number* is $e \approx 2.178$. That $\lim_{n \rightarrow \infty} (1 - 1/n)^n = 1/e$ can be verified by taking the natural logarithm and applying L'Hopital's rule; the non-negativity of the derivative of $(1 - 1/n)^n$ implies it is monotone non-decreasing; therefore, $1/e$ is an upper bound on $(1 - 1/n)^n$ for any finite n .

is $\mathbf{E}[v \mid v \geq \hat{v}]$. Therefore, the expected surplus of uniform pricing is,

$$\text{APX} \geq (1 - 1/e)\mathbf{E}[v \mid v \geq \hat{v}]. \quad (1.3)$$

Combining equations (1.1), (1.2), and (1.3) it is apparent that $\text{APX} \geq (1 - 1/e)\text{REF}$. \square

Question 1.5. When are simple, practical mechanisms like posted pricing a good approximation to the optimal mechanism?

1.1.3 General Routing Mechanisms

Finally we are ready to propose a mechanism for congestion control and routing in general networks. The main idea in the construction is the notion of critical values that was central to showing that the second-price auction has truth-telling as a dominant strategy (Theorem 1.3). In fact, that proof generalizes to any auction wherein each agent faces a critical value (that is not a function of her bid), the agent wins and pays the critical value if her bid exceeds it, and otherwise she loses.

Definition 1.8. The *second-price routing mechanism* is:

- (i) solicit sealed bids,
- (ii) find the set of messages that can be routed simultaneously with the largest total value, and
- (iii) charge the agents of each routed message their critical values.

Theorem 1.6. *The second-price routing mechanism has truthful bidding as a dominant strategy.*

Corollary 1.7. *The second-price routing mechanism maximizes the social surplus.*

The proof of the theorem is similar to the analogous result for the second-price single-item auction, but we will defer its proof to Chapter 3. The corollary follows because the bids are equal to the agents' values, the mechanism is defined to be optimal for the reported bids, and the payments cancel.

Unfortunately, this is far from the end of the story. Step (ii) of the mechanism is known as *winner determination*. To understand exactly what is happening in this step we must be more clear about our model for routing in general networks. For instance, in the Internet, the route that messages take in the network is predetermined by the Border Gateway Protocol (BGP), which enforces that all messages routed to the same

destination through any given router follow the same path. There are no opportunities for load-balancing, i.e., for sending messages to the same destination across different paths so as to keep the loads on any given path at a minimum. Alternatively, we could be in a novel network where the routing can determine which messages to route and which path to route them on.

Once we fix a model, we need to figure out how to solve the optimization problem implied by winner determination. Namely, how do we find the subset of messages with the highest total value that can be simultaneously routed? In principle, we are searching over subsets that meet some complicated feasibility condition. Purely from the point of optimization, this is a challenging task. The problem is related to the infamous *disjoint paths* problems: given a set of pairs of vertices in a graph, find a subset of pairs that can be connected via disjoint paths. This problem is *NP hard* to solve. Meaning: it is at least as hard as any problem in the equivalence class of *NP-complete* problems for which it is widely believed that finding optimal solutions is computationally intractable.

Theorem 1.8. *The disjoint-paths problem is NP hard.*

If we believe it is impossible for a designer to implement a mechanism for which *winner determination* is computationally intractable, we cannot accept the second-price routing mechanism as a solution to the general network routing problem.

Algorithmic theory has an answer to intractability: if computing the optimal solution is intractable, try instead to compute an approximately optimal solution.

Question 1.6. Can we replace Step (ii) in the mechanism with an approximation algorithm and still retain the dominant-strategy incentive property?

Question 1.7. If not, can we (by some other method) design a computationally tractable approximation mechanism for routing?

Question 1.8. Is there a general theory for designing approximation mechanisms from approximation algorithms?

1.2 Mechanism Design

Mechanism design gives a theory for the design of protocols, services,

laws, or other “rules of interaction” in which selfish behavior leads to good outcomes. “Selfish behavior” means that each participant, hereafter *agent*, individually tries to maximize her own utility. Such behavior we define as rational. “Leads” means *in equilibrium*. A set of agent strategies is in equilibrium if no agent prefers to unilaterally change her strategy. Finally, the “good”-ness of an outcome is assessed with respect to the criteria or goals of the designer. Natural economic criteria are *social surplus*, the sum of the utilities of all parties; and *profit*, the total payments made to the mechanism less any cost for providing the outcome.

A theory for mechanism design should satisfy the following four desiderata:

Informative: It pinpoints salient features of the environment and characteristics of good mechanisms therein.

Prescriptive: It gives concrete suggestions for how a good mechanism should be designed.

Predictive: The mechanisms that the theory predicts should be the same as the ones observed in practice.

Tractable: The theory should not assume super-natural ability for the agents or designer to optimize.

Notice that optimality is not one of the desiderata, nor is suggesting a specific mechanism to a practitioner. Instead, intuition from the theory of mechanism design should help guide the design of good mechanisms in practice. Such guidance is possible through informative observations about what good mechanisms do. Observations that are robust to variations in modeling details are especially important.

Sometimes the theory of *optimal mechanism design* meets the above desiderata. The question of designing an optimal mechanism can be viewed as a standard optimization problem: given incentive constraints, imposed by game theoretic strategizing; feasibility constraints, imposed by the environment; and the distribution of agent preferences, optimize the designer’s given objective. In ideal environments the given constraints may simplify and, for instance, allow the mechanism design problem to be reduced to a natural optimization problem without incentive constraints or distribution. We saw an example of this for routing in general networks: in order to invoke the second-price mechanism we only needed to find the optimal set of messages to route. Unfortunately, there are many environments and objectives where the optimal mechanism design problem not simplify as nicely.

1.3 Approximation

In environments where optimal mechanisms do not meet the desiderata above, approximation can provide a remedy. In the formal definition of an approximation, below, a good mechanism is one with a small approximation factor.

Definition 1.9. For an environment given implicitly, denote an *approximation mechanism* and its performance by APX, and a *reference mechanism* and its performance by REF.

- (i) For any environment, APX is a β approximation to REF if $\text{APX} \geq \frac{1}{\beta} \text{REF}$.
- (ii) For any class of environments, a class of mechanisms is a β approximation to REF if for any environment in the class there is a mechanism APX in the class that is a β approximation to REF.
- (iii) For any class of environments, a mechanism APX is a β approximation to REF if for any environment in the class APX is a β approximation to REF.

In the preceding section we saw each of these types of approximation. For i.i.d. $U[0, 1]$, n -agent, single-item environments, posting a uniform price of $\hat{v} = 1 - 1/n$ is a $e/e-1$ approximation to the second-price auction. More generally, for any i.i.d. single-item environment, uniform pricing is a $e/e-1$ approximation to the second-price auction. Finally, for any single-item environment the lottery gives an n approximation to the social surplus of the second-price auction.

Usually we will employ the approximation framework with REF representing the optimal mechanism. For instance, in the preceding section we compared a posted-pricing mechanism to the surplus-optimal second-price auction for i.i.d., single-item environments. For such a comparison, clearly $\text{REF} \geq \text{APX}$, and therefore the approximation factor is at least one. It is often instructive to compare the approximation ability of one class of mechanisms to another. For instance, in the preceding section we compared the surplus of a lottery, as the optimal mechanism without payments, to the surplus of the second-price auction, the optimal mechanism (in general). This kind of apples-to-oranges comparison is useful for understanding the relative importance of various features of a mechanism or environment.

1.3.1 Philosophy of Approximation

While it is, no doubt, a compelling success of the theory of mechanism design that its mechanisms are so prevalent in practice, optimal mechanism design cannot claim the entirety of the credit. These mechanisms are employed by practitioners well beyond the environments for which they are optimal. Approximation can explain why: the mechanisms that are optimal in ideal environments may continue to be approximately optimal much more broadly. It is important for the theory to describe how broadly these mechanisms are approximately optimal and how close to optimal they are. Thus, the theory of approximation can complement the theory of optimality and justify the wide prevalence of certain mechanisms. For instance, in Chapter 4 and ?? we describe how the widely prevalent reserve-price-based mechanisms and posted pricings are corroborated by their approximate optimality.

There are natural environments for mechanism design wherein every “undominated” mechanism is optimal. If we consider only optimal mechanisms we are stuck with the full class from which we can make no observations about what makes a mechanism good; on the other hand, if we relax optimality, we may be able to identify a small subclass of mechanisms that are approximately optimal, i.e., for any environment there is a mechanism in the subclass that approximates the optimal mechanism. This subclass is important in theory as we can potentially observe salient characteristics of it. It is important in practice because, while it is unlikely for a real mechanism designer to be able to optimize over all mechanisms, optimizing over a small class of, hopefully, natural mechanisms may be possible. For instance, a conclusion that we will make precise in Chapter 4 and ?? is that reserve-price-based mechanisms and posted pricings are approximately optimal in a wide range of environments including those with multi-dimensional agent preferences.

Approximation provides a lens with which to explore the salient features of an environment or mechanism. Suppose we wish to determine whether a particular feature of a mechanism is important. If there exists a subclass of mechanisms without that feature that gives a good approximation to the optimal mechanism, then the feature is perhaps not that important. If, on the other hand, there is no such subclass then the feature is quite important. For instance, previously in this chapter we saw that mechanisms without transfers cannot obtain better than a linear approximation to the optimal social surplus in single-item environments. This result suggests that transfers are very important for mechanism

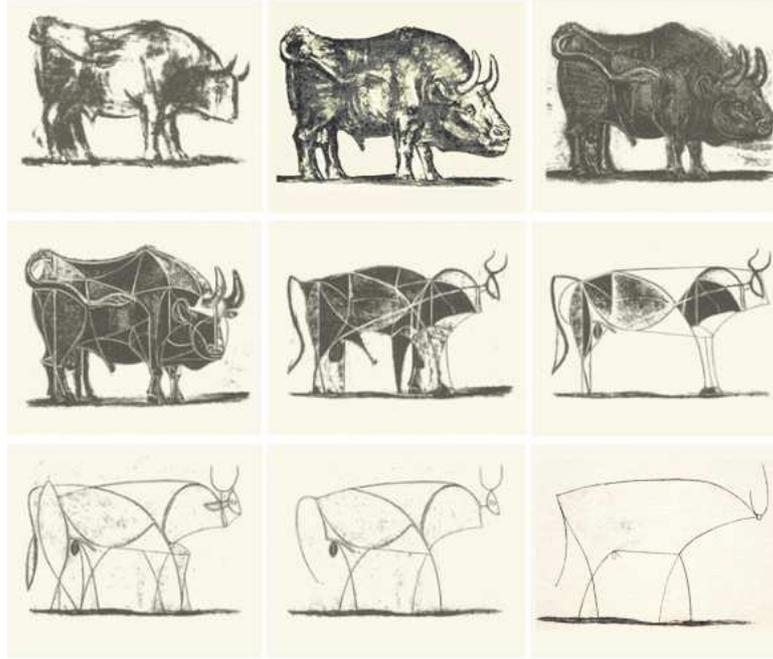


Figure 1.3. Picasso's December, 1945 to January, 1946 abstractionist study of a bull highlights one of the main points of approximation: identifying the salient features of the object of study. Picasso drew these in order from left to right, top to bottom.

design. On the other hand, we also saw that posted-pricing mechanism could obtain an $e/e-1$ approximation to the surplus-optimal mechanism. Posted pricings do not make use of competition between agents, therefore, we can conclude that competition between agents is not that important. Essentially, approximation provides a means to determine which aspect of an environment are details and which are not details. The approximation factor quantifies the relative importance on the spectrum between unimportant details to salient characteristics. Approximation, then allows for design of mechanisms that are not so dependent on details of the setting and therefore more robust. See Figure 1.3 for an illustration of this principle. In particular, in Chapter 4 we will formally observe that revenue-optimal auctions when agent values are drawn from a distribution can be approximated by a mechanism in which the only distributional dependence is a single number; moreover, in Chapter 5 we

will observe that some environments permit a single (prior-independent) mechanism to approximate the revenue-optimal mechanism under any distributional assumption.

Suppose the seller of an item is worried about collusion, risk attitudes, after-market effects, or other economic phenomena that are usually not included in standard ideal models for mechanism design. One option would be to explicitly model these effects and study optimal mechanisms in the augmented model. These complicated models are difficult to analyze and optimal mechanisms may be overly influenced by insignificant-seeming modeling choices. Optimal mechanisms are precisely tuned to details in the model and these details may drive the form of the optimal mechanism. On the other hand, we can consider approximations that are robust to various out-of-model phenomena. In such an environment the comparison between the approximation and the optimal mechanism is unfair because the optimal mechanism may suffer from out-of-model phenomena that the approximation is robust to. In fact, this “optimal mechanism” may perform much worse than our approximation when the phenomena are explicitly modeled. For example, Chapter 4 and ?? describe posted pricing mechanisms that are approximately optimal and robust to timing effects; for this reason an online auction house, such as eBay, may prefer its sellers to use “buy it now” posted pricings instead of auctions.

Finally, there is an issue of non-robustness that is inherent in any optimization over a complex set of objects, such as mechanisms. Suppose the designer does not know the distribution of agent preferences exactly but can learn about it through, e.g., market analysis. Such a market analysis is certainly going to be noisy; exactly optimizing a mechanism to the market analysis may “over fit” to this noise. Both statistics and machine learning theory have techniques for addressing this sort of overfitting. Approximation mechanisms also provide such a robustness. Since the class of approximation mechanisms is restricted from the full set, for these mechanisms to be good, they must pay less attention to details and therefore are robust to sampling noise. Importantly, approximation allows for design and analysis mechanisms for small (a.k.a., *thin*) markets where statistical and machine learning methods are less applicable.

1.3.2 Approximation Factors

Depending on the problem and the approximation mechanism, approximation factors can range from $(1 + \epsilon)$, i.e., arbitrarily close approxima-

tions, to linear factor approximations (or sometimes even worse). Notice a linear factor approximation is one where, as some parameter in the environment grows, i.e., more agents or more resources, the approximation factor gets worse. As examples, we saw earlier an environment in which uniform pricing is a *constant approximation* and the lottery is a *linear approximation*.⁴

In this text we take constant versus super-constant approximation as the separation between good and bad. We will view a proof that a mechanism is a constant approximation as a positive result and a proof that no mechanism (in a certain class) is a constant approximation as a negative result. Constant approximations tend to represent a tradeoff between simplicity and optimality. Properties of constant approximation mechanisms can, thus, be quite informative. Of course, there are many non-mechanism-design environments where super-constant approximations are both useful and informative; however, for mechanism design super-constant approximations tend to be indicative of (a) a bad mechanism, (b) failure to appropriately characterize optimal mechanisms, or (c) an imposition of incompatible modeling assumptions or constraints.

If you were approached by a seller (henceforth: principal) to design a mechanism and you returned to triumphantly reveal an elegant mechanism that gives her a two approximation to the optimal profit, you would probably find her a bit discouraged. After all, your mechanism leaves half of her profit on the table. In the context of this critique we outline the main points of constant, e.g., two, approximations for the practitioner. First, a two approximation provides informative conclusions that can guide the design of even better mechanisms for specific environments. Second, the approximation factor of two is a theoretical result that holds in a large range of environments, in specific environments the mechanism may perform better. It is easy, via simulation, to evaluate the mechanism performance on specific settings to see how close to optimal it actually is. Third, in many environments the optimal mechanism is not understood at all, meaning the principal's alternative to your two approximation is an ad hoc mechanism with no performance guarantee. This principal is of course free to simulate your mechanism and her mechanism in her given environment and decide to use the bet-

⁴ Recall that the approximation factor for uniform pricing bounded by $e/e-1$, an absolute constant that does not increase with various parameters of the auction such as the number of agents. In contrast the approximation factor of the lottery could be as bad as n , the number of agents. As the number of agents increases, so does the approximation bound guaranteed by the lottery.

ter of the two. In this fashion the principal's ad hoc mechanism, if used, is provably a two approximation as well. Fourth, mechanisms that are two approximations in theory arise in practice. In fact, that it is a two approximation explains why the mechanism arises. Even though it is not optimal, it is close enough. If it was far from being optimal the principal (hopefully) would have figured this out and adopted a different approach.

Sometimes it is possible to obtain schemas for approximating the optimal mechanism to within a $(1 + \epsilon)$ factor for any ϵ . These schemas tend to be computational approaches that are useful for addressing potential computational intractability of the optimal mechanism design problem. While they do not tend to yield simple mechanisms, they are relevant in complex environments. Often these approximation schemes are based on (a) identifying a restricted class of mechanisms wherein a near-optimal mechanism can be found and (b) conducting a brute-force search over this restricted class. While very little is learned from such a brute-force search, properties of the restricted class of mechanisms can be informative. Many of the optimal mechanisms we describe can in practice only be implemented as approximation schemes.

Chapter Notes

Routing and congestion control are a central problems in computer systems such as the Internet; see Leiner et al. (1997) for a discussion of design criteria. Demers et al. (1989) analyze "fair queuing" which is a lottery-based mechanism for congestion control. Griffin et al. (2002) discuss the Border Gateway Protocol (BGP) which determines the routes messages take in the Internet. The NP-completeness of the disjoint paths problem (and the related problem of integral multi-commodity flow) was established by Even et al. (1976).

William Vickrey's 1961 analysis of the second-price auction is one of the pillars of mechanism design theory. The second-price routing mechanism is a special case of the more general Vickrey-Clarke-Groves (VCG) mechanism which is attributed additionally to Edward Clarke (1971) and Theodore Groves (1973).

Computational payments were proposed as means for fighting unsolicited electronic mail by Dwork and Naor (1992). Hartline and Roughgarden (2008) consider mechanism design with the objective of residual surplus and describe distributional assumptions under which the lottery

is optimal, the second-price auction is optimal, and when neither are optimal. They also give a single mechanism that approximates the optimal mechanism for any distribution of agent values.

Vincent and Manelli (2007) showed that there are environments for mechanism design wherein every “undominated” mechanism is optimal for some distribution of agent preferences. This result implies that optimality cannot be used to identify properties of good mechanisms. Robert Wilson (1987) suggested that mechanisms that are less dependent on the details of the environment are likely to be more relevant. This suggestion is known as the “Wilson doctrine.”

The $\epsilon/e-1$ approximation via a uniform pricing (Theorem 1.5) is a consequence of Chawla et al. (2010b). Wang et al. (2008) and Reynolds and Wooders (2009) discuss why the “buy it now” (i.e., posted-pricing) mechanism is replacing the second-price auction format in eBay.

2

Equilibrium

The theory of *equilibrium* attempts to predict what happens in a game when players behave strategically. This is a central concept to this text as, in mechanism design, we are optimizing over games to find games with good equilibria. Here, we review the most fundamental notions of equilibrium. They will all be static notions in that players are assumed to understand the game and will play once in the game. While such foreknowledge is certainly questionable, some justification can be derived from imagining the game in a dynamic setting where players can learn from past play.

This chapter reviews equilibrium in both complete and incomplete information games. As games of incomplete information are the most central to mechanism design, special attention will be paid to them. In particular, we will characterize equilibrium when the private information of each agent is single-dimensional and corresponds, for instance, to a value for receiving a good or service. We will show that auctions with the same equilibrium outcome have the same expected revenue. Using this so-called *revenue equivalence* we will describe how to solve for the equilibrium strategies of standard auctions in symmetric environments.

Our emphasis will be on demonstrating the central theories of equilibrium and not on providing the most comprehensive or general results. For that readers are recommended to consult a game theory textbook.

2.1 Complete Information Games

In games of complete information all players are assumed to know precisely the payoff structure of all other players for all possible outcomes

of the game. A classic example of such a game is the *prisoner's dilemma*, the story for which is as follows.

Two prisoners, Bonnie and Clyde, have jointly committed a crime and are being interrogated in separate quarters. Unfortunately, the interrogators are unable to prosecute either prisoner without a confession. Bonnie is offered the following deal: If she confesses and Clyde does not, she will be released and Clyde will serve the full sentence of ten years in prison. If they both confess, she will share the sentence and serve five years. If neither confesses, she will be prosecuted for a minimal offense and receive a year of prison. Clyde is offered the same deal.

This story can be expressed as the following *bimatrix game* where entry (a, b) represents row player's payoff a and column player's payoff b .

	silent	confess
silent	(-1,-1)	(-10,0)
confess	(0,-10)	(-5,-5)

A simple thought experiment enables prediction of what will happen in the prisoners' dilemma. Suppose the Clyde is silent. What should Bonnie do? Remaining silent as well results in one year of prison while confessing results in immediate release. Clearly confessing is better. Now suppose that Clyde confesses. Now what should Bonnie do? Remaining silent results in ten years of prison while confessing as well results in only five. Clearly confessing is better. In other words, no matter what Clyde does, Bonnie is better off by confessing. The prisoners dilemma is hardly a dilemma at all: the *strategy profile* (confess, confess) is a *dominant strategy equilibrium*.

Definition 2.1. A *dominant strategy equilibrium* (DSE) in a complete information game is a strategy profile in which each player's strategy is as least as good as all other strategies regardless of the strategies of all other players.

Dominant strategy equilibrium is a strong notion of equilibrium and is therefore unsurprisingly rare. For an equilibrium notion to be complete it should identify equilibrium in every game. Another well studied game is *chicken*.

James Dean and Buzz (in the movie *Rebel without a Cause*) face off at opposite ends of the street. On the signal they race their cars on a collision course towards each other. The options each have are to swerve or to stay their course. Clearly if they both stay their course they crash. If they both swerve (opposite directions) they escape with their lives but the match is a draw.

Finally, if one swerves and the other stays, the one that stays is the victor and the other the loses.¹

A reasonable bimatrix game depicting this story is the following.

	stay	swerve
stay	(-10,-10)	(1,-1)
swerve	(-1,1)	(0,0)

Again, a simple thought experiment enables us to predict how the players might play. Suppose James Dean is going to stay, what should Buzz do? If Buzz stays they crash and Buzz's payoff is -10 , but if Buzz swerves his payoff is only -1 . Clearly, of these two options Buzz prefers to swerve. Suppose now that Buzz is going to swerve, what should James Dean do? If James Dean stays he wins and his payoff is one, but if he swerves it is a draw and his payoff is zero. Clearly, of these two options James Dean prefers to stay. What we have shown is that the strategy profile (stay, swerve) is a mutual best response, a.k.a., a *Nash equilibrium*. Of course, the game is symmetric so the opposite strategy profile (swerve, stay) is also an equilibrium.

Definition 2.2. A *Nash equilibrium* in a game of complete information is a strategy profile where each player's strategy is a best response to the strategies of the other players as given by the strategy profile.

In the examples above, the strategies of the players correspond directly to actions in the game, a.k.a., *pure strategies*. In general, Nash equilibrium strategies can be randomizations over actions in the game, a.k.a., *mixed strategies* (see Exercise 2.1).

2.2 Incomplete Information Games

Now we turn to the case where the payoff structure of the game is not completely known. We will assume that each agent has some private information and this information affects the payoff of this agent in the game. We will refer to this information as the agent's type and denote it by t_i for agent i . The profile of types for the n agents in the game is $\mathbf{t} = (t_1, \dots, t_n)$.

A *strategy* in a game of incomplete information is a function that maps

¹ The actual chicken game depicted in *Rebel without a Cause* is slightly different from the one described here.

an agent's type to any of the agent's possible actions in the game (or a distribution over actions for mixed strategies). We will denote by $s_i(\cdot)$ the strategy of agent i and $\mathbf{s} = (s_1, \dots, s_n)$ a *strategy profile*.

The auctions described in Chapter 1 were games of incomplete information where an agent's private type was her value for receiving the item, i.e., $t_i = v_i$. As we described, strategies in the ascending-price auction were $s_i(v_i) = \text{"drop out when the price exceeds } v_i\text{"}$ and strategies in the second-price auction were $s_i(v_i) = \text{"bid } b_i = v_i\text{"}$. We refer to this latter strategy as *truth-telling*. Both of these strategy profiles are in *dominant strategy equilibrium* for their respective games.

Definition 2.3. A *dominant strategy equilibrium* (DSE) is a strategy profile \mathbf{s} such that for all i , t_i , and \mathbf{b}_{-i} (where \mathbf{b}_{-i} generically refers to the actions of all players but i), agent i 's utility is maximized by following strategy $s_i(t_i)$.

Notice that aside from strategies being defined as a map from types to actions, this definition of DSE is identical to the definition of DSE for games of complete information.

2.3 Bayes-Nash Equilibrium

Naturally, many games of incomplete information do not have dominant strategy equilibria. Therefore, we will also need to generalize Nash equilibrium to this setting. Recall that equilibrium is a property of a strategy profile. It is in equilibrium if each agent does not want to change her strategy given the other agents' strategies. For an agent i , we want to fix other agent strategies and let i optimize her strategy (meaning: calculate her best response for all possible types t_i she may have). This is an ill specified optimization as just knowing the other agents' strategies is not enough to calculate a best response. Additionally, i 's best response depends on i 's beliefs on the types of the other agents. The standard economic treatment addresses this by assuming a common prior.

Definition 2.4. Under the *common prior assumption*, the agent types \mathbf{t} are drawn at random from a *prior distribution* \mathbf{F} (a joint probability distribution over type profiles) and this prior distribution is *common knowledge*.

The distribution \mathbf{F} over \mathbf{t} may generally be correlated. Which means that an agent with knowledge of her own type must do *Bayesian updating*

to determine the distribution over the types of the remaining bidders. We denote this conditional distribution as $\mathbf{F}_{-i}|_{t_i}$. Of course, when the distribution of types is independent, i.e., \mathbf{F} is the *product distribution* $F_1 \times \cdots \times F_n$, then $\mathbf{F}_{-i}|_{t_i} = \mathbf{F}_{-i}$.

Notice that a prior \mathbf{F} and strategies \mathbf{s} induces a distribution over the actions of each of the agents. With such a distribution over actions, the problem each agent faces of optimizing her own action is fully specified.

Definition 2.5. A *Bayes-Nash equilibrium (BNE)* for a game G and common prior \mathbf{F} is a strategy profile \mathbf{s} such that for all i and t_i , $s_i(t_i)$ is a best response when other agents play $\mathbf{s}_{-i}(t_{-i})$ when $t_{-i} \sim \mathbf{F}_{-i}|_{t_i}$.

To illustrate Bayes-Nash equilibrium, consider using the first-price auction to sell a single item to one of two agents, each with valuation drawn independently and identically from the uniform distribution on $[0, 1]$, i.e., the common prior distribution is $\mathbf{F} = F \times F$ with $F(z) = \Pr_{v \sim F}[v < z] = z$. Here each agent's type is her valuation. We will calculate the BNE of this game by the "guess and verify" technique. First, we guess that there is a symmetric BNE with $s_i(z) = z/2$ for $i \in \{1, 2\}$. Second, we calculate agent 1's expected utility with value v_1 and bid b_1 under the standard assumption that the agent's utility u_i is her value less her payment (when she wins). In this calculation v_1 and b_1 are fixed and $b_2 = v_2/2$ is random. By the definition of the first-price auction:

$$\mathbf{E}[u_1] = (v_1 - b_1) \times \Pr[1 \text{ wins with bid } b_1].$$

Calculate $\Pr[1 \text{ wins with } b_1]$ as

$$\begin{aligned} \Pr[b_2 \leq b_1] &= \Pr[v_2/2 \leq b_1] = \Pr[v_2 \leq 2b_1] = F(2b_1) \\ &= 2b_1. \end{aligned}$$

Thus,

$$\begin{aligned} \mathbf{E}[u_1] &= (v_1 - b_1) \times 2b_1 \\ &= 2v_1b_1 - 2b_1^2. \end{aligned}$$

Third, we optimize agent 1's bid. Agent 1 with value v_1 should maximize $2v_1b_1 - 2b_1^2$ as a function of b_1 , and to do so, can differentiate the function and set its derivative equal to zero. The result is $\frac{d}{db_1}(2v_1b_1 - 2b_1^2) = 2v_1 - 4b_1 = 0$ and we can conclude that the optimal bid is $b_1 = v_1/2$. This proves that agent 1 should bid as prescribed if agent 2 does; and vice versa. Thus, we conclude that the guessed strategy profile is in BNE.

In Bayesian games it is useful to distinguish between stages of the game in terms of the knowledge sets of the agents. The three stages of a Bayesian game are *ex ante*, *interim*, and *ex post*. The *ex ante* stage is before values are drawn from the distribution. *Ex ante*, the agents know this distribution but not their own types. The *interim* stage is immediately after each agent learns her own type, but before playing in the game. In the *interim*, an agent assumes the other agent types are drawn from the prior distribution conditioned on her own type, i.e., via *Bayesian updating*. In the *ex post* stage, the game is played and the actions of all agents are known.

2.4 Single-dimensional Games

We will focus on a conceptually simple class of single-dimensional games that is relevant to the auction problems we have already discussed. In a single-dimensional game, each agent's private type is her value for receiving an abstract service, i.e., $t_i = v_i$. The distribution over types is independent (i.e., a product distribution). A game has an outcome $\mathbf{x} = (x_1, \dots, x_n)$ and payments $\mathbf{p} = (p_1, \dots, p_n)$ where x_i is an indicator for whether agent i indeed received their desired service, i.e., $x_i = 1$ if i is served and 0 otherwise. Price p_i will denote the payment i makes to the mechanism. An agent's value can be positive or negative and an agent's payment can be positive or negative. An agent's utility is linear in her value and payment and specified by $u_i = v_i x_i - p_i$. Agents are risk-neutral expected utility maximizers.

Definition 2.6. A *single-dimensional linear utility* is defined as having utility $u = vx - p$ for service-payment outcomes (x, p) and private value v ; a *single-dimensional linear agent* possesses such a utility function.

A game G maps actions \mathbf{b} of agents to an outcome and payment. Formally we will specify these outcomes and payments as:

- $x_i^G(\mathbf{b}) =$ outcome to i when actions are \mathbf{b} , and
- $p_i^G(\mathbf{b}) =$ payment from i when actions are \mathbf{b} .

Given a game G and a strategy profile \mathbf{s} we can express the outcome and payments of the game as a function of the valuation profile. From the point of view of analysis this description of the the game outcome is much more relevant. Define

- $x_i(\mathbf{v}) = x_i^G(\mathbf{s}(\mathbf{v}))$, and
- $p_i(\mathbf{v}) = p_i^G(\mathbf{s}(\mathbf{v}))$.

We refer to the former as the *allocation rule* and the latter as the *payment rule* for G and \mathbf{s} (implicit). Consider an agent i 's interim perspective. She knows her own value v_i and believes the other agents values to be drawn from the distribution \mathbf{F} (conditioned on her value). For G , \mathbf{s} , and \mathbf{F} taken implicitly we can specify agent i 's interim allocation and payment rules as functions of v_i .

- $x_i(v_i) = \Pr[x_i(v_i) = 1 \mid v_i] = \mathbf{E}[x_i(\mathbf{v}) \mid v_i]$, and
- $p_i(v_i) = \mathbf{E}[p_i(\mathbf{v}) \mid v_i]$.

With linearity of expectation we can combine these with the agent's utility function to write

- $u_i(v_i) = v_i x_i(v_i) - p_i(v_i)$.

Finally, we say that a strategy $s_i(\cdot)$ is *onto* if every action b_i agent i could play in the game is prescribed by s_i for some value v_i , i.e., $\forall b_i \exists v_i s_i(v_i) = b_i$. We say that a strategy profile is *onto* if the strategy of every agent is onto. For instance, the truthtelling strategy in the second-price auction is onto. When the strategies of the agents are onto, the interim allocation and payment rules defined above completely specify whether the strategies are in equilibrium or not. In particular, BNE requires that each agent (weakly) prefers playing the action corresponding (via their strategy) to her value than the action corresponding to any other value.

Proposition 2.1. *When values are drawn from a product distribution \mathbf{F} ; single-dimensional game G and strategy profile \mathbf{s} is in BNE only if for all i , v_i , and z ,*

$$v_i x_i(v_i) - p_i(v_i) \geq v_i x_i(z) - p_i(z).$$

If the strategy profile is onto then the converse also holds.

Notice that in Proposition 2.1 the distribution \mathbf{F} is required to be a product distribution. If \mathbf{F} is not a product distribution, then when agent i 's value is v_i then $x_i(z)$ is not generally the probability that she will win when she follows her designated strategy for value z . This distinction arises because the conditional distribution of the other agents values need not be the same when i 's value is v_i or z .

2.5 Characterization of Bayes-Nash Equilibrium

We now discuss what Bayes-Nash equilibria look like. For instance, when given G , \mathbf{s} , and \mathbf{F} we can calculate the interim allocation and payment rules $x_i(v_i)$ and $p_i(v_i)$ of each agent. We want to succinctly describe properties of these allocation and payment rules that can arise as BNE.

Theorem 2.2. *When values are drawn from a continuous product distribution \mathbf{F} ; single dimensional G and strategy profile \mathbf{s} are in BNE only if for all i ,*

- (i) (monotonicity) $x_i(v_i)$ is monotone non-decreasing, and
- (ii) (payment identity) $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z) dz + p_i(0)$,

where often $p_i(0) = 0$. If the strategy profile is onto then the converse also holds.

Proof. We will prove the theorem in the special case where the support of each agent i 's distribution is $[0, \infty]$. Focusing on a single agent i , who we will refer to as Alice, we drop subscripts i from all notations.

We break this proof into three pieces. First, we show, by picture, that the game is in BNE if the characterization holds and the strategy profile is onto. Next, we will prove that a game is in BNE only if the monotonicity condition holds. Finally, we will prove that a game is in BNE only if the payment identity holds.

Note that if Alice with value v deviates from the equilibrium and takes action $s(v^\dagger)$ instead of $s(v)$ then she will receive outcome and payment $x(v^\dagger)$ and $p(v^\dagger)$. This motivates the definition,

$$u(v, v^\dagger) = vx(v^\dagger) - p(v^\dagger),$$

which corresponds to Alice utility when she makes this deviation. For Alice's strategy to be in equilibrium it must be that for all v , and v^\dagger , $u(v, v) \geq u(v, v^\dagger)$, i.e., Alice derives no increased utility by deviating. The strategy profile \mathbf{s} is in equilibrium if and only if the same condition holds for all agents. (The "if" direction here follows from the assumption that strategies map values onto actions. Meaning: for any action in the game there exists a value v^\dagger such that $s(v^\dagger)$ is that action.)

- (i) G , \mathbf{s} , and \mathbf{F} are in BNE if \mathbf{s} is onto and monotonicity and the payment identity hold.

We prove this by picture. Though the formulaic proof is simple, the pictures provide useful intuition. We consider two possible values

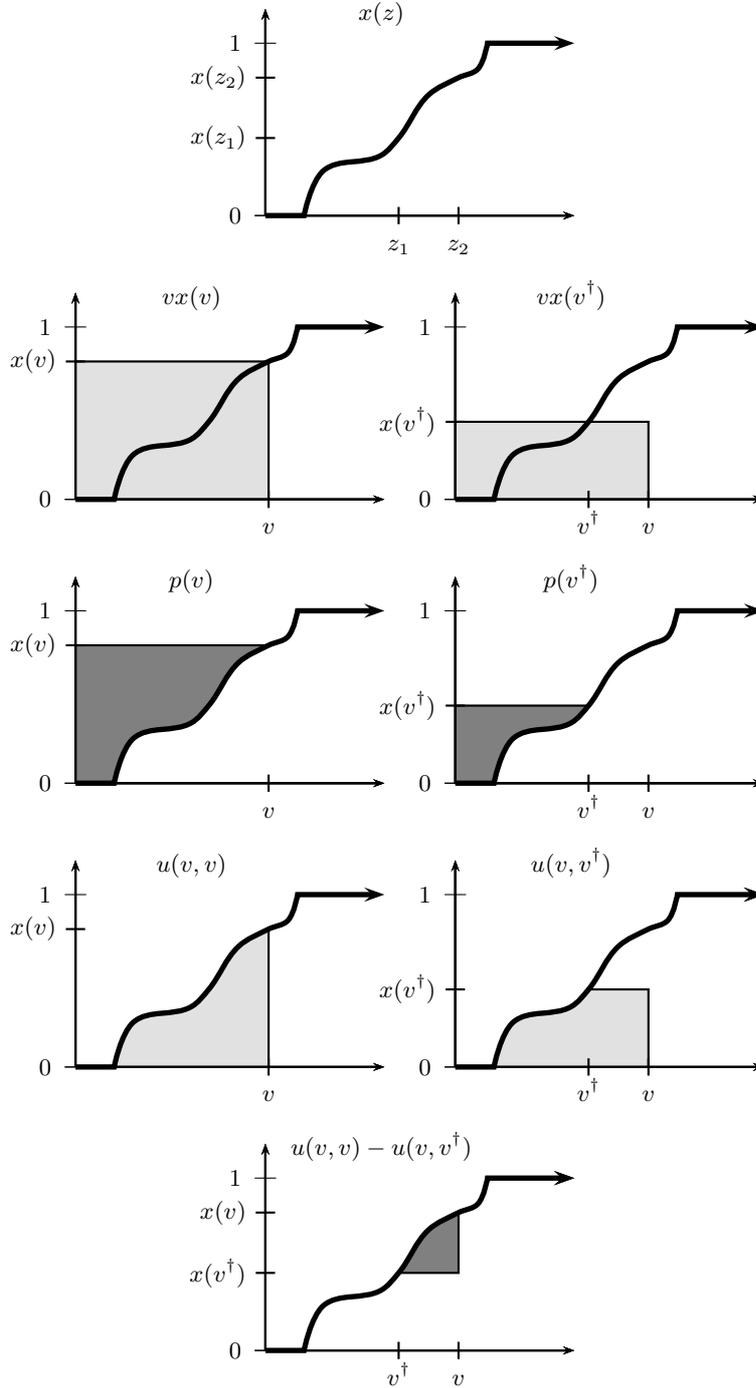


Figure 2.1. The left column shows (shaded) the surplus, payment, and utility of Alice playing action $s(v = z_2)$. The right column shows (shaded) the same for Alice playing action $s(v^\dagger = z_1)$. The final diagram shows (shaded) the difference between Alice's utility for these strategies. Monotonicity implies this difference is non-negative.

z_1 and z_2 with $z_1 < z_2$. Supposing Alice has the high value, $v = z_2$, we argue that Alice does not benefit by simulating her strategy for the lower value, $v^\dagger = z_1$, i.e., by playing $s(v^\dagger)$ to obtain outcome $x(v^\dagger)$ and payment $p(v^\dagger)$. We leave the proof of the opposite, that when $v = z_1$ and Alice is considering simulating the higher strategy $v^\dagger = z_2$, as an exercise for the reader.

To start with this proof, we assume that $x(v)$ is monotone and that $p(v) = vx(v) - \int_0^v x(z) dz$.

Consider the diagrams in Figure 2.1. The first diagram (top, center) shows $x(\cdot)$ which is indeed monotone as per our assumption. The column on the left shows Alice's surplus, $vx(v)$; payment, $p(v)$, and utility, $u(v) = vx(v) - p(v)$, assuming that she follows the BNE strategy $s(v = z_2)$. The column on the right shows the analogous quantities when Alice follows strategy $s(v^\dagger = z_1)$ but has value $v = z_2$. The final diagram (bottom, center) shows the difference in the Alice's utility for the outcome and payments of these two strategies. Note that as the picture shows, the monotonicity of the allocation function implies that this difference is always non-negative. Therefore, there is no incentive for Alice to simulate the strategy of a lower value.

As mentioned, a similar proof shows that Alice has no incentive to simulate her strategy for a higher value. We conclude that she (weakly) prefers to play the action given by the BNE $s(\cdot)$ over any other action in the range of her strategy function; since $s(\cdot)$ is onto this range includes all actions.

- (ii) G , \mathbf{s} , and \mathbf{F} are in BNE only if the allocation rule is monotone.

If we are in BNE then for all valuations, v and v^\dagger , $u(v, v) \geq u(v, v^\dagger)$. Expanding we require

$$vx(v) - p(v) \geq vx(v^\dagger) - p(v^\dagger).$$

We now consider z_1 and z_2 with $z_1 < z_2$ and take turns setting $v = z_1$, $v^\dagger = z_2$, and $v^\dagger = z_1$, $v = z_2$. This yields the following two inequalities:

$$v = z_2, v^\dagger = z_1 \implies z_2x(z_2) - p(z_2) \geq z_2x(z_1) - p(z_1), \text{ and} \quad (2.1)$$

$$v = z_1, v^\dagger = z_2 \implies z_1x(z_1) - p(z_1) \geq z_1x(z_2) - p(z_2). \quad (2.2)$$

Adding these inequalities and canceling the payment terms we have,

$$z_2x(z_2) + z_1x(z_1) \geq z_2x(z_1) + z_1x(z_2).$$

Rearranging,

$$(z_2 - z_1)(x(z_2) - x(z_1)) \geq 0.$$

For $z_2 - z_1 > 0$ it must be that $x(z_2) - x(z_1) \geq 0$, i.e., $x(\cdot)$ is monotone non-decreasing.

- (iii) G , \mathbf{s} , and \mathbf{F} are in BNE only if the payment rule satisfies the payment identity.

We will give two proofs that payment rule must satisfy $p(v) = vx(v) - \int_0^v x(z) dz + p(0)$; the first is a calculus-based proof under the assumption that each of $x(\cdot)$ and $p(\cdot)$ are differentiable and the second is a picture-based proof that requires no assumption.

Calculus-based proof: Fix v and recall that $u(v, z) = vx(z) - p(z)$. Let $u'(v, z)$ be the partial derivative of $u(v, z)$ with respect to z . Thus, $u'(v, z) = vx'(z) - p'(z)$, where $x'(\cdot)$ and $p'(\cdot)$ are the derivatives of $x(\cdot)$ and $p(\cdot)$, respectively. Since BNE implies that $u(v, z)$ is maximized at $z = v$. It must be that

$$u'(v, v) = vx'(v) - p'(v) = 0.$$

This formula must hold true for all values of v . For remainder of the proof, we treat this identity formulaically. To emphasize this, substitute $z = v$:

$$zx'(z) - p'(z) = 0.$$

Solving for $p'(z)$ and then integrating both sides of the equality from 0 to v we have,

$$\begin{aligned} p'(z) &= zx'(z), \text{ so} \\ \int_0^v p'(z) dz &= \int_0^v zx'(z) dz. \end{aligned}$$

Simplifying the left-hand side and adding $p(0)$ to both sides,

$$p(v) = \int_0^v zx'(z) dz + p(0).$$

Finally, we obtained the desired formula by integrating the right-hand side by parts,

$$\begin{aligned} p(v) &= \left[zx(z) \right]_0^v - \int_0^v x(z) dz + p(0) \\ &= vx(v) - \int_0^v x(z) dz + p(0). \end{aligned}$$

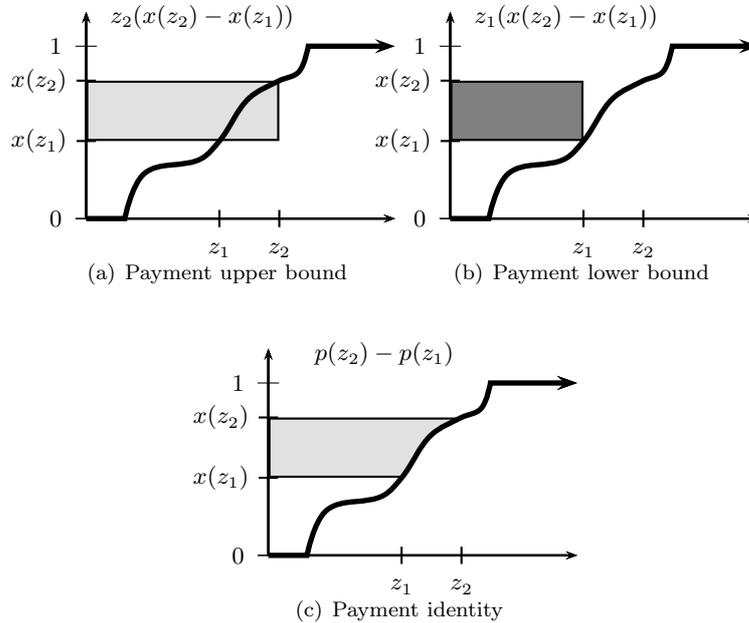


Figure 2.2. Upper (top, left) and lower bounds (top, right) for the difference in payments for two strategies z_1 and z_2 imply that the difference in payments (bottom) must satisfy the payment identity.

Picture-based proof: Consider equations (2.1) and (2.2) and solve for $p(z_2) - p(z_1)$ in each:

$$z_2(x(z_2) - x(z_1)) \geq p(z_2) - p(z_1) \geq z_1(x(z_2) - x(z_1)).$$

The first inequality gives an upper bound on the difference in payments for two types z_2 and z_1 and the second inequality gives a lower bound. It is easy to see that the only payment rule that satisfies these upper and lower bounds for all pairs of types z_2 and z_1 has payment difference exactly equal to the area to the left of the allocation rule between $x(z_1)$ and $x(z_2)$. See Figure 2.2. The payment identity follows by taking $z_1 = 0$ and $z_2 = v$. \square

As we conclude the proof of the BNE characterization theorem, it is important to note how little we have assumed of the underlying game. We did not assume it was a single-round, sealed-bid auction. We did not assume that only a winner will make payments. Therefore, we conclude

for any potentially wacky, multi-round game the outcomes of all Bayes-Nash equilibria have a nice form.

2.6 Characterization of Dominant Strategy Equilibrium

Dominant strategy equilibrium is a stronger equilibrium concept than Bayes-Nash equilibrium. All dominant strategy equilibria are Bayes-Nash equilibria, but as we have seen, the opposite is not true; for instance, there is no DSE in the first-price auction. Recall that a strategy profile is in DSE if each agent's strategy is optimal for her regardless of what other agents are doing. The DSE characterization theorem below follows from the BNE characterization theorem.

Theorem 2.3. *G and \mathbf{s} are in DSE only if for all i and \mathbf{v} ,*

- (i) (monotonicity) $x_i(v_i, \mathbf{v}_{-i})$ is monotone non-decreasing in v_i , and
- (ii) (payment identity) $p_i(v_i, \mathbf{v}_{-i}) = v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) dz + p_i(0, \mathbf{v}_{-i})$,

where (z, \mathbf{v}_{-i}) denotes the valuation profile with the i th coordinate replaced with z . If the strategy profile is onto then the converse also holds.

It was important when discussing BNE to explicitly refer to $x_i(v_i)$ and $p_i(v_i)$ as the probability of allocation and the expected payments because a game played by agents with values drawn from a distribution will inherently, from agent i 's perspective, have a randomized outcome and payment. In contrast, for games with DSE we can consider outcomes and payments in a non-probabilistic sense. A deterministic game, i.e., one with no internal randomization, will result in deterministic outcomes and payments. For our single-dimensional game where an agent is either served or not served we will have $x_i(\mathbf{v}) \in \{0, 1\}$. This specification along with the monotonicity condition implied by DSE implies that the function $x_i(v_i, \mathbf{v}_{-i})$ is a step function in v_i . The reader can easily verify that the payment required for such a step function is exactly the critical value, i.e., \hat{v}_i at which $x_i(\cdot, \mathbf{v}_{-i})$ changes from 0 to 1. This gives the following corollary.

Corollary 2.4. *A deterministic game G and deterministic strategies \mathbf{s} are in DSE only if for all i and \mathbf{v} ,*

- (i) (step-function) $x_i(v_i, \mathbf{v}_{-i})$ steps from 0 to 1 at some $\hat{v}_i(\mathbf{v}_{-i})$, and

$$(ii) \text{ (critical value) } p_i(v_i, \mathbf{v}_{-i}) = \begin{cases} \hat{v}_i(\mathbf{v}_{-i}) & \text{if } x_i(v_i, \mathbf{v}_{-i}) = 1 \\ 0 & \text{otherwise} \end{cases} + p_i(0, \mathbf{v}_{-i}).$$

If the strategy profile is onto then the converse also holds.

Notice that the above theorem deliberately skirts around a subtle tie-breaking issue. Consider the truth-telling DSE of the second-price auction on two agents. What happens when $v_1 = v_2$? One agent should win and pay the other's value. As this results in a utility of zero, from the perspective of utility maximization, both agents are indifferent as to which of them it is. One natural tie-breaking rule is the lexicographical one, i.e., in favor of agent 1 winning. For this rule, agent 1 wins when $v_1 \in [v_2, \infty)$ and agent 2 wins when $v_2 \in (v_1, \infty)$. The critical values are $t_1 = v_2$ and $t_2 = v_1$. We will usually prefer the randomized tie-breaking rule because of its symmetry.

2.7 Revenue Equivalence

We are now ready to make one of the most significant observations in auction theory. Namely, mechanisms with the same outcome in BNE have the same expected revenue. In fact, not only do they have the same expected revenue, but each agent has the same expected payment in each mechanism. This result is in fact a direct corollary of Theorem 2.2. The payment identity means that the payment rule is precisely determined by the allocation rule and the payment of the lowest type, i.e., $p_i(0)$.

Corollary 2.5. *For any two mechanisms where 0-valued agents pay nothing, if the mechanisms have the same BNE outcome then they have same expected revenue.*

We can now quantitatively compare the second-price and first-price auctions from a revenue standpoint. Consider the case where the agent's values are distributed independently and identically. What is the equilibrium outcome of the second-price auction? The agent with the highest valuation wins. What is the equilibrium outcome of the first-price auction? This question requires a little more thought. Since the distributions are identical, it is reasonable to expect that there is a symmetric equilibrium, i.e., one where $s_i = s_{i'}$ for all i and i' . Furthermore, it is reasonable to expect that the strategies are monotone, i.e., an agent with a higher value will out bid an agent with a lower value. Under these assumptions,

the agent with the highest value wins. Of course, in both auctions a 0-valued agent will pay nothing. Therefore, we can conclude that the two auctions obtain the same expected revenue.

As an example of revenue equivalence consider first-price and second-price auctions for selling a single item to two agents with values drawn from $U[0, 1]$. The expected revenue of the second-price auction is $\mathbf{E}[v_{(2)}]$. In Section 2.3 we saw that the symmetric strategy of the first-price auction in this environment is for each agent to bid half her value. The expected revenue of first-price auction is therefore $\mathbf{E}[v_{(1)}/2]$. An important fact about uniform random variables is that in expectation they evenly divide the interval they are over, i.e., $\mathbf{E}[v_{(1)}] = 2/3$ and $\mathbf{E}[v_{(2)}] = 1/3$. How do the revenues of these two auctions compare? Their revenues are identically $1/3$.

Corollary 2.6. *When agents' values are independent and identically distributed according to a continuous distribution, the second-price and first-price auction have the same expected revenue.*

Of course, much more bizarre auctions are governed by revenue equivalence. As an exercise the reader is encouraged to verify that the *all-pay auction*; where agents submit bids, the highest bidder wins, and all agents pay their bids; is revenue equivalent to the first- and second-price auctions.

2.8 Solving for Bayes-Nash Equilibrium

While it is quite important to know what outcomes are possible in BNE, it is also often important to be able to solve for the BNE strategies. For instance, suppose you were a bidder bidding in an auction. How would you bid? In this section we describe an elegant technique for calculating BNE strategies in symmetric environments using revenue equivalence. Actually, we use something a little stronger than revenue equivalence: *interim payment equivalence*. This is the fact that if two mechanisms have the same allocation rule, they must have the same payment rule (because the payment rules satisfy the payment identity). As described previously, the interim payment of agent i with value v_i is $p_i(v_i)$.

Suppose we are to solve for the BNE strategies of mechanism M . The approach is to express an agent's payment in M as a function of the agent's action, then to calculate the agent's expected payment in a strategically-simple mechanism M' that is revenue equivalent to M

(usually a “second-price implementation” of M). Setting these terms equal and solving for the agents action gives the equilibrium strategy.

We give the high level the procedure below. As a running example we will calculate the equilibrium strategies in the first-price auction with two $U[0, 1]$ agents, in doing so we will use a calculation of expected payments in the strategically-simple second-price auction in the same environment.

- (i) *Guess* what the outcome might be in Bayes-Nash equilibrium.

E.g., in the BNE of the first-price auction with two agents with values $U[0, 1]$, we expect the agent with the highest value to win. Thus, guess that the highest-valued agent always wins.

- (ii) *Calculate* the interim payment of an agent in the auction in terms of the strategy function.

E.g., we calculate below the payment of agent 1 in the first-price auction when her bid is $s_1(v_1)$ in expectation when agent 2’s value v_2 is drawn from the uniform distribution.

$$p_1^{\text{FP}}(v_1) = \mathbf{E}[p_1^{\text{FP}}(v_1, v_2) \mid 1 \text{ wins}] \mathbf{Pr}[1 \text{ wins}] \\ + \mathbf{E}[p_1^{\text{FP}}(v_1, v_2) \mid 1 \text{ loses}] \mathbf{Pr}[1 \text{ loses}].$$

Calculate each of these components for the first-price auction where agent 1 follows strategy $s_1(v_1)$:

$$\mathbf{E}[p_1^{\text{FP}}(v_1, v_2) \mid 1 \text{ wins}] = s_1(v_1).$$

This by the definition of the first-price auction: if you win you pay your bid.

$$\mathbf{Pr}[1 \text{ wins}] = \mathbf{Pr}[v_2 < v_1] = v_1.$$

The first equality follows from the guess that the highest-valued agent wins. The second equality is because v_2 is uniform on $[0, 1]$.

$$\mathbf{E}[p_1^{\text{FP}}(v_1) \mid 1 \text{ loses}] = 0.$$

This is because a loser pays nothing in the first-price auction. This means that we do not need to calculate $\mathbf{Pr}[1 \text{ loses}]$. Plug these into the equation above to obtain:

$$p_1^{\text{FP}}(v_1) = s_1(v_1) \cdot v_1.$$

- (iii) *Calculate* the interim payment of an agent in a strategically-simple auction with the same equilibrium outcome.

E.g., recall that it is a dominant strategy equilibrium (a special case of Bayes-Nash equilibrium) in the second-price auction for each agent to bid her value. I.e., $b_1 = v_1$ and $b_2 = v_2$. Thus, in the second-price auction the agent with the highest value to wins. We calculate below the payment of agent 1 in the second-price auction when her value is v_1 in expectation when agent 2's value v_2 is drawn from the uniform distribution.

$$p_1^{\text{SP}}(v_1) = \mathbf{E}[p_1^{\text{SP}}(v_1, v_2) \mid 1 \text{ wins}] \mathbf{Pr}[1 \text{ wins}] \\ + \mathbf{E}[p_1^{\text{SP}}(v_1, v_2) \mid 1 \text{ loses}] \mathbf{Pr}[1 \text{ loses}].$$

Calculate each of these components for the second-price auction:

$$\mathbf{E}[p_1^{\text{SP}}(v_1, v_2) \mid 1 \text{ wins}] = \mathbf{E}[v_2 \mid v_2 < v_1] \\ = v_1/2.$$

The first equality follows by the definition of the second-price auction and its dominant strategy equilibrium (i.e., $b_2 = v_2$). The second equality follows because in expectation a uniform random variable evenly divides the interval it is over, and once we condition on $v_2 < v_1$, v_2 is $U[0, v_1]$.

$$\mathbf{Pr}[1 \text{ wins}] = \mathbf{Pr}[v_2 < v_1] = v_1.$$

The first equality follows from the definition of the second-price auction and its dominant strategy equilibrium. The second equality is because v_2 is uniform on $[0, 1]$.

$$\mathbf{E}[p_1(v_1) \mid 1 \text{ loses}] = 0.$$

This is because a loser pays nothing in the second-price auction. This means that we do not need to calculate $\mathbf{Pr}[1 \text{ loses}]$. Plug these into the equation above to obtain:

$$\mathbf{E}[p_1^{\text{SP}}(v_1)] = v_1^2/2.$$

- (iv) *Solve* for bidding strategies from expected payments.

E.g., the interim payments calculated in the previous steps must be equal, implying:

$$p_1^{\text{FP}}(v_1) = s_1(v_1) \cdot v_1 = v_1^2/2 = p_1^{\text{SP}}(v_1).$$

We can solve for $s_1(v_1)$ and get

$$s_1(v_1) = v_1/2.$$

- (v) Verify initial guess was correct. If the strategy function derived is not onto, verify that actions out of the range of the strategy function are dominated.

E.g., if agents follow symmetric strategies $s_1(z) = s_2(z) = z/2$ then the agent with the highest value wins. With this strategy function, bids are in $[0, 1/2]$ and any bid above $s_1(1) = 1/2$ is dominated by bidding $s_1(1)$. All such bids win with certainty, but of these the bid $s_1(1) = 1/2$ gives the lowest payment.

In the above first-price auction example it should be readily apparent that we did slightly more work than we had to. In this case it would have been enough to note that in both the first- and second-price auction a loser pays nothing. We could therefore simply equate the expected payments conditioned on winning:

$$\mathbf{E}[p_1(v_1) \mid 1 \text{ wins}] = \underbrace{v_1/2}_{\text{second-price}} = \underbrace{s_1(v_i)}_{\text{first-price}}.$$

We can also work through the above framework for the *all-pay* auction where the agents submit bids, the highest bid wins, but all agents pay their bid. The all-pay auction is also revenue equivalent to the second-price auction. However, now we compare the total expected payment (regardless of winning) to conclude:

$$\mathbf{E}[p_1(v_1)] = \underbrace{v_1^2/2}_{\text{second-price}} = \underbrace{s_1(v_i)}_{\text{all-pay}}.$$

I.e., the BNE strategies for the all-pay auction are $s_i(z) = z^2/2$. Remember, of course, that the equilibrium strategies solved for above are for single-item auctions and two agents with values uniform on $[0, 1]$. For different distributions or numbers of agents the equilibrium strategies will generally be different.

We conclude by observing that if we fail to exhibit a Bayes-Nash equilibrium via this approach then our original guess is contracted and there is no equilibrium of the given mechanism that corresponds to the guess. Conversely, if the approach succeeds then the equilibrium found

is the only equilibrium consistent with the guess. As an example, we can conclude the following for first-price auctions.²

Proposition 2.7. *When agents' values are independent and identically distributed from a continuous distribution, the first-price auction has a unique Bayes-Nash equilibrium for which the highest-valued agent always wins.*

2.9 Uniqueness of Equilibria

As equilibrium attempts to make a prediction of what will happen in a game or mechanism, the uniqueness of equilibrium is important. If there are multiple equilibria then the prediction is to a set of outcomes not a single outcome. In terms of mechanism design, some of these outcomes could be good and some could be bad. There are also questions of how the players coordinate on an equilibrium.

As an example, in the second-price auction for two agents with values uniformly distributed on $[0, 1]$ there is the dominant strategy equilibrium where agents truthfully report their values. This outcome is good from the perspective of social surplus in that the item is awarded to the highest-valued agent. There are, however, other Bayes-Nash equilibria. For instance, it is also a BNE for agent 1 (Alice) to bid one and agent 2 (Bob) to bid zero (regardless of their values). Alice is happy to win and pay zero (Bob's bid); Bob with any value $v_2 \leq 1$ is at least as happy to lose and pay zero versus winning and paying one (Alice's bid). Via examples like this the social surplus of the worst BNE in the second-price auction can be arbitrarily worse than the social surplus of the best BNE (Exercise 2.8). This latter equilibrium is not dominant strategy as if Bob were to bid his value (a dominant strategy), then Alice would no longer prefer to bid one. Because of this non-robustness of non-DSE in games that possess DSE, we can assume that agents follow DSE if there exists one.

In contrast, the first-price auction for independent and identical prior distributions does not suffer from multiplicity of Bayes-Nash equilibria. Specifically, the method described in the previous section for solving for the symmetric equilibrium in symmetric auction-like games gives the

² In the next section we will strengthen Proposition 2.7 and show that for the first-price auction (with independent, identical, and continuous distributions) there are no equilibria where the highest-valued agent does not win. Thus, the equilibrium solved for is the unique Bayes-Nash equilibrium.

unique BNE. We describe this result as two parts. First, we exclude the possibility of multiple symmetric equilibria. Second, we exclude the existence of asymmetric equilibria.

Lemma 2.8. *For agents with values drawn independently and identically from a continuous distribution, the first-price auction admits exactly one symmetric Bayes-Nash equilibrium.*

Proof. Consider a symmetric strategy profile $\mathbf{s} = (s, \dots, s)$. First, the common strategy $s(\cdot)$ must be non-decreasing (otherwise BNE is contradicted by Theorem 2.2).

Second, if the strategy is non-strictly increasing then there is a point mass some bid b in the bid distribution. Symmetry with respect to this strategy implies that all agents will make a bid equal to this point mass with some measurable (i.e., strictly positive) probability. All but one of these bidders must lose (perhaps via random tie-breaking). Winning, however, must be strictly preferred to losing for some of the values in the interval (as an agent with value v is only indifferent to winning or losing when $v = b$). Such a losing agent has a deviation of bidding $b + \epsilon$, and for ϵ approaching zero this deviation is strictly better than bidding b . This is a contradiction to the existence of such a non-strictly increasing equilibrium.

Finally, for a strictly increasing strategy s the highest-valued agent must always win; therefore, Proposition 2.7 implies that there is only one such equilibrium. \square

We now make much the same argument as we did in solving for equilibrium (Section 2.8) to exclude the possibility of asymmetric equilibria in the first-price auction. The main idea in this argument is that there are two formulas for the interim utility of an agent in the first-price auction in terms of the allocation rule $x(\cdot)$. The first formula is from the payment identity of Theorem 2.2, the second formula is from the definition of the first-price auction (i.e., in terms of the agent's strategy). They are,

$$u(v) = \int_0^v x(z) dz, \text{ and} \tag{2.3}$$

$$u(v) = (v - s(v)) \cdot x(v). \tag{2.4}$$

The uniqueness of the symmetric Bayes-Nash equilibrium in the first-price auction follows from the following lemma.

Lemma 2.9. *For $n = 2$ agents with values drawn independently and*

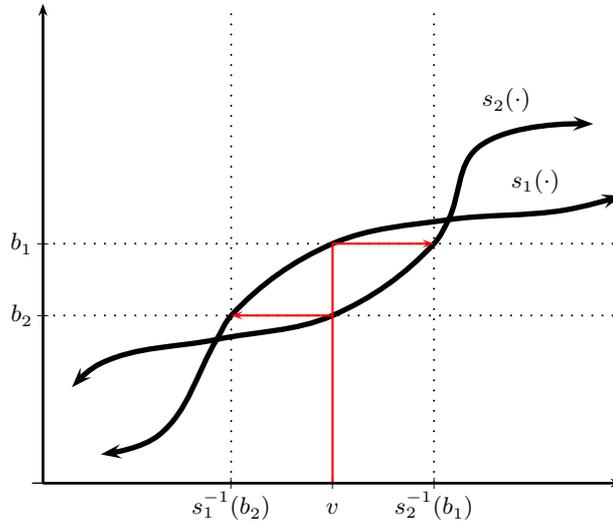


Figure 2.3. Graphical depiction of the first claim in the proof of Lemma 2.9 with $b_i = s_i(v)$. Clearly, $s_2^{-1}(b_1) > s_1^{-1}(b_2)$. Strict monotonicity of the distribution function $F(\cdot)$ then implies that $F(s_2^{-1}(b_1)) > F(s_1^{-1}(b_2))$.

identically from a continuous distribution F , the first-price auction with an unknown random reserve from known distribution G admits no asymmetric Bayes-Nash equilibrium.

Theorem 2.10. For $n \geq 2$ agents with values drawn independently and identically from a continuous distribution F , the first-price auction there is a unique Bayes-Nash equilibrium that is symmetric.

Proof. By Lemma 2.8 there is exactly one symmetric Bayes-Nash equilibrium of an n -agent first-price auction. If there is an asymmetric equilibrium there must be two agents whose strategies are distinct. We can view the n -agent first-price auction in BNE, from the perspective of this pair of agents, as a two-agent first-price auction with a random reserve drawn from the distribution of BNE bids of the other $n - 2$ agents. Lemma 2.9 then contradicts the distinctness of these two strategies. \square

Proof of Lemma 2.9. We will prove this lemma for the special case of strictly-increasing and continuous strategies (for the general argument, see Exercise 2.12). Agent 1 is Alice and agent 2 is Bob.

If the BNE utilities of the agents are the same at all values, i.e., $u_1(v) = u_2(v)$ for all v in the distribution's range, then the payment identity of Theorem 2.2 implies that the strategies are the same at all values. For a contradiction then, fix a strictly-increasing continuous strategy profile $\mathbf{s} = (s_1, s_2)$ for which $u_1(v) > u_2(v)$ at some v . By equation (2.3) there must be a measurable interval of values $I = (a, b)$, i.e., with $\mathbf{Pr}[v \in I] > 0$, containing this value v and for which $x_1(v) \geq x_2(v)$ (assume I is the maximal such interval).

A first claim for strictly-increasing continuous strategies is that $s_1(v) > s_2(v)$ if and only if $x_1(v) > x_2(v)$. See Figure 2.3 for a graphical representation of the following argument. Since the strategies are continuous and strictly increasing, the inverses of the strategies are well defined. Calculate Alice's interim allocation probability x_1 at value v , for Bob's value $v_2 \sim F$ and reserve bid $\hat{b} \sim G$, as:

$$\begin{aligned} x_1(v) &= \mathbf{Pr} \left[s_1(v) > s_2(v_2) \wedge s_1(v) > \hat{b} \right] \\ &= \mathbf{Pr} \left[s_2^{-1}(s_1(v)) > v_2 \wedge s_1(v) > \hat{b} \right] \\ &= F(s_2^{-1}(s_1(v))) \cdot G(s_1(v)). \end{aligned}$$

Likewise, Bob's interim allocation probability is

$$x_2(v) = F(s_1^{-1}(s_2(v))) \cdot G(s_2(v)).$$

For $s_1(v) \geq s_2(v)$ then the last term in the allocation probabilities satisfies $G(s_1(v)) \geq G(s_2(v))$ (as the distribution function $G(\cdot)$ is non-decreasing). Similarly, strict monotonicity of the strategy functions and distribution function imply that for $s_1(v) \geq s_2(v)$ the first term in the allocation probabilities satisfies $F(s_2^{-1}(s_1(v))) \geq F(s_1^{-1}(s_2(v)))$; moreover, either both inequalities are strict or both are equality.

A second claim is that the low-bidding Bob on the interval $I = (a, b)$ obtains (weakly) at most the utility of high-bidding Alice at the endpoint a and (weakly) at least the utility of the high-bidding Alice at the endpoint b . We argue the claim for b , the case of a is similar. The key to this claim is that there are not higher values $v > b$ where $s_2(v) < s_1(b)$. This is either because $s_1(b) = s_2(b)$ (and the strategies are monotonically increasing) or because b is the maximum value in the support of the value distribution F . In the first case, by the above claim $x_1(b) = x_2(b)$ so by (2.4) the agents' utilities are equal. In the second case, Bob with value b could deviate and bid $s_1(b)$ and obtain the same allocation probability as Alice with the same value. By equation (2.4) such a deviation

would give Bob (with value b) the same utility as Alice (with value b). Existence of such a deviation gives a lower bound on Bob's utility.

Finally, we complete the lemma by writing the difference in utilities of each of Alice and Bob with values a and a . By the second claim, above, this difference is (weakly) greater for Bob than Alice (relative to Alice's utility, Bob's utility is no higher at a and no lower at b).

$$u_1(b) - u_1(a) \leq u_2(b) - u_2(a)$$

However, by the first claim and equation (2.3), Alice has a strictly higher allocation rule on I and therefore strictly higher change in utility.

$$\int_a^b x_1(z) dz > \int_a^b x_2(z) dz$$

These observations give a contradiction. \square

2.10 The Revelation Principle

We are interested in designing mechanisms and, while the characterization of Bayes-Nash equilibrium is elegant, solving for equilibrium is still generally quite challenging. The final piece of the puzzle, and the one that has enabled much of modern mechanism design is the *revelation principle*. The revelation principle states, informally, that if we are searching among mechanisms for one with a desirable equilibrium we may restrict our search to single-round, sealed-bid mechanisms in which truthtelling is an equilibrium.

Definition 2.7. A *direct revelation* mechanism is single-round, sealed bid, and has action space equal to the type space, (i.e., an agent can bid any type she might have)

Definition 2.8. A direct revelation mechanism is *Bayesian incentive compatible* (BIC) if truthtelling is a Bayes-Nash equilibrium.

Definition 2.9. A direct revelation mechanism is *dominant strategy incentive compatible* (DSIC) if truthtelling is a dominant strategy equilibrium.

Theorem 2.11. Any mechanism \mathcal{M} with good BNE (resp. DSE) can be converted into a BIC (resp. DSIC) mechanism \mathcal{M}' with the same BNE (resp. DSE) outcome.

Proof. We will prove the BNE variant of the theorem. Let \mathbf{s} , \mathbf{F} , and \mathcal{M} be in BNE. Define single-round, sealed-bid mechanism \mathcal{M}' as follows:

- (i) Accept sealed bids \mathbf{b} .
- (ii) Simulate $\mathbf{s}(\mathbf{b})$ in \mathcal{M} .
- (iii) Output the outcome of the simulation.

We now claim that \mathbf{s} being a BNE of \mathcal{M} implies truthtelling is a BNE of \mathcal{M}' (for distribution \mathbf{F}). Let \mathbf{s}' denote the truthtelling strategy. In \mathcal{M}' , consider agent i and suppose all other agents are truthtelling. This means that the actions of the other players in \mathcal{M} are distributed as $\mathbf{s}_{-i}(\mathbf{s}'_{-i}(\mathbf{v}_{-i})) = \mathbf{s}_{-i}(\mathbf{v}_{-i})$ for $\mathbf{v}_{-i} \sim \mathbf{F}_{-i}|_{v_i}$. Of course, in \mathcal{M} if other players are playing $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ then since \mathbf{s} is a BNE, i 's best response is to play $s_i(v_i)$ as well. Agent i can play this action in the simulation of \mathcal{M} is by playing the truthtelling strategy $s'_i(v_i) = v_i$ in \mathcal{M}' . \square

Notice that we already, in Chapter 1, saw the revelation principle in action. The second-price auction is the revelation principle applied to the ascending-price auction.

Because of the revelation principle, for many of the mechanism design problems we consider, we will look first for Bayesian or dominant-strategy incentive compatible mechanisms. The revelation principle guarantees that, in our search for optimal BNE mechanisms, it suffices to search only those that are BIC (and likewise for DSE and DSIC). The following are corollaries of our BNE and DSE characterization theorems.

We defined the allocation and payment rules $\mathbf{x}(\cdot)$ and $\mathbf{p}(\cdot)$ as functions of the valuation profile for an implicit game G and strategy profile \mathbf{s} . When the strategy profile is truthtelling, the allocation and payment rules are identical the original mappings of the game from actions to allocations and prices, denoted $\mathbf{x}^G(\cdot)$ and $\mathbf{p}^G(\cdot)$. Additionally, let $x_i^G(v_i) = \mathbf{E}[x_i^G(\mathbf{v}) | v_i]$ and $p_i^G(v_i) = \mathbf{E}[p_i^G(\mathbf{v}) | v_i]$ for $\mathbf{v} \sim \mathbf{F}$. Furthermore, the truthtelling strategy profile in a direct-revelation game is onto.

Corollary 2.12. *A direct mechanism \mathcal{M} is BIC for distribution \mathbf{F} if and only if for all i ,*

- (i) (monotonicity) $x_i^{\mathcal{M}}(v_i)$ is monotone non-decreasing, and
- (ii) (payment identity) $p_i^{\mathcal{M}}(v_i) = v_i x_i^{\mathcal{M}}(v_i) - \int_0^{v_i} x_i^{\mathcal{M}}(z) dz + p_i^{\mathcal{M}}(0)$.

Corollary 2.13. *A direct mechanism \mathcal{M} is DSIC if and only if for all i and \mathbf{v} ,*

- (i) (*monotonicity*) $x_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i})$ is monotone non-decreasing in v_i , and
(ii) (*payment identity*) $p_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i}) = v_i x_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i^{\mathcal{M}}(z, \mathbf{v}_{-i}) dz + p_i^{\mathcal{M}}(0, \mathbf{v}_{-i})$.

Corollary 2.14. *A direct, deterministic mechanism \mathcal{M} is DSIC if and only if for all i and \mathbf{v} ,*

- (i) (*step-function*) $x_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i})$ steps from 0 to 1 at some $\hat{v}_i(\mathbf{v}_{-i})$, and
(ii) (*critical value*) $p_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i}) = \begin{cases} \hat{v}_i(\mathbf{v}_{-i}) & \text{if } x_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i}) = 1 \\ 0 & \text{otherwise} \end{cases} + p_i^{\mathcal{M}}(0, \mathbf{v}_{-i})$.

When we construct mechanisms we will use the “if” directions of these theorems. When discussing incentive compatible mechanisms we will assume that agents follow their equilibrium strategies and, therefore, each agent’s bid is equal to her valuation.

Between DSIC and BIC clearly DSIC is a stronger incentive constraint and we should prefer it over BIC if possible. Importantly, DSIC requires fewer assumptions on the agents. For a DSIC mechanisms, each agent must only know her own value; while for a BIC mechanism, each agent must also know the distribution over other agent values. Unfortunately, there will be some environments where we derive BIC mechanisms where no analogous DSIC mechanism is known.

The revelation principle fails to hold in some environments of interest. We will take special care to point these out. Two such environments, for instance, are where agents only learn their values over time, or where the designer does not know the prior distribution (and hence cannot simulate the agent strategies).

Exercises

- 2.1 Find a symmetric mixed strategy equilibrium in the chicken game described in Section 2.1. I.e., find a probability ρ such that if James Dean stays with probability ρ and swerves with probability $1 - \rho$ then Buzz is happy to do the same.
- 2.2 Give a characterization of Bayes-Nash equilibrium for discrete single-dimensional type spaces for agents with linear utility. Assume that $\mathcal{T} = \{v^0, \dots, v^N\}$ with the probability that an agent’s value is $v \in \mathcal{T}$ given by probability mass function $f(v)$. Assume $v^0 = 0$. You will not get a payment identity; instead characterize for any BNE allocation rule, the maximum payments.

- (a) Give a characterization for the special case where the values are uniform, i.e., $v^j = j$ for all j .
- (b) Give a characterization for the special case where the probabilities are uniform, i.e., $f(v^j) = 1/N$ for all j .
- (c) Give a characterization for the general case.

(Hint: You should end up with a very similar characterization to that for continuous type spaces.)

- 2.3 In Section 2.3 we characterized outcomes and payments for BNE in single-dimensional games. This characterization explains what happens when agents behave strategically.

Suppose instead of strategic interaction, we care about fairness. Consider a valuation profile, $\mathbf{v} = (v_1, \dots, v_n)$, an allocation vector, $\mathbf{x} = (x_1, \dots, x_n)$, and payments, $\mathbf{p} = (p_1, \dots, p_n)$. Here x_i is the probability that i is served and p_i is the expected payment of i regardless of whether i is served or not.

Allocation \mathbf{x} and payments \mathbf{p} are *envy-free* for valuation profile \mathbf{v} if no agent wants to unilaterally swap allocation and payment with another agent. I.e., for all i and j ,

$$v_i x_i - p_i \geq v_i x_j - p_j.$$

Characterize envy-free allocations and payments (and prove your characterization correct). Unlike the BNE characterization, your characterization of payments will not be unique. Instead, characterize the minimum payments that are envy-free. Draw a diagram illustrating your payment characterization. (Hint: You should end up with a very similar characterization to that of BNE.)

- 2.4 AdWords is a Google Inc. product in which the company sells the placement of advertisements along side the search results on its search results page. Consider the following *position auction* environment which provides a simplified model of AdWords. There are m advertisement slots that appear along side search results and n advertisers. Advertiser i has value v_i for a click. Slot j has *click-through rate* w_j , meaning, if an advertiser is assigned slot j the advertiser will receive a click with probability w_j . Each advertiser can be assigned at most one slot and each slot can be assigned at most one advertiser. If a slot is left empty, all subsequent slots must be left empty, i.e., slots cannot be skipped. Assume that the slots are ordered from highest click-through rate to lowest, i.e., $w_j \geq w_{j+1}$ for all j .

- (a) Find the envy-free (see Exercise 2.3) outcome and payments with the maximum social surplus. Give a description and formula for the envy-free outcome and payments for each advertiser. (Feel free to specify your payment formula with a comprehensive picture.)
- (b) In the real AdWords problem, advertisers only pay if they receive a click, whereas the payments calculated, i.e., \mathbf{p} , are in expected over all outcomes, click or no click. If we are going to charge advertisers only if they are clicked on, give a formula for calculating these payments \mathbf{p}' from \mathbf{p} .
- (c) The real AdWords problem is solved by auction. Design an auction that maximizes the social surplus in dominant strategy equilibrium. Give a formula for the payment rule of your auction (again, a comprehensive picture is fine). Compare your DSE payment rule to the envy-free payment rule. Draw some informal conclusions.
- 2.5 Consider the first-price auction for selling a single item to two agents whose values are independent but not identical. In each of the settings below prove or disprove the claim that there is a Bayes-Nash equilibrium wherein the item is always allocated to the agent with the highest value.
- (a) Agent 1 has value $U[0, 1]$ and agent 2 has value $U[0, 1/2]$.
- (b) Agent 1 has value $U[0, 1]$ and agent 2 has value $U[1/2, 1]$.
- 2.6 Consider a first-price position auction (see Exercise 2.4) with n agents, n positions, and decreasing weights $\mathbf{w} = (w_1, w_2)$. In such an auction agents are assigned to positions in the order of their bids. The agent assigned to position i is served with probability w_i and she pays her bid when served. Use revenue equivalence to solve for symmetric Bayes-Nash equilibrium strategies \mathbf{s} when the values of the agents are drawn independent and identically from $U[0, 1]$.
- 2.7 Consider a first-price position auction (see Exercise 2.4) with n agents, n positions, and position weights \mathbf{w} defined by $w_i = 1 - i^{-1/n-1}$ for $i \in \{1, \dots, n\}$. In such an auction agents are assigned to positions in the order of their bids. The agent assigned to position i is served with probability w_i and she pays her bid when served. Use revenue equivalence to prove that the symmetric strategy profile $\mathbf{s} = (s, \dots, s)$ with $s(v) = v/2$ is a Bayes-Nash equilibrium when

the values of the agents are drawn independent and identically from $U[0, 1]$.

- 2.8 Prove that in a two-agent second-price auction for a single-item, that the best Bayes-Nash equilibrium can have a social surplus (i.e., the expected value of the winner) that is arbitrarily larger than the worst Bayes-Nash equilibrium. (Hint: Show that for any fixed β that there is a value distribution F and two BNE where the social surplus in one BNE is strictly larger than a β fraction of the social surplus of the other BNE.)
- 2.9 Show that with independent, identical, and continuously distributed values, the two-agent all-pay auction (where agents bid, the highest-bidder wins, and all agents pay their bids) admits exactly one strictly continuous Bayes-Nash equilibrium.
- 2.10 Show that with independent, identical, and continuously distributed values, the two-agent first-price position auction (cf. Exercise 2.4; where agents bid, the highest bidder is served with given probability w_1 , the second-highest bidder is served with given probability $w_2 \leq w_1$, and all agents pay their bids when they are served) admits exactly one strictly continuous Bayes-Nash equilibrium.
- 2.11 Consider the following auction with first-price payment semantics. Agents bid, any agent whose bid is (weakly) higher than all other bids wins, all winners are charged their bids. Notice that in the case of a tie in the highest bid, all of the tied agents win. Prove that there are multiple Bayes-Nash equilibria when agents have values that are independently, identically, and continuously distributed.
- 2.12 Prove Lemma 2.9: For two agents with values drawn independently and identically from a continuous distribution F with support $[0, 1]$, the first-price auction with an unknown random reserve from known distribution G admits no asymmetric Bayes-Nash equilibrium. I.e., remove the assumption of strictly-increasing and continuous strategies from the proof given in the text.

Chapter Notes

The formulation of Bayesian games is due to Harsanyi (1967). The characterization of Bayes-Nash equilibrium, revenue equivalence, and the revelation principle come from Myerson (1981). Parts of the BNE characterization proof presented here come from Archer and Tardos (2001). Amann and Leininger (1996), Bajari (2001), Maskin and Riley (2003),

and Lebrun (2006) studied the uniqueness of equilibrium in the first-price and all-pay auctions. The revenue-equivalence-based uniqueness proof presented here is from Chawla and Hartline (2013).

The position auction was formulated by Edelman et al. (2007) and Varian (2007); see Jansen and Mullen (2008) for the history of auctions for advertisements on search engines. Envy freedom has been considered in algorithmic (e.g., Guruswami et al., 2005) and economic (e.g., Jackson and Kremer, 2007) contexts. Hartline and Yan (2011) characterized envy-free outcomes for single-dimensional agents.

3

Optimal Mechanisms

In this chapter we discuss the objectives of social surplus and profit. As we will see, the economics of designing mechanisms to maximize social surplus is relatively simple. The optimal mechanism is a simple generalization of the second-price auction that we have already discussed. Furthermore, it is dominant strategy incentive compatible and prior-free, i.e., it is not dependent on distributional assumptions. Social surplus maximization is unique among economic objectives in this regard.

The objective of profit maximization, on the other hand, adds significant new challenge: for profit there is no single optimal mechanism. For any mechanism, there is a distribution over agent preferences and another mechanism where this new mechanism has strictly larger profit than the first one.

This non-existence of an absolutely optimal mechanism requires a relaxation of what we consider a good mechanism. To address this challenge, this chapter follows the traditional economics approach of Bayesian optimization. We will assume that the distribution of the agents' preferences is common knowledge, even to the mechanism designer. This designer should then search for the mechanism that maximizes her expected profit when preferences are indeed drawn from the distribution.

As an example, consider two agents with values drawn independently and identically from $U[0, 1]$. The second-price auction obtains revenue equal to the expected second-highest value, $\mathbf{E}[v_{(2)}] = 1/3$. A natural question is whether more revenue can be had. As a first step, it is similarly easy to calculate that the second-price auction with reserve $1/2$ obtains an expected revenue of $5/12$ (which is higher than $1/3$).¹ Above,

¹ There are three cases: (i) $1/2 > v_{(1)} > v_{(2)}$, (ii) $v_{(1)} > 1/2 > v_{(2)}$, and (iii), $v_{(1)} > v_{(2)} > 1/2$. Case (i) happens with probability $1/4$ and has no revenue; case

perhaps surprisingly, a seller makes more money by sometimes not selling the item even when there is a buyer willing to pay. In this chapter we show that the second-price auction with reserve $1/2$ is indeed optimal for this two agent example and furthermore we give a concise characterization of the revenue-optimal auction for any single-dimensional agent environment.

3.1 Single-dimensional Environments

In our previous discussion of Bayes-Nash equilibrium we focused on the agents' incentives. Single-dimensional linear agents each have a single private value for receiving some abstract service and linear utility, i.e., the agent's utility is her value for the service less her payment (Definition 2.6). Recall that the outcome of a single-dimensional game is an allocation $\mathbf{x} = (x_1, \dots, x_n)$, where x_i is an indicator for whether agent i is served, and payments $\mathbf{p} = (p_1, \dots, p_n)$, where p_i is the payment made by agent i . Here we formalize the designer's constraints and objectives.

Definition 3.1. A *general cost environment* is one where the designer must pay a service cost $c(\mathbf{x})$ for the allocation \mathbf{x} produced. A *general feasibility environment* is one where there is a feasibility constraint over the set of agents that can be simultaneously served. A *downward-closed feasibility constraint* is one where subsets of feasible sets are feasible.

Of course, downward-closed environments are a special case of general feasibility environments which are a special case of general cost environments. We can express general feasibility environments as general costs environments where $c(\cdot) \in \{0, \infty\}$. We can similarly express downward-closed feasibility environments as the further restriction where $\mathbf{x}^\dagger \leq \mathbf{x}$ (i.e., for all i , $x_i^\dagger \leq x_i$) and $c(\mathbf{x}) = 0$ and implies that $c(\mathbf{x}^\dagger) = 0$. We will be aiming for general mechanism design results and the most general results will be the ones that hold in the most general environments. We will pay special attention to restrictions on the environment that enable illuminating observations about optimal mechanisms.

(ii) happens with probability $1/2$ and has revenue $1/2$; and case (iii) happens with probability $1/4$ and has expected revenue $\mathbf{E}[v_{(2)} \mid \text{case (iii) occurs}] = 2/3$. The calculation of the expected revenue in case (iii) follows from the conditional values being $U[1/2, 1]$ and the fact that, in expectation, uniform random variables evenly divide the interval they are over. The total expected revenue can then be calculated as $5/12$.

The two most fundamental designer objectives are social surplus, a.k.a., social welfare,² and profit.

Definition 3.2. The *social surplus* of an allocation is the cumulative value of the agents served less the service cost:

$$\text{Surplus}(\mathbf{v}, \mathbf{x}) = \sum_i v_i \cdot x_i - c(\mathbf{x}).$$

The *profit* of allocation and payments is the cumulative payment of the agents less the service cost:

$$\text{Profit}(\mathbf{p}, \mathbf{x}) = \sum_i p_i - c(\mathbf{x}).$$

Implicit in the definition of social surplus is the fact that the payments from the agents are transferred to the service provider and therefore do not affect the objective.³

The single-item and routing environments that were discussed in Chapter 1 are special cases of downward-closed environments. Single-item environments have

$$c(\mathbf{x}) = \begin{cases} 0 & \text{if } \sum_i x_i \leq 1, \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

In routing environments, recall, each agent has a message to send between a source and destination in the network.

$$c(\mathbf{x}) = \begin{cases} 0 & \text{if messages with } x_i = 1 \text{ can be simultaneously routed, and} \\ \infty & \text{otherwise.} \end{cases}$$

We have yet to see any examples of general cost environments. One natural one is that of a *multicast auction*. The story for this problem comes from live video streaming. Suppose we wish to stream live video to viewers (agents) in a computer network. Because of the high-bandwidth nature of video streaming the content provider must lease the network links. Each link has a publicly known cost. To serve a set of agents, the designer must pay the cost of network links that connect each agent, located at different nodes in the network, to the “root”, i.e., the origin of the multicast. The nature of multicast is that the messages need only

² A mechanism that optimizes social surplus is said to be *economically efficient*; though, we will not use this terminology because of possible confusion with *computational efficiency*. A mechanism is computationally efficient if it computes its outcome quickly (see ??).

³ An alternative notion would be to consider only the total value derived by the agents, i.e., the surplus less the total payments. This *residual surplus* was discussed in detail in Chapter 1; mechanisms for optimizing residual surplus are the subject of Exercise 3.1.

be transmitted once on each edge to reach the agents. Therefore, the total cost to serve these agents is the minimum cost of the *multicast tree* that connects them.⁴

3.2 Social Surplus

We now derive the optimal mechanism for social surplus. To do this we walk through a standard approach in mechanism design. We completely relax the Bayes-Nash equilibrium incentive constraints and ask and solve the remaining non-game-theoretic optimization question. We then verify that this solution does not violate the incentive constraints. We conclude that the resulting mechanism is optimal.

The non-game-theoretic optimization problem of maximizing surplus for input $\mathbf{v} = (v_1, \dots, v_n)$ is that of finding \mathbf{x} to maximize $\text{Surplus}(\mathbf{v}, \mathbf{x}) = \sum_i v_i x_i - c(\mathbf{x})$. Let OPT be an optimal algorithm for solving this problem. We will care about both the allocation that OPT selects, i.e., $\text{argmax}_{\mathbf{x}} \text{Surplus}(\mathbf{v}, \mathbf{x})$ and its surplus $\max_{\mathbf{x}} \text{Surplus}(\mathbf{v}, \mathbf{x})$. Where it is unambiguous we will use notation $\text{OPT}(\mathbf{v})$ to denote either of these quantities. Notice that the formulation of OPT has no mention of Bayes-Nash equilibrium incentive constraints.

We know from our characterization that the allocation rule of any BNE is monotone, and that any monotone allocation rule can be implemented in BNE with the appropriate payment rule. Thus, relative to the non-game-theoretic optimization, the mechanism design problem of finding a BIC mechanism to maximize surplus has an added monotonicity constraint. As it turns out, even though we did not impose a monotonicity constraint on OPT, it is satisfied anyway.

Lemma 3.1. *For each agent i and all values of other agents \mathbf{v}_{-i} , the allocation rule of OPT for agent i is a step function.*

Proof. Consider any agent i . There are two situations of interest. Either i is served by $\text{OPT}(\mathbf{v})$ or i is not served by $\text{OPT}(\mathbf{v})$. We write out the surplus of OPT in both of these cases. Below, notation (z, \mathbf{v}_{-i}) denotes the vector \mathbf{v} with the i th coordinate replaced with z .

⁴ In combinatorial optimization this problem is known as the *weighted Steiner tree* problem. It is a computationally challenging variant of the *minimum spanning tree* problem.

Case 1 ($i \in \text{OPT}$):

$$\begin{aligned}\text{OPT}(\mathbf{v}) &= \max_{\mathbf{x}} \text{Surplus}(\mathbf{v}, \mathbf{x}) \\ &= v_i + \max_{\mathbf{x}_{-i}} \text{Surplus}((0, \mathbf{v}_{-i}), (1, \mathbf{x}_{-i})).\end{aligned}$$

Define $\text{OPT}_{-i}(\infty, \mathbf{v}_{-i})$, the optimal surplus from agents other than i assuming that i is served, as the second term on the right hand side. Thus,

$$\text{OPT}(\mathbf{v}) = v_i + \text{OPT}_{-i}(\infty, \mathbf{v}_{-i}).$$

Notice that $\text{OPT}_{-i}(\infty, \mathbf{v}_{-i})$ is not a function of v_i .

Case 2 ($i \notin \text{OPT}$):

$$\begin{aligned}\text{OPT}(\mathbf{v}) &= \max_{\mathbf{x}} \text{Surplus}(\mathbf{v}, \mathbf{x}) \\ &= \max_{\mathbf{x}_{-i}} \text{Surplus}((0, \mathbf{v}_{-i}), (0, \mathbf{x}_{-i})).\end{aligned}$$

Define $\text{OPT}(0, \mathbf{v}_{-i})$, the optimal surplus from agents other than i assuming that i is not served, as the term on the right hand side. Thus,

$$\text{OPT}(\mathbf{v}) = \text{OPT}(0, \mathbf{v}_{-i}).$$

Notice that $\text{OPT}(0, \mathbf{v}_{-i})$ is not a function of v_i .

OPT chooses whether or not to allocate to agent i , and thus which of these cases we are in, so as to optimize the surplus. Therefore, OPT allocates to i whenever the surplus from Case 1 is greater than the surplus from Case 2. I.e., when

$$v_i + \text{OPT}_{-i}(\infty, \mathbf{v}_{-i}) \geq \text{OPT}(0, \mathbf{v}_{-i}).$$

Solving for v_i we conclude that OPT allocates to i whenever

$$v_i \geq \text{OPT}(0, \mathbf{v}_{-i}) - \text{OPT}_{-i}(\infty, \mathbf{v}_{-i}).$$

Notice that neither of the terms on the right hand side contain v_i . Therefore, the allocation rule for i is a step function with critical value $\hat{v}_i = \text{OPT}(0, \mathbf{v}_{-i}) - \text{OPT}_{-i}(\infty, \mathbf{v}_{-i})$. \square

Since the allocation rule induced by OPT is a step function, it satisfies our strongest incentive constraint: with the appropriate payments (i.e., the “critical values”) truthtelling is a dominant strategy equilibrium (Corollary 2.14). The resulting surplus maximization mechanism is often referred to as the *Vickrey-Clarke-Groves* (VCG) mechanism, named after William Vickrey, Edward Clarke, and Theodore Groves.

Definition 3.3. The *surplus maximization* (SM) mechanism is:

- (i) Solicit and accept sealed bids \mathbf{b} .
- (ii) find the optimal outcome $\mathbf{x} \leftarrow \text{OPT}(\mathbf{b})$, and
- (iii) set prices \mathbf{p} as

$$p_i \leftarrow \begin{cases} \text{OPT}(0, \mathbf{b}_{-i}) - \text{OPT}_{-i}(\infty, \mathbf{b}) & \text{if } i \text{ is served} \\ 0 & \text{otherwise.} \end{cases}$$

An intuitive description of the critical value $\hat{v}_i = \text{OPT}(0, \mathbf{v}_{-i}) - \text{OPT}_{-i}(\infty, \mathbf{v}_{-i})$ is the *externality* that agent i imposes on the other agents by being served. In other words, because i is served the other agents obtain total surplus $\text{OPT}_{-i}(\infty, \mathbf{v}_{-i})$ instead of the surplus $\text{OPT}(0, \mathbf{v}_{-i})$ that they would have received if i was not served. We can similarly write $p_i = \text{OPT}(0, \mathbf{v}_{-i}) - \text{OPT}_{-i}(\mathbf{v})$ as the externality agent i imposes by being present in the mechanism (regardless of whether she is served or not). Note that if she is not served then the second term is equal to the first and the externality she imposes is zero. Hence, we can interpret the surplus maximization mechanism as serving agents to maximize the social surplus and charging each agent the externality she imposes on the others.

By Corollary 2.14 and Lemma 3.1 we have the following theorem, and by the optimality of OPT and the assumption that agents follow the dominant truth-telling strategy we have the following corollary.

Theorem 3.2. *The surplus maximization mechanism is dominant strategy incentive compatible.*

Corollary 3.3. *The surplus maximization mechanism optimizes social surplus in dominant strategy equilibrium.*

Example 3.1. The second-price routing auction from Chapter 1 is an instantiation of the surplus maximization mechanism where feasible outcomes are subsets of agents whose messages can be simultaneously routed.

It is useful to view the surplus maximization mechanism as a reduction from the mechanism design problem to the non-game-theoretic optimization problem. Given an algorithm that solves the non-game-theoretic optimization problem, i.e., OPT, we can construct the surplus maximization mechanism from it.

Surplus maximization is singular among objectives in that there is a

single mechanism that is optimal regardless of distributional assumptions. Essentially: the agents' incentives are already aligned with the designer's objective and one only needs to derive the appropriate payments, i.e., the critical values. For general objectives, e.g., in the next section we will discuss profit maximization, the optimal mechanism is distribution dependent.

There are other ways to implement surplus maximization besides that of Definition 3.3. By revenue equivalence, the payment rule of the surplus maximization mechanism is unique up to the payments each agent would make if her value was zero, i.e., $p_i(0, \mathbf{v}_{-i})$ for agent i . For instance $p_i = \text{OPT}_{-i}(\mathbf{v})$ is an DSIC payment rule as well with $p_i(0, \mathbf{v}_{-i}) = \text{OPT}(0, \mathbf{v}_{-i})$. This payment rule does not satisfy the natural *no-positive-transfers* condition which requires that agents not be paid to participate. It is also possible to design BNE mechanisms, e.g., with first-price semantics, that implement the same outcome in equilibrium as the surplus maximization mechanism (see Exercise 3.2), though unlike the surplus maximization mechanism given above, design of such a BNE mechanism requires distributional knowledge.

3.3 Profit

A non-game-theoretic optimization problem looks to maximize some objective subject to feasibility. Given the input, we can search over feasible outcomes for the one with the highest objective value for this input. The outcome produced on one input need not bear any relation to the outcome produced on an (even slightly) different input. Mechanisms, on the other hand, additionally must address agent incentives which impose constraints over the outcomes that the mechanism produces across all possible misreports of the agents. In other words, the mechanism's outcome on one input is constrained by its outcome on similar inputs. Therefore, a mechanism may need to tradeoff its objective performance across inputs.

When the distribution of agent values is specified, e.g., by a common prior (Definition 2.5) and the designer has knowledge of this prior, such a tradeoff can be optimized. In particular, the prior assigns a probability to each input and the designer can then optimize expected objective value over this probability distribution. The mechanism that results from such an optimization is said to be *Bayesian optimal*. In this section we derive the Bayesian optimal mechanism for the objective of profit. Other

objectives that are linear in social surplus and payments can be similarly considered (e.g., residual surplus, see Exercise 3.1).

We will use agents with values drawn from the following distributions as examples.

Mathematical Note. At various points in the remainder of this chapter it will be convenient to write the expectations of discontinuous distributions via the integral of their density function which is, at their discontinuity, not well defined. We will then reinterpret the expectation via integration by parts. This notational convenience can be made precise via the Dirac delta function which integrates to a step function; however, we will not describe these details formally.

Consider, as an example, the following which is taken from the construction of Proposition 3.9 on page 73. Draw a random variable $\hat{q} \in [0, 1]$ from a distribution G with distribution function $G(q)$. If G is continuous then its density $g(q) = \frac{d}{dq}G(q)$ is well defined and we can write the expectation of some function $P(\cdot)$ of \hat{q} as $\mathbf{E}_{\hat{q} \sim G}[P(\hat{q})] = \int_0^1 P(q)g(q) dq$. If G is discontinuous (i.e., it possesses point masses) the same formula is correct when the density g contains the appropriate Dirac delta function.

A change of variables allows any integral over $[0, 1]$ to be reinterpreted as the expectation of a function of a uniform random variable. From the above example,

$$\mathbf{E}_{\hat{q} \sim G}[P(\hat{q})] = \mathbf{E}_{q \sim U[0,1]}[P(q)g(q)].$$

Finally, integration by parts gives, for example, the following formula for rearranging an integral, with $\frac{d}{dq}P(q)$ denoted by $p(q)$,

$$\int_0^1 P(q)g(q) dq = \left[P(q)G(q) \right]_0^1 - \int_0^1 p(q)G(q) dq.$$

When $P(0) = P(1) = 0$ the first term on the right-hand side is identically zero. If not, we can set $P(0) = P(1) = 0$ which will introduce a discontinuity in to $P(\cdot)$ which we can express in $p(\cdot)$ via the Dirac delta function as described above. Formulaically, this modification allows the first term of the right-hand side to be accounted for by the integral. We can, as above, write these integrals as expectations of functions of a uniform random variable. Integration by parts can be thus expressed for $q \sim U[0, 1]$ as:

$$\mathbf{E}[P(q)g(q)] = \mathbf{E}[-p(q)G(q)].$$

Example 3.2. A *uniform agent* has single-dimensional linear utility with value v drawn uniformly from $[0, 1]$, i.e., $F(z) = z$ and $f(z) = 1$.

Example 3.3. A *bimodal agent* has single-dimensional linear utility with value v drawn uniformly from $[0, 3]$ with probability $3/4$ and uniformly from $(3, 8]$ with probability $1/4$, i.e., the distribution defined by density function $f(v) = 1/4$ for $v \in [0, 3]$ and $f(v) = 1/20$ for $v \in (3, 8]$ (see Figure 3.4, page 71).

3.3.1 High-level Approach: Amortized Analysis

The profit of a mechanism is given by the sum of the agents' payments (minus the cost of serving them) which, via the payment identity of Theorem 2.2, namely

$$p(v) = v \cdot x(v) - \int_0^v x(v^\dagger) dv^\dagger, \quad (3.1)$$

depends on the allocation rule of each agent (in particular, on $x(v^\dagger)$ for $v^\dagger \leq v$ for an agent with value v). In other words, what the mechanism chooses to do when the agent's value is $v^\dagger < v$ affects the revenue the mechanism obtains when her value is v .

This dependence of the payment on the allocation that the agent would receive if she had a lower value implies that there is no pointwise optimal mechanism (as there was for social surplus maximization, cf. Section 3.2). Consider selling an item to a single agent with value v drawn uniformly from $[0, 1]$ (Example 3.2). If her value is 0.2, then it is pointwise optimal to offer her the item at price 0.2. This corresponds to the allocation rule which steps from zero to one at 0.2. Similarly if her value is 0.7, then it is pointwise optimal to offer her the item at price 0.7. Of course, offering a 0.7-valued agent a price of 0.2 or a 0.2-valued agent a price of 0.7 is not optimal. There is no single mechanism that is pointwise optimal on both of these inputs. On the other hand, given a distribution over the agent's value, we can easily optimize for the price with maximum expected revenue: post the price \hat{v} that maximizes $\hat{v} \cdot (1 - F(\hat{v}))$. For the uniform agent where $F(z) = z$, this optimal price is $\hat{v}^* = 1/2$.⁵

The payment identity (3.1) gives a formula for the expected payment

⁵ Set $\frac{d}{d\hat{v}} [\hat{v} \cdot (1 - \hat{v})] = 1 - 2\hat{v} = 0$ and solve for \hat{v} to get the optimal price to post of $\hat{v}^* = 1/2$.

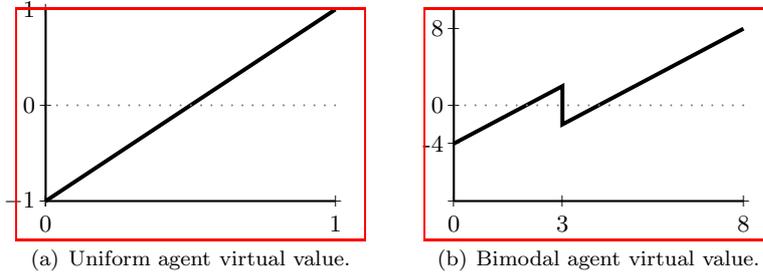


Figure 3.1. Depicted are virtual value functions $\phi(v) = v - \frac{1-F(v)}{f(v)}$ for the uniform and bimodal agent examples (Example 3.2 and Example 3.3). Notice that the virtual value function in the uniform example is monotone non-decreasing in value while in the bimodal example it is not. For reference, the line $v_2 = v_1$ is depicted (grey dotted line).

that a v -valued agent makes in terms of her allocation rule. As is evident from the integral form of the payment identity, an agent's payment at a given value depends on the allocation probability she would have obtained with a lower value. In fact, her payment is highest when the allocation to lower values is the lowest. Our approach to optimizing profit will be via an *amortized analysis* where we charge the loss in revenue from high values due to high allocation probability at low values to the low values themselves. Via such an approach, the amortized benefit from serving an agent with a given value is her value less a deduction that accounts for the lowered the payment for higher values. We will refer to this amortized benefit as *virtual value* and we will show that the problem of optimizing profit in expectation over the distribution of values reduces to the problem of maximizing *virtual surplus* pointwise.

A straightforward approach to such an amortized analysis (given subsequently in Section 3.3.4) will give virtual value function

$$\phi(v) = v - \frac{1 - F(v)}{f(v)}. \quad (3.2)$$

In equation (3.2), v is the revenue from serving the agent with value v (at a price of v) and $\frac{1-F(v)}{f(v)}$ represents the loss of revenue from serving higher values. We will see that such a formulation satisfies

$$\mathbf{E}_{v \sim F}[p(v)] = \mathbf{E}_{v \sim F}[\phi(v) \cdot x(v)] \quad (3.3)$$

for any allocation and payment rules (x, p) that satisfy the Bayes-Nash equilibrium characterization (Theorem 2.2; i.e., monotonicity of x and

the payment identity (3.1)). Equation (3.3) can be derived simply by applying the definition of expectation (as an integral) to the payment identity and simplifying (see Exercise 3.3); we will give a less direct but more economically intuitive construction subsequently in Section 3.3.4.

From equation (3.2) the virtual value function for the uniform agent example is $\phi(v) = 2v - 1$; for the bimodal agent example it is depicted in Figure 3.1. Notice that $\phi(0) < 0$ as there is no value from serving an agent with value zero but serving such an agent lowers the price that she could be charged if her value were higher. Notice that the highest virtual value is always equal to the highest value as there is no amortized deduction necessary to account for lower prices obtained by higher values as no higher values exist, e.g., the uniform agent with values on interval $[0, 1]$ has $\phi(1) = 1$ and the bimodal agent with values on interval $[0, 8]$ has $\phi(8) = 8$.

The importance of equation (3.3) is that it enables the non-pointwise optimization of expected payments to be recast as a pointwise optimization of virtual surplus. The non-game-theoretic optimization problem of maximizing virtual surplus is that of finding \mathbf{x} to maximize $\text{Surplus}(\phi(\mathbf{v}), \mathbf{x}) = \sum_i \phi_i(v_i) \cdot x_i - c(\mathbf{x})$.⁶ Let OPT again be the surplus maximizing algorithm. We will care about both the allocation that $\text{OPT}(\phi(\mathbf{v}))$ selects, i.e., $\text{argmax}_{\mathbf{x}} \text{Surplus}(\phi(\mathbf{v}), \mathbf{x})$ and its virtual surplus $\max_{\mathbf{x}} \text{Surplus}(\phi(\mathbf{v}), \mathbf{x})$. Where it is unambiguous we will use notation $\text{OPT}(\phi(\mathbf{v}))$ to denote either of these quantities. Note that this formulation of OPT has no mention of the incentive constraints.

We now give the first part of the derivation of the optimal mechanism for virtual surplus (and, hence, for profit). To do this we again walk through a standard approach in mechanism design. We completely relax the incentive constraints and solve the remaining non-game-theoretic optimization problem. Since expected profit equals expected virtual surplus, this non-game-theoretic optimization problem is to optimize virtual surplus. We then verify that this solution does not violate the incentive constraints (under some conditions). We conclude that (under the same conditions) the resulting mechanism is optimal.

We know from the BIC characterization (Corollary 2.12) that incentive constraints require that the allocation rule be monotone. Thus, the mechanism design problem of finding a BIC mechanism to maximize virtual surplus has an added monotonicity constraint. Notice that, even though we did not impose a monotonicity constraint on $\text{OPT}(\phi(\cdot))$, if

⁶ Here, $\phi(\mathbf{v})$ denotes the profile of virtual values $(\phi_1(v_1), \dots, \phi_n(v_n))$.

the virtual valuation functions $\phi_i(\cdot)$ for each agent i are monotone then $\text{OPT}(\phi(\cdot))$ is monotone.

Lemma 3.4. *For any profile of virtual value functions ϕ , monotonicity of $\phi_i(\cdot)$ implies the monotonicity of the allocation to agent i of $\text{OPT}(\phi(z, \mathbf{v}_{-i}))$ with respect to z .*

Proof. Let $\mathbf{x}(\cdot)$ be the allocation rules of OPT , i.e., $\mathbf{x}(\mathbf{v}) = \text{argmax}_{\mathbf{x}^\dagger} \text{Surplus}(\mathbf{v}, \mathbf{x}^\dagger)$. Recall from Lemma 3.1 that maximizing surplus is monotone in that $x_i(z, \mathbf{v}_{-i})$ is monotone in z . Therefore $x_i(\phi_i(z), \phi_{-i}(\mathbf{v}_{-i}))$ is monotone in $\phi_i(z)$, i.e., increasing $\phi_i(z)$ does decrease x_i . By assumption $\phi_i(z)$ is monotone in z ; therefore, increasing z cannot decrease $\phi_i(z)$ which cannot decrease $x_i(\phi_i(z), \phi_{-i}(\mathbf{v}_{-i}))$. \square

For many distributions the virtual value function $v - \frac{1-F(v)}{f(v)}$ of equation (3.2) is monotone, e.g., uniform (Example 3.2), normal, and exponential distributions. We refer to these as regular distributions. For regular distributions the approach suggested above is sufficient for describing the optimal mechanism.

Definition 3.4. A distribution F is *regular* if $v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing.

On the other hand, many relevant distributions are irregular, e.g., bimodal (Example 3.3; Figure 3.1(b)). For irregular distributions a more sophisticated amortized analysis is needed to derive the appropriate virtual values. To obtain a mechanism that optimizes non-monotone virtual value functions we cannot initially relax the monotonicity constraint; instead we must optimize virtual surplus subject to monotonicity. In Section 3.3.5 we will describe a generic procedure for *ironing* a non-monotone virtual value function to obtain a monotone (ironed) virtual value function. For ironed virtual values from this procedure, pointwise optimization of the ironed virtual surplus is equivalent to optimization of the original virtual surplus subject to monotonicity. We conclude that, even for irregular distributions, the design of optimal mechanisms in expectation for a known distribution on values is equivalent to the pointwise optimization of a virtual surplus that is given by monotone virtual value functions.

3.3.2 The Virtual Surplus Maximization Mechanism

As revenue-optimal mechanism are virtual surplus maximizers, we now give a generic and formal description of this sort of mechanism. For

monotone virtual value functions, Lemma 3.4 implies that virtual surplus maximization gives a monotone allocation rule for each agent and any fixed values of the other agents; therefore, it satisfies our strongest incentive constraint. With the appropriate payments (i.e., the “critical values”) truth-telling is a dominant strategy equilibrium (recall Corollary 2.14). One way to view the suggested virtual surplus maximization mechanism is as a reduction to surplus maximization, which is solved by the SM mechanism (Definition 3.3; also known as VCG).

Definition 3.5. The *virtual surplus maximization* (VSM) mechanism for single-dimensional linear agents and monotone virtual value functions ϕ is:

- (i) Solicit and accept sealed bids \mathbf{b} ,
- (ii) simulate the surplus maximization mechanism on virtual bids

$$(\mathbf{x}, \mathbf{p}^\dagger) \leftarrow \text{SM}(\phi(\mathbf{b})),$$

- (iii) set prices \mathbf{p} from critical values as

$$p_i \leftarrow \begin{cases} \phi_i^{-1}(p_i^\dagger) & \text{if } i \text{ is served,} \\ 0 & \text{otherwise, and} \end{cases}$$

- (iv) output outcome (\mathbf{x}, \mathbf{p}) .

Notice that the payments \mathbf{p} calculated by VSM can be viewed as follows. SM on virtual values outputs virtual prices \mathbf{p}^\dagger . For winners these correspond to the minimum virtual value that the agent must have to win. The price an agent pays is the minimum value that she must have to win, this can be calculated from these virtual prices via the inverse virtual valuation function. (For virtual value functions $\phi(\cdot)$ that are discontinuous or not strictly increasing this inverse virtual value function is defined as $\phi^{-1}(z) = \inf\{v^\dagger : \phi(v^\dagger) \geq z\}$.)

Theorem 3.5. *For monotone virtual value functions $\phi = (\phi_1, \dots, \phi_n)$, the virtual surplus maximization mechanism VSM is dominant strategy incentive compatible.*

Proof. The theorem follows from Lemma 3.4 applied to each agent, the definition of VSM, and Corollary 2.14. \square

Corollary 3.6. *For monotone virtual value functions ϕ , the virtual surplus maximization mechanism optimizes virtual surplus in dominant strategy equilibrium.*

Notice that the approach above was for optimization of an objective in expectation in Bayes-Nash equilibrium. The mechanism we obtained, in fact, satisfies the stronger dominant strategy incentive compatibility condition. Moreover, even though possibly randomized mechanisms were optimized over, the optimal mechanism is deterministic. When there are ties in virtual surplus, i.e., by multiple distinct outcomes each of which gives the same virtual surplus, these ties can be broken arbitrarily; we may, however, prefer the symmetry of random tie breaking.

To employ Corollary 3.6 for optimizing a given objective, it remains to find a virtual value function for which pointwise optimization of virtual surplus corresponds to optimization of the expected objective value.

Definition 3.6. A *virtual value function* $\phi(\cdot)$ for a given objective is a weakly monotone function that maps a value to a virtual value for which expected optimal virtual surplus is equal to the optimal expected objective value.

3.3.3 Single-item Environments

The above description of the virtual surplus maximization mechanisms does not offer much in the way of intuition. To get a clearer picture, we consider optimal mechanisms the special case of single-item environments, i.e., where the feasible outcomes serve at most one agent. We will consider here four special cases: a single agent, multiple (generally asymmetric) agents, multiple agents with a symmetric strictly-increasing virtual value function, and multiple agents with a symmetric (not strictly) increasing virtual value function.

For a single agent with a monotone virtual value function $\phi(\cdot)$, there is some value $\hat{v}^* = \phi^{-1}(0)$ where the function crosses zero. For example, for the uniform agent this value is $\hat{v}^* = 1/2$, see Figure 3.1(a). Maximizing virtual surplus is simple: if $v \geq \hat{v}^*$ then serve the agent; otherwise, do not serve the agent. In other words, the agent has a critical value of \hat{v}^* and the outcome is identical to that from posting a take-it-or-leave-it price of \hat{v}^* .

Definition 3.7. For an agent with value v drawn from distribution F and virtual value function ϕ , the *monopoly price* $\hat{v}^* = \phi^{-1}(0)$ is the posted price that obtains the highest expected virtual surplus.

Now consider a single-item auction environment and the virtual surplus maximization mechanism for the profile of virtual value functions

ϕ . The mechanism will serve the agent with the highest positive virtual value, or nobody if all virtual values are negative. To see what the critical value of an agent i in this auction is we can write out the condition that must hold for the agent to win. In particular, $\phi_i(v_i) \geq \max(\phi_j(v_j), 0)$ for all $j \neq i$, so i 's critical value is

$$\hat{v}_i = \max(\phi_i^{-1}(\phi_j(v_j)), \phi_i^{-1}(0)) \quad (3.4)$$

for j with the highest virtual value of the other agents. Notice that the auction depends on the precise details of the virtual value functions (see Example 3.4 below). Notice that the second term in the maximization is the monopoly price $\hat{v}_i^* = \phi_i^{-1}(0)$. If the other agents are not competitive, i.e., all agents j have $\phi_j(v_j) < 0$, then the optimization problem reduces to the single-agent case and agent i should see a reserve price of \hat{v}_i^* .

Corollary 3.7. *For single-item environments and monotone virtual value functions, the auction that allocates to the agent with the highest non-negative virtual value maximizes virtual surplus in dominant strategy equilibrium.*

Example 3.4. Consider a two-agent single-item environment with agent 1's (Alice) value from $U[0, 1]$ (as in Example 3.2) and agent 2's (Bob) value from $U[0, 2]$ (with distribution function $F_2(z) = z/2$). The virtual values for revenue from equation (3.2) are $\phi_1(v_1) = 2v_1 - 1$ and $\phi_2(v_2) = 2v_2 - 2$. The virtual surplus maximization mechanism serves Alice whenever $\phi_1(v_1) > \max(\phi_2(v_2), 0)$, i.e., when $v_1 > \max(v_2 - 1/2, 1/2)$. Note that in this revenue-optimal auction Alice may have a lower value than Bob and still win.

Now suppose the virtual value functions are monotone, strictly increasing, identical, and denoted by ϕ . This happens when the agents are independent and identically distributed and, as discussed above, the function $v - \frac{1-F(v)}{f(v)}$ is strictly monotone. In such a scenario, $\phi_i^{-1}(\phi_j(v_j)) = \phi^{-1}(\phi(v_j)) = v_j$, and equation (3.4) for agent i 's critical value simplifies to $\hat{v}_i = \max(v_j, \hat{v}^*)$ where j is the highest valued of the other agents. The virtual surplus maximizing auction thus serves the agent with the highest value that is at least $\hat{v}^* = \phi^{-1}(0)$, a.k.a., the monopoly price. What auction has this equilibrium outcome? The second-price auction with monopoly reserve \hat{v}^* .

Definition 3.8. The *second-price auction with reservation price \hat{v}* , sells the item if any agent bids above \hat{v} . The price the winning agent pays

the maximum of the second highest bid and \hat{v} . The *monopoly-reserve auction* sets $\hat{v} = \hat{v}^*$.

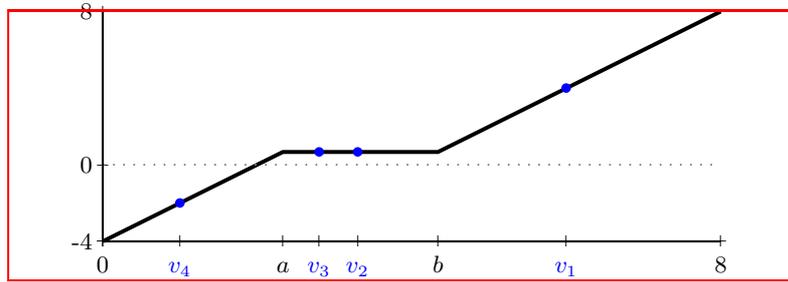
Corollary 3.8. *In single-item environments with identical strictly-increasing virtual value function ϕ , the virtual surplus maximizing mechanism is the second-price auction with monopoly reserve $\hat{v}^* = \phi^{-1}(0)$.*

Example 3.5. Consider a two-agent single-item environment with i.i.d. uniform agents (as in Example 3.2). As we have calculated, $\phi(v) = 2v - 1$ is monotone and strictly increasing, the monopoly price is $\hat{v}^* = \phi^{-1}(0) = 1/2$, and the revenue-optimal auction is the second-price auction with reserve price $1/2$. Our calculation at the introduction of this chapter showed its expected revenue to be $5/12$. Now we see that this revenue is optimal among all mechanisms for this scenario.

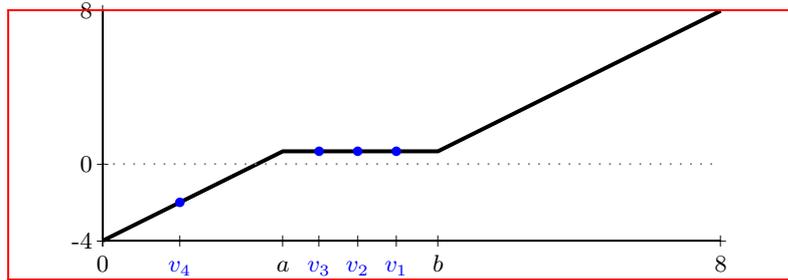
Notice that the optimal reserve price is not a function of the number of agents. For more intuition for why the reserve price is invariant to the number of agents, notice the following. Either the other agents are competitive and the reserve is irrelevant or the other agents are irrelevant and the designer faces the same revenue tradeoffs as in the single-agent example. This single-agent tradeoff is optimized by a reserve equal to the monopoly price. Furthermore, the result can easily be extended to single-item multi-unit auctions where the optimal reserve price is also not a function of the number of units that are for sale (and beyond, see Proposition 4.23 in Chapter 4).

We conclude this section by considering the case of symmetric virtual value functions that are increasing but not strictly so. Notice that, with strictly increasing virtual value functions and values drawn from a continuous distribution, ties in virtual value are a measure zero event, i.e., for any two agents i and j , $\Pr[\phi_i(v_i) = \phi_j(v_j)] = 0$. On the other hand, when virtual value functions are constant on an interval $[a, b]$ and the distribution assigns some non-zero probability to values in this interval, there is a measurable, i.e., non-zero, probability of ties. The virtual surplus maximization mechanism can break these ties arbitrarily or randomly. Especially in symmetric environments we will prefer the symmetric tie-breaking rule by, e.g., for single-unit environments, choosing the winner of the tie uniformly at random.

It is instructive to see exactly what the virtual surplus maximization mechanism does when there are ties in virtual values. Figure 3.2 depicts such a virtual valuation function (which corresponds to the ironed virtual value for revenue for the bimodal agent that will be derived subsequently



(a) Unique highest virtual value.



(b) Non-unique highest virtual value.

Figure 3.2. The weakly monotone virtual valuation function $\phi(v)$ under two realizations of four agent values depicting both the case where the highest virtual value is unique and the case where it is not unique.

in Section 3.3.5). Instantiating the agents' values corresponds to picking points on the horizontal axis. The agents' virtual valuations can then be read off the plot. The optimal auction assigns the item to the agent with the highest virtual value. If there is a tie, it picks a random tied agent to win.

Figure 3.2(a) depicts a realization of values for $n = 4$ agents where the highest virtual value is unique. What does the virtual surplus maximization do here? It allocates the item to the highest-valued agent, i.e., agent 1 in the figure. Figure 3.2(b) depicts a second realization of values where the highest virtual value is not unique. With uniform random tie breaking, a random tied agent is selected as the winner, i.e., one of agents 1, 2, and 3 in the figure. In general if the highest virtual value has a k -agent tie then each of these tied agents wins with probability $1/k$.

The payment an agent must make in expectation over the random tie-breaking rule can be calculated as follows. Consider the case where there is a unique highest virtual value. The agent with this virtual value

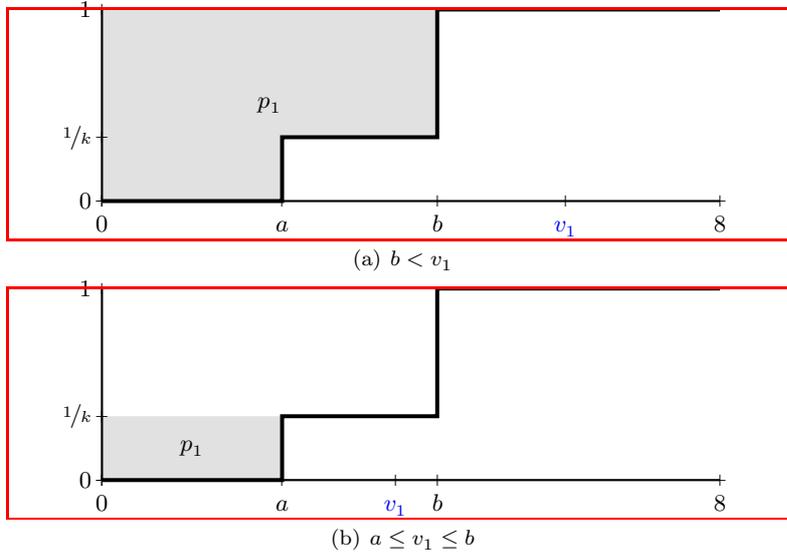


Figure 3.3. The allocation (black line) and payment rule (gray region) for agent 1 given fixed \mathbf{v}_{-1} with $k - 1$ of the other agents tied for having the highest virtual value, i.e., with values in $[a, b]$ (e.g., from virtual valuation function of Figure 3.2). For $v_1 \in [a, b]$, agent 1 would be in a k -agent tie for the highest virtual value; for $v_1 > b$ agent 1 would win outright.

wins, assume it is agent 1 (Alice). To calculate her payment we need to consider her allocation rule for fixed values \mathbf{v}_{-1} of the other agents. This allocation rule is

$$x_1(z, \mathbf{v}_{-1}) = \begin{cases} 1 & \text{if } z > b \\ 1/k & \text{if } z \in [a, b] \\ 0 & \text{if } z < a. \end{cases}$$

when \mathbf{v}_{-1} has a $k - 1$ agents in interval $[a, b]$. The $1/k$ probability of winning for $z \in [a, b]$ arises from our analysis of what happens in a k -agent tie. When Alice has the unique highest virtual value, i.e., $v_1 > b$, then $p_1 = b - a/k$, see Figure 3.3(a). On the other hand, when Alice is tied for the highest virtual value with $k - 1$ other agents with values in interval $[a, b]$, as depicted in Figure 3.3(b), her expected payment is $p_1 = a/k$. Of course, $x_1 = 1/k$ so such an expected payment can be implemented by charging a to the tied agent that wins and zero to the losers.

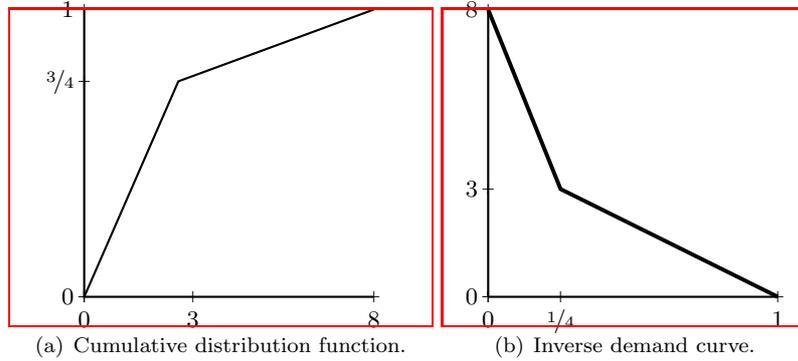


Figure 3.4. Depicted are the cumulative distribution function $F(v)$ and inverse demand curve $V(q)$ corresponding to the bimodal agent of Example 3.3. The inverse demand curve is obtained from the cumulative distribution function by rotating it 90 degrees counterclockwise.

3.3.4 Quantile Space, Price-posting Revenue, and Derivation of Virtual Values

In this section we give an economically intuitive derivation of virtual value functions for revenue maximization.

Consider an agent Alice with a single-dimensional linear preference (Definition 2.6). Alice's preference is described by her value v which is drawn from distribution F . There is a one-to-one mapping between Alice's value and her strength relative to the distribution. For instance, Alice with value $v = 0.9$ drawn from $U[0, 1]$ is stronger than 90% and weaker than 10% of values drawn from the same distribution. Denote by *quantile* quantile q the relative strength of a value where $q = 0$ is the strongest and $q = 1$ is the weakest, and by $V(\cdot)$ the *inverse demand curve* that maps quantiles to values. Importantly, the distribution of an agent's quantile is always $U[0, 1]$ as the probability that an agent's quantile q is below a given \hat{q} is exactly \hat{q} .

Definition 3.9. The *quantile* of a single-dimensional agent with value $v \sim F$ is the measure with respect to F of stronger values, i.e., $q = 1 - F(v)$; the *inverse demand curve* maps an agent's quantile to her value, i.e., $V(q) = F^{-1}(1 - q)$.

Example 3.6. For the example of a uniform agent (Example 3.2) where $F(z) = z$, the inverse demand curve is $V(q) = 1 - q$; for the example of

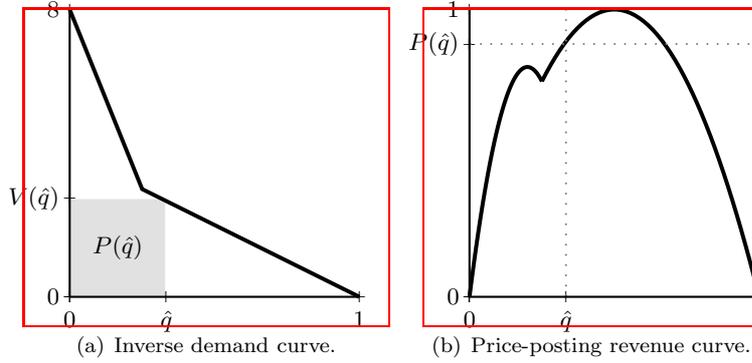


Figure 3.5. Depicted are the inverse demand curve $V(q)$ and price-posting revenue curve $P(q)$ corresponding to the bimodal agent of Example 3.3. The price-posting revenue curve is given by $P(\hat{q}) = \hat{q} \cdot V(\hat{q})$, i.e., the area of the rectangle of width \hat{q} and height $V(\hat{q})$ that fits under the inverse demand curve.

a bimodal agent (Example 3.3), the inverse demand curve is depicted in Figure 3.4.

In Section 2.4 we defined the allocation rule for an agent as a function of her value as $x(\cdot)$ and characterized the allocation rules that can arise in Bayes-Nash equilibrium as the class of monotone non-decreasing functions (of value). The *allocation rule in quantile space* is denoted $y(q) = x(V(q))$. Since quantile and value are indexed in the opposite direction, $y(\cdot)$ will be monotone non-increasing in quantile.

Consider posting a take-it-or-leave-it price of $V(\hat{q})$ for some quantile \hat{q} . By the definition of the inverse demand curve $V(\cdot)$, such a price is accepted with probability \hat{q} . In other words, the ex ante sale probability of posting price $V(\hat{q})$ is \hat{q} . Notice that the allocation rule of this price-posting mechanism is simply the reverse step function that starts at one and steps from one to zero at \hat{q} . We can define a revenue curve by considering the revenue from this *price-posting* approach as a function of the ex ante service probability \hat{q} . For the uniform example, the price-posting revenue curve is $P(\hat{q}) = \hat{q} - \hat{q}^2$; for the bimodal example, it is depicted in Figure 3.5(b).

Definition 3.10. The *price-posting revenue curve* of a single-dimensional linear agent specified by inverse demand curve $V(\cdot)$ is $P(\hat{q}) = \hat{q} \cdot V(\hat{q})$ for any $\hat{q} \in [0, 1]$.

We can use revenue equivalence (via the payment identity) to express

the revenue of any allocation rule in terms of the price-posting revenue curve. The main idea is the following. By revenue equivalence, any two mechanisms with the same allocation rule have the same revenue. Given an allocation rule y we can construct a mechanism with that allocation rule by taking the appropriate convex combination of price-posting mechanisms. Below we walk through this approach in detail.

An allocation rule y is a monotone non-increasing function from $[0, 1]$ to $[0, 1]$. The allocation rules for price postings are reverse step functions. The class of reverse step functions are a basis for the class of monotone non-increasing functions from $[0, 1]$ to $[0, 1]$: any such monotone non-increasing function can be expressed as a convex combination of (a.k.a., distribution over) reverse step functions. Consider the distribution $G^y(z) = 1 - y(z)$ and the mechanism that draws $\hat{q} \sim G^y$ and posts price $V(\hat{q})$. Notice, that the probability that Alice with fixed quantile q and value $V(q)$ is allocated by this mechanism is:

$$\Pr_{\hat{q} \sim G^y}[V(\hat{q}) < V(q)] = \Pr_{\hat{q} \sim G^y}[\hat{q} > q] = 1 - G^y(q) = y(q).$$

The mechanism resulting from the above convex combination of price postings has allocation rule exactly $y(\cdot)$ and Alice's expected payment (i.e., the expected revenue) is equal to the same convex combination of revenues $P(\hat{q})$ from posting price $V(\hat{q})$ with $\hat{q} \sim G^y$. This revenue is as follows, via a change of variables from $\hat{q} \sim G^y$ to $q \sim U[0, 1]$ according to G^y 's density function $g^y(z) = \frac{d}{dz}G^y(z) = \frac{d}{dz}(1 - y(z)) = -y'(z)$, integration by parts, and the assumption that $P(0) = P(1) = 0$ (there is no revenue from always selling or never selling; Mathematical Note on page 60).

$$\begin{aligned} \mathbf{E}_{\hat{q} \sim G^y}[P(\hat{q})] &= \mathbf{E}_{q \sim U[0,1]}[-y'(q) \cdot P(q)] \\ &= \mathbf{E}_{q \sim U[0,1]}[P'(q) \cdot y(q)], \end{aligned}$$

where $P'(q) = \frac{d}{dq}P(q)$ is the marginal increase in price-posting revenue for an increase in ex ante allocation probability, a.k.a., the *marginal price-posting revenue* at q . Notice that the calculation of Alice's expected payment for allocation rule y above is implicitly taking the expectation over Alice's quantile $q \sim U[0, 1]$ via the definition of the price-posting revenue curve $P(\cdot)$. Of course, by revenue equivalence (Theorem 2.2), any mechanism with the same allocation rule generates the same revenue.

Proposition 3.9. *A single-agent mechanism with allocation rule y has expected revenue equal to the surplus of marginal price-posting revenue $\mathbf{E}_q[P'(q) \cdot y(q)]$.*

The above rephrasing of the expected revenue in terms of marginal revenue is an amortized analysis. Notice that if we serve Alice with quantile q with some probability then, were her quantile lower (i.e., stronger), she would be served with no lower a probability. Therefore, the contribution to the revenue from all quantiles above quantile q can be credited to the change in service probability at q . The marginal price-posting revenue is precisely this reamortizing of revenues across the different agent quantiles.

The marginal price-posting revenues are exactly the virtual values described previously by equation (3.2).

$$P'(q) = \frac{d}{dq}(q \cdot V(q)) = V(q) + qV'(q) = v - \frac{1-F(v)}{f(v)}, \quad (3.5)$$

where the first equality follows from the definition of price-posting revenue (Definition 3.10) and the last equality follows from the definition of the inverse demand curve $V(\cdot)$ whereby $v = V(q)$ satisfies $F(v) = 1 - q$ and $1/f(v) = -\frac{d}{dq}V(q) = -V'(q)$. Recall that a distribution is regular if $v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing or, equivalently, the marginal price-posting revenue is monotone non-increasing, or equivalently the price-posting revenue curve is concave.

Proposition 3.10. *A distribution F is regular if and only if its corresponding price-posting revenue curve is concave.*

Proposition 3.9 shows the expected revenue of a mechanism is equal to its surplus of marginal price-posting revenue. For regular distributions, the marginal price posting revenue derived above is monotone; therefore, we can conclude that the virtual surplus maximization mechanism with virtual value function defined by the marginal price-posting revenue curve (Definition 3.5) is dominant strategy incentive compatible and profit optimal (Corollary 3.6).

Theorem 3.11. *For agents with values drawn from regular distributions the marginal price-posting revenue curves are virtual value functions for revenue and the virtual surplus maximization mechanism optimizes expected profit in dominant strategy equilibrium.*

The price-posting revenue curve $P(\hat{q})$ is defined by the revenue obtained by posting a price that is accepted with probability \hat{q} . Consider instead the single-agent optimization of optimizing revenue subject to an ex ante constraint \hat{q} . This optimization problem is not generally solved by a price posting; however, for regular distributions it is. Subsequently in Section 3.4 we will consider this more general problem and define

from it (optimal) revenue curves. For regular distributions price-posting revenue curves and (optimal) revenue curves are equal.

3.3.5 Virtual Surplus Maximization Subject to Monotonicity

We now turn our attention to the case where the non-game-theoretic problem of optimization of marginal price-posting revenue is not itself inherently monotone. An *irregular* distribution is one for which the price-posting revenue curve is non-concave (in quantile). The marginal price-posting revenue curves (and virtual value functions defined from them) are non-monotone; therefore, a higher value might result in a lower virtual value. As $\text{OPT}(\phi(\cdot))$ is non-monotone for such a virtual value function, there is no payment rule with which its outcome is incentive compatible (by the only-if direction of Corollary 2.12). We must instead optimize this virtual surplus subject to monotonicity.

Recall that virtual values, e.g., $v - \frac{1-F(v)}{f(v)}$, correspond to an amortized analysis where we “charge” the value v if it is served for the lower price its service implies for higher values. When this direct approach to an amortized analysis gives a non-monotone virtual value function, the following generic *ironing procedure* gives an ironed virtual value function which is monotone and for which pointwise optimization is equivalent to the optimization of expected virtual surplus subject to monotonicity of the allocation rule.

There are two key ideas to this ironing procedure. First, if there is some interval $[a, b]$ of quantiles that all receive the same allocation probability, then the virtual values of these quantiles can be reamortized arbitrarily and the expected virtual value of the allocation rule is unchanged. Second, if we reamortize by simple averaging then we get “ironed” virtual values that are constant on the $[a, b]$ interval and optimization of the ironed virtual surplus will give the same allocation probability to quantiles within the interval. Therefore, the approach of the second part implies the assumption of the first part. Moreover, in terms of fixing non-monotonicities, after ironing the virtual value are constant (and therefore weakly monotone) on the interval $[a, b]$.

As in previous sections, the geometry of this reamortization is more transparent in quantile space rather than value space. This is because quantiles are drawn from a uniform distribution so reamortizing by moving virtual value from one quantile to another is balanced with respect to the distribution. If we were to do such a shift of virtual value in value

space then we would need to normalize by the density function of the distribution. We therefore proceed by considering a virtual value function $\phi(\cdot)$ in quantile space. We denote the cumulative virtual value for quantiles at most \hat{q} as $\Phi(\hat{q}) = \int_0^{\hat{q}} \phi(q) dq$. For profit maximization, the virtual value functions correspond to marginal price-posting revenue curves and cumulative virtual value functions correspond to price-posting revenue curves, i.e., $\phi(q) = P'(q)$ and $\Phi(q) = P(q)$. The ironing procedure we will describe, however, can be applied to any non-monotone virtual value function.

The goal of ironing is arrive at a monotone (ironed) virtual value function, equivalently, a concave cumulative virtual value function, without any loss in virtual surplus for monotone allocation rules. We now investigate the consequences of the ironing procedure proposed above on the virtual value and cumulative virtual value functions. The averaging of virtual value over an interval $[a, b]$ in quantile space replaces the function on that interval with a constant equal to the original function's average. We can then integrate to see what the effect on the cumulative virtual value is. Notice that on $q \in [0, a]$ and $q \in [b, 1]$ this integral is identically $\Phi(q)$; while for $q \in [a, b]$ it is the integral of a constant function and therefore linearly connects $(a, \Phi(a))$ to $(b, \Phi(b))$ with a line segment. For the bimodal agent of Example 3.3 these quantities are depicted in Figure 3.6 with an arbitrary choice of a and b .

If we iron the virtual value functions and then optimize with ironed virtual values as virtual values, then the revenue is again the virtual surplus (by the correctness of ironing construction, e.g., as proven by Theorem 3.12, below). It remains to choose the appropriate intervals on which to iron so that the ironed virtual value functions are monotone (equivalently, the ironed revenue curve is concave) and the optimization of ironed virtual surplus also optimizes the virtual surplus. Intuitively, higher revenue curves produce higher revenues. As the ironing procedure operates on the cumulative virtual value functions by replacing an interval with a line segment, we can construct the concave hull, i.e., the smallest concave upper-bound, of the cumulative virtual value function by ironing. Notice that this ironed cumulative virtual value function has two advantages over the original cumulative virtual value function: it is pointwise higher and it is concave.

Definition 3.11. The *ironing procedure* for (non-monotone) virtual value function ϕ (in quantile space)⁷ is:

⁷ The ironing procedure can also be expressed in value space by first mapping

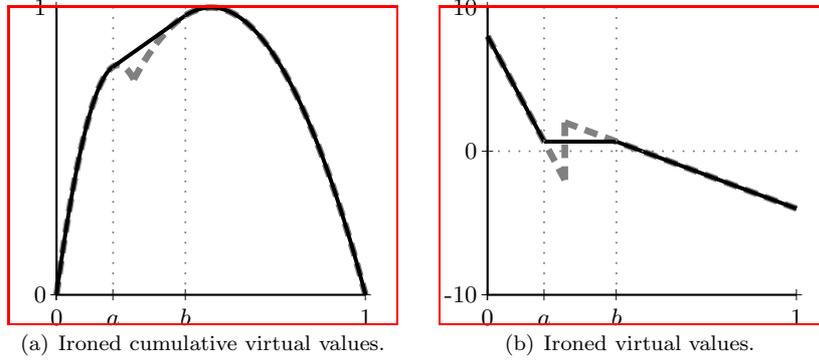


Figure 3.6. Consider the bimodal agent of Example 3.3 and virtual value function equal to the marginal price-posting revenue curve. The cumulative virtual value and virtual value functions in quantile space are depicted (thick, gray, dashed lines) in the left and right diagram, respectively. After ironing on an arbitrarily selected interval $[a, b]$, the resulting cumulative virtual value and virtual value functions are depicted (thin, black, solid lines).

- (i) Define the *cumulative virtual value* function as $\Phi(\hat{q}) = \int_0^{\hat{q}} \phi(q) dq$.
- (ii) Define *ironed cumulative virtual value* function as $\bar{\Phi}(\cdot)$ as the concave hull of $\Phi(\cdot)$.
- (iii) Define the *ironed virtual value* function as $\bar{\phi}(q) = \frac{d}{dq} \bar{\Phi}(q) = \bar{\Phi}'(q)$.

Theorem 3.12. *For any monotone allocation rule $y(\cdot)$ and any virtual value function $\phi(\cdot)$, the expected virtual surplus of an agent is upper-bounded by her expected ironed virtual surplus, i.e.,*

$$\mathbf{E}[\phi(q) \cdot y(q)] \leq \mathbf{E}[\bar{\phi}(q) \cdot y(q)].$$

Furthermore, this inequality holds with equality if the allocation rule y satisfies $y'(q) = 0$ for all q where $\bar{\Phi}(q) > \Phi(q)$.

Proof. By integration by parts for any virtual value function $\phi^\dagger(\cdot)$ and monotone allocation rule $y(\cdot)$ (Mathematical Note on page 60),

$$\mathbf{E}[\phi^\dagger(q) \cdot y(q)] = \mathbf{E}[-y'(q) \cdot \Phi^\dagger(q)]. \quad (3.6)$$

Notice that the (non-increasing) monotonicity of the allocation rule $y(\cdot)$

values to quantiles via the cumulative distribution function or inverse demand curve, executing the ironing procedure in quantile space, and then mapping ironed virtual value functions back into value space.

implies the non-negativity of $-y'(q)$. With the left-hand side of equation (3.6) as the expected virtual surplus, it is clear that a higher cumulative virtual value implies no lower expected virtual surplus. By definition of $\bar{\Phi}(\cdot)$ as the concave hull of $\Phi(\cdot)$, $\bar{\Phi}(q) \geq \Phi(q)$ and, therefore, for any monotone allocation rule, in expectation, the ironed virtual surplus is at least the virtual surplus. I.e., $\mathbf{E}[-y(q) \cdot \bar{\Phi}(q)] \geq \mathbf{E}[-y(q) \cdot \Phi(q)]$.

To see the equality under the assumption that $y'(q) = 0$ for all q where $\bar{\Phi}(q) > \Phi(q)$, rewrite the difference between the ironed virtual surplus and the virtual surplus via equation (3.6) as,

$$\mathbf{E}[\bar{\phi}(q) \cdot y(q)] - \mathbf{E}[\phi(q) \cdot y(q)] = \mathbf{E}[-y'(q) \cdot [\bar{\Phi}(q) - \Phi(q)]].$$

The assumption implies the term inside the expectation on the left-hand side is zero for all q . \square

Corollary 3.13. *For any virtual value function $\phi(\cdot)$ with ironed virtual value $\bar{\phi}(\cdot)$ from the ironing procedure (Definition 3.11), the optimization of virtual surplus subject to monotonicity of the allocation rule is equivalent to optimization of ironed virtual surplus pointwise.*

We now conclude this section by summarizing the consequences of ironing for virtual surplus maximization. First, we can define the *ironed virtual surplus maximization* mechanism for virtual value functions ϕ as the virtual surplus maximization mechanism applied to the ironed virtual value functions $\bar{\phi}$. This profile $\bar{\phi}$ of ironed virtual value functions is constructed from the profile ϕ of virtual value functions by applying the ironing procedure individually to each virtual value function.

Theorem 3.14. *For any (non-monotone) virtual value functions ϕ , the ironed virtual surplus maximization mechanism maximizes expected virtual surplus in dominant strategy equilibrium.*

Corollary 3.15. *For (irregular) single-dimensional linear agents, the ironed marginal price-posting revenue curves are virtual value functions for revenue and the virtual surplus maximization mechanism optimizes expected profit in dominant strategy equilibrium.*

The ironing procedure above results in virtual value functions that are not strictly monotone. See Section 3.3.3 for a discussion of the virtual surplus maximization mechanism with non-strictly monotone virtual value functions in single-item environments.

3.4 Multi- to Single-agent Reduction

While the previous sections gave a complete approach to profit maximization for single-dimensional linear agents, here we give an alternative derivation that comes to the same conclusion but provides more conceptual understanding, especially for irregular distributions. The approach will be to reduce the problem of solving a multi-agent mechanism design problem to that of solving a collection of simple single-agent pricing problems. It observes and makes use of a *revenue-linearity* property that is satisfied by single-dimensional agents with linear utility. In ?? this reduction is extended to multi-dimensional non-linear agents.

A mechanism for a single agent is simply a menu of outcomes where, after the agent realizes her value from the distribution, she chooses the outcome she most prefers. This observation is known as the *taxation principle* and is a simple consequence of the revelation principle (Theorem 2.11). It can be seen as follows: The agent's actions in the mechanism induce a set of (possibly randomized) outcomes; for a fully rational agent, these probabilistic outcomes may as well be listed on a menu from which the agent just chooses her favorite. Each of these probabilistic outcomes can be summarized by its allocation probability and expected payment (as far as the preferences of a single-dimensional linear agent is concerned). We call such a probabilistic allocation a lottery, and the menu of lotteries and their accompanying prices a *lottery pricing*. The allocation and payment rules $(x(\cdot), p(\cdot))$ described in Section 3.1 precisely define such a menu where the outcomes are indexed so that the agent with value v prefers outcome $(x(v), p(v))$ over all other outcomes.

Below we will look at two optimization problems. The first will be an *ex ante pricing* problem where we look for the lottery pricing with the optimal revenue subject to a constraint on the ex ante service probability $\mathbf{E}_v[x(v)]$. The revenue of the optimal ex ante pricings induce a concave *revenue curve*. We will then look at an *interim pricing* problem where we have a constraint on the allocation rule $x(\cdot)$ and we again wish to optimize revenue subject to that constraint. The main conclusion will be that we can express the optimal interim pricing as a convex combination of optimal ex ante pricings. The decomposition will enable the expected payments to be expressed in terms of a monotone *marginal revenue curve* (cf. Section 3.3.4). Pointwise optimization of the surplus of marginal revenue then gives the optimal revenue.

3.4.1 Revenue Curves

It will be more economically intuitive to study lottery pricings in quantile space. Alice has her quantile q drawn from the uniform distribution $U[0, 1]$ and value $V(\cdot)$ according to the inverse demand curve. Upon realizing her quantile, she will choose her preferred outcome from a lottery pricing. This two step process induces an allocation rule $y(q) = x(V(q))$ and an ex ante probability $\mathbf{E}_q[y(q)]$ that Alice is served. Recall that the allocation rule is taken in expectation with respect to the randomization in the outcome of the lottery that Alice buys, and the ex ante service probability is taken additionally in expectation with respect to the randomization of Alice's quantile.

Definition 3.12. With equality constraint \hat{q} on the ex ante allocation probability, the single-agent *ex ante pricing problem* is to find the revenue-optimal lottery pricing. The optimal ex ante revenue, as a function of \hat{q} , is denoted by the *revenue curve* $R(\hat{q})$.

It will be important to contrast the revenue-optimal lottery pricing for an ex ante constraint \hat{q} with the price posting that satisfies the same constraint. The revenues of these two pricings are given by the revenue curve $R(\hat{q})$ and price-posting revenue curve $P(\hat{q})$ (from Section 3.3.4). First, recall that the difficulty with deriving optimal mechanisms directly from the price-posting revenue curve $P(\cdot)$ is that it may not be concave. On the other hand the revenue curve $R(\cdot)$ is always concave.⁸ Second, notice that the allocation rule for price posting, which serves all values that are at least $V(\hat{q})$, is the strongest allocation rule with ex ante service probability \hat{q} in the following sense. Any other allocation rule can shift allocation probability from stronger (lower) quantiles to weaker (higher) quantiles but cannot allocate with any greater probability to the strongest \hat{q} measure of quantiles. Therefore, for the ex ante probability \hat{q} , the allocation rule of the optimal ex ante pricing is no stronger than that of price posting. Third, the optimal ex ante pricing for constraint \hat{q} obtains at least the revenue of price posting. This observation is immediate from the fact that it is optimizing over lottery pricings that include the posting price $V(\hat{q})$. We summarize these observations as the

⁸ This observation follows from the fact that the space of lottery pricings is convex: randomizing between two lottery pricings gives a lottery pricing that corresponds to the lotteries' convex combination and gives ex ante allocation probability and expected revenue according to the same convex combination. In contrast, the space of price postings is not convex: the convex combination of two price postings cannot be expressed as a price posting. Consequently and as we have already observed, the price-posting revenue curve is not generally concave.

following proposition which, with Proposition 3.9 (essentially, revenue equivalence), will be sufficient for proving the optimality of marginal revenue maximization; we defer precise characterization of the optimal ex ante lottery pricing to later in this section.

Proposition 3.16. *The optimal ex ante pricing problems induce a concave revenue curve and, for any ex ante service probability, the optimal lottery has no stronger an allocation rule and no lower a revenue than price posting.*

3.4.2 Optimal and Marginal Revenue

We now formulate an interim lottery pricing problem that takes an allocation rule as a constraint and asks for the optimal lottery pricing with an allocation rule that is no stronger than the one given. To do so we must first generalize the definition of strength (as discussed previously when comparing price posting with optimal lotteries). Recall that with the same ex ante allocation probability the difference between the price posting and an optimal lottery is that the optimal lottery may have service probability shifted from strong (low) quantiles to weak (high) quantiles. This condition generalizes naturally.

The ex ante probability that allocation rule $y(\cdot)$ allocates to the strongest \hat{q} measure of quantiles is $Y(\hat{q}) = \int_0^{\hat{q}} y(q) dq$; we refer to $Y(\cdot)$ as the *cumulative allocation rule* for $y(\cdot)$. The (non-increasing) monotonicity of allocation rules implies that cumulative allocation rules are concave. As follows, we can view an allocation rule $\hat{y}(\cdot)$ as a constraint via its cumulative allocation rule \hat{Y} .

Definition 3.13. Given an allocation constraint \hat{y} with cumulative constraint \hat{Y} , the allocation rule y with cumulative allocation rule Y is *weaker* (resp. \hat{y} is *stronger*) if and only if it satisfies $Y(\hat{q}) \leq \hat{Y}(\hat{q})$ for all \hat{q} ; denote this relationship by $y \preceq \hat{y}$.

A strong allocation rule as a constraint corresponds to a weak constraint as it permits the most flexibility in allocation rules that satisfy it. The ex ante pricing problem for constraint \hat{q} is a special case of the interim pricing problem. The strongest allocation rule that serves with probability \hat{q} is the reverse step function that steps from one to zero at \hat{q} ; therefore, the allocation constraint $\hat{y}^{\hat{q}}$ is the weakest constraint that allows service probability at most \hat{q} . In comparison, a general allocation constraint \hat{y} (e.g., with total allocation probability $\mathbf{E}[\hat{y}(q)] = \hat{q}$) allows

more fine-grained control by giving a constraint, for all \hat{q}^\dagger , on the cumulative service probability of any $[0, \hat{q}^\dagger]$ measure of quantiles by $\hat{Y}(\hat{q}^\dagger)$. Of course, given an allocation constraint \hat{y} , the strongest allocation rule that satisfies the constraint is the constraint itself, i.e., $y = \hat{y}$. From this notion of strength we can take an allocation rule as a constraint and consider the optimization question of finding an allocation rule that is no stronger and with the highest possible revenue.

Definition 3.14. The optimal revenue subject to an allocation constraint $\hat{y}(\cdot)$ is $\mathbf{Rev}[\hat{y}]$ and it is attained by the *optimal interim pricing* for \hat{y} .

An important property of this definition of the strength of an allocation rule is that it closed under convex combination, i.e., if $\hat{y} = \hat{y}^\dagger + \hat{y}^\ddagger$, $y^\dagger \preceq \hat{y}^\dagger$, and $y^\ddagger \preceq \hat{y}^\ddagger$ then $y \preceq \hat{y}$ for $y = y^\dagger + y^\ddagger$. This means that one approach to construct an allocation rule y that satisfies the allocation constraint \hat{y} is to express y as a convex combination of ex ante constraints, and to implement each with the optimal ex ante pricing. Relative to the construction of Proposition 3.9, using optimal lottery pricings improves on price postings in that for each \hat{q} the optimal ex ante revenue $R(\hat{q})$ may exceed the price-posting revenue $P(\hat{q})$. Consider the mechanism that draws \hat{q} from the distribution $G^{\hat{y}}(z) = 1 - \hat{y}(z)$ and offers Alice the optimal ex ante pricing for \hat{q} . The optimal revenue for allocation constraint \hat{y} must be at least the revenue of this mechanism. By the Mathematical Note on page 60, we have:

$$\begin{aligned} \mathbf{Rev}[\hat{y}] &\geq \mathbf{E}_{\hat{q} \sim G^{\hat{y}}} [R(\hat{q})] \\ &= \mathbf{E}_q [-\hat{y}'(q) \cdot R(q)] \\ &= \mathbf{E}_q [R'(q) \cdot \hat{y}(q)], \end{aligned}$$

where $R'(q) = \frac{d}{dq}R(q)$ is the *marginal revenue* at q .

Definition 3.15. The *surplus of marginal revenue* of an allocation constraint \hat{y} is $\mathbf{MargRev}[\hat{y}] = \mathbf{E}_q [R'(q) \cdot \hat{y}(q)]$.

3.4.3 Downward Closure and Pricing

We now make a brief aside to discuss downward closure of the environment and its relationship to the previously defined single-agent lottery pricing problems. Recall that a downward closure environment is one

where from any feasible outcome it is always feasible to additionally reject and agent who was previously being served. Our definition of the optimal ex ante pricing problem is not downward closed as we required that the ex ante constraint be met with equality. On the other hand, our definition of the optimal interim pricing problem was downward closed as it was allowed that $Y(1) < \hat{Y}(1)$. These definitions were given above as they are the most informative.

It is possible to consider a downward-closed variant of the ex ante pricing problem where a lottery pricing is sought with ex ante probability at most \hat{q} . Obviously, adding downward closure results in a revenue curve that is monotone non-decreasing. From the non-downward-closed revenue curve, the downward-closed revenue curve is given as a function of \hat{q} by $\max_{q \leq \hat{q}} R(q)$. Thus, the downward-closed revenue curve after the monopoly quantile is constant. Importantly, the downward-closed marginal revenue curve is always non-negative. It is similarly possible to consider a non-downward-closed variant of the interim pricing problem where it is additionally required that $Y(1) = \hat{Y}(1)$.

In our discussion of revenue linearity in the subsequent section, it will be important not to mix-and-match with respect to downward closure.

3.4.4 Revenue Linearity

The above derivation says the surplus of marginal revenue of an allocation constraint is a lower bound on its optimal revenue. A central dichotomy in optimal mechanism design is given by the partitioning of single-agent problems into those for which this inequality is tight and those when it is not. Notice that linearity of the revenue operator $\mathbf{Rev}[\cdot]$ implies by the above derivation that for any allocation constraint the optimal revenue and surplus of marginal revenue are equal.

Definition 3.16. A agent (with implicit utility function, type space, and distribution over types) is *revenue linear* if $\mathbf{Rev}[\cdot]$ is linear, i.e., if when $\hat{y} = \hat{y}^\dagger + \hat{y}^\ddagger$ then $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$.⁹

Proposition 3.17. For a revenue-linear agent and any allocation constraint \hat{y} , the optimal revenue is equal to the surplus of marginal revenue, i.e., $\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}]$.

We now show that single-dimensional linear agents are revenue linear.

⁹ It is assumed that the ex ante and interim problem are consistent with respect to downward closure, see Section 3.4.3.

This result is a consequence of three main ingredients: the concavity of the revenue curve $R(\cdot)$, that the optimal ex ante pricings which define the revenue curve gives more revenue with a weaker allocation rule than the price postings which define price-posting revenue curves (Proposition 3.16), and that revenue equivalence allows revenue to be expressed in terms of price-posting revenue curves (Proposition 3.9). Optimal revenue equaling surplus of marginal revenue for single-dimensional linear agents, then, is an immediate corollary of this revenue linearity and Proposition 3.17.

Theorem 3.18. *A single-dimensional linear agent is revenue linear.*

Proof. Before we begin, notice that for any revenue curve $R(\cdot)$ and allocation rule $y(\cdot)$ the surplus of marginal revenue $\mathbf{MargRev}[y]$ can be equivalently expressed as

$$\mathbf{E}_q[-y'(q)R(q)] = \mathbf{E}_q[R'(q)y(q)] = \mathbf{E}_q[-R''(q)Y(q)] + R'(1)Y(1)$$

via integration by parts (with $R(1) = R(0) = Y(0) = 0$; Mathematical Note on page 60). The same equations also govern the surplus of marginal price-posting revenue in terms of revenue curve $P(\cdot)$. Two observations:

- (i) The left-hand side shows that a pointwise higher revenue curve gives a no lower revenue (as $-y'(\cdot)$ is non-negative). In particular, the allocated marginal revenue exceeds the surplus of marginal price-posting revenue as $R(q) \geq P(q)$ for all q (by Proposition 3.16).
- (ii) The right-hand side shows that for concave revenue curves, i.e., where $-R''(\cdot)$ is non-negative, e.g., $R(\cdot)$ not $P(\cdot)$; a stronger allocation rule gives higher revenue. In particular, the allocation rule y obtained by optimizing for \hat{y} has no higher surplus of marginal revenue than does \hat{y} .¹⁰

We have already concluded that the surplus of marginal revenue lower bounds the optimal revenue; so to prove the theorem it suffices to upper bound the optimal revenue by the surplus of marginal revenue. Suppose we optimize for \hat{y} and get some weaker allocation rule y , then y is a fixed

¹⁰ Consistency with respect to downward-closure (see Section 3.4.3) implies the inequality on the $R'(1)Y(1)$ term. For the downward-closed case: the marginal revenue $R'(1)$ is non-negative and thus $R'(1)\hat{Y}(1) \geq R'(1)Y(1)$. For the non-downward-closed case: it is required that $\hat{Y}(1) = Y(1)$ and thus $R'(1)\hat{Y}(1) = R'(1)Y(1)$.

point of $\mathbf{Rev}[\cdot]$ (optimizing with y as an allocation constraint gives back allocation rule y); therefore,

$$\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[y].$$

By revenue equivalence (Proposition 3.9), the revenue of any allocation rule is equal to its surplus of marginal price-posting revenue, so

$$\mathbf{Rev}[y] = \mathbf{E}[P'(q) \cdot y(q)].$$

By observation (i), for allocation rule y , the surplus of marginal revenue is at least the surplus of marginal price-posting revenue,

$$\mathbf{E}[-y'(q) \cdot P(q)] \leq \mathbf{E}[-y'(q) \cdot R(q)].$$

By observation (ii), the surplus of marginal revenue for \hat{y} is at least that of y ,

$$\mathbf{E}[-R''(q) \cdot Y(q)] \leq \mathbf{E}[-R''(q) \cdot \hat{Y}(q)] = \mathbf{MargRev}[\hat{y}].$$

The above sequence of inequalities implies that the surplus of marginal revenue is at least the optimal revenue for \hat{y} ,

$$\mathbf{Rev}[\hat{y}] \leq \mathbf{MargRev}[\hat{y}]. \quad \square$$

Corollary 3.19. *For an agent with single-dimensional, linear utility, the optimal revenue equals the marginal revenue, i.e.,*

$$\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}] = \mathbf{E}[R'(q)\hat{y}(q)].$$

Observe that Corollary 3.19 implies that the marginal revenue curve is a virtual value function for revenue. The virtual surplus maximization mechanism for these virtual values maximizes expected profit.

Theorem 3.20. *For linear single-dimensional agents, the marginal revenue curves are a virtual value functions for revenue and the virtual surplus maximization mechanism optimizes expected profit in dominant strategy equilibrium.*

3.4.5 Optimal Ex Ante Pricings, Revisited

We now return to the question of characterizing the optimal ex ante pricings that define the revenue curve (Definition 3.12). Given an ex ante constraint \hat{q} , what is the optimal lottery pricing? We saw previously that price posting $V(\hat{q})$ is a simple way to serve an agent with ex ante probability \hat{q} . When the distribution is regular, it is easy to see that price posting is optimal. By monotonicity of the marginal price-posting

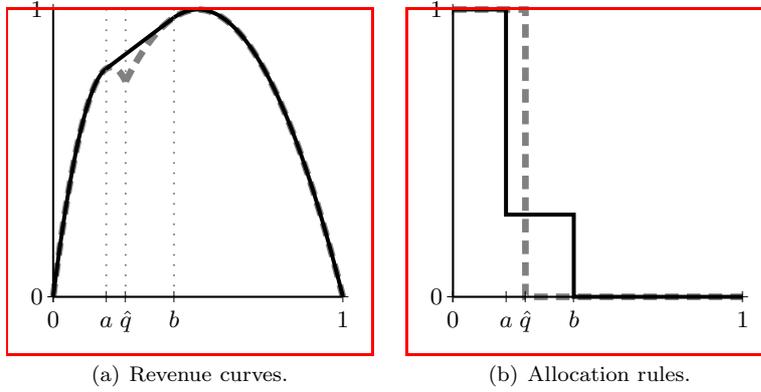


Figure 3.7. Depicted are the revenue curve $R(q)$, price-posting revenue curve $P(q)$, and the allocation rules corresponding to ex ante allocation constraint \hat{q} for the bimodal agent of Example 3.3. For this agent the revenue curve $R(\cdot)$ (thin, black, solid line) is obtained from the price-posting revenue curve $P(\cdot)$ (thick, grey, striped line) by replacing the curve on interval $[a, b]$ with a line segment. The allocation rule $\hat{y}(\cdot)$ for posting price $V(\hat{q})$ is the reverse step function at \hat{q} (thick, grey, striped line). For $\hat{q} \in [a, b]$ as depicted, the allocation rule $y(\cdot)$ for the \hat{q} optimal ex ante pricing (thin, black, solid line) is the appropriate convex combination of the reverse step functions at a and b . Notice that the area under both allocation rules is equal to the ex ante service probability \hat{q} .

revenue curve, the \hat{q} measure of types with the highest marginal revenues is precisely those with quantile in $[0, \hat{q}]$. The mechanism that serves only these types is the $V(\hat{q})$ price posting. Therefore, for regular distributions $R(\cdot) = P(\cdot)$. The following is a restatement of Proposition 3.9 in terms of the revenue curve for the regular case.

Corollary 3.21. *For regular single-agent environments, allocation rule y has expected revenue equal to the surplus of marginal revenue $\mathbf{E}_q[R'(q) \cdot y(q)]$.*

To solving the ex ante pricing problem for irregular distributions we will define a very natural class of lottery pricings which directly resolve the problematic non-concavity of the price-posting revenue curves. Suppose the price-posting revenue is non-concave at some \hat{q} , instead of posting price $V(\hat{q})$ another method for serving with ex ante probability \hat{q} would be to pick any interval $[a, b]$ that contains \hat{q} and take the appropriate convex combination of posting prices $V(a)$, which serves with probability $a < \hat{q}$, and $V(b)$, which serves with probability $b > \hat{q}$, so that the combined service probability is exactly \hat{q} . The revenue from

this convex combination is the same convex combination of the revenues; the allocation rule is given by the same convex combination of the two reverse step functions. Figure 3.7(b) depicts these allocation rules. Formulaically,

$$y^{\hat{q}}(q) = \begin{cases} 1 & \text{if } q < a, \\ \frac{\hat{q}-a}{b-a} & \text{if } q \in [a, b], \text{ and} \\ 0 & \text{if } b < q. \end{cases}$$

It is easy to see that via *two-price lotteries* of this form we can obtain an ex ante revenue for every \hat{q} that corresponds to the convex hull of $P(\cdot)$. See Figure 3.7(a).

This class of two-price lotteries satisfies all the conditions that the optimal pricings satisfies with respect to Proposition 3.16. Optimal two-price lotteries (a) induce a concave revenue curve, (b) have at least the revenue of price posting, and (c) have allocation rules no stronger than those of price posting. Consequently, via the exact same proof as Theorem 3.18 (and Corollary 3.19) the optimal revenue is given by convex combination of ex ante pricings from this class. Applying this revenue-optimality result to the allocation constraint $\hat{y}^{\hat{q}}(\cdot)$, for which the aforementioned convex combination places probability one on \hat{q} , we see that the optimal two-price lottery for ex ante constraint \hat{q} is in fact optimal among all lottery pricings.

Theorem 3.22. *For a single-dimensional linear agent and ex ante constraint \hat{q} , the optimal ex ante pricing is a two-priced lottery and the optimal ex ante revenue $R(\hat{q})$ is given by the concave hull of the price-posting revenue curve $P(\cdot)$ at \hat{q} .*

3.4.6 Optimal Interim Pricings, Revisited

We now reconsider the problem of finding the optimal interim pricing (with allocation rule y) for allocation constraint \hat{y} , i.e., solving $\mathbf{Rev}[\hat{y}]$. Recall that \hat{y} is a constraint, but the allocation rule y of the optimal mechanism subject to \hat{y} may be generally weaker than \hat{y} , i.e., $y \preceq \hat{y}$. Just as we can view the ironing of the price-posting revenue curve on interval I as averaging marginal price-posting revenue on this interval, we can so view the optimization of y subject to \hat{y} . To optimize a weakly monotone function $R'(\cdot)$ subject to \hat{y} we should greedily assign low quantiles to high probabilities of service except on ironed intervals, i.e., $[a, b]$

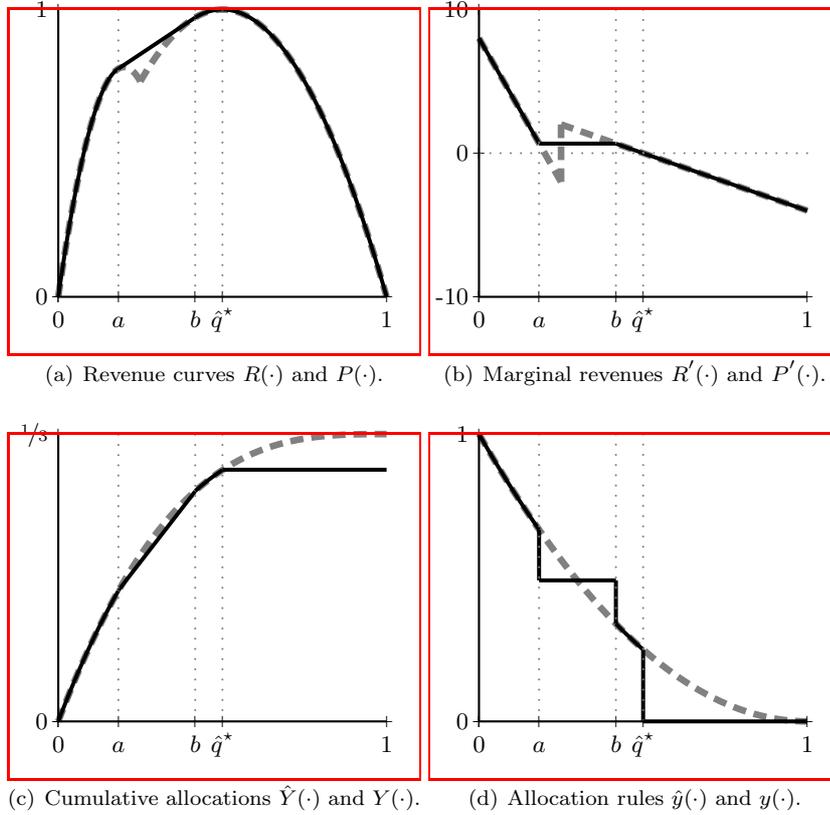


Figure 3.8. The optimal single-item auction is depicted for three bimodal agents (Example 3.3). The price-posting revenue curve $P(\cdot)$ is depicted by a thick, grey, dashed line in Figure 3.8(a). The revenue curve (thin, black, solid line) is its concave hull. The ironed interval (a, b) where $R(q) > P(q)$ is depicted. The allocation constraint $\hat{y}(q) = (1 - q)^2$ (Figure 3.8(d), thick, grey, dashed line) corresponds to lowest-quantile-wins for three agents; the allocation rule $y(q)$ (thin, black, solid line) results from optimizing $\mathbf{Rev}[\hat{y}]$. Simply, ironing corresponds to a line-segment for revenue curves and cumulative allocation rules and to averaging for marginal revenues and allocation rules.

where $q \in [a, b]$ satisfies $R''(q) = 0$. Quantiles on ironed intervals are assigned to the average probability of service for the ironed interval. One way to obtain such an allocation rule is via a *resampling transformation* σ that, for quantile q in some ironed interval $[a, b]$, resamples the quantile from this interval, i.e., as $y(q) = \mathbf{E}_\sigma[\hat{y}(\sigma(q))]$. The cumulative allocation rule Y is exactly equal to the cumulative allocation constraint

\hat{Y} except every ironed interval is replaced with a line segment. In other words, the revenue optimization of $\mathbf{Rev}[\cdot]$ can be effectively solved by superimposing the revenue curve and the allocation constraint on the same quantile axis and then ironing the allocation constraint where the revenue curve is ironed. Figure 3.8 illustrates this construction.

We will typically be in environments that are downward-closed where optimizing revenue allows the exclusion of any agent with negative virtual value. Thus, the optimal allocation rule y drops to zero after the quantile \hat{q}^* of the monopoly price; equivalently Y is flat after \hat{q}^* . For non-downward-closed environments the definition of $\mathbf{Rev}[\cdot]$ can be modified so that the total ex ante allocation probability of the constraint is met with equality, i.e., $\hat{Y}(1) = Y(1)$. See Section 3.4.3.

3.5 Social Surplus with a Balanced Budget

In this section we explore the role that the designer's budget constraint plays on mechanism design for the objective of social surplus. Assume that the mechanism designer would like to maximize social surplus, but cannot subsidize the transaction, i.e., she is constrained to mechanisms with non-negative profit. Notice that such a constraint introduces a non-linearity into the designer's objective; however, this particular non-linearity instead can be instead represented as a constraint on total payments which, because revenue is linear (Theorem 3.18), is a linear constraint.

Recall that with outcome (\mathbf{x}, \mathbf{p}) the social surplus of a mechanism is $\sum_i v_i x_i - c(\mathbf{x})$ and its profit is $\sum_i p_i - c(\mathbf{x})$. There are two standard environments where budget balance is a crucial issue. First, in an *exchange* the mechanism designer is the mediator between a buyer and seller. The feasibility constraint is *all or none* in that either the trade occurs, in which case both agents are "served," or the trade does not occur, in which case neither agent is served. Second, in a *non-excludable public project* there is a fixed cost for producing a public good, e.g., for building a bridge, and if the good is produced then all agents can make use of the good. Again, the feasibility constraint is all or none.

The surplus maximization mechanism (Definition 3.3) has a deficit, i.e., negative profit, in non-trivial all-or-none environments. For instance, to maximize surplus in an exchange, the good should be traded when the buyer's value exceeds the seller's value for the good. The critical value for the buyer is the seller's value; the critical value for the seller is the

buyer's value. When the good is sold the buyer pays the seller's value, the seller is paid the buyer's value, and the mechanism has a deficit of the difference between the two values. This difference is positive as otherwise the trade would not have occurred.

Here we address the question of maximizing social surplus subject to budget balance (taking both quantities in expectation). As with profit maximization, there is no mechanism that optimizes surplus subject to budget balance pointwise. E.g., in an exchange, if the values were known then the buyer and seller would be happy to trade at any price between their values; this is budget balanced. This approach, however, requires knowledge of a price that is between the buyer and seller's values, and this knowledge is not generally available in Bayesian mechanism design.

Our objective is surplus:

$$\text{Surplus}(\mathbf{v}, \mathbf{x}) = \sum_i v_i x_i - c(\mathbf{x});$$

in addition to the feasibility constraint (which is given by $c(\cdot)$), incentive constraints (i.e., monotonicity of each agent's allocation rule), and individual rationality constraints we have a budget-balance constraint

$$\text{Profit}(\mathbf{p}, \mathbf{x}) = \sum_i p_i - c(\mathbf{x}) \geq 0.$$

To optimize this objective in expectation subject to budget balanced in expectation we obtain the mathematical program

$$\begin{aligned} \max_{\mathbf{x}(\cdot), \mathbf{p}(\cdot)} \quad & \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i x_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v})) \right] & (3.7) \\ \text{s.t.} \quad & \mathbf{x}(\cdot) \text{ and } \mathbf{p}(\cdot) \text{ are IC and IR} \\ & \mathbf{E}_{\mathbf{v}} \left[\sum_i p_i - c(\mathbf{x}) \right] \geq 0 \end{aligned}$$

where expectations are simply integrals with respect to the density function of the valuation profile.

3.5.1 Lagrangian Relaxation

We will make two transformations of mathematical program (3.7) so as to be able to describe its solution. First, we will employ Proposition 3.9 to write expected payments in terms of the allocation rule (and the marginal price-posting revenue curve). Second, we will employ the method of Lagrangian relaxation on the budget-balance constraint to move it into the objective. Intuitively, Lagrangian relaxation allows the

constraint to be violated but places a linear cost on violating the constraint. This cost is parameterized by the Lagrangian parameter λ , for high values of λ there is a high cost for violating the constraint (and a high benefit for slack in the constraint, i.e., the margin by which the constraint is satisfied), for low values of λ there is a low cost for violating the constraint. E.g., $\lambda = 0$ the optimization is the original problem without the budget-balance constraint; with $\lambda = \infty$ the optimization is entirely one of maximizing the slack in the constraint. In our case the slack in the constraint is the profit of the mechanism. Therefore, the $\lambda = \infty$ case is to maximize profit and the $\lambda = 0$ case is to maximize social surplus (without budget balance). Adjusting the Lagrangian parameter λ traces out the *Pareto frontier* between the two objectives of social surplus and profit (see Figure 3.9(a)). From this Pareto frontier we can see how to optimize social surplus subject to a constraint on profit (such as budget balance) or optimize profit subject to a constraint on social surplus. Notice that when the constraint that is Lagrangian relaxed is met with equality then it drops from the objective entirely and the objective value obtained is the optimal value of the original program.

In quantile space with payments expressed in terms of the allocation rule, the Lagrangian relaxation of our program is as follows.

$$\begin{aligned} \max_{\hat{\mathbf{y}}(\cdot)} \mathbf{E}_{\mathbf{q}} \left[\sum_i V_i(q_i) \hat{y}_i(\mathbf{q}) - c(\hat{\mathbf{y}}(\mathbf{q})) \right] & \quad (3.8) \\ + \lambda \mathbf{E}_{\mathbf{q}} \left[\sum_i P'(q_i) \hat{y}_i(\mathbf{v}) - c(\hat{\mathbf{y}}(\mathbf{q})) \right] & \\ \text{s.t. } \mathbf{y}(\cdot) \text{ is monotone.} & \end{aligned}$$

Simplifying the objective with the identity (3.5) of $P'(q) = \frac{d}{dq}(q \cdot V(q)) = V(q) - q \cdot V'(q)$, we have

$$\mathbf{E}_{\mathbf{q}} \left[\sum_i [(1 + \lambda) \cdot V_i(q_i) + \lambda q \cdot V'_i(q_i)] \cdot \hat{y}_i(\mathbf{q}) - (1 + \lambda) \cdot c(\mathbf{y}(\mathbf{q})) \right].$$

This is simply a (Lagrangian) virtual surplus optimization where agent i 's virtual value is

$$\phi_i^\lambda(q) = (1 + \lambda) \cdot V_i(q_i) + \lambda q \cdot V'_i(q_i). \quad (3.9)$$

and with (Lagrangian) cost $(1 + \lambda)c(\cdot)$, subject to monotonicity of each agent's the allocation rule.

If our original non-game-theoretic problem (without incentive and budget-balance constraints) is solvable, the same solution can be applied to solve this Lagrangian optimization. First, we can normalize the

objective by dividing by $(1 + \lambda)$, the result is a virtual surplus optimization with the same cost function as the original problem. Second, the budget-balance constrained optimization problem be effectively solved to an arbitrary degree of precision, e.g., by binary searching for the Lagrangian parameter λ for which solutions to the Lagrangian optimization are just barely budget balanced. The details of this search are described below.

3.5.2 Monotone Lagrangian Virtual Values

For any Lagrangian parameter λ , the optimal mechanism for the Lagrangian objective is the one that maximizes Lagrangian virtual surplus subject to monotonicity of each agent's the allocation rule. When the Lagrangian virtual value $\phi_i^\lambda(\cdot)$ is monotone non-increasing in q_i for each i the virtual surplus maximization mechanism for these Lagrangian virtual values and Lagrangian cost optimizes the Lagrangian objective in dominant strategy equilibrium (Corollary 3.6).

Lemma 3.23. *For a regular distribution (Definition 3.4 on page 64) given by inverse demand function $V(\cdot)$ and any non-negative Lagrangian parameter λ , the Lagrangian virtual value function $\phi^\lambda(q) = (1 + \lambda) \cdot V(q) + \lambda q \cdot V'(q)$ is monotonically decreasing.*

Proof. The Lagrangian virtual value function of equation (3.9) is a convex combination of the inverse demand curve $V(\cdot)$ and the marginal price-posting revenue curve $P'(q) = V(q) - q \cdot V'(q)$, i.e., virtual values for revenue. The inverse demand curve is strictly decreasing by definition (Definition 3.9) and the marginal price-posting revenue curve is non-increasing by the regularity assumption (Proposition 3.10). The convex combination of two monotone functions is monotone; if one of the functions is strictly monotone then so is any non-trivial convex combination of them. The lemma follows. \square

To optimize expected social surplus subject to budget balance we need to tune the Lagrangian parameter so that the budget-balance constraint is met with equality. So tuned, the mechanism's expected profit will be zero and the expected Lagrangian objective will be equal to the true objective (expected social surplus). Expected profit is, as described above, a monotone function of the Lagrangian parameter. When expected profit is continuous in the Lagrangian parameter λ , this tuning of λ is straightforward. Recall that for surplus maximization subject to budget balance, the slack in the Lagrangian constraint is equal to the expected profit.

Lemma 3.24. *For Lagrangian virtual value functions that are continuous in the Lagrangian parameter, the slack in the Lagrangian constraint for expected Lagrangian virtual surplus maximization is continuously non-decreasing in the Lagrangian parameter.*

Proof. The distribution of quantiles and a fixed Lagrangian parameter induce a distribution on profiles of Lagrangian virtual values. Continuity of Lagrangian virtual values with respect to the Lagrangian parameter implies that the joint density function on profiles of Lagrangian virtual values is continuous in the Lagrangian parameter. For any fixed profile of Lagrangian virtual values, Lagrangian virtual surplus maximization finds a (deterministic) pointwise optimal solution, the slack of this solution is also fixed and deterministic. As the distribution over these profiles is continuous in the Lagrangian parameter so is the expected slack. \square

Theorem 3.25. *For regular general-costs environments, an Lagrangian virtual values from equation (3.9), there exists a Lagrangian parameter for which the virtual surplus maximization mechanism has zero expected profit and with this parameter the mechanism maximizes expected social surplus subject to budget balance in dominant strategy equilibrium.*

Example 3.7. Consider two agents with uniformly distributed values and a non-excludable public project with cost one, i.e.,

$$c(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} = (1, 1), \\ 0 & \text{if } \mathbf{x} = (0, 0), \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

The Lagrangian virtual values in value space are $\phi(v) = (2\lambda + 1) \cdot v - \lambda$. The Lagrangian virtual surplus mechanism serves both agents when $(2\lambda + 1)(v_1 + v_2) - 2\lambda > 1 + \lambda$ (for allocation $\mathbf{x} = (1, 1)$, the left-hand side is the Lagrangian virtual surplus, the right-hand side is the Lagrangian cost), i.e., when

$$v_1 + v_2 \geq \frac{3\lambda + 1}{2\lambda + 1}. \quad (3.10)$$

For $\lambda = 0$ we serve if $v_1 + v_2 \geq 1$ (clearly this maximizes surplus) and for $\lambda = \infty$ we serve if $v_1 + v_2 \geq 3/2$ (this maximizes profit). In equation (3.10) we see that (for the uniform distribution), for any Lagrangian parameter λ , the form of the optimal mechanism is a threshold rule on the sum of the agent values. It is easy then to solve for the threshold satisfies the budget-balance constraint with equality. The optimal threshold is $5/4$,

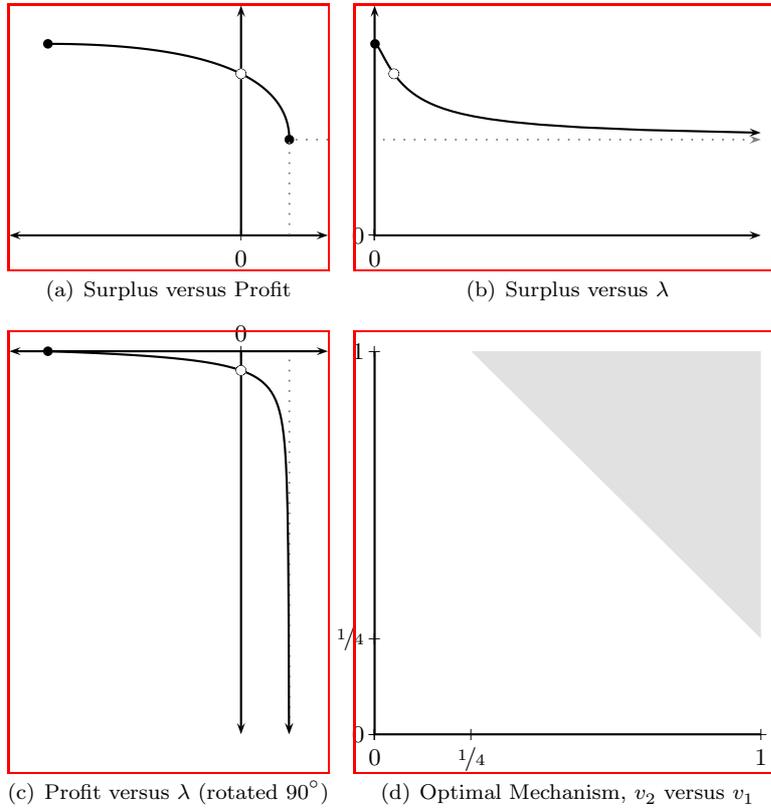


Figure 3.9. Depiction of the Pareto frontier for surplus (vertical axis) and profit (horizontal axis). On the Pareto frontier, the surplus maximizing point is profit minimizing (with negative profit) and the profit maximizing point is surplus minimizing. The surplus optimal point subject to budget balance is denoted by “o”. The surplus and profit versus the Lagrangian parameter λ are depicted along with their asymptote (grey, dotted line) as $\lambda \rightarrow \infty$. The profit versus λ plot has been rotated 90° clockwise so as to line up with the profit axis of the Pareto frontier plot. The optimal mechanism is depicted by plotting v_2 versus v_1 where the region of valuation profiles for which the project is provided is shaded.

the optimal Lagrangian parameter is $\lambda^* = 1/2$, and the social surplus is $9/64 \approx 0.14$. This example is depicted in Figure 3.9.

3.5.3 Non-monotone Lagrangian Virtual Values and Partial Ironing

When the Lagrangian virtual value functions are non-monotone then the ironing procedure (Definition 3.11) can be applied and the virtual surplus maximization mechanism with the resulting ironed Lagrangian ironed virtual values is optimal for the Lagrangian objective. After ironing, however, the slack in the Lagrangian constraint, e.g., expected profit, is generally discontinuous in the Lagrangian parameter. In such case there is a point λ^* such that for $\lambda < \lambda^*$ the expected profit of any solution is negative and for $\lambda > \lambda^*$ the expected profit of any solution is positive. At $\lambda = \lambda^*$ there are multiple solutions to the Lagrangian objective. These solutions vary in the contribution to the relaxed objective from the original objective and from the slack in the Lagrangian constraint (which is part of the relaxed objective); the expected profits of these solution span the gap between the negative profit solutions and the positive profit solutions. In particular, a convex combination of the supremum (with respect to expected profit) of solutions with negative profit with infimum of solutions with positive profit will optimize ironed Lagrangian virtual surplus and meet the budget-balance constraint with equality.

This convex combination of mechanisms can be interpreted as an ironed virtual surplus optimizer with a non-standard tie-breaking rule. Consider virtual value function $\phi(\cdot)$ and ironed virtual value function $\bar{\phi}(\cdot)$ constructed for $\phi(\cdot)$ for distribution F via the ironing procedure (Definition 3.11). By the definition of the ironing procedure, the cumulative ironed virtual value function $\bar{\Phi}(\cdot)$ is the smallest concave upper bound on the cumulative virtual value function $\Phi(\cdot)$. Define $[a, b]$ to be an *ironed interval* if $\bar{\Phi}(q) > \Phi(q)$ for $q \in (a, b)$ and $\bar{\Phi}(q) = \Phi(q)$ for $q \in \{a, b\}$. The ironing procedure gives ironed virtual values that are equal to virtual values in expectation under the assumption that all quantiles within the same ironed interval have the same allocation probability (Theorem 3.12). Such an outcome is always obtained for outcomes selected solely based on ironed virtual values (ignoring actual values).

For Lagrangian ironed virtual value functions, it may be that two adjacent ironed intervals have the same ironed virtual value. In such a case outcomes selected solely based on ironed virtual values will produce the same allocation probability for quantiles in the union of the adjacent ironed intervals. Notice that the equality of ironed virtual values across adjacent ironed intervals is sensitive to small changes in the Lagrangian

parameter. With a slightly higher Lagrangian parameter these ironed intervals will be strictly merged; with a slightly lower Lagrangian parameter these ironed intervals will be strictly distinct. Thus, infimum mechanism is the one that tie-breaks to merge adjacent ironed intervals with the same ironed virtual value and the supremum mechanism is the one that tie-breaks to keep adjacent ironed intervals distinct. We refer to the mixing over two tie-breaking rule for maximizing ironed virtual surplus as *partial ironing*.

Theorem 3.26. *For general-cost environments, and Lagrangian virtual values from equation (3.9), there exists a Lagrangian parameter and partial-ironing parameter for which the partially-ironed Lagrangian virtual surplus maximization mechanism optimizes social surplus subject to budget balance in dominant strategy equilibrium.*

Exercises

- 3.1 In computer networks such as the Internet it is often not possible to use monetary payments to ensure the allocation of resources to those who value them the most. Computational payments, e.g., in the form of “proofs of work”, however, are often possible. One important difference between monetary payments and computational payments is that computational payments can be used to align incentives but do not transfer utility from the agents to the seller. I.e., the seller has no direct value from an agent performing a proof-of-work computation. Define the *residual surplus* as the social surplus less the payments, i.e., $\sum_i (v_i \cdot x_i - p_i) - c(\mathbf{x})$. (For more details, see the discussion of non-monetary payments in Chapter 1.)

Describe the mechanism that maximizes residual surplus when the distribution on agents’ values satisfy the *monotone hazard rate* assumption, i.e., $f(v)/1-F(v)$ is monotone non-decreasing. Your description should first include a description in terms of virtual values and then you should interpret the implication of the monotone hazard rate assumption to give a simple description of the optimal mechanism. In particular, consider monotone hazard rate distributions in the following environments:

- (a) a single-item auction with i.i.d. values,
- (b) a single-item auction with non-identical values, and

- (c) an environment with general costs specified by $c(\cdot)$ and non-identical values.
- 3.2 Give a mechanism with first-price payment semantics that implements the social surplus maximizing outcome in equilibrium for any single-dimensional agent environment. Hint: Your mechanism may be parameterized by the distribution.
- 3.3 Derive equation (3.3),

$$\mathbf{E}_{v \sim F}[p(v)] = \mathbf{E}_{v \sim F}[\phi(v) \cdot x(v)] \quad (3.3)$$

by taking expectation of the payment identity (3.1),

$$p(v) = v \cdot x(v) - \int_0^v x(z) \, dz, \quad (3.1)$$

for $v \sim F$ and simplifying.

- 3.4 Consider the non-downward closed environment of *public projects*: either every agent can be served or none of them. I.e., the cost structure satisfies:

$$c(\mathbf{x}) = \begin{cases} 0 & \text{if } \sum_i x_i = 0, \\ 0 & \text{if } \sum_i x_i = n, \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

- (a) Describe the revenue-optimal mechanism for general distributions.
- (b) Describe the revenue-optimal mechanism when agents' values are i.i.d. from $U[0, 1]$.
- (c) Give an asymptotic, in terms of the number n of agents, analysis of the expected revenue of the revenue-optimal public project mechanism when agents' values are i.i.d. from $U[0, 1]$.
- 3.5 Consider a two unit auction to four agents and a virtual value function that is strictly monotone except for an interval $[a, b]$ where it is a positive constant (e.g., Figure 3.2 on 69). Suppose the valuation profile \mathbf{v} satisfies $v_1 > b$, $v_2, v_3 \in [a, b]$, and $v_4 < a$. Calculate the probability of winning and expected payments of all agents (in terms of a and b).
- 3.6 Consider profit maximization with values drawn from a discrete distribution. Derive virtual values for revenue for discrete single-dimensional type spaces for agents with linear utility. Assume that $\mathcal{T} = \{v^0, \dots, v^N\}$ with the probability that an agent's value is $v \in \mathcal{T}$ given by probability mass function $f(v)$. Assume $v^0 = 0$. Note: You must first solve Exercise 2.2 to characterize BNE equilibrium.

- (a) Derive virtual values for the special case where the values are uniform, i.e., $v^j = j$ for all j .
- (b) Derive virtual values for the special case where the probabilities are uniform, i.e., $f(v^j) = 1/N$ for all j .
- (c) Give virtual values for the general case.

(Hint: You should end up with a very similar formulation to that for continuous type spaces.)

- 3.7 The text has focused on *forward auctions* where the auctioneer is a seller and the agents are buyers. The same theory can be applied to *reverse auctions* (or *procurement*) where the auctioneer is a buyer and the agents are sellers. It is possible to consider reverse auctions within the framework described in this chapter where an agent's value for service is negative, i.e., in order to provide the service they must pay a cost. It is more intuitive, however, to think in terms of positive costs instead of negative values.
- (a) Derive a notion analogous to revenue curves for an agent (as a seller) with private cost drawn from a distribution F .
 - (b) Derive a notion of *virtual cost functions* analogous to virtual value functions.
 - (c) Suppose the auctioneer has a value of v for procuring a service from one of several sellers with costs distributed i.i.d. and uniformly on $[0, 1]$. Describe the auction that optimizes the seller's profit (value for procurement less payments made to agents).
- 3.8 Consider a profit-maximizing broker mediating the exchange between a buyer and a seller. The broker's profit is the difference between payment made by the buyer and payment made to the seller. Use the derivation of virtual values for revenue (from Section 3.3.4) and virtual costs (from Exercise 3.7).
- (a) Derive the optimal exchange mechanism for regular distributions for the buyer and seller.
 - (b) Solve for the optimal exchange mechanism in the special case where the buyer's and seller's values are both distributed uniformly on $[0, 1]$.
- 3.9 In Example 3.7 it was shown that for two agents with uniform values on interval $[0, 1]$ and a cost of one for serving both of them together, the surplus maximizing mechanism with a balanced budget in expectation serves the agents when the sum of their values is at least $4/3$. There is a natural dominant strategy "second-price" implementation of this mechanism; instead give a "first-price" (a.k.a.,

pay-your-bid) implementation. Your mechanism should solicit bids, decide based on the bids whether to serve the agents, and charge each agent her bid if they are served.

Chapter Notes

The surplus-optimal Vickrey-Clarke-Groves (VCG) mechanism is credited to Vickrey (1961), Clarke (1971), and Groves (1973).

The characterization of revenue-optimal single-item auctions as virtual value maximizers (for regular distributions) and ironed virtual value maximizers (for irregular distributions) was derived by Roger Myerson (1981). Its generalization to single-dimensional agent environments is an obvious extension. The relationship between revenue-optimal auctions, price-posting revenue curves, and marginal price-posting revenue (equivalent to virtual values) is due to Bulow and Roberts (1989). The revenue-linearity-based approach is from Alaei et al. (2013).

Myerson and Satterthwaite (1983) characterized mechanisms that maximize social surplus subject to budget balance via Lagrangian relaxation of the budget-balance constraint. The discussion of partial ironing for Lagrangian virtual surplus maximizers given here is from Devanur et al. (2013). This partial ironing suggests that the optimal mechanism is not deterministic, the problem of finding a deterministic mechanism to maximize social surplus subject to budget balance is much more complex as the space of deterministic mechanisms is not convex (Diakonikolas et al., 2012).

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.

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* $\hat{v} = (\hat{v}_1, \dots, \hat{v}_n)$ 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 opti-

Technical 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. Recall that a regular distribution can be equivalently specified by its distribution function or its revenue curve. The *equal revenue distribution* has constant revenue curve, $R^{\text{EQR}}(q) = 1$, and therefore any price on $[1, \infty)$ is optimal. A sufficient perturbation to make unique monopoly price $\hat{v}^* = 1$ is given by revenue curve $R^{\text{EQR}}(q) = 1 - \epsilon(1 - q)$.

In the previous two chapters, with the characterization of Bayes-Nash equilibrium (Theorem 2.2) and the characterization of profit-optimal mechanisms (Corollary 3.15), 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 be 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 re-derive Theorem 2.2 and Corollary 3.15 (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 $[v, v + \epsilon]$.

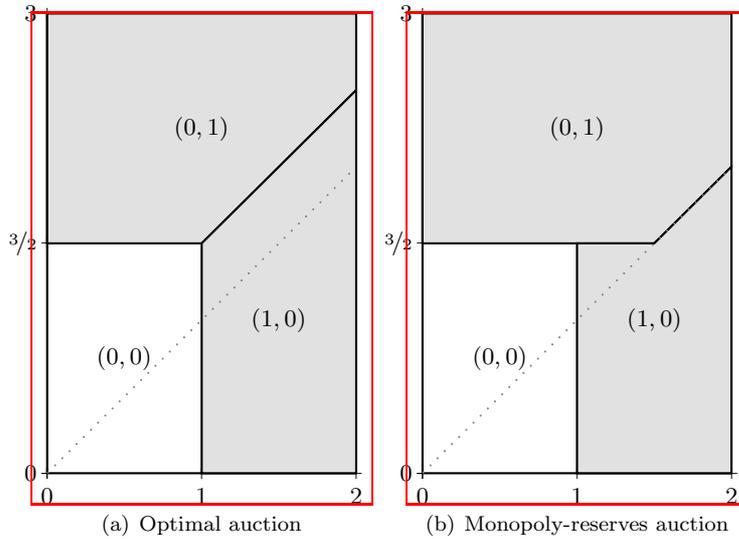


Figure 4.1. In Example 4.1 Agent 1 has value $v_1 \sim U[0, 2]$; agent 2 has value $v_2 \sim U[0, 3]$. In the space of valuation profiles $\mathbf{v} \in [0, 2] \times [0, 3]$, with agent 1's value on the horizontal axis and agent 2's value on the vertical axis, the allocation $\mathbf{x} = (x_1, x_2)$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

mal 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 3.4, 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 4.1. *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.*

The proof of Theorem 4.1 is enabled by the following three properties of regular distributions and virtual value functions. First, Corollary 3.21 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 4.2. *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). \square

Lemma 4.3. *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 4.1) 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. \square

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 4.1. 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, $\text{REF} \geq \text{APX}$; our goal is to give an approximate inequality in the opposite direction by showing that $2 \text{APX} \geq \text{REF}$. 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 3.21, $\text{REF} = \mathbf{E}[\phi_I(v_I)]$ and $\text{APX} = \mathbf{E}[\phi_J(v_J)]$.

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

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

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

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

The inequality in the above calculation follows from Lemma 4.2 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:

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

The first inequality in the above calculation follow from values upper bounding virtual values (Lemma 4.3). 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 4.1 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 $F^{\text{EQR}}(z) = 1 - 1/z$ and density function $f^{\text{EQR}}(z) = 1/z^2$ 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 $\hat{v} \cdot (1 - F^{\text{EQR}}(\hat{v})) = 1$. As the price-posting revenue curve is the constant function $P^{\text{EQR}}(\hat{q}) = 1$, 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 4.4. *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 $\hat{v}_2^* = 1$ (Technical Note on page 102). Thus the monopoly prices are $\hat{v}^* = (1, 1)$ 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 4.1 relied on regularity crucially when it invoked Corollary 3.21 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

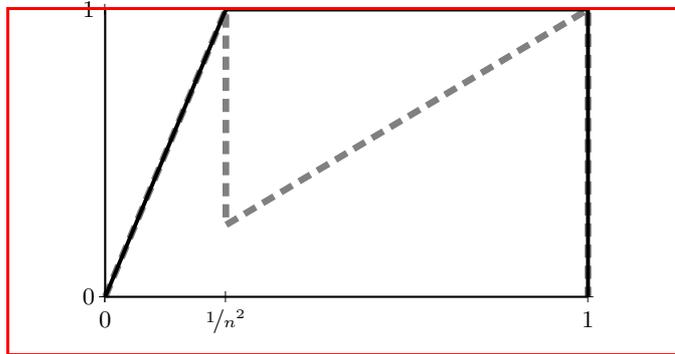


Figure 4.2. The revenue curve (thin, solid, black) and price-posting revenue curve (gray, thick, dashed) for the discrete two-point equal revenue distribution from the proof of Proposition 4.5 with $h = 2$. As usual for revenue curves, the horizontal axis is quantile.

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 4.5), and that this lower bound is tight as the monopoly-reserves auction for (non-identical) irregular distributions is, trivially, an n approximation (Proposition 4.6).

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 4.5. *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 (Technical Note on page 102). With a monopoly reserve of $\hat{v}^* = 1$ 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 follow-

ing limit result which implies the proposition.

$$\text{APX} = \lim_{h \rightarrow \infty} \text{APX}(h) = 1, \text{ and} \quad (4.1)$$

$$\text{REF} = \lim_{h \rightarrow \infty} \text{REF}(h) = n. \quad (4.2)$$

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:

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

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,

$$\lim_{h \rightarrow \infty} h^k \cdot \Pr[\text{exactly } k \text{ are high valued}] = \binom{n}{k}, \text{ and} \quad (4.3)$$

$$\lim_{h \rightarrow \infty} h^k \cdot \Pr[\text{at least } k \text{ are high valued}] = \binom{n}{k}. \quad (4.4)$$

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 $\text{REF} = \lim_{h \rightarrow \infty} \text{REF}(h)$ 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. $\text{REF} = 0 + \binom{n}{1} = n$.

We can similarly calculate $\text{APX} = \lim_{h \rightarrow \infty} \text{APX}(h)$ 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, $\text{APX} = 1$. \square

Proposition 4.6. *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, $\text{UB} \geq \text{REF}$ 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 $\mathbf{R}^* = (R_0^*, \dots, R_n^*)$ the profile of monopoly revenues. The revenue of the optimal n -unit auction is:

$$\text{UB} = \sum_i R_i^*.$$

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^* with the highest monopoly revenue and offers this agent her monopoly price $\hat{v}_{i^*}^*$. LB obtains revenue

$$\text{LB} = \max_i R_i^*.$$

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

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

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 4.7. *For any product distribution on prize values $\mathbf{F} = F_1 \times \dots \times F_n$, 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 4.7 is a *prophet inequality*: it suggests 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 4.7. 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 $\hat{q}_i = 1 - F_i(\hat{v}) = \Pr[v_i \geq \hat{v}]$ 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 \mathbf{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 $i \in \{0, \dots, n\}$.

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

The prophet's expected payoff is

$$\begin{aligned} \text{REF} &= \mathbf{E}[\max_i v_i] = \hat{v} + \mathbf{E}[\max_i (v_i - \hat{v})] \\ &\leq \hat{v} + \mathbf{E}[\max_i (v_i - \hat{v})^+] \\ &\leq \hat{v} + \sum_i \mathbf{E}[(v_i - \hat{v})^+]. \end{aligned} \quad (4.5)$$

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

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

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 $\text{APX}_1 = (1 - \chi) \cdot \hat{v}$. 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:

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

The second line follows because $\Pr[\mathcal{E}_i] = \prod_{j \neq i} (1 - \hat{q}_j) \geq \prod_j (1 - \hat{q}_j) = \chi$ 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:

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

(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} , $\text{APX} \geq \text{REF} / 2$.

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. \square

The prophet inequality is tight in the sense that a better approximation bound cannot generally be 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 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 use 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 4.8. *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* $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$ 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 4.9. *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) $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$. 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 3.12. 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 4.7), 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 func-

tions. 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 a 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 4.10. *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 $\hat{\mathbf{q}} = (\hat{q}_1, \dots, \hat{q}_n)$ 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 $\hat{\mathbf{q}}$ 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 fraction of the supply allocated to it. Both of these optimization problem are given by the following convex program:

$$\begin{aligned} \max_{\hat{\mathbf{q}}} \quad & \sum_i R_i(\hat{q}_i) \\ \text{s.t.} \quad & \sum_i \hat{q}_i \leq 1. \end{aligned} \tag{4.7}$$

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 4.11. *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 $N = \{1, \dots, n\}$ and a distribution over subsets given by \mathcal{D} . Let \hat{q}_i be the ex ante¹ 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{v} the *maximum-weight-element* set function $g^{\text{MWE}}(S) = \max_{i \in S} \hat{v}_i$ has a correlation gap of $e/e-1$.

Lemma 4.12. *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 $g^{\text{MWE}}(S) = \max_{i \in S} \hat{v}_i$. 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 $\mathbf{E}_{S \sim \mathcal{D}}[g^{\text{MWE}}(S)]$ 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 $\hat{\mathbf{q}} = (\hat{q}_1, \dots, \hat{q}_n)$

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 $\mathbf{E}_{S \sim \mathcal{D}^I} [g^{\text{MWE}}(S)]$ 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 $g^{\text{MWE}}(S) = \max_{i \in S} \hat{v}_i$, the correlated value simplifies to $\mathbf{E}_{S \sim \mathcal{D}} [g^{\text{MWE}}(S)] = \sum_i \hat{q}_i \hat{v}_i$. Scaling the weights $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$ 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 $\hat{\mathbf{v}}$, 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, $g^{\text{OOM}}(S) = 1$ if $|S| \geq 1$ and otherwise $g^{\text{OOM}}(S) = 0$.

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

$$\sum_i c_i \hat{q}_i \hat{v}_i = \sum_i c_i \cdot (\hat{q}_i + \delta_i) \geq \sum_i c_i \hat{q}_i. \quad (4.8)$$

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 $S \sim \mathcal{D}^I$,

$$\begin{aligned} \mathbf{E} [g^{\text{OOM}}(S)] &= \mathbf{Pr}[|S| \geq 1] = 1 - \mathbf{Pr}[|S| = 0] = 1 - \prod_i (1 - \hat{q}_i) \\ &\geq 1 - (1 - 1/n)^n \geq 1 - 1/e, \end{aligned}$$

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 $\hat{\mathbf{q}}$ 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., $g^{\text{MWE}}(S) = \max_{i \in S} \hat{v}_i$. 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 $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$ 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 $g^{\text{MWE}}(S)$ on $S \sim \mathcal{D}^I$.

For regular distributions, the translation from the solution to the optimal ex ante mechanism which is given by $\hat{\mathbf{q}}$ 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 3.22, 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 $\hat{\mathbf{q}}$ 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 4.13. *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 4.10 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 4.12). The prices correspond to a uniform virtual pricing by the characterization of the optimal ex ante mechanism (Proposition 4.11). \square

The construction and analysis of Theorem 4.13 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, \dots, 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 4.14. *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} . \square

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 3.8). For anonymous posted pricing, Theorem 4.13 implies a $e/e-1 \approx 1.58$ approximation for regular distributions and Theorem 4.9 implies a two approximation for irregular distributions. Notice that while Theorem 4.13 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 4.15. *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 4.13 shows that this anonymous posted pricing is a $e/e-1$ approximation. For tightness, see Exercise 4.12. \square

Corollary 4.16. *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 4.9 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 4.14, 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

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

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 $\text{REF} = 2n - 1$ and $\text{APX} = n$; 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 4.5; 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, $\text{REF} = 2n - 1$.

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, $\text{APX} = n$. \square

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 4.17. *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. \square

Theorem 4.18. *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 $\text{REF} = \sum_i \hat{v}_i \hat{q}_i$ 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. \square

We now give a tight inapproximation bound for anonymous reserves

and pricings with irregular distributions. Recall the proof of Proposition 4.6 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 4.19. *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 4.6 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 4.27. For irregular distributions the tight approximation bound for single-unit environments of Proposition 4.5 and Proposition 4.6 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 4.20. *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 approximation³ is $(1 - 1/\sqrt{2\pi k})^{-1}$ (see Exercise 4.17).

Proposition 4.21. *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 4.21, 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 4.20, above) implies the same bound by an anonymous pricing (and consequently for the $(k+1)$ st price auction with an anonymous

³ Stirling's approximation is $k! = (k/e)^k \sqrt{2\pi k}$. This approximation is obtained by approximating the natural logarithm as $\ln(k!) = \ln(1) + \dots + \ln(k)$ 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 4.22, below). For irregular, non-identical distributions the n -approximation bound of Theorem 4.19 for single-item environments generalizes and is tight.

Lemma 4.22. *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 k th 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 $\{1, \dots, k\}$. The value distribution that gives this bound is the one where F_i is a pointmass at $1/i$ for agents $i \in \{1, \dots, k\}$ 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 $i \in \{1, \dots, k\}$, 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 k th harmonic number $\sum_{i=1}^k 1/i = \mathcal{H}_k$. The k th harmonic number can be approximated by the integral $\int_1^k 1/i \, di$ and satisfies $\ln k - 1 \leq \mathcal{H}_k \leq \ln k$. \square

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) $\mathbf{x} \leftarrow \mathbf{0}$ (the null assignment).
- (iii) For each agent i (in sorted order),
 - if $(1, \mathbf{x}_{-i})$ is feasible, $x_i \leftarrow 1$.
 - (I.e., serve i if i can be served alongside previously served agents.)
- (iv) Output \mathbf{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 3.8). 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 4.23. *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{v} is:

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

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

- (iii) set prices \mathbf{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 4.24. *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 \mathbf{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 \mathbf{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 start adding elements $i \in J$. As J is not feasible, it must fail to add at

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.

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 \mathbf{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, i.e., a *single-dimensional constrained matching market*.
- 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 4.25. *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 4.24. The “if” direction is as follows. Let r be the *rank* of the matroid. Let $I = \{i_1, \dots, i_r\}$ be the set of agents selected in the surplus maximizing assignment, and let $J = \{j_1, \dots, j_r\}$ 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 $I_k = \{i_1, \dots, i_k\}$ and let $J_{k-1} = \{j_1, \dots, j_{k-1}\}$. 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. \square

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 4.26. *For matching agents $N = \{1, \dots, n\}$ to items $K = \{1, \dots, k\}$ 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. \square

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 4.27. *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 4.27.

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 4.28 and Theorem 4.29, below. We will prove Lemma 4.28 and leave the formal proofs of Theorem 4.27 and Theorem 4.29 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 $j = \operatorname{argmax}_{\{k: I \cup \{k\} \in \mathcal{I}\}} v_k$.

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

Lemma 4.28. *In matroid environments, the surplus maximization mechanism on valuation profile \mathbf{v} has the critical values $\hat{\mathbf{v}}$ 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 $\epsilon \rightarrow 0$ 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 $\operatorname{OPT}(0, \mathbf{v}_{-i})$ and $\operatorname{OPT}_{-i}(\infty, \mathbf{v}_{-i})$ are optimal surplus from agents other than i with i is not served and served, respectively.

$$\hat{v}_i = \operatorname{OPT}(0, \mathbf{v}_{-i}) - \operatorname{OPT}_{-i}(\infty, \mathbf{v}_{-i}). \quad (4.9)$$

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,

$$\hat{v}_i = v_j. \quad (4.10)$$

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 $\text{OPT}(0, \mathbf{v}_{-i})$. Let $J \cup \{j\}$ be the agents served in $\text{OPT}(0, \mathbf{v}_{-i})$ and let $I \cup \{i\}$ be the agents served in $\text{OPT}(\infty, \mathbf{v}_{-i})$. We can deduce from equations (4.9) and (4.10) that,

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

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 $\text{OPT}(0, \mathbf{v}_{-i})$, j must be the best, i.e., highest valued, replacement. \square

Theorem 4.29. *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 4.30. *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 $\text{rank}(S)$ denote this matroid rank function.

A profile of ex ante probabilities $\hat{\mathbf{q}} = (\hat{q}_1, \dots, \hat{q}_n)$ 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 $S \subset N$,

$$\sum_{i \in S} \hat{q}_i \leq \text{rank}(S). \quad (4.11)$$

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

$$\begin{aligned} \max_{\hat{\mathbf{q}}} \quad & \sum_i R(\hat{q}_i) \\ \text{s.t.} \quad & \sum_{i \in S} \hat{q}_i \leq \text{rank}(S), \quad \forall S \subset N. \end{aligned} \quad (4.12)$$

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 $\text{rank}_{\hat{v}}(\cdot)$ for weights \hat{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{q} is ex ante feasible, there exists a correlated distribution \mathcal{D} over feasible sets which induces ex ante probabilities \hat{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 \mathbf{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 $\text{rank}_{\hat{v}}(S)$.

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 ℓ evenly sized pieces. Consider the ℓ -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., $1/\ell$ 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 $g : 2^N \rightarrow \mathbb{R}$. 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 4.32. *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 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 4.12, Section 4.3).

Theorem 4.33. *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 4.34. *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 4.23 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 agent's 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 4.28). 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 4.35. *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 4.12). Thus, the optimal revenue is at least $\text{REF} \geq h^2 \cdot e^{-1/e}$.

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 $\text{APX} \leq 2h$. \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 4.2. The expected value the equal revenue distribution on $[1, h]$ is $1 + \ln h$ (for the unbounded equal revenue distribution it is infinite). This can be calculated from the formula $\mathbf{E}[v] = \int_0^\infty (1 - F^{\text{EQR}}(z)) dz$ with $F^{\text{EQR}}(z) = 1 - 1/z$. On the other hand, the monopoly revenue for the equal revenue distribution is one. Therefore, the optimal social sur-

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) $h = \Theta(\log n / \log \log n)$.

plus 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 4.36. *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 $2, \dots, n + 1$ (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 4.2, 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 $F^{\text{EXP}}(z) = 1 - e^{-z}$, $f^{\text{EXP}}(z) = e^{-z}$, and $h^{\text{EXP}}(z) = 1$ 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

$$\phi(v) = v - 1/h(v). \quad (4.13)$$

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 4.3), 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 4.2, above, which shows that for non-monotone-hazard-rate distributions, the ratio of surplus to revenue can be unbounded.

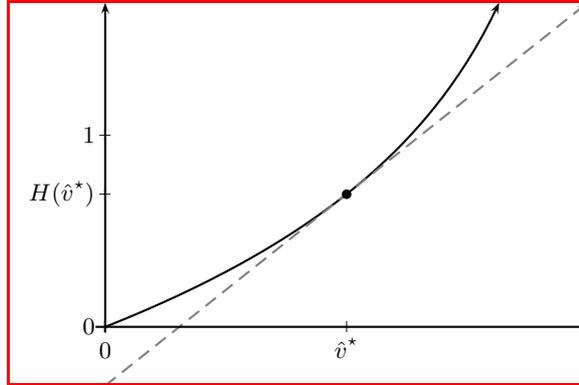


Figure 4.4. The cumulative hazard rate function (solid, black) for the uniform distribution is $H(v) = -\ln(1 - v)$ and it is lower bounded by its tangent (dashed, gray) at $\hat{v}^* = 1/2$.

Example 4.3. The expected value of the exponential distribution (with rate one) is one. This can be calculated from the formula $\mathbf{E}[v] = \int_0^\infty (1 - F^{\text{EXP}}(z)) dz$ with $F^{\text{EXP}}(z) = 1 - e^{-z}$. Since the exponential distribution has hazard rate $h^{\text{EXP}}(z) = 1$, the virtual valuation formula for the exponential distribution is $\phi^{\text{EXP}}(v) = v - 1$. The monopoly price is one. The probability that the agent accepts the monopoly price is $1 - F^{\text{EXP}}(1) = 1/e$ so its expected revenue is $1/e$. The ratio of the expected surplus to expected revenue is e .

Theorem 4.37. 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 4.38. For any monotone-hazard-rate distribution its expected value is at most e times more than the expected monopoly revenue.

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

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

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 $\frac{d}{dv} \ln(1 - F(v)) = \frac{-f(v)}{1 - F(v)}$ and $\frac{d}{dv} \ln(e^{-H(v)}) = -h(v)$.

of $v \sim \mathbf{F}$ is $\int_0^\infty (1 - F(z)) dz$. 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 $\hat{v}^* = \phi^{-1}(0)$ 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}^* . See Figure 4.4. I.e.,

$$\begin{aligned} H(v) &\geq H(\hat{v}^*) + h(\hat{v}^*)(v - \hat{v}^*) \\ &= H(\hat{v}^*) + \frac{v - \hat{v}^*}{\hat{v}^*}. \end{aligned} \quad (4.15)$$

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

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

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

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

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

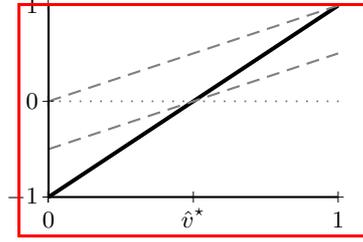


Figure 4.5. The virtual value for the uniform distribution is depicted. For $v \geq \hat{v}^*$ the virtual value $\phi(v)$ (solid, black) is sandwiched between the value v (dashed, gray) and value less the monopoly price $v - \hat{v}^*$ (dashed, gray).

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

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

Theorem 4.40. *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.

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

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

$$\begin{aligned} 2 \cdot \text{APX} &\geq \mathbf{E}[\text{APX's winners' virtual values} + \text{reserve prices}] \\ &= \mathbf{E} \left[\sum_i (\phi_i(v_i) + \hat{v}_i^*) \cdot x_i(v) \right] \\ &\geq \mathbf{E} \left[\sum_i v_i \cdot x_i(v) \right] = \mathbf{E}[\text{APX's surplus}] \\ &\geq \mathbf{E}[\text{REF's surplus}] \geq \mathbf{E}[\text{REF's revenue}] = \text{REF}. \end{aligned}$$

The second inequality follows from Lemma 4.39. 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 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 3.4 which are depicted in Figure 4.1.
- 4.3 Finish the proof of Lemma 4.3 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 $[0, \hat{q}^*]$, 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 4.1 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 $F^{\text{EQR}}(z) = 1 - 1/z$). 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 -nearby 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 -nearby 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 4.13 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 4.18 by adapting the analysis of the prophet inequality (Theorem 4.7) 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 4.19). 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 4.5 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 $i \in \{1, \dots, n\}$) 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 $i \in \{1, \dots, n\}$) 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 4.6 and Proposition 4.5 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 4.21, 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 $\hat{v} = (\hat{v}_1, \dots, \hat{v}_n)$ the value of $g^{\text{kMWE}}(S)$ 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 $g^{\text{kCC}}(S) = \min(k, |S|)$ (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 4.27. Hint: This result can be proved using Lemma 4.28 and Theorem 4.29 and a similar argument to the proof of Theorem 4.1.
- 4.20 A mechanism \mathcal{M} is *revenue monotone* if for all pairs of valuation profiles \mathbf{v} and \mathbf{v}^\dagger such that for all i , $v_i \geq v_i^\dagger$, the revenue of \mathcal{M} on \mathbf{v} is at least its revenue on \mathbf{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 4.32.

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{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{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 4.36 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 $[0, h]$:

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

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^* :

- (a) $(\mathbf{x}^\dagger, \mathbf{p}^\dagger) \leftarrow \text{SM}(\mathbf{v})$,
- (b) $x_i = \begin{cases} x_i^\dagger & \text{if } v_i \geq \hat{v}_i^* \\ 0 & \text{otherwise, and} \end{cases}$
- (c) $p_i = \max(\hat{v}_i^*, p_i^\dagger)$.

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 4.37 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 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. (2010b) 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. (2010b) 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. (2010b) 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. ??). 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.

5

Prior-independent Approximation

In the last two chapters we discussed mechanism that performed well for a given Bayesian prior distribution. Assuming the existence of such a Bayesian prior is natural when deriving mechanisms for games of incomplete information as the Bayes-Nash equilibrium concept requires a prior distribution that is common knowledge. In this chapter we will relax the assumption the designer has knowledge of the prior distribution and is able to tune the parameters of her mechanism with it. The goal of *prior-independent* mechanism design is to identify a single mechanism that has good performance for all distributions in a large family of relevant distributions, e.g., the family of i.i.d. regular distributions.

As is evident from our analysis of Bayesian optimal auctions, e.g., for profit maximization, for any auction that one might consider good for one prior, there is another prior for which another auction performs strictly better. This consequence is obvious because optimal auctions for distinct distributions are generally distinct. So, while no single auction is optimal for all value distributions, there may be a single auction that is approximately optimal across a wide range of distributions.

In this chapter we will take two approaches to prior-independent mechanism design. The first approach considers “resource” augmentation. We will show that in some environments the (prior-independent) surplus maximization mechanism with increasing competition, e.g., by recruiting more agents, earns more revenue than the revenue-optimal mechanism without the increased competition. The second approach is to design mechanisms that do a little market analysis on the fly. Via this second approach, we will show that for a large class of environments there is a single mechanism that approximates the revenue of the optimal mechanism.

5.1 Motivation

Since prior-independence is not without loss it is important to consider its motivation; however, before doing so recall the original justification for the common prior assumption (see Section 2.3). Auctions and mechanisms are games of incomplete information and in such games, in order to argue about strategic choice, we needed to formalize how players deal with uncertainty. We did this by assuming a Bayesian prior. In a *Stackelberg game*, instead of moving simultaneously, players make actions in a prespecified order. We can view mechanism design as a two stage Stackelberg game where the designer is a player who moves first and the agents are players who (simultaneously) move second. To analyze the Bayes-Nash equilibrium in such a Stackelberg game, the designer bases her strategy on the common prior. Without such prior knowledge, the problem of predicting the designer's strategy is ill posed. Thus, in so far as the theory of mechanism design should describe (or predict) the outcome of a game, within the standard equilibrium concept for games of incomplete information, a prior assumption is necessary.

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

It is important to consider the incentives of information acquisition within the mechanism design problem; even if the designer has knowledge of a prior distribution, it may be problematic to employ this knowledge in a mechanism. Suppose the designer obtained her prior knowledge from previous market experience. The problem with designing the mechanism with this knowledge is that the earlier agents may strategize so that information about their preferences is not exploited by the designer later. For example, a monopolist who cannot commit not to lower the sale price in the future cannot sell at a high price now (see Exercise 5.1).

It is similarly important to consider the losses due to inaccuracies in in-

formation acquisition within the mechanism design problem. To learn the prior a designer could perform a *market analysis*, for example, by hiring a marketing firm to survey the market and provide distributional estimates of agent preferences. This mode of operation is quite reasonable in large markets. However, in large markets mechanism design is not such an interesting topic; each agent will have little impact on the others and therefore the designer may as well stick to posted-pricing mechanisms. Indeed, for commodity markets posted prices are standard in practice. Mechanisms, on the other hand, are most interesting in small, a.k.a., *thin*, markets. Contrast the large market for automobiles to the thin market for spacecrafts. There may be five organizations in the world in the market for spacecrafts; how would a designer optimize a mechanism for selling them? First, even if the agents' values do come from a distribution, the only way to sample from the distribution is to interview the agents themselves. Second, even if we did interview the agents, we could obtain at most five data points. This sample size is hardly enough for statistical approaches to be able to estimate the distribution of agent values. A motivating question this perspective raises, and one that is closely tied to prior-independent mechanism design, is: How many samples from a distribution are sufficient for the design of an approximately optimal mechanism?

There are other reasons to consider prior-independent mechanism design besides the questionable origin of prior information. Perhaps the most striking of which is the frequent inability of a designer to redesign a new mechanism for each scenario in which she wishes to run a mechanism. This is not just a concern; in many settings, it is a principle. Consider the standard Internet routing protocol IP. This is the protocol responsible for sending emails, browsing web pages, streaming video, etc. Notice that the workloads for each of these tasks is quite different. Emails are small and can be delivered with several minutes delay without issue. Web pages are small, but must be delivered immediately. Comparably, video streaming permits high latency but requires continuous bandwidth. It would be impractical to install new protocols in Internet routers each time a new network usage pattern arises. Instead, a protocol for computer networks, such as IP, should work pretty well in any setting, even ones well beyond the imaginations of the original designers.

5.2 “Resource” Augmentation

In this section we describe a classical result from auction theory that shows that a little more competition in a surplus maximizing mechanism revenue dominates the revenue maximizing mechanism without the increased competition. From an economic point of view this result questions the *exogenous-participation* assumption, i.e., that there a certain number of agents, say n , that will participate in the mechanism. If, for instance, agents only participate in the mechanism when their utility from doing so is large enough, i.e., with *endogenous participation*, then running an optimal mechanism may decrease participation and thus result in a lower revenue than the surplus maximizing mechanism.

On the other hand, the suggestion of this result, that slightly increasing competition can ensure good revenue, is inherently prior independent. The designer does not need to know the prior distribution to market her service so as to attract more agent participation.

5.2.1 Single-item Environments

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

Theorem 5.1. *For i.i.d. regular single-item environments, the expected revenue of the second-price auction with $n + 1$ agents is at least that of the optimal auction with n agents.*

Proof. First consider the following question. What is the optimal single-item auction for $n + 1$ agents that always sells the item? The requirement that the item always be sold implies that, even if all virtual values are negative, a winner must still be selected. Clearly the optimal such auction is the one that assigns the item to the agent with the highest virtual value (cf. Corollary 3.8). Since the distribution is i.i.d. and regular, the agent with the highest virtual value is the agent with the highest value. Therefore, this optimal auction that always sells the item is the second-price auction.

Now consider an $(n + 1)$ -agent mechanism LB that runs the optimal auction on agents $1, \dots, n$ and if this auction fails to sell the item, it gives the item away for free to agent $n + 1$. Obviously, LB’s expected revenue is equal to the expected revenue of the optimal n -agent auction. It is, however, an $(n + 1)$ -agent auction that always sells the item. Therefore,

its revenue is a lower bound on that of the optimal $(n + 1)$ -agent auction that always sells.

We conclude that the expected revenue of the second-price auction with $n + 1$ agents is at least that of LB which is equal to that of the optimal auction for n agents. \square

This resource augmentation result provides the beginning of a prior-independent theory for mechanism design. For instance, we can easily obtain a prior-independent approximation result as a corollary to Theorem 5.1 and Theorem 5.2, below.

Theorem 5.2. *For i.i.d. single-item environments the optimal $(n - 1)$ -agent auction is an $n/n-1$ approximation to the optimal n -agent auction.*

Proof. See Exercise 5.2. \square

Corollary 5.3. *For i.i.d. regular single-item environments with $n \geq 2$ agents, the second-price auction is an $n/n-1$ approximation to the optimal auction revenue.*

5.2.2 Multi-unit and Matroid Environments

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

Example 5.1. For large n consider an n -unit environment and agents with uniformly distributed values on $[0, 1]$. With $n + 1$ agents, the expected revenue of the $(n + 1)$ st-price auction on $n + 1$ agents is about one as there are n winners and the $(n + 1)$ st value is $1/n+2 \approx 1/n$ in expectation.¹ On the other hand, the optimal auction with n agents will post a price of $1/2$ to each agent and achieve an expected revenue of $n/4$.

The resource augmentation result does extend, and in a very natural way, but more than a single agent must be recruited. For k -unit environments we have to recruit k additional agents. Notice that to extend the proof of Theorem 5.1 to a k -unit environment we can define the auction LB for $n + k$ agents to run the optimal n -agent auction on agents $1, \dots, n$

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

and to give any remaining units to agents $n + 1, \dots, n + k$. The desired conclusion follows. In fact, this argument can be extended to matroid environments. Of course matroid set systems are generally asymmetric, so we have to be specific as to the role with respect to the feasibility constraint of the added agents. The result is more intuitive when stated in terms of removing agents from the optimal mechanism instead of adding agents to surplus maximization mechanism, though the consequence is analogous. Recall from Section 4.6 that a base of a matroid is a feasible set of maximal cardinality.

Theorem 5.4. *For any i.i.d. regular matroid environment the expected revenue of the surplus maximization mechanism is at least that of the optimal mechanism in the environment obtained by removing any set of agents that corresponds to a base of the matroid.*

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

5.3 Single-sample Mechanisms

While the assumption that it is possible to recruit an additional agent seems not to be too severe, once we have to recruit k new agents in k -unit environments or a new base for matroid environments, the approach seems less actionable. In this section we will show that a single additional agent is enough to obtain a good approximation to the optimal auction revenue. We will not, however, just add this agent to the market; instead, we will use this agent for market analysis.

In the opening of this chapter we discussed the need to connect the size of the sample for market analysis with the size of the actual market. In this context, the assumption that the prior distribution is known is tantamount to assuming that an infinitely large sample is available for market analysis. In this section we show that this impossibly large sample can be approximated by a single sample from the distribution.

Definition 5.1. The *surplus maximization mechanism with lazy reserves* \hat{v} is the following:

- (i) simulate the surplus maximization mechanism on the bids,

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

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

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

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

where SM denotes the surplus maximization mechanism.

The *lazy single-sample-reserve* mechanism sets $\hat{\mathbf{v}} = (\hat{v}, \dots, \hat{v})$ for $\hat{v} \sim F$. The *lazy monopoly-reserve* mechanism sets $\hat{\mathbf{v}} = \hat{\mathbf{v}}^*$.

Proposition 5.5. *The surplus maximization mechanism with lazy reserves is dominant strategy incentive compatible.*

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

5.3.1 The Geometric Interpretation

Consider a single-agent environment. The optimal auction in such an environment is simply to post the monopoly price as a take-it-or-leave-it offer. In comparison, the single-sample-reserve mechanism would post a random price that is drawn from the same distribution as the agent's value is drawn. We will give a geometric proof that shows that for regular distributions, the revenue from posting such a random price is within a factor of two of that of the (optimal) monopoly price.

This statement can be viewed as the $n = 1$ special case of the Theorem 5.1, i.e., that the two-agent second-price auction obtains at least the (one-agent) monopoly revenue. In a two-agent second-price auction each

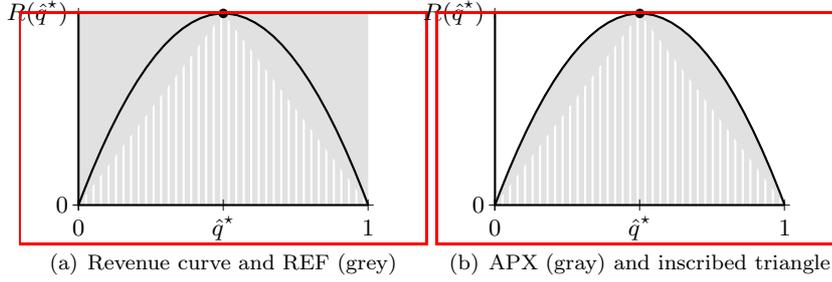


Figure 5.1. The revenue curve (black line) for the uniform distribution is depicted. REF is the area of the rectangle (gray); by geometry the area of the inscribed triangle (white striped) is $1/2$ REF. APX is the area under the revenue curve (gray); by convexity it is lower bounded by the area of the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

agent is offered the a price equal to the value of the other, i.e., a random price from the distribution. Therefore, the two-agent second-price auction obtains twice the revenue of a single sample reserve. The result showing that the single-sample revenue is at least half of the monopoly revenue then implies that the two-agent second-price auction obtains at least the (one-agent) monopoly revenue.

Lemma 5.6. *For a regular single-agent environment, posting a random price from the agent's value distribution obtains at least half the revenue from posting the (optimal) monopoly price.*

Proof. Let $R(\cdot)$ be the agent's revenue curve. Let \hat{q}^* be the quantile corresponding to the monopoly price, i.e., $\hat{q}^* = \operatorname{argmax}_{\hat{q}} R(\hat{q})$. The expected revenue from (optimal) monopoly pricing is $\text{REF} = R(\hat{q}^*)$; this revenue is represented in Figure 5.1(a) by the area of the rectangle (grey) of width one and height $R(\hat{q}^*)$. Recall that drawing a random value from the distribution is equivalent to drawing a uniform quantile. The expected revenue from the corresponding random price is $\text{APX} = \mathbf{E}_{\hat{q}}[R(\hat{q})] = \int_0^1 R(\hat{q}) d\hat{q}$; this revenue is depicted in Figure 5.1(b) by the area below the revenue curve (grey). This area is convex because the revenue curve is concave; therefore, by geometry it contains an inscribed triangle with vertices corresponding to 0, \hat{q}^* , and 1 on the revenue curve (Figure 5.1, white striped). This triangle has width one, height $\text{REF} = R(\hat{q}^*)$, and therefore its area is equal to $1/2 \text{REF}$. Thus, $\text{APX} \geq 1/2 \text{REF}$. \square

Example 5.2. For the uniform distribution where $R(\hat{q}) = \hat{q} - \hat{q}^2$, the

quantities in the proof of Lemma 5.6 can be easily calculated:

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

5.3.2 Monopoly versus Single-sample Reserves

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

Theorem 5.7. *For i.i.d. regular downward-closed environments, the expected revenue of the lazy single-sample-reserve mechanism is at least half of that of the lazy monopoly-reserve mechanism.*

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

Denote the monopoly quantile by \hat{q}^* , denote the critical quantile for Alice in the surplus maximization mechanism with no reserve by \hat{q}_i^{SM} , and denote the quantile of a lazy reserve by \hat{q} . Alice's wins in the surplus maximization mechanism with this lazy reserve when her quantile is below $\min(\hat{q}, \hat{q}_i^{\text{SM}})$. For a fixed \hat{q}_i^{SM} , the revenue from Alice, in expectation over her own quantile and as a function of the lazy reserve quantile \hat{q} , induces the revenue curve $R^\dagger(\hat{q}) = R(\min(\hat{q}, \hat{q}_i^{\text{SM}}))$. Figure 5.2 depicts Alice's original revenue curve $R(\cdot)$ and this induced revenue curve $R^\dagger(\cdot)$ in the cases that $\hat{q}_i^{\text{SM}} \leq \hat{q}^*$ and $\hat{q}_i^{\text{SM}} \geq \hat{q}^*$.

Alice's expected payment in the lazy monopoly-reserve mechanism is $\text{REF} = R^\dagger(\hat{q}^*)$ which is geometrically the maximum height of the revenue curve R^\dagger ; and her expected payment in the lazy single-sample-reserve mechanism, where $\hat{q} \sim U[0, 1]$, is $\text{APX} = \mathbf{E}_{\hat{q}}[R^\dagger(\hat{q})]$. We conclude with the same geometric argument as in Lemma 5.6 that relates REF

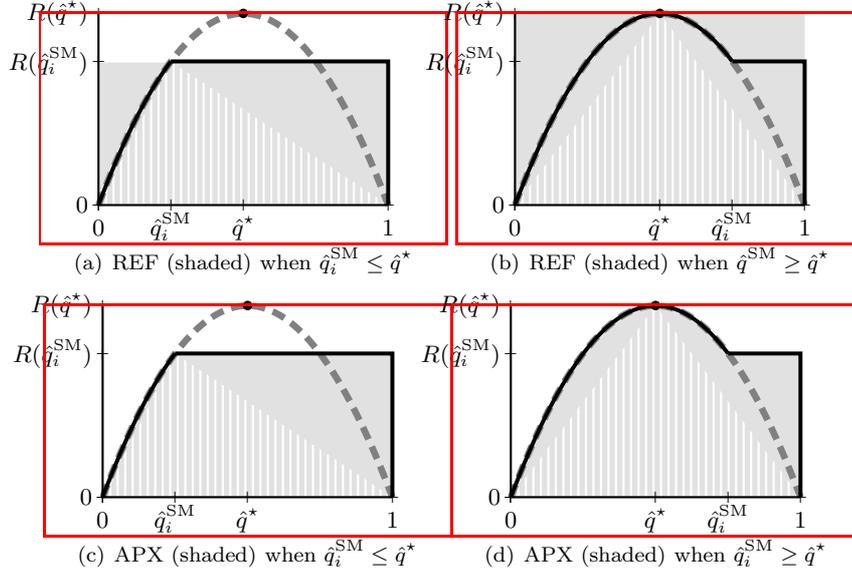


Figure 5.2. In each diagram, the revenue curve $R(\cdot)$ (thick, dashed, grey line) of the uniform distribution and the induced revenue curve $R^\dagger(\cdot) = R(\max(\cdot, \hat{q}_i^{\text{SM}}))$ (thin, solid, black line). On the left is the case that $\hat{q}_i^{\text{SM}} \leq \hat{q}^*$; on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

to a rectangle, APX to the area under the induced revenue curve, and $1/2 \text{REF}$ to the area of an inscribed triangle (see Figure 5.2). \square

5.3.3 Optimal versus Lazy Single-sample-reserve Mechanism

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

For i.i.d. matroid environments, as discussed above, lazy monopoly reserve pricing is identical to (eager) monopoly reserve pricing. Moreover, surplus maximization with the monopoly reserve is revenue optimal (Proposition 4.23). We conclude the following corollary. Recall that matroid environments include multi-unit environments as a special case.

Corollary 5.8. *For any i.i.d. regular matroid environment, the revenue of the single-sample-reserve mechanism is a two approximation to that of the revenue-optimal mechanism.*

Theorem 4.40 shows that for monotone-hazard-rate distributions the surplus maximization mechanism with (eager) monopoly reserves is a two approximation to the optimal mechanism; however, as in downward-closed environments eager and lazy reserve pricing are not identical (see Exercise 5.3), we have slightly more work to do. Recall Theorem 4.37 which states that for MHR distributions the optimal revenue and optimal social surplus are within an ϵ factor of each other. One way to prove this theorem is, in fact, by showing that the revenue of the surplus maximization mechanism with lazy monopoly reserve prices is an ϵ approximation to the optimal social surplus and hence so is the optimal revenue (see Exercise 4.25). Combining this observation with Theorem 5.7 it is evident that the lazy single-sample-reserve mechanism is a 2ϵ approximation. The approximation bound can be improved to four via a more careful analysis that we omit.

Theorem 5.9. *For any i.i.d. monotone-hazard-rate downward-closed environment, the revenue of the lazy single-sample-reserve mechanism is a four approximation to that of the revenue-optimal mechanism.*

5.4 Prior-independent Mechanisms

We now turn to mechanisms that are completely prior independent. Unlike the mechanisms of the preceding section, these mechanisms will not require any distributional information, not even a single sample from the distribution. We will, however, still assume that there is a distribution.

Definition 5.2. A mechanism APX is a *prior-independent* β *approximation* if

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

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

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

Consider the following k -unit auction:

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

This auction is clearly incentive compatible. Furthermore, it is easy to see that it is a $2^{n/n-1}$ approximation for $n \geq 2$ agents with values drawn i.i.d. from a regular distribution. This follows from the fact that rejecting a random agent loses at most a $1/n$ fraction of the optimal revenue (Theorem 5.2), and from the previous single-sample-reserve result (Corollary 5.8). This approximation bound is clearly worst for $n = 2$ where it guarantees a four approximation. The same approach can be applied to matroid and downward-closed environments as well; instead, we will discuss a slightly more sophisticated approach.

5.4.1 Digital Good Environments

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

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

Definition 5.3. For digital-good environments,

- the (digital good) *pairing auction* arbitrarily pairs agents and runs the second-price auction on each pair (assuming n is even), and
- the (digital good) *circuit auction* orders the agents arbitrarily (e.g., lexicographically) and offers each agent a price equal to the value of the preceding agent in the order (the first agent is offered the last agent's value).

The *random pairing auction* and the *random circuit auction* are the variants where the pairing or circuit is selected randomly.

Theorem 5.10. *For i.i.d. regular digital-good environments, any auction wherein each agent is offered the price of another random or arbitrary (but not value dependent) agent is a two approximation to the optimal auction revenue.*

The proof of this theorem follows directly from the geometric analysis of single-sample pricing (Lemma 5.6). Clearly, the pairing and circuit auctions satisfy the conditions of the above theorem. In conclusion, in i.i.d. environments it is relatively easy to obtain samples from the distribution while running a mechanism.

5.4.2 General Environments

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

Consider the following definition of mechanism composition (cf. Exercise 5.9). Notice that the mechanisms we have been discussing can all be interpreted as calculating a critical value for each agent, serving each agent whose value exceeds her critical value, and charging each served agent her critical value. In fact, by Corollary 2.14, any randomization over deterministic dominant strategy incentive compatible mechanisms admits such an interpretation.

Definition 5.4. The *parallel composite* \mathcal{M} of two (randomizations over) deterministic DSIC mechanisms, \mathcal{M}^\dagger and \mathcal{M}^\ddagger is as follows:

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

Notice that in the parallel composite, \mathcal{M} , the set of agents served is the intersection of those served by \mathcal{M}^\dagger and \mathcal{M}^\ddagger . By downward closure, then, the outcome of the composition is feasible as long as the outcome of one

of \mathcal{M}^\dagger or \mathcal{M}^\ddagger is feasible. The mechanism is dominant strategy incentive compatible by its definition via critical values and Corollary 2.14.

Proposition 5.11. *The parallel composite of two (randomizations over) deterministic dominant strategy incentive compatible mechanisms is dominant strategy incentive compatible and, if one of the mechanisms is feasible, feasible.*

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

Definition 5.5. For downward-closed environments,

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

Theorem 5.12. *For i.i.d. regular matroid environments, the revenues of the pairing and circuit mechanisms are two approximations to the optimal mechanism revenue.*

Theorem 5.13. *For i.i.d. monotone-hazard-rate downward-closed environments, the revenues of the pairing and circuit mechanisms are four approximations to the optimal mechanism revenue.*

The results presented in this chapter are representative of the techniques for the design and analysis of prior-independent approximation mechanisms; however, a number of extensions are possible. If we use more than one samples from the distribution, bounds for regular distributions can be improved and bounds for irregular distributions can be obtained. Both of these directions will be taken up during our discussion of prior-free mechanisms in Chapter 7. Finally, the i.i.d. assumption can be relaxed, either by assuming that agents are partitioned by demographic (see Exercise 5.10) or by an ordering assumption.

Exercises

- 5.1 Consider the sale of a magazine subscription over two periods to a single agent who has a linear uniform additive value for each period's issue of the magazine. Her value v is drawn from a regular distribution F and if x_1 , x_2 , p_1 , and p_2 denote her allocation and payments in each period then her utility is $v(x_1 + x_2) - p_1 - p_2$. In each period, the designer publishes her mechanism and then the agent bids for receiving that period's issue of the magazine.
- Suppose that the designer can commit to the mechanism to be used in period two before the agent bids in period one, describe the revenue optimal mechanisms and the equilibrium behavior of the agent.
 - Suppose that the designer cannot commit to the mechanism to be used in period two before the agent bids in period one, describe the revenue optimal mechanisms and the equilibrium behavior of the agent.
 - Compare the revenues from the previous steps for the uniform distribution.
- 5.2 Prove Theorem 5.2: For i.i.d. single-item environments the optimal auction with $n - 1$ agents auction is an $n/n-1$ approximation to the optimal auction with n agents.
- 5.3 Consider the surplus maximization mechanism with an anonymous reserve that is either lazy or eager.
- Find a valuation profile, downward-closed feasibility constraint, and anonymous reserve price such that different outcomes result from lazy and eager reserve pricing.
 - Prove that for anonymous reserve pricing in matroid environments, lazy and eager reserve pricing give the same outcome.
- 5.4 Consider a regular single-agent environment. Show that posting the median price from the agent's value distribution obtains at least half the revenue from posting the monopoly price. The median price for an agent with inverse demand function $V(\cdot)$ is $\hat{v} = V(1/2)$.
- 5.5 In Example 5.2 it is apparent that the approximation bound of a sample reserve to the monopoly reserve for a uniform distribution is $3/2$. Use this bound to derive better bounds for the lazy single-sample-reserve mechanism versus the lazy monopoly-reserve mechanism. In particular, show that if the single-agent approximation of sample reserve to monopoly reserves is β then the the same

bound holds in general for the lazy single-sample-reserve and lazy monopoly reserve mechanism.

- 5.6 Consider the surplus maximization mechanism with lazy monopoly reserve prices in downward-closed monotone-hazard-rate environments.
- Show that in a single-agent environment, that its expected surplus is at most twice its expected revenue.
 - Show that in a downward-closed environment, that its expected surplus is at most twice its expected revenue.
- 5.7 Suppose we are in a non-identical environment, i.e., agent i 's value is drawn from independently from distribution F_i , and suppose the mechanism can draw one sample from each agent's distribution.
- Give a constant approximation mechanism for regular, matroid environments (and give the constant).
 - Give a constant approximation mechanism for monotone-hazard-rate, downward-closed environments (and give the constant).
- 5.8 This chapter has been mostly concerned with the profit objective. Suppose we wished to have a single mechanism that obtained good surplus and good profit.
- Show that surplus maximization with monopoly reserves is not generally a constant approximation to the optimal social surplus in regular, single-item environments.
 - Show that the lazy single sample mechanism is a constant approximation to the optimal social surplus in i.i.d., regular, matroid environments.
 - Investigate the Pareto frontier between prior-independent approximation of surplus and revenue. I.e., if a mechanism is an α approximation to the optimal surplus and a β approximation to the optimal revenue, plot it as point $(1/\alpha, 1/\beta)$ in the positive quadrant.
- 5.9 Define the *sequential composite* \mathcal{M} of two mechanism \mathcal{M}^\dagger and \mathcal{M}^\ddagger as first simulating \mathcal{M}^\dagger , second simulating \mathcal{M}^\ddagger on the winners of \mathcal{M}^\dagger , and serving the agents served by the second mechanism at the maximum of their prices in the two mechanisms.
- Give an example of deterministic DSIC mechanisms \mathcal{M}^\dagger and \mathcal{M}^\ddagger such that the sequential composite \mathcal{M} is not DSIC.
 - Show that if \mathcal{M}^\dagger is the surplus maximizing mechanism (and \mathcal{M}^\ddagger is any randomization over DSIC mechanisms) then the composition is DSIC.

- (c) Describe a property of the surplus maximizing mechanism as \mathcal{M}^\dagger that enables the incentive compatibility of the sequential composite \mathcal{M} .
- 5.10 Suppose the agents are divided into k markets where the value of agents in the same market are identically distributed, e.g., by demographic. Assume that the partitioning of agents into markets is known, but not the distributions of the markets. Assume there are at least two agents in each market. Unrelated to the markets, assume the environment has a downward-closed feasibility constraint.
- (a) Give a prior-independent constant approximation to the revenue-optimal mechanism for regular matroid environments.
 - (b) Give a prior-independent constant approximation to the revenue-optimal mechanism for monotone-hazard-rate downward-closed environments.

Chapter Notes

The resource augmentation result that shows that recruiting one more agent to a single-item auction raises more revenue than setting the optimal reserve price is due to Bulow and Klemperer (1996). The proof of the Bulow-Klemperer Theorem that was presented in this text is due to René Kirkegaard (2006). A generalization of the Bulow-Klemperer Theorem to non-identical distributions was given by Hartline and Roughgarden (2009).

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

6

Bayes-Nash Approximation

This text primarily focuses on the design of incentive compatible mechanisms, i.e., ones where truth telling is an equilibrium. This focus is justified in theory by the revelation principle (Section 2.10 on page 46) which suggests that if there is a mechanism with a good equilibrium then there is one where truth telling is a good equilibrium. Thus, nothing “good” is lost by the restriction. In practice, though, designed mechanisms are rarely incentive compatible, and undoing the revelation principle is not straightforward. It is not always an easy task to identify a practical mechanism with the same Bayes-Nash equilibrium outcome as a designed Bayesian incentive compatible mechanisms. This chapter focuses on the analysis of mechanisms that are not incentive compatible, and in design criteria for them.

In the design of Bayes-Nash (i.e., non-incentive-compatible) mechanisms there will be less fine grained control over the exact equilibrium selected by the mechanism, instead we will look to identify properties of mechanisms from which we can guarantee that any equilibrium is approximately optimal.

Our motivating example is the first-price auction with agents with independent but non-identically distributed values. Recall that with identically distributed values the first-price auction possesses a unique symmetric equilibrium in which the highest valued agent always wins the item (see Section 2.9 on page 42). This outcome is optimal from the perspective of social surplus. Moreover, the first-price auction with the monopoly reserve price, for values drawn i.i.d. from a regular distribution, is revenue optimal in equilibrium. For asymmetric distributions the first-price auction is neither optimal for surplus nor revenue. We will show that the first-price auction is an $e/e-1 \approx 1.58$ approximation

for social surplus, and the first price auction with asymmetric monopoly reserves is a $2e/e-1 \approx 3.16$ approximation for revenue.

One of the reasons analysis of Bayes-Nash mechanisms is important is that the ideal setting of incentive compatible mechanism design, where a mechanism is being run in a closed system, is rare. In many practical applications of mechanism design, agents may have the option to participate in many mechanisms, simultaneously or in sequence. Incentive compatibility of these individual mechanisms does not imply incentive compatibility of the composition of mechanisms. An important development of this chapter is a theory of composition for mechanisms. Via this theory we will show that simultaneous first-price auctions for multiple items (albeit for single-dimensional agents) have the same performance guarantees stated above for the first-price auction in isolation.

The conventional approach to the analysis of Bayes-Nash equilibrium, as a first step, explicitly solves for the Bayes-Nash equilibrium. For asymmetric environments such an analysis would require the solution to analytically intractable differential equations. The approximation-based approach presented herein circumvents solving for BNE by decomposing the analysis into the following two parts. The first part isolates the best-response property of Bayes-Nash equilibrium and formalizes the intuition that either an agent gets good utility or must be facing fierce competition. The second part identifies a *revenue covering* property, that revenue exceeds an aggregate measure of the competition faced by each agent, as a criteria to be approximated. With bounds on utility and revenue, we get approximation bounds on the social surplus (the sum of utility and revenue).

The bounds we derive on the social surplus and revenue of auctions in Bayes-Nash equilibrium are parameterized by the extent to which revenue covering approximately holds. This observation then gives clear direction for optimization in mechanism design. A Bayes-Nash mechanism's performance is proportional to its approximation with respect to revenue covering. Bayes-Nash mechanisms should be designed to minimize this approximation.

6.1 Social Surplus of Winner-pays-bid Mechanisms

The first-price auction with asymmetric value distributions does not maximize social surplus in Bayes-Nash equilibrium. For two agents and the uniform distribution on distinct supports, the differential equations

that govern Bayes-Nash equilibrium can be solved; Example 6.1, below, gives the solution one special case. For more general distributions, more than two agents, and more complex auction formats, equilibrium is analytically intractable. A main result of this section will be the following bound on the social surplus of any Bayes-Nash equilibrium of the first-price auction for any product distribution on agent values.

Theorem 6.1. *For any product distribution and Bayes-Nash equilibrium of the first-price auction, the expected social surplus is an $\frac{e}{e-1} \approx 1.58$ approximation to the expected social surplus of the optimal outcome.*

Example 6.1. The equilibrium s of Alice (agent 1) and Bob (agent 2) in the first-price auction with $v_1 \sim U[0, 1]$ and $v_2 \sim U[0, 2]$ is

- $s_1(v) = \frac{2}{3v} \left(2 - \sqrt{4 - 3v^2} \right)$ and
- $s_2(v) = \frac{2}{3v} \left(\sqrt{4 + 3v^2} - 2 \right)$.

The asymmetry of strategies implies that the highest-valued agent does not always win, i.e., the auction is inefficient. Bob wins when $v_2 > (v_1^{-2} + 3/4)^{-1/2} > v_1$. See Figure 6.1.

In this section we will consider a generalizations of the first-price auction where the mechanism selects an allocation based on the bids and all the winners pay their bid. As per the following definition, the first-price auction is the winner-pays-bid highest-bids-win mechanism for single-item environments.

Definition 6.1. A *winner-pays-bid* mechanism

- (i) solicits bids,
- (ii) selects a feasible set of winners, and
- (iii) charges each winner her bid.

In the winner-pays-bid *highest-bids-win* mechanism, the winners selected in Step (ii) are the feasible set of agents with the highest sum of bids.

Chapter 6: Topics Covered.

- The geometry of best response for single-dimensional agents.
- Revenue covering, a criterion for Bayes-Nash optimization.
- Analysis of welfare and revenue in Bayes-Nash equilibrium.
- Analysis of the simultaneous composition of mechanisms.
- Reserve prices in Bayes-Nash mechanisms.

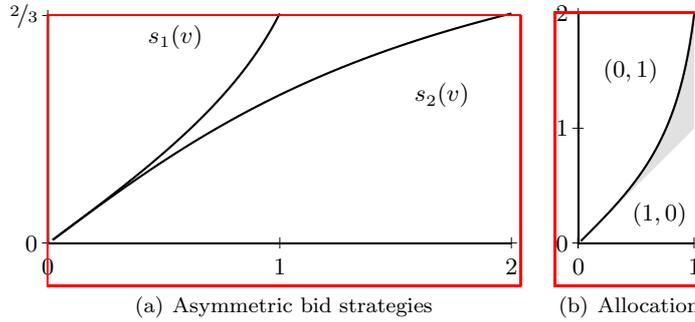


Figure 6.1. In (a) the Bayes-Nash bid strategies in the first-price auction with the asymmetric value distribution of Example 6.1 are depicted. In (b) the Bayes-Nash allocation $\mathbf{x}(v)$ is depicted with v_1 on the horizontal axis and v_2 on the vertical axis. In the shaded gray area the BNE allocates to agent 1 while $v_2 > v_1$ and this allocation is inefficient.

6.1.1 The Geometry of Best Response

Consider an agent, Alice, in a winner-pays-bid mechanism. Alice wins the auction when her bid b exceeds a *critical bid* \hat{b} which is given by the bids of others and the rules of the mechanism. For the first-price auction this critical bid is the maximum of the other agents' bids. Denote the interim *bid allocation rule*, which maps Alice's bid to her probability of winning, as given the distribution of other agents' bids and the mechanism's rules, by $\tilde{x}(b) = \Pr_{\hat{b}}[b > \hat{b}]$. Notice that this bid allocation rule is precisely to the cumulative distribution function of Alice's critical bid. The expected critical bid Alice faces is a measure of the level of competition in the auction, and is given by the area above its cumulative distribution function (equivalently, the area above the bid allocation rule).¹ By the definition of the auction, Alice's utility with value v for any bid b is given by $u(v, b) = (v - b) \tilde{x}(b)$. Our analysis will relate Alice's utility, her expected critical bid, and her value; each of which can be compared geometrically (Figure 6.2).

- Alice's expected utility for any bid, denoted $u(v, b) = (v - b) \tilde{x}(b)$, is given by a rectangle below the bid allocation rule. Alice's utility in Bayes-Nash equilibrium, denoted $u(v)$, is the largest such rectangle.
- Alice's expected critical bid, denoted $\hat{B} = \mathbf{E}[\hat{b}]$, is given by the area above the bid allocation rule, i.e., $\hat{B} = \int_0^\infty (1 - \tilde{x}(b)) db$.

¹ Recall, the expected value of a non-negative random variable $v \sim F$ is given by $\mathbf{E}[v] = \int_0^\infty (1 - F(z)) dz$, cf. Section A.3 on page 330.

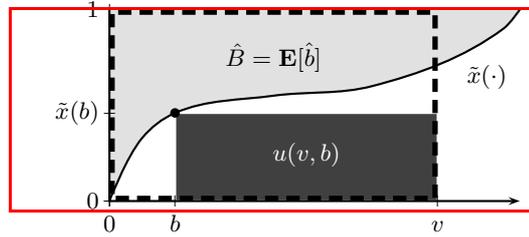


Figure 6.2. Geometry of best-response in the first price auction. The expected critical bid is the area (light gray) above the bid allocation rule (thin solid line). The utility from a bid is given by the a rectangle (dark gray) below the bid allocation rule. The value of the agent can be depicted by the area of a rectangle (thick dashed outline).

- Alice's value v can be compared geometrically to the above quantities as the area of the rectangle of width v and height 1.

Intuitively, either Alice's utility or her expected critical bid is a large fraction of her value.

6.1.2 Utility Approximates Value

We formalize the geometric intuition that either Alice's utility or expected critical bid is large compared to her value in the following theorem.

Theorem 6.2. *In any Bayes-Nash equilibrium of any winner-pays-bid mechanism and for any agent, the expected sum of her utility and her critical bid is an $e/e-1 \approx 1.58$ approximation to her value; i.e.,*

$$u(v) + \hat{B} \geq e^{-1/e} v.$$

One way to prove Theorem 6.2 is via a best-response argument. In particular, BNE utility is at least the utility $u(v, b)$ for any value v and any deviation bid b ; a careful selection of deviation gives the desired bound. As a warm up, consider deviating to $b = v/2$ and observe that

$$u(v, v/2) + \hat{B} \geq 1/2 v. \quad (6.1)$$

Fix any critical bid \hat{b} . If $\hat{b} \geq v/2$, non-negativity of BNE utility implies inequality (6.1). On the other hand, if $\hat{b} \leq v/2$, then Alice wins by bidding $b = v/2$ and her utility is $v - v/2 = v/2$ as required by (6.1). Taking expectations of these inequalities over \hat{b} gives equation (6.1). This argument is depicted geometrically in Figure 6.3.

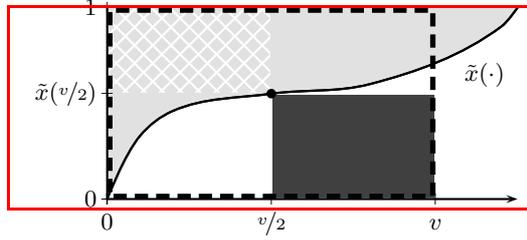


Figure 6.3. Geometry of equation (6.1): $u(v, v/2) + \hat{B} \geq 1/2 v$. The utility $u(v, b)$ from deviating to $b = v/2$ is depicted by the dark gray area; the expected critical bid \hat{B} is depicted by the light gray area; the lower bound on \hat{B} is crosshatched.

Proof of Theorem 6.2. Fix the critical bid \hat{b} , and consider the utility from deviating to a random bid b drawn from the distribution G on support $[0, e^{-1/e} v]$ with density function $g(z) = 1/v-z$. If $\hat{b} \geq e^{-1/e} v$ the inequality of the theorem holds. Otherwise, the utility from such a deviation is $v - b$ when $b \geq \hat{b}$, zero otherwise.

$$\begin{aligned} u(v) &\geq u(v, b) \\ &\geq \int_{\hat{b}}^{e^{-1/e} v} (v - b) g(b) db \geq \int_{\hat{b}}^{e^{-1/e} v} 1 db \\ &= e^{-1/e} v - \hat{b}. \end{aligned}$$

Thus, $u(v) + \hat{b} \geq e^{-1/e} v$. The deviation strategy is independent of \hat{b} , so taking expectation over \hat{b} yields the theorem. \square

6.1.3 Revenue Covering Approximation

The next part of the analysis is to bound the sum of the expected critical bids for any feasible subset of the agents by the expected revenue of the auction. This analysis is performed for any distribution of bids and does not, in particular, assume that these bids are in equilibrium. For pay-your-bid mechanisms, we may as well perform this bound pointwise for all bid profiles. Such a non-equilibrium pointwise analysis is both easy and versatile.

Definition 6.2. A pay-your-bid mechanism \mathcal{M} has *revenue covering approximation* μ if, for any profile of bids \mathbf{b} and any feasible allocation \mathbf{y} ,

$$\text{Revenue}(\mathbf{b}) \geq \frac{1}{\mu} \sum_i \hat{b}_i y_i, \tag{6.2}$$

where, for bid profile \mathbf{b} and mechanism \mathcal{M} , $\text{Revenue}(\mathbf{b})$ is the revenue and \hat{b}_i is the critical bid of agent i . Mechanism \mathcal{M} is *revenue covered* if $\mu = 1$.

Of course by taking expectations of both sides of inequality (6.2) for and distribution of bids $\mathbf{b} \sim \mathbf{G}$, the expected revenue is at least the expected critical bid of any feasible set of agents \mathbf{y} , i.e.,

$$\mathbf{E}_{\mathbf{b}}[\text{Revenue}(\mathbf{b})] \geq \frac{1}{\mu} \sum_i \hat{B}_i y_i,$$

where $\hat{B}_i = \mathbf{E}_{\mathbf{b}}[\hat{b}_i]$ for \mathcal{M} as previously discussed.

Notice that in the definition of revenue covering the revenue of the auction and the critical bids are given by the bid profile and the definition of the rules of the mechanism. The allocation \mathbf{y} is unrelated to the bid profile and the rules of the mechanisms, it is only constrained by the feasibility constraint of the single-dimensional allocation problem, as defined in Section 3.1 on page 54. As alluded to above, analysis of the approximation μ of revenue covering of any given mechanism in any given environment is generally straightforward.

Theorem 6.3. *The first-price auction in single-item environments is revenue covered, i.e., $\mu = 1$.*

Proof. First, recall that for the first-price auction the critical bid faced by an agent is equal to the highest of the other agents' bids. The revenue of the auction is the highest bid over all. Thus, the revenue of the auction $\text{Revenue}(\mathbf{b}) = \max_j b_j$ is at least the critical bid $\hat{b}_i = \max_{j \neq i} b_j$ of any agent j . Second, feasibility of \mathbf{y} in single-item environments requires that $\sum_i y_i \leq 1$. Combining these observations,

$$\begin{aligned} \sum_i \hat{b}_i y_i &\leq \sum_i \text{Revenue}(\mathbf{b}) y_i \leq \text{Revenue}(\mathbf{b}) \sum_i y_i \\ &\leq \text{Revenue}(\mathbf{b}). \end{aligned} \quad \square$$

Similarly, it is a relatively easy exercise to show that the winner-pays-bid highest-bids-win mechanism for matroid environments (see Section 4.6 on page 129) is revenue covered ($\mu = 1$). Not all mechanisms are revenue covered, in fact, the winner-pays-bid highest-bids-win mechanism for the single-minded combinatorial auction environment, one of the canonical examples of a downward-closed environment that is not a matroid, is not revenue covered for any $\mu < m$; neither is the winner-pays-bid highest-bids-win mechanism for the routing environment discussed in Section 1.1.3 on page 14 (see Exercise 6.2).

Theorem 6.4. *The winner-pays-bid highest-bids-win mechanism for matroid environments is revenue covered, i.e., $\mu = 1$.*

Proof. See Exercise 6.3. □

Example 6.2. In the *single-minded combinatorial auction* environment there are n agents and m items. Agent i has value v_i for bundle $S_i \subseteq [m]$ (and no value for any other bundle of items). Item j may be sold to at most one agent. It is assumed that the bundles are known and the values are each agent's private information. An allocation \mathbf{x} is feasible for a single-minded combinatorial environment no items are allocated to multiple agents, i.e., $x_i = x_{i^\dagger} = 1$ only if $S_i \cup S_{i^\dagger} = \emptyset$.

Consider winner-pays-bid highest-bids-win mechanism for the single-minded combinatorial auction environment. This mechanism does not generally have good surplus in BNE; moreover it does not have revenue covering approximation μ for any $\mu < m$. To show this we need to exhibit an environment (given by bundles $\mathbf{S} = (S_1, \dots, S_n)$, a bid profile \mathbf{b} , and feasible allocation \mathbf{y} such that the inequality of Definition 6.2 is only satisfied if $\mu = m$.

Consider the environment with $n = m + 1$ agents with:

- agent n demanding $S_n = \{1, \dots, m\}$, the grand bundle, and
- agent $i \neq n$ demanding $S_i = \{i\}$, a singleton bundle.

Consider the bid profile $\mathbf{b} = (0, \dots, 0, 1)$ and feasible (but not highest-bids-win) allocation $\mathbf{y} = (1, \dots, 1, 0)$. In other words, the agent n demanding the grand bundle bids $b_n = 1$ and is not served ($y_n = 0$), while singleton agents $i \neq n$ bid $b_i = 0$ and are served ($y_i = 1$).

The highest-bids-win mechanism's revenue for bids \mathbf{b} is only 1 as it selects feasible outcome $\mathbf{x} = \mathbf{b} = (0, \dots, 0, 1)$ that serves only the grand-bundle agent n . Each singleton agent $i \neq n$ faces a critical bid of $\hat{b}_i = 1$ as she must beat agent n . As the sum of the feasible critical bids is $\sum_i \hat{b}_i y_i = m$, the auction is not a revenue covering approximation for any $\mu < m$.

In the next sections we will see that the approximation of social surplus (and revenue, with monopoly reserves) of a Bayes-Nash mechanism is proportional to its revenue covering approximation μ . Thus, to design a good Bayes-Nash mechanism it suffices to design a mechanism with small revenue covering approximation.

6.1.4 Social Surplus in Bayes-Nash Equilibrium

Theorem 6.2, which shows that utility approximates value, and revenue covering approximation combine to give a bound on the Bayes-Nash equilibrium surplus relative to any feasible outcome, including the one that maximizes social surplus.

Theorem 6.5. *For any winner-pays-bid mechanism that has a revenue covering approximation of $\mu \geq 1$, the expected social surplus in Bayes-Nash equilibrium is an $\mu^e/e-1$ approximation to the optimal social surplus.*

Proof. Consider a valuation profile \mathbf{v} and agent i . By Theorem 6.2 in BNE,

$$u_i(v_i) + \hat{B}_i \geq e^{-1/e} v_i.$$

Denote by $y^*(\mathbf{v}) = \operatorname{argmax}_y \sum_i v_i y_i$ the surplus optimizing allocation. Thus, $\sum_i v_i y_i^*(\mathbf{v}) = \operatorname{REF}(\mathbf{v})$, the optimal social surplus. Notice that $y_i^*(\mathbf{v}) \in [0, 1]$; thus,

$$u_i(v_i) + \hat{B}_i y_i^*(\mathbf{v}) \geq e^{-1/e} v_i y_i^*(\mathbf{v}).$$

Sum over all agents i and invoke μ revenue covering:

$$\sum_i u_i(v_i) + \mu \mathbf{E}[\text{Revenue}] \geq e^{-1/e} \operatorname{REF}(\mathbf{v}).$$

Take expectation over values \mathbf{v} from the distribution \mathbf{F} and use $\mu \geq 1$:

$$\mu (\mathbf{E}[\text{Utilities}] + \mathbf{E}[\text{Revenue}]) \geq e^{-1/e} \mathbf{E}[\operatorname{REF}(\mathbf{v})].$$

The theorem follows from observing that the surplus of the mechanism APX is equal to sum of the utilities of the agents and the mechanism's revenue. \square

As is evident from the statement of Theorem 6.5, to show that a winner-pays-bid auction has good welfare in Bayes-Nash equilibrium it suffices to show that it has a revenue covering approximation. As we saw above, the first-price auction has revenue covering approximation of $\mu = 1$ (Theorem 6.3); thus, it is a $e/e-1 \approx 1.58$ approximation to social surplus in any Bayes-Nash equilibrium. In other words, Theorem 6.1 is proved. More generally, winner-pays-bid highest-bids-win matroid mechanisms are also a 1.58 approximation to social surplus (by Theorem 6.4 and Theorem 6.5).

6.2 Beyond Winner-pays-bid Mechanisms

In this section we will extend the analysis of the preceding section to mechanisms that do not have winner-pays-bid semantics. This extension will allow straightforward generalization to all-pay mechanisms and mechanisms with complex action spaces. A motivating example of a mechanism with a complex action space is the simultaneous first-price auction for single-dimensional constrained matching markets, cf. Section 4.6.1 on page 131.

In constrained matching markets there are n agents and m items. Each agent i has a value v_i for any of the items in bundle $S_i \subset [m]$. In the simultaneous first-price auction, each agent selects which items to bid on and how much to bid, the agents submit bids simultaneously to the auctions, and each item is sold at the highest bid to the highest bidder. Importantly, an agent who bids in more than one auction may win more than one item even though she only has value for one item.

While the simultaneous first-price auction allows multi-dimensional bids, it is still a single-dimensional game, see Section 2.4 on page 29. Just as the Bayes-Nash equilibrium characterization for single-dimensional games is expressed by allocation and payment rules in terms of each agent's valuation (Theorem 2.2 on page 31), revenue equivalence suggests that we can equally well express the allocation and payment rule of the BNE of any mechanism in terms of the winner-pays-bid implementation of the BNE.

Definition 6.3. Consider any mechanism \mathcal{M} , an agent with action space A in \mathcal{M} , and any distribution of other agents' actions.

- The (interim) *action allocation rule* $x^{\mathcal{M}} : A \rightarrow [0, 1]$ maps any action $a \in A$ to a probability of allocation.
- The *action payment rule* $p^{\mathcal{M}} : A \rightarrow [0, 1]$ maps any action to an expected payment.
- The (interim, effective) *winner-pays-bid allocation rule*, denoted $\tilde{x}(\cdot)$, is the smallest monotone function that upper bounds the pointset given by:²

$$\{(p^{\mathcal{M}}(a)/x^{\mathcal{M}}(a), x^{\mathcal{M}}(a)) : a \in A\}.$$

- The (interim, effective) *expected critical bid* \hat{B} is the area above $\tilde{x}(\cdot)$.

Especially for auctions like the all-pay auction, is is not well-defined to

² Note, if the mechanism \mathcal{M} is individually rational then $(0, 0)$ is always in the pointset.

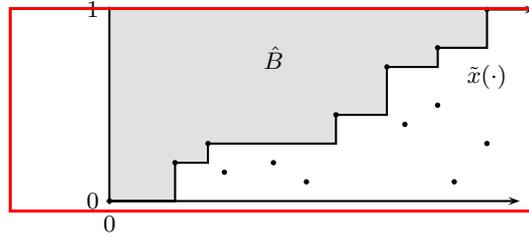


Figure 6.4. The pointset of equation 6.3 is depicted. The effective winner-pays-bid allocation rule $\tilde{x}(\cdot)$ is the smallest monotone function that upper bounds the point set. The points strictly below $\tilde{x}(\cdot)$ are dominated and correspond to actions that will never be taken. The expected critical bid \hat{B} is given by the shaded (light gray) area.

talk about the (effective) winner-pays-bid allocation rule without imposing a distribution on actions. The following definition generalizes revenue covering approximation to mechanisms that do not have winner-pays-bid semantics.

Definition 6.4. A mechanism \mathcal{M} has *revenue covering approximation* μ if, for any product distribution on action profiles $\mathbf{a} \sim \mathbf{G}$ and any feasible allocation \mathbf{y} ,

$$\mathbf{E}_{\mathbf{a}}[\text{Revenue}(\mathbf{a})] \geq \frac{1}{\mu} \sum_i \hat{B}_i y_i,$$

where, for action profile $\mathbf{a} \sim \mathbf{G}$ and mechanism \mathcal{M} , $\text{Revenue}(\mathbf{a}) = \sum_i p_i^{\mathcal{M}}(\mathbf{a})$ is the mechanism's revenue and \hat{B}_i is the expected (effective) critical bid of agent i (from her effective winner-pays-bid allocation rule; Definition 6.3).

The developments of the previous section; specifically Theorem 6.2 and Theorem 6.5; extend without modification to non-winner-pays-bid mechanisms via Definition 6.3 and Definition 6.4. The following theorem summarizes.

Theorem 6.6. *For any individually-rational mechanism that has a revenue covering approximation of $\mu \geq 1$, the expected social surplus in Bayes-Nash equilibrium is an $\mu^e/e-1$ approximation to the optimal social surplus.*

For example, the following theorem can be shown. From it and Theo-

rem 6.6 we conclude that the all-pay auction is a $2e/e-1 \approx 3.16$ approximation to social welfare.³

Theorem 6.7. *In single-item environments the all-pay auction is 2 revenue covered.*

Proof. See Exercise 6.5. □

6.3 Simultaneous Composition

In this section we consider the simultaneous composition of revenue covered mechanisms and show that the composite mechanism is itself revenue covered. An example to have in mind is the simultaneous first-price auction for single-dimensional constrained matching markets that was described at the onset of the preceding section. We impose three assumptions on the environment of these mechanisms:

- (i) The agents are *unit-demand* with respect to simultaneous allocation across several mechanisms. In other words, an agent is considered served if she is served by any of the individual mechanisms in the composition, and she has no additional value for being served by multiple mechanisms over being served in a single mechanism. She must pay for each mechanisms in which she is served.
- (ii) Each mechanism is *individually rational*. This assumption requires that each agent has an action that gives non-negative utility. In particular, an agent with value zero must have an action with zero (expected) payment; we may as well assume that such an agent will also not be served. This action effectively enables an agent to abstain from participation in each mechanism.
- (iii) The individual environments in the composition are *downward closed* and the composite environment is their *union environment*. In other words, if $\mathbf{x}^1, \dots, \mathbf{x}^m \in \{0, 1\}^n$ are deterministic feasible outcomes for $\mathcal{M}^1, \dots, \mathcal{M}^m$, respectively; then \mathbf{x} with $x_i = \max_j x_i^j$ is feasible for \mathcal{M} .

Definition 6.5. Given m mechanisms $\mathcal{M}^1, \dots, \mathcal{M}^j$; the *simultaneous composite mechanism* \mathcal{M} for unit-demand agents is the following:

³ An improved analysis of the surplus of the all-pay auction is available by proving a version of Theorem 6.2 for all-pay-style payment rules. See Exercise 6.1.

- Agent i 's action space in \mathcal{M} is $A_i = A_i^1 \times \dots \times A_i^m$ where A_i^j is agent i 's action space for mechanism \mathcal{M}^j .
- On action profile $\mathbf{a} = (\mathbf{a}^1, \dots, \mathbf{a}^m)$ with $\mathbf{a}^j = (a_1^j, \dots, a_n^j)$, the outcome of the mechanism is $\mathcal{M}(\mathbf{a}) = (\mathcal{M}^1(\mathbf{a}^1), \dots, \mathcal{M}^m(\mathbf{a}^m))$.
- The action allocation rule is $\mathbf{x}^{\mathcal{M}}(\mathbf{a})$ with $x_i^{\mathcal{M}}(\mathbf{a}) = \max_j x_i^{\mathcal{M}^j}(\mathbf{a}^j)$.
- The action payment rule is $\mathbf{p}^{\mathcal{M}}(\mathbf{a})$ with $p_i^{\mathcal{M}}(\mathbf{a}) = \sum_j p_i^{\mathcal{M}^j}(\mathbf{a}^j)$.

Theorem 6.8. *Revenue covering approximation is closed under simultaneous composition; i.e., if mechanisms $\mathcal{M}^1, \dots, \mathcal{M}^m$ are downward closed, individually rational, and have revenue covering approximation μ ; then their simultaneous composite mechanism \mathcal{M} has revenue covering approximation μ .*

The following two lemmas, implied by downward closure and individual rationality, respectively, enable the proof of Theorem 6.8.

Lemma 6.9. *For the union environment of m downward-closed environments, allocation \mathbf{x} is feasible if and only if there exists $\mathbf{x}^1, \dots, \mathbf{x}^m$ feasible for the individual environments that satisfy $x_i = \sum_j x_i^j$ for all i and j .*

Proof. By definition of feasibility in the union environment, if $\mathbf{x}^1, \dots, \mathbf{x}^m$ are feasible for the environment of $\mathcal{M}^1, \dots, \mathcal{M}^m$, respectively, then

$$x_i = \max_j x_i^j \quad (6.3)$$

is feasible for the union environment of \mathcal{M} . Moreover, by downward closure of each individual mechanism \mathcal{M}^j if \mathbf{x} is feasible, then there exists $\mathbf{x}^1, \dots, \mathbf{x}^m$ with each \mathbf{x}^j feasible for \mathcal{M}^j and

$$x_i = \sum_j x_i^j \quad (6.4)$$

for all i and j . We are able to replace the maximization in equation (6.3) with the summation in equation (6.4) because downward closure allows the summation to be reduced to the maximum by removing service from an agent in all but at most one of the individual mechanisms. \square

Lemma 6.10. *For the simultaneous composite mechanism \mathcal{M} of m individually rational mechanisms $\mathcal{M}^1, \dots, \mathcal{M}^m$, any agent i , and any effective winner-pays-bid $b \in \mathbb{R}$,*

- Agent i 's allocation probability with effective winner-pays-bid b is greater in \mathcal{M} than in \mathcal{M}^j for any j , i.e., $\tilde{x}_i(b) \geq \tilde{x}_i^j(b)$.

- (ii) Agent i 's expected critical bid is smaller in \mathcal{M} than in \mathcal{M}^j for any j , i.e., $\hat{B}_i \leq \hat{B}_i^j$.

Proof. Fix any agent i . The pointset of equation (6.3) that defines the winner-pays-bid allocation rule for i in \mathcal{M} contains that of \mathcal{M}^j for all j as one allowable bid in \mathcal{M} is to bid only in \mathcal{M}^j (by individual rationality of the other mechanisms). As such, the smallest monotone function that contains this pointset is higher for \mathcal{M} than for \mathcal{M}^j , i.e., $\tilde{x}_i(b) \geq \tilde{x}_i^j(b)$ for all b . As \hat{B} and \hat{B}^j are defined as the area above winner-pays-bid allocation rules \tilde{x} and \tilde{x}^j , the former is smaller than the latter. \square

Proof of Theorem 6.8. Consider feasible allocation \mathbf{y} for the composite mechanism and the following sequence of inequalities with explanation below.

$$\begin{aligned} \mu \mathbf{E}[\text{Revenue}] &\geq \sum_j \mu \mathbf{E}[\text{Revenue}_j] \\ &\geq \sum_j \sum_i \hat{B}_i^j y_i^j \\ &\geq \sum_i \hat{B}_i \sum_j y_i^j \\ &= \sum_i \hat{B}_i y_i. \end{aligned}$$

The first line follows from the definition of revenue as the sum of payments from all agents in all mechanisms. By Lemma 6.9 and the feasibility of \mathbf{y} there exists $\mathbf{y}^1, \dots, \mathbf{y}^m$ which are feasible for $\mathcal{M}^1, \dots, \mathcal{M}^m$, respectively, and satisfy $y_i = \sum_j y_i^j$. The second line follows from revenue covering of \mathcal{M}^j for each j with respect to \mathbf{y}^j . Swapping the order of summation and employing the lower bound of $\hat{B}_i \leq \hat{B}_i^j$ from Lemma 6.10 for all i and j gives the third line. The fourth line is from the definition of $\mathbf{y}^1, \dots, \mathbf{y}^m$ in terms of \mathbf{y} . We are left with the inequality that shows that \mathcal{M} has revenue covering approximation μ . \square

6.4 Reserve Prices

We will shortly be analyzing the revenue of Bayes-Nash mechanisms like the first-price auction. As we understand from Chapter 3 and Chapter 4, reserve prices play an important role in revenue maximization. According to the previous definition of revenue covering approximation (Definition 6.4), auctions with reserve prices are not generally approximately revenue covered. Revenue covering arguments stem from relating the critical bid of an agent to potential payments of other agents. For

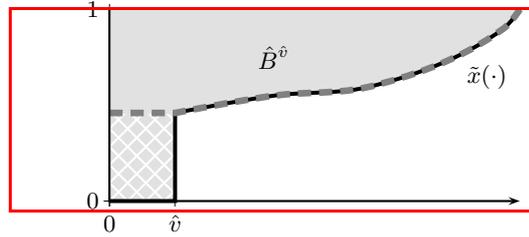


Figure 6.5. Geometry of best-response in a winner-pays-bid auction with reserve \hat{v} . The expected critical bid \hat{B} is the area (light gray solid and crosshatched) above the bid allocation rule $\tilde{x}(\cdot)$ (thin solid line), the expected critical bid with discounted reserve is the area (light gray solid) above its cumulative distribution function $\hat{x}^{\hat{v}}(\cdot)$ (thick dashed line).

example, in the first-price auction the critical bid of an agent is the maximum bid of the other agents, and if this agent does not bid above this critical bid then this maximum bid of the others is equal to the auction revenue. With a reserve price, an agent's critical bid may come from either bids of others or the reserve price. When the agent does not bid above her reserve price, the reserve price does not translate into auction revenue. See Figure 6.5.

In this section we alter the framework of analysis to account for reserve prices. As is evident from Figure 6.5, the critical bid \hat{B} as the area above the bid allocation rule $\tilde{x}(\cdot)$ over counts the contribution to revenue from the agent's critical bid. One resolution to this over counting is to explicitly discount the contribution to \hat{B} from the reserve. The following definition captures this idea. Recall the bid allocation rule is equivalently the distribution function for the critical bid; thus, to discount the reserve is to assume the critical bid is zero whenever it would otherwise be the reserve.

Definition 6.6. The *critical bid with discounted reserve* is

$$\hat{b}^{\hat{v}} = \begin{cases} 0 & \text{if } \hat{b} \leq \hat{v}, \text{ and} \\ \hat{b} & \text{otherwise.} \end{cases}$$

The cumulative distribution function for the *critical bid with discounted reserve* \hat{v} is $\hat{x}^{\hat{v}}(b) = \tilde{x}(\max(b, \hat{v}))$; see Figure 6.5; its expected value is:

$$\hat{B}^{\hat{v}} = \mathbf{E}[\hat{b}^{\hat{v}}] = \int_0^{\infty} (1 - \tilde{x}(\max(b, \hat{v}))) db.$$

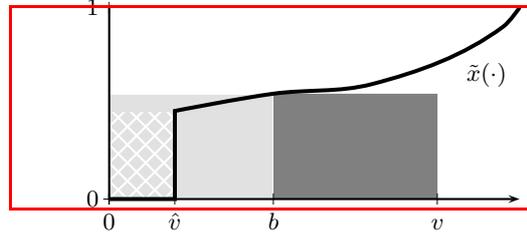


Figure 6.6. Geometric demonstration of equation (6.7). The expected discounted reserve $\hat{B} - \hat{B}^{\hat{v}}$ (light gray crosshatched) is at most the expected payment from bid b (light gray solid and crosshatched). The surplus from bid b is the utility (dark gray area) plus the expected payment.

6.4.1 Surplus Approximates Value

We now lift the utility approximation of value result of Theorem 6.2 for mechanisms without reserves to mechanisms with reserves.

Theorem 6.11. *In any Bayes-Nash equilibrium of any mechanism and for any agent with value v exceeding her reserve \hat{v} , the expected sum of her surplus and her critical bid with discounted reserve is an $\frac{e}{e-1} \approx 1.58$ approximation to her value; i.e.,*

$$v x(v) + \hat{B}^{\hat{v}} \geq e^{-1/e} v. \quad (6.5)$$

Proof. Theorem 6.2 states

$$u(v) + \hat{B} \geq e^{-1/e} v. \quad (6.6)$$

In BNE, an agent with $v \geq \hat{v}$ will bid $b \geq \hat{v}$ as any lower bid results in zero utility. Recall that the expected payment of an agent with equilibrium bid b is $p(v) = b \tilde{x}(b)$; geometrically as $b \geq \hat{v}$ this payment exceeds the amount of \hat{B} discounted by the reserve (Figure 6.6). Thus,

$$p(v) + \hat{B}^{\hat{v}} \geq \hat{B}. \quad (6.7)$$

Recall that surplus is utility plus payment, i.e., $v x(v) = u(v) + p(v)$. The proof concludes by adding equation (6.6) to (6.7). \square

6.4.2 Revenue Covering Approximation

For the appropriate definition of revenue covering approximation with reserves, revenue covering without reserves implies revenue covering with reserves.

Definition 6.7. A mechanism with reserves has revenue covering approximation μ if the revenue covering approximation condition (Definition 6.4) holds with respect to expected critical bids with discounted reserves.

Theorem 6.12. *Revenue covering approximation is closed under reserve pricing; i.e., if a mechanism \mathcal{M} without reserves has revenue covering approximation μ , then with reserves it has revenue covered approximation μ .*

Proof. The revenue covering condition with reserves is only weaker as $\hat{B}_i^{\hat{v}} \leq \hat{B}_i$ for all agents i . \square

6.4.3 Social Surplus in Bayes-Nash Equilibrium

Theorem 6.13. *For any individually-rational mechanism with reserves that has revenue covering approximation $\mu \geq 1$, the expected social surplus in Bayes-Nash equilibrium is an $e/e-1(1 + \mu)$ approximation to the optimal social surplus with the same reserves.*

Proof. Denote the reserves by $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$. Consider a valuation profile \mathbf{v} . By Theorem 6.11, in BNE any agent i with $v_i \geq \hat{v}_i$ satisfies,

$$v_i x_i(v_i) + \hat{B}_i^{\hat{v}_i} \geq e^{-1/e} v_i.$$

Denote by $\mathbf{y}^*(\mathbf{v}) = \operatorname{argmax}_{\mathbf{y}} \sum_{i: v_i \geq \hat{v}_i} v_i y_i$ the surplus optimizing allocation with reserves $\hat{\mathbf{v}}$ (with $y_i^*(\mathbf{v}) = 0$ for i with $v_i < \hat{v}_i$). Thus, $\sum_i v_i y_i^*(\mathbf{v}) = \operatorname{REF}(\mathbf{v})$, the optimal social surplus with reserves $\hat{\mathbf{v}}$. Notice that $y_i^*(\mathbf{v}) \in [0, 1]$; thus,

$$v_i x_i(v_i) + \hat{B}_i^{\hat{v}_i} y_i^*(\mathbf{v}) \geq e^{-1/e} v_i y_i^*(\mathbf{v}).$$

The above equation was derived for agent i with $v_i \geq \hat{v}_i$; however, it holds trivially for i with $v_i < \hat{v}_i$ as $y_i^*(\mathbf{v}) = 0$ for such agents. Sum over all agents i and invoke μ revenue covering,

$$\sum_i v_i x_i(v_i) + \mu \mathbf{E}[\operatorname{Revenue}] \geq e^{-1/e} \operatorname{REF}(\mathbf{v}).$$

Take expectation over values \mathbf{v} from the distribution \mathbf{F} ,

$$\mathbf{E}[\operatorname{Surplus}] + \mu \mathbf{E}[\operatorname{Revenue}] \geq e^{-1/e} \mathbf{E}[\operatorname{REF}(\mathbf{v})].$$

The surplus of an individually-rational mechanism always exceeds its revenue; the theorem follows. \square

An example consequence of Theorem 6.13 is the following. Moreover, analogous corollaries hold for the winner-pays-bid highest-bids-win matroid mechanism and the simultaneous composition of revenue covered mechanisms.

Corollary 6.14. *For any product distribution on values, first-price auction with reserves has Bayes-Nash equilibrium surplus that is an $2e/e-1 \approx 3.16$ approximation to the optimal surplus with the same reserves.*

This approach for treating reserves applies to any mechanism that can be interpreted as having a reserve price. Importantly, our definition of reserves is in value space; while reserves, in the definition of a mechanism, bind in bid space. For the first-price auction and the simultaneous composition thereof, these are the same thing. For all-pay auctions, however, the value at which a bid-based reserve binds is endogenous to the equilibrium. For all-pay auctions, any bid-based reserves and BNE induce value-based reserves for which Theorem 6.13 holds.

6.5 Analysis of Revenue

We will adapt the framework for Bayes-Nash analysis of the surplus of mechanisms with reserves to analyze the revenue of Bayes-Nash equilibrium in mechanisms with monopoly reserves. Recall from Chapter 3 that the expected payment in BNE (and thus revenue) from an agent with value $v \sim F$ satisfies $\mathbf{E}_v[p(v)] = \mathbf{E}_v[\phi(v)x(v)]$ with virtual value function given by $\phi(v) = v - \frac{1-F(v)}{f(v)}$ (see Section 3.3.1 on page 61). The approach will be to adapt Theorem 6.11, which bounds an agent's BNE surplus in terms of her value, to bound an agent's BNE virtual surplus in terms of her virtual value. Our analysis is necessarily restricted to regular distributions where the virtual value function $\phi(\cdot)$ given above is monotone non-decreasing (see Definition 3.4 on page 64)

Theorem 6.15. *In any Bayes-Nash equilibrium of any mechanism and for any agent with value v exceeding her reserve \hat{v} and with non-negative virtual value $\phi(v)$, the expected sum of her virtual surplus and her critical bid with discounted reserve is an $e/e-1 \approx 1.58$ approximation to her virtual value; i.e.,*

$$\phi(v)x(v) + \hat{B}^{\hat{v}} \geq e^{-1/e}\phi(v). \quad (6.8)$$

Proof. The definition of virtual values for revenue as $\phi(v) = v - \frac{1-F(v)}{f(v)}$

implies that $v \geq \phi(v)$ or, in other words, $\phi(v)/v \leq 1$. Thus, relative to the surplus and value terms of inequality (6.5) of Theorem 6.11, the virtual-surplus and virtual-value terms of (6.8) are scaled downward. Equivalently, the expected-critical-bid term on the right-hand side is relatively scaled upward. Thus, the inequality (6.8) of the present theorem is implied by Theorem 6.11. \square

The following theorem is proved as was Theorem 6.13 but with the following key differences. The proof begins with the virtual surplus approximation of virtual value bound of Theorem 6.15 instead of the analogous bound of Theorem 6.11. It finishes by observing, as virtual surplus and revenue are equal in expectation, that expected virtual surplus plus expected revenue is exactly twice the expected revenue. Additionally, the theorem is stated for monopoly reserves and agents with regular distributions which necessarily excludes from analysis agents with negative virtual value.

Theorem 6.16. *For agents with regularly distributed values and any mechanism with monopoly reserves that has revenue covering approximation $\mu \geq 1$, the expected revenue in Bayes-Nash equilibrium is an $e/e-1 (1 + \mu)$ approximation to the optimal revenue.*

Again, this theorem can be applied to any of the revenue covered mechanisms previously discussed. The following corollary is for the first-price auction, there are similar corollaries for the winner-pays-bid highest-bids-win matroid mechanism and the simultaneous composition of mechanisms.

Corollary 6.17. *For any regular product distribution on values, the first-price auction with monopoly reserves has Bayes-Nash equilibrium revenue that is an $2e/e-1 \approx 3.16$ approximation to the optimal revenue.*

In Section 5.2 on page 158 we saw that with sufficient competition the surplus maximizing mechanism (without reserves) approximates the revenue optimal mechanism (e.g., Theorem 5.4). Similar sufficient competition results extend to revenue covered mechanisms. One such definition of sufficient competition is that there are at least two agents from each distribution that are in direct competition with each other. The following theorem is an example.

Theorem 6.18. *For any regular product distribution on values with at least two agents with values drawn from each distinct distribution,*

the first-price auction has Bayes-Nash equilibrium revenue that is an $3e/e-1 \approx 4.75$ approximation to the optimal revenue.

Proof. See Exercise 6.6. □

6.6 Revenue Covering Optimization

We have seen that the revenue covering approximation of a mechanism governs its Bayes-Nash approximation with respect to both social surplus and revenue. We now consider the problem of optimizing the rules of a mechanism to minimize its revenue covering approximation. The motivating example will be that of single-minded combinatorial auctions. We saw that the winner-pays-bid highest-bids-win mechanism for m -item single-minded combinatorial auctions is not a revenue covering approximation of μ for any $\mu < m$ (Example 6.2). Faced with this negative result, the question remains to identify a winner-pays-bid mechanism that obtains a non-trivial revenue covering approximation. Importantly, such a mechanism will have to choose a suboptimal, in terms of sum of bids, set of winners.

The running example for this section will be a single-minded combinatorial auction environment for n agents and m items. Each agent i has value v_i for obtaining bundle $S_i \subset [m]$. Two agents that desire the same item, i.e., i and i^\dagger with $S_i \cup S_{i^\dagger} \neq \emptyset$, cannot simultaneously be served. The section culminates by showing that a winner-pays-bid mechanism based on a simple greedy heuristic has a revenue covering approximation of \sqrt{m} .

6.6.1 Non-bossiness, Approximation, and Greedy Algorithms

The difficulty of single-minded combinatorial auctions is that one agent can block many other agents that could be simultaneously served. It could be optimal to serve the blocked agents, but in equilibrium the blocking agent bids enough to dissuade any of the blocked agents from individually deviating to win. In Example 6.2 this situation was exhibited with one agent demanding the grand bundle $[m]$ and many agents each demanding a single item; the grand-bundle agent then blocked all the singleton agents. When the grand-bundle agent bids 1, and the singleton agents bid 0, then the deviation bid that any singleton agent

must make to win is 1. Since their values in the example are 1, this deviation does not improve the singleton agent's utility. Of course, the singleton agents would win if the sum of their bids exceeds the grand-bundle agent's bid of 1. Thus, as one of the singleton agent increases her bid — though, all other bids unchanged, she continues to lose — the critical bids of all other singleton agents are reduced. This bad property is precisely what inhibits revenue covering approximation. The following definition formalizes the non-exhibition of this property.

Definition 6.8. A mechanism is *subcritically non-bossy* if for any bid profile \mathbf{b} , critical bids $\hat{\mathbf{b}}$, and any other bid profile where losers may increase their bids up to their critical bids, i.e., \mathbf{b}^\dagger with $b_i^\dagger \in [b_i, \max(b_i, \hat{b}_i)]$, the same set of agents win under \mathbf{b} and \mathbf{b}^\dagger .⁴

To solve the combinatorial auction problem we are going to have to replace the highest-bids-win allocation rule with an allocation rule that does not maximize the sum of the bids of the agents served. There are two potential losses from such an allocation rule. First, there is the direct loss from the fact that the allocation rule chooses a suboptimal set of bids. Even if there is a feasible set of agents with high bid sum, its revenue could be low. Second, there is the indirect loss from strategization on the part of the agents. The highest-bids-win allocation rule suffers no direct losses, but prohibitively in indirect losses. On the other hand, the first-price auction for the grand bundle, i.e., where only one agent ever wins her desired bundle, suffers prohibitive direct losses but, as the first-price auction is revenue covered, suffers no indirect losses with respect to the optimal mechanism that only serves one agent). Ideally both direct and indirect losses should be kept small. The following definition formalizes a bound on the direct loss in terms of approximation.

Definition 6.9. A mechanism (APX) with ex post bid allocation rule $\tilde{\mathbf{x}}(\mathbf{b})$, which maps a profile of bids to an allocation, is a β *approximation* to highest-bids-win (REF) if

$$\text{APX}(\mathbf{b}) = \sum_i b_i \tilde{x}_i(\mathbf{b}) \geq 1/\beta \max_{\mathbf{x}} b_i x_i = \text{REF}(\mathbf{b}).$$

We now show that in a subcritically non-bossy mechanism the only

⁴ This definition adopts the convention that ties in the bid allocation rule, when any loser increases her bid to equal her critical bid, are broken in favor of the current winners. The arguments below can be made without this tie-breaking convention by considering \mathbf{b}^\dagger with losers bidding $b_i^\dagger \in [b_i, \max(b_i, \hat{b}_i - \epsilon)]$ for an arbitrarily small ϵ .

loss in surplus is the direct loss from the non-optimality of the bid allocation rule, i.e., there is no indirect loss.

Theorem 6.19. *A winner-pays-bid subcritically non-bossy mechanism that is a β approximation to highest-bids-win has a revenue covering approximation of $\mu = \beta$.*

Proof. Fix a bid profile \mathbf{b} , the critical bid profile $\hat{\mathbf{b}}$, and any feasible allocation \mathbf{y} . Denote the bid allocation rule by $\tilde{\mathbf{x}}(\mathbf{b}) = (\tilde{x}_1(\mathbf{b}), \dots, \tilde{x}_n(\mathbf{b})) \in \{0, 1\}^n$. Denote the maximum subcritical bid profile \mathbf{b}^\dagger with $b_i^\dagger = \max(b_i, \hat{b}_i)$. Subcritical non-bossiness requires allocation to be unchanged if all losers increase their bids to their critical values, i.e., $\tilde{\mathbf{x}}(\mathbf{b}^\dagger) = \tilde{\mathbf{x}}(\mathbf{b})$.

The following sequence of equations implies that the mechanism has revenue covering approximation $\mu = \beta$; formal justification for each equation follows.

$$\begin{aligned} \text{Revenue}(\mathbf{b}) &= \sum_i b_i \tilde{x}_i(\mathbf{b}) \\ &= \sum_i b_i^\dagger \tilde{x}_i(\mathbf{b}) \\ &= \sum_i b_i^\dagger \tilde{x}_i(\mathbf{b}^\dagger) \\ &\geq \frac{1}{\beta} \sum_i b_i^\dagger y_i \\ &\geq \frac{1}{\beta} \sum_i \hat{b}_i y_i. \end{aligned}$$

The first equation is by definition of winner-pays-bid mechanisms. The second equation is the equality of b_i and $b_i^\dagger = \max(b_i, \hat{b}_i)$ where winning ($\tilde{x}_i(\mathbf{b}) = 1$) implies $b_i \geq \hat{b}_i$. The third equation is by subcritical non-bossiness, as discussed above. The fourth equation follows by the β -approximation optimality of $\tilde{\mathbf{x}}(\cdot)$ on \mathbf{b}^\dagger . The fifth and final equation follows from the definition of $b_i^\dagger = \max(b_i, \hat{b}_i) \geq \hat{b}_i$. We conclude that the mechanisms has a revenue covering approximation of $\mu = \beta$. \square

Theorem 6.19 shows that to find winner-pays-bid mechanisms that are revenue covered it suffices to find a subcritically non-bossy mechanism that is a good approximation to highest-bids-win. A greedy algorithm is one that sort the agents by some priority and then serve each agent if it is feasible to do so given the agents previously served by the algorithm. Greedy algorithms are a standard design methodology in the field of approximation algorithms and they have important consequences for mechanism design. For example, we saw in Section 4.6 on page 129 that greedy algorithms are optimal in ordinal environments such as those

given by a matroid set system. Subsequently, we will see that greedy algorithms are approximately optimal in some environments and mechanisms based on them are winning-bids non-bossy.

Definition 6.10. For any downward-closed environment and any profile of priority functions $\vartheta = (\vartheta_1, \dots, \vartheta_n)$, the *greedy-by-priority algorithm*:

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

Theorem 6.20. *The greedy-by-priority bid allocation rule is subcritically non-bossy.*

Proof. Fix a profile of bids \mathbf{b} , critical bids $\hat{\mathbf{b}}$, and maximum subcritical bid profile \mathbf{b}^\dagger with $b_i^\dagger = \max(b_i, \hat{b}_i)$. Consider varying a single losing bid i on the range $[0, \hat{b}_i]$ and simulating the algorithm. Wherever this bid arises in the sorted order of agents by priority, since $b_i \leq \hat{b}_i$, it must be infeasible to serve the agent. Thus, this agent is discarded and all decisions by the algorithm to serve or not to serve any other agents are unaffected. The same holds for all losing agents simultaneously. For any winning i , $b_i^\dagger = b_i$ which is unchanged; for any losing agent i , $b_i^\dagger = \hat{b}_i$ which is unchanged. Thus, the bid allocation rule is subcritically non-bossy. \square

6.6.2 Single-minded Combinatorial Auctions

We now instantiate the approach of the preceding section to design a single-minded combinatorial auction that has a non-trivial revenue covering approximation. Theorem 6.19 and Theorem 6.20 imply that to find a winner-pays-bid mechanism that is revenue covered, it suffices to identify a profile of priority functions ϑ such that the greedy-by-priority algorithm obtains a good approximation to highest-bids-win. We now consider this task and identify a priority for which greedy-by-priority is a $\beta = \sqrt{m}$ approximation and, thus, has a revenue covering approximation of $\mu = \sqrt{m}$.

We begin by considering two extremal approaches, both of which yield only m approximation, and then look at trading off these extremes to get the desired \sqrt{m} approximation. The first failed approach to consider

is *greedy by bid*, i.e., the prespecified sorting criterion in the static greedy template above is by agent bids, i.e., the priority function is the identity $\vartheta_i(b_i) = b_i$. This algorithm is bad because it is an m approximation on the following $n = m + 1$ agent input. Agents i , for $0 \leq i \leq m$, have $S_i = \{i\}$ and $b_i = 1$; agent $m+1$ has $b_{m+1} = 1+\epsilon$ and demands the grand bundle $S_{m+1} = \{1, \dots, m\}$ (for some small $\epsilon > 0$). See Figure 6.7(a) with $A = 1$ and $B = 1 + \epsilon$. Greedy-by-bid orders agent $m + 1$ first, this agent is feasible and therefore served. All remaining agents are infeasible after agent $m + 1$ is served. Therefore, the algorithm serves only this one agent and has surplus $1 + \epsilon$. Of course highest-bids-win serves the m small agents for a total surplus of m . The approximation factor of greedy-by-bid is the ratio of these two performances, i.e., m .

Obviously what went wrong in greedy-by-bid is that we gave preference to an agent with large demand who then blocked a large number of mutually-compatible small-demand agents. We can compensate for this by instead sorting by bid-per-size, i.e., $\vartheta(b_i) = b_i/|S_i|$. *Greedy by bid-per-size* also fails on the following $n = 2$ agent input. Agent 1 has $S_1 = \{1\}$ and $b_1 = 1 + \epsilon$ and agent 2 has $b_2 = m$ demands the grand bundle $S_2 = \{1, \dots, m\}$. See Figure 6.7(b) with $A = 1 + \epsilon$ and $B = m$. Greedy-by-bid-per-item orders agent 1 first, this agent is feasible and therefore served. Agent 2 is infeasible after agent 1 is served. Therefore, the algorithm serves only agent 1 and has surplus $1+\epsilon$. Of course highest-bids-win serves agent 2 and has surplus of m . The approximation factor of greedy-by-bid-per-item is the ratio of these two performances, i.e., m .

The flaw with this second algorithm is that it makes the opposite mistake of the first algorithm; it undervalues large-demand agents. While we correctly realized that we need to trade off bid for size, we have only considered extremal examples of this trade-off. To get a better idea for this trade-off, consider the cases of a single large-demand agent and either m small-demand agents or 1 small-demand agent. We will leave the bids of the two kinds of agents as variables A for the small-demand agent(s) and B for the large-demand agent. Assume, as in our previous examples, that $mA > B > A$. These settings are depicted in ??.

Notice that any greedy algorithm that orders by some function of bid and size will either prefer A -bidding or B -bidding agents in both cases. The A -preferred algorithm has surplus Am in the m -small-agent case and surplus A in the 1-small-agent case. The B -preferred algorithm has surplus B in both cases. The Highest-bids-win outcome, on the other hand, has surplus mA in the m -small-agent case and surplus B in the 1-small-agent case. Therefore, the worst-case approximation for A -preferred is

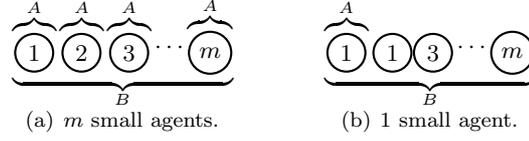


Figure 6.7. Challenge cases for greedy orderings as a function of bid and bundle size.

B/A (achieved in the 1-small-agent case), and the worst-case approximation for B -preferred is mA/B (achieved in the m -small-agent case). These performances and worst-case ratios are summarized in Figure 6.8.

	m small agents	1 small agent	approximation
highest-bids-win	mA	B	1
A -preferred	mA	A	B/A
B -preferred	B	B	mA/B

Figure 6.8. Performances of A - and B -preferred greedy algorithms and their approximation to highest-bids-win in worst-case over the two cases.

If we are to use the greedy algorithm design paradigm we need to minimize the worst-case ratio. The approach suggested by the analysis of the above cases would be trade off A versus B to equalize the worst-case approximation, i.e., when $B/A = mA/B$. Here m was a stand-in for the size of the large-demand agent. The suggested algorithm is greedy by bid-per-square-root-size which orders the agents by the priority $\vartheta(b_i) = b_i/\sqrt{|S_i|}$. The tradeoff above can be observed explicitly in the in the proof of Theorem 6.21, below.

Theorem 6.21. *For m -item single-minded combinatorial auctions environments, the greedy by bid-per-square-root-size algorithm is a $\beta = \sqrt{m}$ approximation to highest-bids-win.*

Proof. Let APX represent the greedy by bid-per-square-root-size algorithm and its surplus; let REF represent the optimal algorithm and its surplus. Let I be the set selected by APX and I^* be the set selected by REF. We will proceed with a *charging argument* to show that if $i \in I$ blocks some set of agents $C_i \subset I^*$ then the sum of bids of the blocked agents is not too large relative to the bid of agent i .

Consider the agents sorted (as in APX) by $b_i/\sqrt{|S_i|}$. For an agent $i^* \in I^*$ not to be served by APX, it must be that at the time it is considered by the greedy algorithm, another agent i has already been selected that

blocks i^* , i.e., the bundles S_i and S_{i^*} have non-empty intersection. Intuitively we will charge one such agent i with the loss from not accepting agent i^* . We define C_i as the set of all $i^* \in I^*$ that are charged to i as described above. Of special note, if $i^* \in I$, i.e., it was not yet blocked when considered by APX, we charge it to itself, i.e., $C_{i^*} = \{i^*\}$. Notice that the sets C_i for winners $i \in I$ of APX partition the winners I^* of REF.

The theorem follows from the inequalities below. Explanations of each non-trivial step are given afterwards.

$$\text{REF} = \sum_{i^* \in I^*} b_{i^*} = \sum_{i \in I} \sum_{i^* \in C_i} b_{i^*} \quad (6.9)$$

$$\leq \sum_{i \in I} \frac{b_i}{\sqrt{|S_i|}} \sum_{i^* \in C_i} \sqrt{|S_{i^*}|} \quad (6.10)$$

$$\leq \sum_{i \in I} \frac{b_i}{\sqrt{|S_i|}} \sum_{i^* \in C_i} \sqrt{m/|C_i|} \quad (6.11)$$

$$= \sum_{i \in I} \frac{b_i}{\sqrt{|S_i|}} \sqrt{m|C_i|} \quad (6.12)$$

$$\leq \sum_{i \in I} b_i \sqrt{m} = \sqrt{m} \cdot \text{APX}. \quad (6.13)$$

Line (6.9) follows because C_i partition I^* . Line (6.10) follows because $i^* \in C_i$ implies that i precedes i^* in the greedy ordering and therefore $b_{i^*} \leq b_i \sqrt{|S_{i^*}|}/\sqrt{|S_i|}$. The demand sets S_{i^*} of $i^* \in C_i$ are disjoint (because they are a subset of I^* which is feasible and therefore disjoint). Thus, we can bound $\sum_{i^* \in C_i} |S_{i^*}| \leq m$. The square-root function is concave and the sum of a concave function is maximized when each term is equal, i.e., when $|S_{i^*}| = m/|C_i|$. Therefore, $\sum_{i^* \in C_i} \sqrt{|S_{i^*}|} \leq \sum_{i^* \in C_i} \sqrt{m/|C_i|}$ and line (6.11) follows. Line (6.12) follows from independence of the inner summand on i^* . Finally, line (6.13) follows because the bundle S_{i^*} of each agent $i^* \in C_i$ is disjoint but contain some demanded item in S_i and, therefore, $|C_i| \leq |S_i|$. \square

We conclude the section with the following corollary. The first part is a consequence of Theorem 6.19, Theorem 6.20, and Theorem 6.21. The second part is a consequence of the first part and Theorem 6.5. The third part is a consequence of the first part and Theorem 6.16.

Corollary 6.22. *For m -item single-minded combinatorial auction environments, the winner-pays-bid greedy-by-value-per-square-root-size mechanism has revenue covering approximation $\mu = \sqrt{m}$; its surplus in Bayes-Nash equilibrium is an $e/(e-1)\sqrt{m}$ approximation to the optimal surplus; and with monopoly reserves and regular distributions its revenue is an $e/(e-1)(1 + \sqrt{m})$ approximation to the optimal revenue.*

Exercises

- 6.1 Consider an all-pay auction and show an analogous utility value covering to Theorem 6.2. Specifically, in BNE,

$$u(v) + \hat{B} \geq 1/2 v,$$

where \hat{B} is the expected critical bid of the agent. Combine this result with revenue covering (with respect to the all-pay-bid allocation rule) to show that the expected social surplus of the all pay auction is a two approximation to the optimal social surplus.

- 6.2 Consider the single-dimensional routing environment discussed in Section 1.1.3 on page 14 where there is a graph $G = (V, E)$, each agent i has a message to send from source vertex $s_i \in V$ to target vertex $t_i \in V$ (public knowledge) and a private value v_i for sending such a message. A feasible outcome is given by an edge disjoint collection of paths in the graph. Show that the winner-pays-bid highest-bids-win mechanism is not $\mu \leq d$ revenue covered where d is the *diameter* of the graph, i.e., the maximum over pairs of vertices of the shortest path between the pair.
- 6.3 Show that the winner-pays-bid highest-bids-win auction for matroid environments is 1 revenue covered, i.e., prove Theorem 6.4.
- 6.4 Consider the single-minded combinatorial auction problem of Example 6.2. The optimization problem of selecting the feasible set of agents with the highest sum of bids corresponds to the *weighted set packing* problem which is NP-hard (cf. Section 1.1.3 and ??). The following greedy algorithm is known to be a \sqrt{m} approximation, i.e., it always finds a feasible subset of agents with bids that sum to at least a \sqrt{m} fraction of the sum of the optimal feasible set of bids (see ?? on page ??).
- (a) Sort the bids b_i by $b_i/|S_i|$.
 - (b) Considering the bids in this order, accept a bid if it is feasible with previously accepted bids.

Prove that the mechanism that selects winners with this greedy algorithm and charges each winner her bid has revenue covering approximation \sqrt{m} .

- 6.5 Show that with respect to the effective winner-pays-bid allocation rule (Definition 6.3) the all-pay auction is 2 revenue covered, i.e., prove Theorem 6.7.
- 6.6 Prove Theorem 6.18. Consider a single-item environment with agent values drawn from regular distributions with least two agents with

values drawn from each distinct distribution. Show that the first-price auction has Bayes-Nash equilibrium revenue that is an $3e/e-1 \approx 4.75$ approximation to the optimal revenue.

Chapter Notes

Vickrey (1961) posed the question of solving for the equilibrium in the first-price auction and two agents with values drawn from the uniform distribution with asymmetric supports. The solution when the lower bound of the supports is the same, as in the $U[0, 1]$ and $U[0, 2]$ case of Example 6.1, was given by Griesmer et al. (1967). The general case of two agents with arbitrary uniform distributions was solved by Kaplan and Zamir (2012).

The quantification of the disutility of equilibrium versus the social surplus maximizing outcome is known as the *price of anarchy*. This topic of study was initially proposed by Koutsoupias and Papadimitriou (1999). It was applied to (full information) congestion games by Roughgarden and Tardos (2002), cf. the routing game of Section 1.1 on page 2. Roughgarden (2012a) abstracted the canonical price of anarchy analysis as what is referred to as the *smoothness framework*. Roughgarden (2012b) and Syrgkanis and Tardos (2013) generalize this smoothness framework to games of incomplete information and auctions, respectively. There has been extensive study of the price of anarchy of specific auction games to which detailed reference is omitted. This text focuses on an adaptation of the smoothness paradigm to single-dimensional agents that was given by Hartline et al. (2014).

The proof that the sum of utility and critical bid approximate an agent's value for first-price auctions that is given in this text is from Syrgkanis and Tardos (2013); an alternative geometric argument can be found in Hartline et al. (2014). The improved analysis of the all-pay auction of Exercise 6.1 is based on Syrgkanis and Tardos (2013). A smoothness framework for analyzing the simultaneous composition of auctions was first given by Syrgkanis and Tardos (2013); the analysis given here is the refinement of Hartline et al. (2014) for single-dimensional agents. The analysis of revenue in Bayes-Nash equilibrium is from Hartline et al. (2014).

The relationship between revenue covering approximation and greedy algorithms is a recasting of the main result of Lucier and Borodin (2010) into the analysis framework of Hartline et al. (2014).

The analysis of Syrgkanis and Tardos (2013) is more general than the one presented here primarily in that it allows for multi-dimensional agent preferences. They also give numerous results that are not covered here, one such result is for the sequential composition of mechanisms, i.e., when mechanisms are run one after the other.

7

Prior-free Mechanisms

In Chapter 3 we derived optimal mechanisms for social surplus and profit. For social surplus, the surplus maximization mechanism (Definition 3.3, page 58) is optimal pointwise on all valuation profiles. For profit, the virtual surplus maximization mechanism (Definition 3.5, page 65) is optimal in expectation for values drawn from the given distribution. The difference between the statement of these results is significant: for social surplus there is a pointwise optimal mechanism whereas optimal mechanisms for expected profit are parameterized by the distribution from which values are drawn. The goal of this chapter is to design mechanisms that obtain approximately optimal profit pointwise on all valuation profiles.

As an example, consider a digital good environment with $n = 100$ agents. Consider first the valuation profile where agent i has value $v_i = i$ for all i . How much revenue could a mechanism hope to obtain in such an environment? For example, this valuation profile seems similar to the uniform distribution on $[0, 100]$ for which the Bayesian optimal mechanism would post a price of 50 and obtain an expected revenue of $2500 = 50 \times 50$. Consider second the valuation profile where all agents have value one. This valuation profile seems similar to a pointmass distribution where the Bayesian optimal mechanism post a price of one for a revenue of 100. Can we come up with one mechanism that on the first profile obtains revenue close to 2500 and on the second profile obtains revenue close to 100? Moreover, what is an appropriate target revenue in general and is there an auction that approximates this target? These are the questions we address in this chapter.

The main difficulty in prior-free mechanism design for non-trivial objectives like profit (or, e.g., social surplus with a balanced budget, see Section 3.5) is that there is no pointwise, i.e., for all valuation pro-

files, optimal mechanism. Recall that incentive constraints in mechanism design bind across all valuation profiles. For example, the payment of an agent depends on the what the mechanism would have done had the agent possessed a lower value (Theorem 2.2). Therefore, mechanisms for the profit objective must trade off performance on one input for another. In Chapter 3 this tradeoff was optimized in expectation with respect to the prior distribution from which the agents' values are drawn; without a prior another method for navigating this tradeoff is needed.

This challenge can be resolved with approximation by comparing the performance of a mechanism to an economically meaningful prior-free benchmark. A mechanism approximates a prior-free benchmark if, for all valuation profiles, the mechanism's performance approximates the benchmark performance. A benchmark is economically meaningful if, for a large class of distributions, the expected value of the benchmark is at least the expected performance of the Bayesian optimal mechanism. If a mechanism approximates an economically meaningful benchmark then, as a corollary, the mechanism is also a prior-independent approximation (as defined in Chapter 5). Notice that this approach gives a purely prior-free design and analysis framework, but still requires returning to the Bayesian setting for economic justification of the benchmark.

A final concern is the equilibrium concept. Recall from Chapter 2 that we introduced the common prior assumption (Definition 2.5, page 28) so that strategic choice in games of incomplete information is well defined. Recall also that most of the optimal and approximately optimal mechanisms that we discussed in previous chapters were dominant strategy incentive compatible. In this chapter we resolve the issue of strategic choice absent a common prior by requiring that the designed mechanisms satisfy this stronger dominant-strategy incentive-compatibility condition.

The chapter begins by formalizing the framework for design and analysis of prior-free mechanisms via an economically meaningful prior-free benchmark. This framework is instantiated first in the structurally simple environment of a digital good and then subsequently generalized to environments with richer structure. The prior-free mechanism discussed will all be based a natural market analysis metaphor.

7.1 The Framework for Prior-free Mechanism Design

A main challenge for prior-free mechanism design is in identifying an economically meaningful method for evaluating a mechanism's performance. While the prior-independent mechanisms (of Chapter 5) can be compared to the optimal mechanism for the unknown distribution, absent a prior, there is no optimal mechanism with which to compare. This challenge can be resolved by decomposing the prior-independent analysis into two steps. Fix a large, relevant class of prior distributions. In the first step a prior-free benchmark is identified and normalized so that for all distributions in the class the expected benchmark is at least the Bayesian optimal performance. In the second step an auction is constructed and proven to approximate the benchmark pointwise on all valuation profiles. These steps combine to imply a prior-independent approximation and are formalized below.

Definition 7.1. A *prior-free benchmark* maps valuation profiles to target performances. A prior-free benchmark (APX) is *normalized* for a class of distributions if for all distributions in the class, in expectation the benchmark is at least the performance of the Bayesian optimal mechanism ($\text{REF}_{\mathbf{F}}$) for the distribution. I.e., for all \mathbf{F} in the class, $\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\text{APX}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\text{REF}_{\mathbf{F}}(\mathbf{v})]$.

Definition 7.2. A mechanism (APX) is a *prior-free β approximation* to prior-free benchmark (REF) if for all valuation profiles, its performance is at least a β fraction of the benchmark. I.e., for all \mathbf{v} , $\text{APX}(\mathbf{v}) \geq 1/\beta \text{REF}(\mathbf{v})$.

Proposition 7.1. For any prior-free mechanism, class of distributions, and prior-free benchmark, if the benchmark is normalized for the class of distributions and the mechanism a prior-free β approximation to the

Chapter 7: Topics Covered.

- prior-free benchmarks,
- envy-free optimal pricings,
- random sampling auctions,
- profit extraction as a decision problem for mechanism design, and
- stochastic analysis of random walks and the gamblers ruin.

benchmark, then the mechanism is a prior-independent β approximation for the class of distributions.

We can distinguish good prior-free benchmarks from bad prior-free benchmarks by how much they overestimate the performance. (Note: a normalized prior-free benchmark never underestimates performance.) The extent to which a prior-free benchmark overestimates performance can be quantified by again considering the benchmark relative to a class of prior distributions. As the benchmark is normalized, for any distribution the expected benchmark exceeds the expected performance of the Bayesian optimal mechanism. Of course, the performance of any mechanism is no better than the Bayesian optimal mechanism for the distribution; therefore, the extent to which the Bayesian optimal mechanism approximates the benchmark gives a lower bound on the prior-free approximation of any mechanism to the benchmark. This is formalized in the following definition and proposition.

Definition 7.3. The *resolution* γ of a prior-free benchmark (REF) is the largest ratio of the benchmark to the performance of the Bayesian optimal mechanism (APX $_{\mathbf{F}}$) for any prior-distribution \mathbf{F} . I.e., γ satisfies $\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\text{APX}_{\mathbf{F}}(\mathbf{v})] \geq 1/\gamma \mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\text{REF}(\mathbf{v})]$ for all \mathbf{F} .

Proposition 7.2. For any class of distributions and any prior-free benchmark, the prior-free approximation β of any mechanism is at least the benchmark's resolution γ .

This prior-free design and analysis framework turns the question of approximation into one of optimization. There is some mechanism that obtains the optimal prior-free approximation relative to the benchmark. In most of the cases we will discuss in this chapter the optimal mechanism has an approximation factor that matches the resolution of the benchmark.

Definition 7.4. The *optimal prior-free approximation* β^* for a prior-free benchmark (REF) satisfies

$$\beta^* = \min_{\text{APX}} \max_{\mathbf{v}} \frac{\text{REF}(\mathbf{v})}{\text{APX}(\mathbf{v})}$$

where APX ranges over all dominant strategy incentive compatible mechanisms for the given environment.

In summary, we need a normalized benchmark so that its approximation has economic meaning, and we need a benchmark with fine

resolution as its resolution lower bounds the best prior-free approximation. Intuitively, a benchmark with finer resolution will be better for distinguishing good mechanisms from bad mechanisms. A first and fundamental task in prior-free mechanism design is to identify a benchmark with fine resolution.

Example: Prior-free Monopoly Pricing

We conclude this section by instantiating the framework for design and analysis of prior-free mechanisms for the single-agent monopoly pricing problem. This is the problem of selling a single item to a single agent to maximize revenue. When the agent's value is drawn from a known distribution F , the seller's optimal mechanism, is to post the monopoly price $\hat{v}^* = \operatorname{argmax} \hat{v} (1 - F(\hat{v}))$ for the distribution (see Section 3.3.3).

Consider the class of distributions over a single agent's value with support $[1, h]$. The surplus gives a normalized prior-free benchmark and is defined by the identity function. Notice that (a) for any distribution the expected value of the benchmark exceeds the monopoly revenue and (b) and this inequality is tight for pointmass distributions. The latter observation implies that the surplus is the smallest normalized benchmark (hence, it obtains the finest resolution).

We approach the problem of analyzing the resolution of a benchmark in tandem with its optimal prior-free approximation. First, we give a lower bound on the resolution by considering the expected benchmark on the distribution for which all mechanisms attain the same performance. For the revenue objective, this distribution is the equal-revenue distribution. Second, we give a mechanism with a prior-free approximation factor that matches the lower bound. As, by Proposition 7.2, the optimal prior-free approximation factor is at least the resolution this upper bound implies that the lower bound on resolution is tight.

Lemma 7.3. *For single-agent environments, the class of distributions with support $[0, h]$, and the objective of profit, the surplus benchmark has resolution γ at least $1 + \ln h$.*

Proof. Consider the equal revenue distribution (truncated to the range $[1, h]$ with a pointmass at h with probability $1/h$). The monopoly revenue for the equal-revenue distribution is one and the expected surplus (and therefore the expected benchmark) is $1 + \ln h$ (as also calculated in Example 4.2, page 142); therefore, the resolution γ of the benchmark is at least $1 + \ln h$. \square

Now consider the purely prior-free question of posting a price to obtain a revenue that approximates the surplus benchmark. It should be clear that no deterministic price \hat{v} will do: if $\hat{v} > 1$ the prior-free approximation is infinite for value $v = 1$, and if $\hat{v} = 1$ then the prior-free approximation is h for value $v = h$. On the other hand, picking a randomized price uniformly from the powers of two on the $[1, h]$ interval gives a logarithmic approximation to the surplus. For such a randomized pricing, with probability $1/\log h$ the power of two immediately below v is posted and when this happens the revenue is at least half of the surplus benchmark. This approach and analysis can be tightened to give an approximation ratio that exactly matches the resolution of the benchmark.

Lemma 7.4. *For values in the interval $[1, h]$ there is a prior-free distribution over posted prices with revenue that is a $1 + \ln h$ approximation to the surplus benchmark.*

Proof. Consider the distribution over prices G with cumulative distribution function $G(z) = 1 + \ln z / 1 + \ln h$ and a pointmass at one with probability $1/1 + \ln h$. For any particular value $v \in [1, h]$, the expected revenue from a random price drawn from G is $v/1 + \ln h$. \square

Theorem 7.5. *For single-agent environments, values in $[1, h]$, and the objective of profit, the resolution of the surplus benchmark and the optimal prior-free approximation are $1 + \ln h$.*

Notice that the resolution of the surplus benchmark, which is optimal among all normalized benchmarks, is not constant. In particular, it grows logarithmically with h and, when the agent's value is not bounded within some interval $[1, h]$, it is infinite. We will address this deficiency in the subsequent section where a benchmark with constant resolution and prior-free mechanisms with constant approximation ratios are derived (for $n \geq 2$ agents).

7.2 The Digital-good Environment

Our foray into prior-free mechanism design begins with the benevolent digital-good environment. In a digital-good environment any subset of agents can be simultaneously served. The absence of a feasibility constraint will enable us to focus directly on the main challenge of prior-free mechanism design which is in overcoming the lack of a prior.

We begin by deriving a benchmark with constant resolution. This benchmark is based on a theory of envy-free pricing and we will refer to it as the *envy-free benchmark*. The resolution of the envy-free benchmark and the prior-free optimal approximation are 2.42 (in the limit with n). In the remainder of the section, we will focus on the design of simple mechanisms that approximate this envy-free benchmark (but are not optimal). First, we show that anonymous deterministic auctions cannot give good prior-free approximation. Second, we describe two approaches for designing randomized prior-free auctions for digital goods. The first auction is based on a straightforward market analysis metaphor: use a random sample of the agents to estimate the distribution of values, and run the optimal auction for the estimated distribution on the remaining agents. The approximation ratio of this auction is upper bounded by 4.68. The second auction is based on a standard algorithmic design paradigm: reduction to the a “decision version” of the problem. It gives a four approximation. These mechanisms are randomizations over deterministic dominant strategy incentive compatible (DSIC) mechanisms, the characterization of which is restated from Corollary 2.14 as follows.

Theorem 7.6. *A direct, deterministic mechanism \mathcal{M} is DSIC if and only if for all i and \mathbf{v} ,*

- (i) (*step-function*) $x_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i})$ steps from 0 to 1 at some $\hat{v}_i(\mathbf{v}_{-i})$, and
- (ii) (*critical value*) $p_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i}) = \begin{cases} \hat{v}_i(\mathbf{v}_{-i}) & \text{if } x_i^{\mathcal{M}}(v_i, \mathbf{v}_{-i}) = 1 \\ 0 & \text{otherwise} \end{cases} + p_i^{\mathcal{M}}(0, \mathbf{v}_{-i})$.

7.2.1 The Envy-free Benchmark

Consider the following definition of and motivation for the envy-free benchmark. Recall that, when the agents’ values are drawn from an i.i.d. distribution, the Bayesian optimal digital-good auction would simply post the monopoly price for the distribution as a take-it-or-leave-it offer independently to each agent. For such a posted pricing, the agents with values above the monopoly price would choose to purchase the item and the agents with values below the monopoly price would not. As each agent selects her preferred outcome, this outcome is *envy free* no agent is envious of the outcome obtained by any other agent.

Without a prior, the monopoly price is not well defined. Instead, the *empirical monopoly price* for valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ is the monopoly price of the empirical distribution; it is calculated as $v_{(i^*)}$ with

$i^* = \operatorname{argmax}_i iv_{(i)}$ and $v_{(i)}$ denoting the i th highest value in \mathbf{v} . It is easy to see that the empirical monopoly revenue $\max_i iv_{(i)}$ is an upper bound on the revenue that would be obtained by monopoly pricing if there were a known prior distribution on values. While it is not incentive compatible to inspect the valuation profile, calculate the empirical monopoly price $v_{(i^*)}$, and offer it to each agent; it is envy free. Furthermore, as we will see subsequently, for digital good environments empirical monopoly pricing gives the *envy-free optimal revenue* which we will denote by $\operatorname{EFO}(\mathbf{v})$.

For the class of i.i.d. distributions, the envy-free optimal revenue is a normalized benchmark. Unfortunately, the resolution of the envy-free optimal revenue, as a benchmark, is super constant. When there is $n = 1$ agent the optimal envy-free revenue is the surplus and, from the discussion of the monopoly pricing problem in the preceding section, its resolution is $1 + \ln h$ for values in $[1, h]$ and unbounded in general. The only thing, however, preventing $\operatorname{EFO}(\mathbf{v}) = \max_i iv_{(i)}$ from being a good benchmark is the case where the maximum is obtained at $i^* = 1$ by selling to the highest value agent at her value. This discussion motivates the definition of an envy-free benchmark that explicitly excludes the $i^* = 1$ case.

Definition 7.5. The *envy-free benchmark* $\operatorname{EFO}^{(2)}(\mathbf{v})$ for digital goods is the optimal revenue from posting a uniform price that is bought by two or more agents. I.e., $\operatorname{EFO}^{(2)}(\mathbf{v}) = \max_{i \geq 2} iv_{(i)}$.

Our discussion will distinguish between the envy-free optimal revenue, EFO , and the envy-free benchmark, $\operatorname{EFO}^{(2)}$. The difference between them is that the latter excludes the possibility of selling to just the highest-valued agent. While the envy-free optimal revenue (as a benchmark) is normalized for all i.i.d. distributions, the envy-free benchmark is not. The envy-free benchmark is, however, normalized for a large class of distributions; we omit a precise characterization of this class, though subsequently in Section 7.3, we show that it includes all i.i.d. regular distributions on $n \geq 2$ agents.

Analysis of resolution of the envy-free benchmark is difficult because we must quantify over all distributions. We follow the same high-level approach as for bounding the benchmark resolution in the monopoly pricing problem. First, we analyze the ratio between the expected benchmark and the Bayesian optimal auction revenue for the equal revenue distribution to get a lower bound on the resolution. Second, we observe that an auction exists with prior-free approximation that matches this resolution. Proposition 7.2, which states that any prior-free approxima-

tion is an upper bound on the resolution, implies that the resolution and optimal prior-free approximation are equal. The following theorem summarizes this analysis.

Theorem 7.7. *In digital good environments, the resolution and optimal prior-free approximation of the envy-free benchmark are equal. For $n = 2, 3,$ and $4,$ the resolution and optimal prior-free approximation are $2, 13/6 \approx 2.17,$ and $2^{15}/96 \approx 2.24,$ respectively; in the limit with n it is $2.42.$*

We give the complete two-step proof of the $n = 2$ special case of Theorem 7.7: Lemma 7.8 proves a lower bound of two on the resolution, and Lemma 7.9 proves an upper bound of two on the optimal prior-free approximation. Proposition 7.2, then, implies the equality. The generalization of this proof to $n \geq 3$ agents is technical and the subsequent discussion will treat it only at a high level.

Lemma 7.8. *For two-agent digital-good environments, the resolution of the envy-free benchmark is at least two.*

Proof. We give a lower bound on the resolution by comparing the expected envy-free benchmark (REF) to the expected revenue of the Bayesian optimal auction (APX) for the equal revenue distribution. Recall that the equal-revenue distribution (Definition 4.2, page 106) is given by distribution $F^{\text{EQR}}(z) = 1 - 1/z$ and the revenue from posting any price $\hat{v} \geq 1$ is one. Therefore, the expected revenue of the Bayesian optimal digital-good auction for $n = 2$ agents is $\text{APX} = n = 2.$

It remains to calculate the expected value of the envy-free benchmark $\text{REF} = \mathbf{E}_{\mathbf{v}}[\text{EFO}^{(2)}(\mathbf{v})].$ In the case that $n = 2,$ the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$ simplifies to $2v_{(2)}.$ The expectation of a non-negative random variable X can be calculated as $\mathbf{E}[X] = \int_0^\infty \mathbf{Pr}[X > z] dz;$ to employ this formula we calculate $\mathbf{Pr}[2v_{(2)} > z].$ For $z \geq 2$ we have:

$$\begin{aligned} \mathbf{Pr}_{\mathbf{v}}[2v_{(2)} > z] &= \mathbf{Pr}_{\mathbf{v}}[v_1 > z/2 \wedge v_2 > z/2] \\ &= \mathbf{Pr}_{\mathbf{v}}[v_1 > z/2] \mathbf{Pr}_{\mathbf{v}}[v_2 > z/2] \\ &= 4/z^2. \end{aligned}$$

For $z < 2$ we have: $\mathbf{Pr}[2v_{(2)} > z] = 1.$ The calculation the envy-free benchmark's expected value concludes as follows.

$$\text{REF} = \mathbf{E}_{\mathbf{v}}[2v_{(2)}] = \int_0^\infty \mathbf{Pr}[2v_{(2)} > z] dz = 2 + \int_2^\infty 4/z^2 dz = 4.$$

The resolution of the envy-free benchmark is thus at least $\text{REF}/\text{APX} = 4/2 = 2.$ \square

The generalization of Lemma 7.8 to $n > 2$ follows same proof structure. The main difficulty of the analysis is in calculating the expectation of the benchmark. This is complicated because it becomes the maximum of many terms. E.g., for $n = 3$ agents, $\text{EFO}^{(2)}(\mathbf{v}) = \max(2v_{(2)}, 3v_{(3)})$. Nonetheless, for general n its expectation can be calculated exactly; in the limit with n it is about 2.42.

Lemma 7.9. *For two-agent digital-good environments, the second-price auction is a prior-free two approximation of the envy-free benchmark.*

Proof. For $n = 2$ agents the the envy-free benchmark is $2v_{(2)}$ which is twice the revenue of the second-price auction. Therefore, the second-price auction is a prior-free two approximation to the envy-free benchmark. \square

The generalization of Lemma 7.9 beyond $n = 2$ agents is technical and does not give a natural auction. For example, the $n = 3$ agent optimal auction offers each agent a price drawn from a probability distribution with a pointmass at each of the other two agents' values and continuous density at prices strictly higher than these values. The probabilities depend on the ratio of the two other agents' values. For larger $n \geq 4$ no closed-form expression is known; though, the prior-free optimal auction can be seen to match the lower bound on the resolution by a brute-force construction. This prior-free optimal auction suffers from the main drawback of optimal mechanisms: it is quite complicated. In the next sections, we will derive simple mechanisms that approximate the prior-free optimal digital-good auction.

7.2.2 Deterministic Auctions

The main idea that enables approximation of the envy-free benchmark is that when selecting a price to offer agent i we can use statistics from the values of all other agents as given by their reports \mathbf{v}_{-i} . This motivates the following mechanism which differs from empirical monopoly pricing in that the price to agent i is from the empirical distribution for \mathbf{v}_{-i} not \mathbf{v} .

Definition 7.6. The *deterministic optimal price* auction offers each agent i the take-it-or-leave-it price of \hat{v}_i set as the monopoly price for the profile of other agent values \mathbf{v}_{-i} .

The deterministic optimal price auction is dominant strategy incentive

compatible. It is possible to show that the auction is a prior-independent constant approximation (cf. Chapter 5); however it is not a prior-free approximation. In fact, this deficiency of the deterministic optimal price auction is one that is fundamental to all anonymous (a.k.a., symmetric) deterministic auctions.

Example 7.1. Consider the valuation profile with ten high-valued agents, with value ten, and 90 low-valued agents, with value one. What does the auction do on such a valuation profile? The offer to a high-valued agent is $\hat{v}_h = 1$, as \mathbf{v}_{-h} consists of 90 low-valued agents and 9 high-valued agents. The revenue from the high price is 90; while the revenue from the low price is 99. The offer to a low-valued agent is $\hat{v}_1 = 10$, as \mathbf{v}_{-1} consists of 89 low-valued agents and 10 high-valued agents. The revenue from the high price is 100; while the revenue from the low price is 99. With these offers all high-valued agents will win and pay one, while all low-valued agents will lose. The total revenue of ten is far from the envy-free benchmark revenue of $\text{EFO}^{(2)}(\mathbf{v}) = 100$.

Theorem 7.10. *No n -agent anonymous deterministic dominant-strategy incentive-compatible digital-good auction is better than an n approximation to the envy-free benchmark.*

Proof. Consider valuation profiles \mathbf{v} with values $v_i \in \{1, h\}$. Let $n_h(\mathbf{v})$ and $n_1(\mathbf{v})$ denote the number of h values and 1 values in \mathbf{v} , respectively. By Theorem 7.6, any deterministic and dominant strategy incentive compatible auction APX has a critical value at which each agent is served. That APX is anonymous implies that the critical value for agent i , as a function of the reports of other agents, is independent of the index i and only a function of $n_h(\mathbf{v}_{-i})$ and $n_1(\mathbf{v}_{-i})$. Thus, we can let $\hat{v}(n_h, n_1)$ represent the offer price of APX for any agent i when we plug in $n_h = n_h(\mathbf{v}_{-i})$ and $n_1 = n_1(\mathbf{v}_{-i})$. Finally, we assume that $\hat{v}(n_h, n_1) \in \{1, h\}$ as this restriction cannot hurt the auction profit on the valuation profiles we are considering.

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

- (i) Observe that for any auction that is a good approximation, it must be that for all m , $\hat{v}(m, 0) = h$. Otherwise, on the $n = m + 1$ agent all h 's input, the auction only achieves profit n while the envy-free benchmark is hn . Thus, the auction would be at most an $h \geq n$ approximation on profiles with $h \geq n$.
- (ii) Likewise, observe that for any auction that is a good approximation,

it must be that for all m , $\hat{v}(0, m) = 1$. Otherwise, on the $n = m + 1$ agent all 1's input, the auction achieves no profit and is clearly not an approximation of the envy-free benchmark n .

(iii) We now identify a bad valuation profile for the auction. Take m sufficiently large and consider $\hat{v}(k, m - k)$ as a function of k . As we have argued for $k = 0$, $\hat{v}(k, m - k) = 1$. Consider increasing k until $\hat{v}(k, m - k) = h$. This must occur since at $k = m$ we have $\hat{v}(k, m - k) = h$. Let $k^* = \min\{k : \hat{v}(k, m - k) = h\} > 1$ be this transition point. Now consider an $n = m + 1$ agent valuation profile with $n_h(\mathbf{v}) = k^*$ and $n_1(\mathbf{v}) = m - k^* + 1$.

- For low-valued agents: $\hat{v}(n_h(\mathbf{v}_{-1}), n_1(\mathbf{v}_{-1})) = \hat{v}(k^*, m - k^*) = h$. Thus, all low-valued agents are rejected and contribute nothing to the auction profit.
- For high-valued agents: $\hat{v}(n_h(\mathbf{v}_{-h}), n_1(\mathbf{v}_{-h})) = \hat{v}(k^* - 1, m - k^* + 1) = 1$. Thus, all high-valued agents are offered a price of one which they accept. Thus, the contribution to the auction profit from such agents is $1 \times n_h(\mathbf{v}) = k^*$.

Thus, the total auction profit for this valuation profile is $\text{APX} = k^*$.

(iv) For $h = n$, the envy-free benchmark on this valuation profile is $\text{REF} = nk^*$. There are two cases. If $k^* = 1$ then the benchmark is n (from selling to all agents at price 1); of course, for $k^* = 1$ then $n = nk^*$. If $k^* \geq 2$ the benchmark is also nk^* (from selling to the k^* high-valued agents at price $h = n$).

In conclusion, we have identified a valuation profile where the auction revenue is $\text{APX} = k^*$ and the envy-free benchmark is $\text{REF} = nk^*$; the auction is at best a prior-free n approximation. \square

Theorem 7.10 implies that either randomization or asymmetry is necessary to obtain good prior-free mechanisms. While either approach will permit the design of good mechanisms, all deterministic asymmetric auctions known to date are based on derandomizations of randomized auctions. This text will discuss only these randomized auctions.

7.2.3 The Random Sampling Auction

We now discuss a prior-free auction based on a natural market-analysis metaphor. Notice that the problem with the deterministic optimal price auction in the preceding section is that it may simultaneously offer high-valued agents a low price and low-valued agents a high price. Of course,

either of these prices would have been good if it were offered consistently to all agents. One approach for combating this lack of coordination is to coordinate using random sampling. The idea is roughly to partition the agents into a market and sample and then use the sample to estimate a good price and then offer that price to the agents in the market. With a random partition we expect a fair share of high- and low-valued agents to be in both the market and the sample; therefore, a price that is good for the sample should also be good for the market.

Definition 7.7. The *random sampling (optimal price) auction* works as follows:

- (i) randomly partition the agents into sample S and market M (by flipping a fair coin for each agent),
- (ii) compute (empirical) monopoly prices \hat{v}_S^* and \hat{v}_M^* for S and M respectively, and
- (iii) offer \hat{v}_S^* to M and \hat{v}_M^* to S .

We first, and easily, observe that the random sampling auction is dominant strategy incentive compatible.

Theorem 7.11. *The random sampling auction is dominant strategy incentive compatible.*

Proof. Fix a randomized partition of the agents into a market and sample. For this partitioning, each agent faces a critical value that is a function only of other agent reports. Theorem 7.6 then implies that the auction for this partitioning is dominant strategy incentive compatible. Of course, if it is dominant strategy for any fixed partitioning it is certainly dominant strategy in expectation over the random partitioning. \square

The following example, as a warm up exercise, demonstrates that the random sampling auction is not better than a four approximation to the envy-free benchmark.

Example 7.2. Consider the 2-agent input $\mathbf{v} = (1.1, 1)$ for which the envy-free benchmark is $\text{EFO}^{(2)}(\mathbf{v}) = 2$. To calculate the auction's revenue on this input, notice that these two agents are in the same partition with probability $1/2$ and in different partitions with probability $1/2$. In the former case, the auction's revenue is zero. In the latter case it is the lower value, i.e., one. The auction's expected profit is therefore $1/2$, which is a four approximation to the benchmark.¹

¹ It is natural to think this example could be improved if the auction were to

Theorem 7.12. *For digital good environments and all valuation profiles, the random sampling auction is at most a 4.68 approximation to the envy-free benchmark.*

This theorem is technical and it is generally believed that the bound it provides is loose and the random sampling auction is in fact a worst-case four approximation. Below we will prove the weaker claim that it is at worst a 15 approximation. This weaker claim highlights the main techniques involved in proving that variants and generalizations of the random sampling auction are constant approximations.

Lemma 7.13. *For digital good environments and all valuation profiles, the random sampling auction is at most a 15 approximation to the envy-free benchmark.*

Proof. Assume without loss of generality that the highest-valued agent is in the market M . This terminology comes from the fact that if the highest agent value is sufficiently large then all agents in other partition (in this case S) will be rejected; the role of S is then only as a sample for statistical analysis. There are two main steps in the proof. Step (i) is to show that the optimal envy-free revenue from the sample $\text{EFO}(\mathbf{v}_S)$ is close to the envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v})$. Step (ii) is to show that the revenue from price \hat{v}_S^* on M is close to the envy-free optimal revenue from the sample which is, recall, the revenue from price \hat{v}_S^* on S .

We will use the following definitions. First sort the agents by value so that v_i is the i th largest valued agent. Define y_i as an indicator variable for the event that $i \in S$ (the sample). Notice that $\mathbf{E}[y_i] = 1/2$ except for $i = 1$; $y_1 = 0$ by the assumption that the highest valued agent is in the market. Define $Y_i = \sum_{j \leq i} y_j$ as the number of the i highest-valued agents who are in the sample. Let $\text{EFO}^{(2)}(\mathbf{v}) = i^* \hat{v}^*$ where i^* is the number of winners in the benchmark and $\hat{v}^* = v_{i^*}$ is the benchmark price.

- (i) With good probability, the optimal envy-free revenue for the sample, $\text{EFO}(\mathbf{v}_S)$, is close to the envy-free benchmark, $\text{EFO}^{(2)}(\mathbf{v})$.

Define \mathcal{B} as the event that the sample contains at least half of the i^* highest-valued agents, i.e., $Y_{i^*} \geq i^*/2$. Of course the envy-free optimal

partition half of the agents into the market and half into the sample. However in worst case, this improved partitioning cannot help. Pad the valuation profile with agents who have zero value for the item and then observe that the same analysis on this padded valuation profile gives a lower bound of four on the auction's approximation ratio.

revenue for the sample is at least the revenue from posting price \hat{v}^* (which is envy-free), i.e., $\text{EFO}(\mathbf{v}_S) \geq Y_{i^*} \hat{v}^*$. Event \mathcal{B} then implies that $Y_{i^*} \hat{v}^* \geq 1/2 i^* \hat{v}^*$, or equivalently $\text{EFO}(\mathbf{v}_S) \geq 1/2 \text{EFO}^{(2)}(\mathbf{v})$.

We now show that $\Pr[\mathcal{B}] = 1/2$ when i^* is even. Recall that the highest valued agent is always in the market. Therefore there are $i^* - 1$ (an odd number) of agents which we partition between the market and the sample. One partition receives at least $i^*/2$ of these and half the time it is the sample; therefore, $\Pr[\mathcal{B}] = 1/2$.

When i^* is odd $\Pr[\mathcal{B}] < 1/2$, and a slightly more complicated argument is needed to complete the proof. A sketch of the argument is as follows. Define \mathcal{C} as the event that $Y_{i^*} \geq i^*/2$. When this event occurs, by a similar analysis as in the even case, $\text{EFO}(\mathbf{v}_S) \geq 1/2(1 - 1/i^*) \text{EFO}^{(2)}(\mathbf{v})$. The implied bound is worse than the analogous bound for the even case by an $1 - 1/i^*$ factor. The probability that the event \mathcal{C} holds improves over event \mathcal{B} , however, and this improvement more than compensates for the loss. Notice that strictly more of the top $i^* - 1$ agents are in the sample or market with equal probability but event \mathcal{C} also occurs when the numbers are equal. Thus, $\Pr[\mathcal{C}] > 1/2 = \Pr[\mathcal{B}]$. The intuition that these bounds combine to improve over the even case, above, is that the probability that the $i^* - 1$ top agents are split evenly grows as $\Theta(\sqrt{1/i^*})$ and the loss from the event providing a weaker bound grows as $\Theta(1/i^*)$.

- (ii) With good probability, the revenue from price \hat{v}_S^* on M is close to $\text{EFO}(\mathbf{v}_S)$.

Define \mathcal{E} as the event that for all indices i that the market contains at least a third as many of the i highest-valued agents as the sample, i.e., $\forall i, i - Y_i \geq 1/3 Y_i$. Notice that the left hand side of this equation is the number of agents with value at least v_i in the market, while the right hand side is a third of the number of such agents in the sample. Importantly, this event implies that the partitioning of agents is not too imbalanced in favor of the sample. We refer to this event as the *balanced sample* event; though, note that it is only a one-directional balanced condition.

Let the envy-free optimal revenue for the sample be $\text{EFO}(\mathbf{v}_S) = Y_{i_S^*} \hat{v}_S^*$ where i_S^* is the index of the agent whose value is used as its price, $\hat{v}_S^* = v_{i_S^*}$ is its price, and $Y_{i_S^*}$ is its number of winners. The profit of the random sampling auction is equal to $(i_S^* - Y_{i_S^*}) \hat{v}_S^*$. Under the balanced sample condition this is lower bounded by $1/3 Y_{i_S^*} \hat{v}_S^* = 1/3 \text{EFO}(\mathbf{v}_S)$.

Subsequently, we will prove a *balanced sampling lemma* (Lemma 7.14) that shows that $\Pr[\mathcal{E}] \geq 0.9$.

We combine the two steps, above, as follows. If both good events \mathcal{E} and \mathcal{B} hold, then the expected revenue of random sampling auction is at least $1/6 \text{ EFO}^{(2)}(\mathbf{v})$. By the union bound, the probability of this good fortune is $\Pr[\mathcal{E} \wedge \mathcal{B}] \geq 1 - \Pr[\neg\mathcal{E}] - \Pr[\neg\mathcal{B}] \geq 0.4$.² We conclude that the random sampling auction is a $15 = 6 \times 1/0.4$ approximation to the envy-free benchmark. \square

Lemma 7.14 (Balanced Sampling). *For $y_1 = 0$, y_i for $i \geq 2$ an indicator variable for a independent fair coin flipping to heads, and sum $Y_i = \sum_{j \leq i} y_j$,*

$$\Pr[\forall i, (i - Y_i) \geq 1/3 Y_i] \geq 0.9.$$

Proof. We relate the condition of the lemma to the *probability of ruin* in a *random walk* on the integers. Notice that $(i - Y_i) \geq 1/3 Y_i$ if and only if, for integers i and Y_i , $3i - 4Y_i + 1 > 0$. So let $Z_i = 3i - 4Y_i + 1$ and view Z_i as the position, in step i , of a random walk on the integers. Since $Y_1 = y_1 = 0$ this random walk starts at $Z_i = 4$. Notice that at step i in the random walk with is in position Z_i , so at step $i + 1$ we have

$$Z_{i+1} = \begin{cases} Z_i - 1 & \text{if } y_{i+1} = 1, \text{ and} \\ Z_i + 3 & \text{if } y_{i+1} = 0; \end{cases}$$

i.e., the random walk either takes three steps forward or one step back. We wish to calculate the probability that this random walk never touches zero. This type of calculation is known as a *probability of ruin* analysis in reference to a gambler's fate when playing a game with such a payoff structure.

Let r_k denote the probability of ruin from position k . This is the probability that the random walk eventually takes k steps backwards. Clearly $r_0 = 1$, as at position $k = 0$ we are already ruined, and $r_k = r_1^k$, as taking k steps back is equivalent to stepping back k times. By the

² We denote the event that \mathcal{E} does not occur by $\neg\mathcal{E}$, which should be read as “not \mathcal{E} .” The probabilities of any event \mathcal{E} and its complement $\neg\mathcal{E}$ satisfy $\Pr[\neg\mathcal{E}] = 1 - \Pr[\mathcal{E}]$. A typical approach for bounding the probability of the conjunction (i.e., the “and”) of two events is by the disjunction (i.e., the “or”) of their negations, i.e., $\Pr[\mathcal{E} \wedge \mathcal{B}] = 1 - \Pr[\neg(\mathcal{E} \wedge \mathcal{B})] = 1 - \Pr[\neg\mathcal{E} \vee \neg\mathcal{B}]$. The *union bound* states that the probability of the disjunction of two events is at most the sum of the probabilities of each event. (This bound is tight for disjoint events, while for events that may simultaneously occur, it double counts the probability of outcomes that satisfy both events.) Thus, $\Pr[\mathcal{E} \wedge \mathcal{B}] \geq 1 - \Pr[\neg\mathcal{E}] - \Pr[\neg\mathcal{B}]$.

definition of the random walk, we have the recurrence,

$$r_k = 1/2(r_{k-1} + r_{k+3}).$$

Plugging in the above identities for $k = 1$ we have,

$$r_1 = 1/2(1 + r_1^4).$$

This is a quartic equation that can be solved, e.g., by *Ferarri's formula* (though we omit the details). Since our random walk starts at $Z_1 = 4$ we calculate $r_4 = r_1^4 \leq 0.1$, meaning that the probability that the balanced sampling condition is satisfied is at least 0.9. \square

The proof of Theorem 7.12 follows a very similar structure to that of Lemma 7.13. The main additional idea is that, instead of fixing the level of imbalance to be tolerated, it is a random variable. In Lemma 7.13 the imbalance is fixed to $1/3$. Notice that the performance bound constructed in the lemma scales linearly with the imbalance. Thus, the expected bound can be factored into the expected imbalance times the worst case performance for imbalance one.

7.2.4 Decision Problems for Mechanism Design

Decision problems play a central role in computational complexity and algorithm design. Where as an optimization problem is to find the optimal solution to a problem, a *decision problem* is to decide whether or not there exists a solution that meets a given objective criterion. While it is clear that decision problems are no harder to solve than optimization problems, the opposite is also true, for instance, we can search for the optimal objective value of any feasible solution by making repeated calls to an algorithm that solves the decision problem. This search is single-dimensional and can be effectively solved, e.g., by *binary search*. In this section we develop a similar theory for prior-free mechanism design.

Profit Extraction

For profit maximization in mechanism design, recall, there is no point-wise optimal mechanism. Therefore, we define the mechanism design decision problem in terms of the envy-free optimal revenue EFO. The decision problem for profit target Π is to design a mechanism that obtains profit at least Π on any valuation profile \mathbf{v} with $\text{EFO}(\mathbf{v}) \geq \Pi$. We call a mechanism that solves the decision problem a *profit extractor*.

Definition 7.8. The *digital good profit extractor* for target Π and valuation profile \mathbf{v} finds the largest k such that $v_{(k)} \geq 1/k \Pi$, sells to the top k agents at price $1/k \Pi$, and rejects all other agents. If no such k exists, it rejects all agents.

Lemma 7.15. *The digital good profit extractor is dominant strategy incentive compatible.*

Proof. Consider the following ascending auction. See if all agents can evenly split the target Π . If some agents cannot afford to pay their fair share, reject them. Repeat with the remaining agents. Notice that the number of remaining agents in this process is decreasing, and thus, the fair share of each remaining agent is increasing. Therefore, each agent faces an ascending price until she drops out. It is a dominant strategy for her to drop out when the ascending price exceeds her value (c.f. the single-item ascending-price auction of Definition 1.5, page 5).

The outcome selected by this ascending auction is identical to that of the profit extractor. Therefore, we can interpret the profit extractor as the revelation principle (Theorem 2.11) applied to the ascending auction. The dominant strategy equilibrium of the ascending auction, then, implies that the profit extractor is dominant strategy incentive compatible. \square

Lemma 7.16. *For all valuation profiles \mathbf{v} , the digital good profit extractor for target Π obtains revenue Π if $\text{EFO}(\mathbf{v}) \geq \Pi$ and zero otherwise.*

Proof. Recall, $\text{EFO}(\mathbf{v}) = i^* v_{(i^*)}$. If $\Pi \leq \text{EFO}(\mathbf{v})$ then there exists a k such that $v_{(k)} \geq 1/k \Pi$, e.g., $k = i^*$. In this case its revenue is exactly Π . On the other hand, if $\Pi > \text{EFO}(\mathbf{v}) = \max_k k v_{(k)}$ then there is no such k for which $v_{(k)} \geq 1/k \Pi$ and the mechanism has no winners and no revenue. \square

Approximate Reduction to Decision Problem

We now employ random sampling to approximately reduce the mechanism design problem of optimizing profit to the decision problem. The key observation in this reduction is an analogy. Notice that given a single agent with value v , if we offer this agent a threshold \hat{v} the agent buys and pays \hat{v} if and only if $v \geq \hat{v}$. Analogously a profit extractor with target Π on a subset S of the agents obtains revenue Π if and only if $\text{EFO}(\mathbf{v}_S) \geq \Pi$. We can thus view the subset S of agents like a single “meta agent” with value $\text{EFO}(\mathbf{v}_S)$. The idea then is to randomly partition the agents into two parts, treat each part as a meta agent, and

run the second-price auction on these two meta agents. The last step is accomplished by attempting to profit extract the envy-free optimal revenue for one part from the other part, and vice versa.

Definition 7.9. The *random sampling profit extraction* auction works as follows:

- (i) randomly partition the agents into S and M (by flipping a fair coin for each agent),
- (ii) Calculate $\Pi_M = \text{EFO}(\mathbf{v}_M)$ and $\Pi_S = \text{EFO}(\mathbf{v}_S)$, the benchmark profit for each part.
- (iii) Profit extract Π_S from M and Π_M from S .

Notice that the intuition from the analogy to the second-price auction implies that the revenue of the random sampling profit extraction auction is exactly the minimum of Π_M and Π_S . Since the profit extractor is dominant strategy incentive compatible, so is the random sampling profit extraction auction.

Lemma 7.17. *The random sampling profit extraction auction is dominant strategy incentive compatible.*

Before we prove that the random sampling profit extraction auction is a four approximation to the envy-free benchmark, we give a simple proof of a lemma that will be important in the analysis.

Lemma 7.18. *With $k \geq 2$ fair coin flips, the expected minimum of the number of heads or tails is at least $1/4 k$.*

Proof. Let W_i be a random variable for the minimum number of heads or tails in the first i coin flips. The following calculations are elementary:

$$\begin{aligned}\mathbf{E}[W_1] &= 0, \\ \mathbf{E}[W_2] &= 1/2, \text{ and} \\ \mathbf{E}[W_3] &= 3/4.\end{aligned}$$

We now obtain a general bound on $\mathbf{E}[W_i]$ for $i > 3$. Let $w_i = W_i - W_{i-1}$ representing the change to the minimum number of heads or tails when we flip the i th coin. Notice that linearity of expectation implies that $\mathbf{E}[W_i] = \sum_{i=1}^k \mathbf{E}[w_i]$. Thus, it will suffice to calculate $\mathbf{E}[w_i]$ for all i . We consider this calculation in three cases:

Case 1 (i even): This implies that $i - 1$ is odd, and prior to flipping the i th coin it was not the case that there was a tie. Assume

without loss of generality that there were more tails than heads. Now when we flip the i th coin, there is probability $1/2$ that it is heads and we increase the minimum by one; otherwise, we get tails have no increase to the minimum. Thus, $\mathbf{E}[w_i] = 1/2$.

Case 2 (i odd): Here we use the crude bound that $\mathbf{E}[w_i] \geq 0$. Note that this is actually the best we can claim in worst case since $i - 1$ is even so before flipping the i th coin it could be that there is a tie. If this were the case then regardless of the i th coin flip, $w_i = 0$ and the minimum number of heads or tails would be unchanged.

Case 3 ($i = 3$): This is a special case of Case 2; however we can get a better bound using the calculations of $\mathbf{E}[W_2] = 1/2$ and $\mathbf{E}[W_3] = 3/4$ above to deduce that $\mathbf{E}[w_3] = \mathbf{E}[W_3] - \mathbf{E}[W_2] = 1/4$.

Finally we are ready to calculate a lower bound on $\mathbf{E}[W_k]$.

$$\begin{aligned} \mathbf{E}[W_k] &= \sum_{i=1}^k \mathbf{E}[w_i] \\ &\geq 0 + 1/2 + 1/4 + 1/2 + 0 + 1/2 + 0 + 1/2 + \dots \\ &= 1/4 + 1/2 \lfloor k/2 \rfloor \\ &\geq 1/4 k. \end{aligned} \quad \square$$

Theorem 7.19. *For digital good environments and all valuation profiles, the random sampling profit extraction auction is a four approximation to the envy-free benchmark.*

Proof. For valuation profile \mathbf{v} , let REF be the envy-free benchmark and its revenue and APX be the random sampling profit extraction auction and its expected revenue. From the analogy to the second-price auction on meta-agents, the expected revenue of the auction is $\text{APX} = \mathbf{E}[\min(\Pi_M, \Pi_S)]$ (where the expectation is taken over the randomized of the partitioning of agents).

Assume that the envy-free benchmark sells to $i^* \geq 2$ agents at price \hat{v}^* , i.e., $\text{REF} = i^* \hat{v}^*$. Of the i^* winners in REF, let i_M^* be the number of them that are in M and i_S^* the number that are in S . Since there are i_M^* agents in M above price \hat{v}^* , then $\Pi_M \geq i_M^* \hat{v}^*$. Likewise, $\Pi_S \geq i_S^* \hat{v}^*$.

$$\frac{\text{APX}}{\text{REF}} = \frac{\mathbf{E}[\min(\Pi_M, \Pi_S)]}{i^* \hat{v}^*} \geq \frac{\mathbf{E}[\min(i_M^* \hat{v}^*, i_S^* \hat{v}^*)]}{i^* \hat{v}^*} = \frac{\mathbf{E}[\min(i_M^*, i_S^*)]}{i^*} \geq \frac{1}{4}.$$

The last inequality follows by applying Lemma 7.18 when we consider $i^* \geq 2$ coins and heads as putting an agent in S and a tails as putting the agent in M .

This bound is tight by an adaptation of the analysis of Example 7.2 from which we concluded that the random sampling optimal price auction is at best a four approximation. \square

One question that should seem pertinent at this point is whether partitioning into two groups is optimal. We could alternatively partition into three parts and run a three-agent auction on the benchmark revenue of these parts. Of course, the same could be said for partitioning into ℓ parts for any ℓ . In fact, the optimal partitioning comes from $\ell = 3$, though we omit the proof and full definition of the mechanism.

Theorem 7.20. *For digital good environments and all valuation profiles, the random three-partitioning profit extraction auction is a 3.25 approximation to the envy-free benchmark.*

7.3 The Envy-free Benchmark

The first step in generalizing the framework for prior-free approximation from the preceding sections is to generalize the envy-free benchmark. In this section we consider envy-free optimal pricing in general environments. We will give characterizations of envy-free pricings and envy-free optimal pricings that mirror those of incentive compatibility. These characterizations will promote the viewpoint that envy freedom is a relaxation of incentive compatibility that admits pointwise optimization. The section will conclude with the general definition and discussion of the envy-free benchmark.

Definition 7.10. For valuation profile \mathbf{v} , an outcome with allocation \mathbf{x} and payments \mathbf{p} is *envy free* if no agent prefers the outcome of another agent to her own, i.e.,

$$\forall i, j, v_i x_i - p_i \geq v_i x_j - p_j.$$

Example 7.3. As a running example for this section consider an $n = 90$ agent, $k = 20$ unit environment with a valuation profile \mathbf{v} that consists of ten high-valued agents each with value ten and 80 low-valued agents each with value two. The following three pricings are envy free (and feasible for the environment).

- (i) Post a price of ten. Serve the ten high-valued agents at this price, reject the low-valued agents. This pricing is envy free: the high-valued

agents weakly prefer buying and the low-valued agents prefer not buying. The total revenue is $100 = 10 \times 10$.

- (ii) Post a price of two. Serve the ten high-valued agents and ten of the low-valued agents at this price. This pricing is envy free: the high-valued agents prefer buying and the low-valued agents are indifferent between buying and not buying. The total revenue is $40 = 20 \times 2$.
- (iii) Post a price of nine to buy the item with certainty and a price of $1/4$ to buy the item with probability $1/8$ (equivalently, a probability $1/8$ chance to buy at price of two). Serve the ten high-valued agents with the certainty outcome, and serve the 80 low-valued agents with the probabilistic outcome. By an elementary analysis this pricing is envy free: the high-valued agents weakly prefer to buy the certainty outcome and the low-valued agents weakly prefer to buy the probabilistic outcome (over nothing). The total revenue is $110 = 10 \times 9 + 80 \times 1/4$.

7.3.1 Envy-free Pricing

The definition of envy freedom can be contrasted to definition of incentive compatibility as given by the revelation principle and the defining inequality of Bayes-Nash equilibrium (Proposition 2.1, page 30). Importantly, incentive compatibility constrains the outcome an agent would receive upon a unilateral misreport where as envy freedom constrains the outcome she would receive upon swapping with another agent. The similarity of envy freedom and incentive compatibility enables an analogous characterization (cf. Section 2.5, page 31) and optimization (cf. Section 3.3, page 59) of envy-free pricings. However, unlike the incentive-compatibility constraints, envy-freedom constraints bind pointwise on the given valuation profile; therefore, there is always a pointwise optimal envy-free outcome.

Theorem 7.21. *For valuation profile \mathbf{v} (sorted with $v_1 \geq \dots \geq v_n$), an outcome (\mathbf{x}, \mathbf{p}) is envy free if and only if*

- (monotonicity) $x_1 \geq \dots \geq x_n$.
- (payment correspondence) there exists a p_0 and monotone function $y(\cdot)$ with $y(v_i) = x_i$ such that for all i

$$p_i = v_i x_i - \int_0^{v_i} y(z) dz + p_0,$$

where usually $p_0 = 0$.

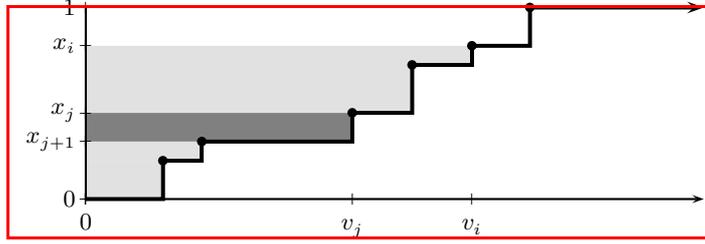


Figure 7.1. The allocation is depicted as points (v_j, x_j) for each agent j . The envy-free payment of agent i is depicted as the total shaded area. The j th term in the sum of equation (7.1) is the dark shaded rectangle. The effective allocation rule y from Theorem 7.21 is the stair function depicted by a solid line.

Notice that the envy-free payments are not pinned down precisely by the allocation; instead, there is a range of appropriate payments. As these payment can be interpreted as the “area above the curve $y(\cdot)$,” the maximum payments are given when $y(\cdot)$ is the smallest monotone function consistent with the allocation. Given our objective of profit maximization, for any monotone allocation rule, we focus on the maximum envy-free payments. These maximum envy-free payments are thus given by the following formula and depicted in Figure 7.1:

$$p_i = \sum_{j \geq i}^n v_j (x_j - x_{j+1}), \quad (7.1)$$

again, with \mathbf{v} sorted as $v_1 \geq \dots \geq v_n$.

Proof of Theorem 7.21. We prove the theorem for the maximum envy-free payments as specified by (7.1) and leave the general payment correspondence as an exercise.

Monotonicity and the payment identity of equation 7.1 imply envy freedom: Suppose \mathbf{x} is swap monotone. Let \mathbf{p} be given as by equation 7.1. We verify that (\mathbf{x}, \mathbf{p}) is envy-free. There are two cases: if $i \leq j$, we have:

$$p_i - p_j = \sum_{k=i}^{j-1} v_k \cdot (x_k - x_{k+1}) \leq v_i \cdot \sum_{k=i}^{j-1} (x_k - x_{k+1}) = v_i \cdot (x_i - x_j),$$

and if $i \geq j$, we have:

$$p_i - p_j = - \sum_{k=j}^{i-1} v_k \cdot (x_k - x_{k+1}) \leq -v_i \cdot \sum_{k=j}^{i-1} (x_k - x_{k+1}) = v_i \cdot (x_i - x_j).$$

Each equation above can be rearranged to give the definition of envy freedom.

Envy freedom implies monotonicity: Suppose \mathbf{x} admits \mathbf{p} such that (\mathbf{x}, \mathbf{p}) is envy-free. By definition, $v_i x_i - p_i \geq v_i x_j - p_j$ and $v_j x_j - p_j \geq v_j x_i - p_i$. By summing these two inequalities and rearranging, $(x_i - x_j) \cdot (v_i - v_j) \geq 0$, and hence \mathbf{x} is monotone.

The maximum envy-free prices satisfy the payment identity of equation 7.1: Agent i does not envy $i + 1$ so $v_i x_i - p_i \geq v_i x_{i+1} - p_{i+1}$, or rearranging: $p_i \leq v_i(x_i - x_{i+1}) + p_{i+1}$. Given p_{i+1} the maximum p_i satisfies this inequality with equality. Letting $p_n = v_n x_n$ (the maximum individually rational payment) and induction gives the payment identity: $p_i = \sum_{j=i}^n v_j \cdot (x_j - x_{j+1})$. \square

7.3.2 Envy-free Optimal Revenue

Definition 7.11. Given any symmetric environment and valuation profile \mathbf{v} , the *envy-free optimal revenue*, denoted $\text{EFO}(\mathbf{v})$, is the maximum revenue attained by a feasible envy-free outcome.

In Section 7.2 we discussed the envy-free optimal revenue for digital good environments and observed that it can be viewed as the revenue from the monopoly pricing of the *empirical distribution* for the valuation profile. The empirical distribution for a valuation profile \mathbf{v} is the discrete distribution with probability $1/n$ at value v_i .

Consider envy-free optimal pricing in multi-unit environments where, unlike digital goods, there is a constraint on the number of agents that can be simultaneously served (see Example 7.3). Recall that for irregular multi-unit auction environments the Bayesian optimal auction is not just the second-price auction with the monopoly reserve (in particular, it may iron). For these environments the envy-free optimal pricing also may iron. In particular, it corresponds to a virtual value optimization for virtual values given by the empirical distribution. Below we define the empirical revenue and empirical marginal revenue from which the envy-free optimal revenue can be calculated (cf. Definition 3.11, Definition 3.12, and Definition 3.15 in Section 3.3, page 59).

Definition 7.12. For valuation profile \mathbf{v} sorted as $v_1 \geq \dots \geq v_n$, the *empirical price-posting revenues* are $\mathbf{P} = (P_0, \dots, P_n)$ with $P_0 = 0$ and $P_i = i v_i$ for all $i \in [n]$. The *empirical price-posting revenue curve* is the piece-wise linear function connecting the points $(0, P_0), \dots, (n, P_n)$. The *empirical revenue curve* is the smallest concave function that upper bounds the empirical price-posting revenue curve; i.e., the empirical

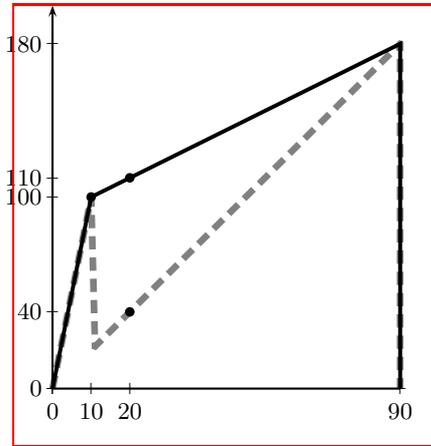


Figure 7.2. The empirical revenue and empirical price-posting revenue curves corresponding to the $n = 90$ agent valuation profile with ten high-valued agents and 80 low-valued agents (Example 7.3). The three envy-free pricings of the example are depicted as P_{10} , P_{20} , and R_{20} .

revenue curve is given by ironing the empirical price-posting revenue curve. The empirical revenues are $\mathbf{R} = (R_0, \dots, R_n)$ with R_i obtained by evaluating the empirical revenue curve at i . *Empirical marginal revenues* and *empirical marginal price-posting revenues* are given by the left slope of their respective empirical revenue curves, or equivalently, as $P'_i = P_i - P_{i-1}$ and $R'_i = R_i - R_{i-1}$.

Example 7.4. The empirical marginal revenues for Example 7.3 ($n = 90$ agents, ten with value ten and 80 with value two). The empirical revenues and price-posting revenues for this valuation profile are given in Figure 7.3.2; The empirical marginal revenues are:

$$R'_i = \begin{cases} 10 & i \in \{1, \dots, 10\}, \text{ and} \\ 1 & i \in \{11, \dots, 90\}. \end{cases}$$

Analogously to the Bayesian optimal incentive compatible auction, the envy-free optimal pricing is a virtual value maximizer for virtual values defined by the empirical marginal revenue. The proofs of Theorem 7.22 and Corollary 7.23, below, are essentially the same as the proofs of Theorem 3.12 and Corollary 3.15.

Theorem 7.22. *The maximal envy-free revenue for monotone alloca-*

tion \mathbf{x} is

$$\sum_i P'_i x_i = \sum_i P_i (x_i - x_{i+1}) \leq \sum_i R_i (x_i - x_{i+1}) = \sum_i R'_i x_i$$

with equality if and only if $R_i \neq P_i \Rightarrow x_i = x_{i+1}$.

Corollary 7.23. *In symmetric environments, with virtual values defined as the empirical marginal revenues, virtual surplus maximization (with random tie-breaking) gives the envy-free outcome with the maximum profit.*

Proof of Theorem 7.22. The inner inequality holds by the following sequence of inequalities:

$$\begin{aligned} \sum_{i=1}^n p_i &= \sum_{i=1}^n \sum_{j=i}^n v_j \cdot (x_j - x_{j+1}) \\ &= \sum_{i=1}^n i v_i \cdot (x_i - x_{i+1}) = \sum_{i=1}^n P_i \cdot (x_i - x_{i+1}) \\ &= \sum_{i=1}^n R_i \cdot (x_i - x_{i+1}) - \sum_{i=1}^n (R_i - P_i) \cdot (x_i - x_{i+1}) \\ &\leq \sum_{i=1}^n R_i \cdot (x_i - x_{i+1}), \end{aligned}$$

where the final inequality follows from the facts that $R_i \geq P_i$ and $x_i \geq x_{i+1}$. Clearly the inequality holds with equality if and only if $x_i = x_{i+1}$ whenever $R_i > P_i$.

The outer equalities hold by collecting like terms in the summation as follows,

$$\sum_{i=1}^n P_i \cdot (x_i - x_{i+1}) = \sum_{i=1}^n (P_i - P_{i-1}) \cdot x_i = \sum_{i=1}^n P'_i \cdot x_i,$$

with the analogous equations relating R_i and R'_i . \square

7.3.3 Envy freedom versus Incentive Compatibility

Optimal envy-free pricing and Bayesian optimal mechanisms are structurally similar; they are both virtual value maximizers. In this section we observe that their optimal revenues are also similar.

An *empirical virtual value function* can be defined from a valuation profile \mathbf{v} with empirical marginal revenues \mathbf{R}' as follows (recall $v_{n+1} = 0$):

$$\phi(v) = \begin{cases} R'_{i+1} & \text{if } v \in [v_{i+1}, v_i) \text{ for some } i \in [n], \text{ and} \\ v & \text{otherwise.} \end{cases} \quad (7.2)$$

This definition is true to the geometric revenue curve interpretation

where the value v can be represented as a diagonal line from the origin with slope v and the marginal revenue for v is the left slope of the revenue curve at its intersection with this line.

For any virtual value function, symmetric environment, and valuation profile; virtual surplus maximization gives an allocation that is monotone, i.e., $v_i > v_j \Rightarrow x_i \geq x_j$, as well as an allocation rule that is monotone, i.e., $z > z^\dagger \Rightarrow x_i(z) \geq x_i(z^\dagger)$. From this allocation and allocation rule the incentive-compatible and envy-free revenues can be calculated and compared. Recall that the maximal envy-free payment of agent i for this allocation comes from equation (7.1) whereas the payment of the incentive compatible mechanism with this allocation rule comes from Corollary 2.13. These payments are related but distinct.

Example 7.5. Compare the envy-free revenue and incentive-compatible revenue corresponding to Example 7.3 ($k = 20$ units, $n = 90$ agents, ten with value ten, and 80 with value two). The virtual value function from equation (7.2) is:

$$\phi(v) = \begin{cases} -180 & v < 2, \\ 1 & v \in [2, 10), \text{ and} \\ v & v \in [10, \infty). \end{cases}$$

We now calculate the revenue of the incentive compatible mechanism that serves the 20 agents with the highest virtual value. In the virtual-surplus-maximizing auction, on the valuation profile \mathbf{v} , the high-valued agents win with probability one and the low-valued agents win with probability $1/8$ (as there are ten remaining units to be allocated randomly among 80 low-valued agents). To calculate payments we must calculate the allocation rule for both high- and low-valued agents. Low-valued agents, by misreporting a high value, win with probability one. The allocation rule for low-valued agents is depicted in Figure 7.3(a). High-valued agents, by misreporting a low value, on the other hand, win with probability $11/81$. Such a misreport leaves only nine high-value-reporting agents and so there are 11 remaining units to allocate randomly to the 81 low-value-reporting agents. The allocation rule for high-valued agents is given in Figure 7.3(b). Payments can be determined from the allocation rules: a high-valued agent pays about 8.9 and a low-valued agent (in expectation) pays $1/4$. The total incentive compatible revenue is about 109. Notice that this revenue is only slightly below the envy-free optimal revenue of 110.

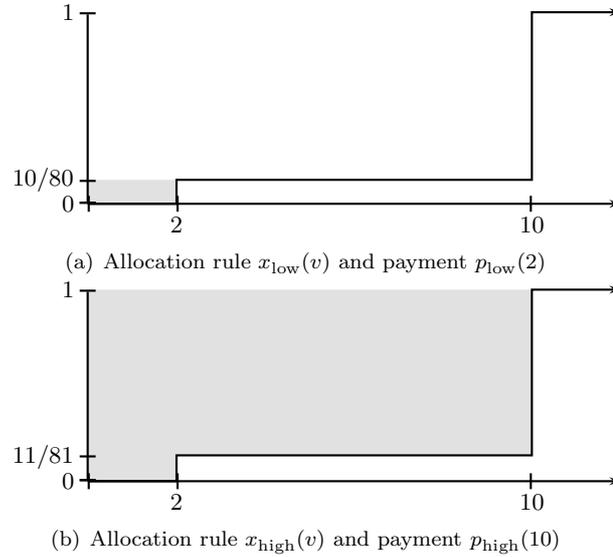


Figure 7.3. The allocation rules for high- and low-valued agents induced by the valuation profile and mechanism with virtual values given in Example 7.5. The incentive compatible payments are given by the area of the shaded region.

The revenue calculation in Example 7.5 is complicated by the fact that when a high-valued agent reports truthfully there are ten remaining units to allocate to the 80 low-valued agents; whereas when misreporting a low value, there are 11 remaining units to allocate to 81 low-value reporting agents. Importantly: the allocation rule for high-valued agents and low-valued agents are not the same (compare Figure 7.3(a) and Figure 7.3(b)). The envy-free payments for both high- and low-valued agents, on the other hand, are calculated from the same “allocation rule” (denoted as $y(\cdot)$ in Theorem 7.21) which is, in fact, identical to the incentive-compatible allocation rule of the low-value agents (Figure 7.3(a)). Thus, the envy-free revenue can be viewed as a relaxation of the incentive-compatible revenue that is simpler and, therefore, more analytically tractable.

We now formalize the fact that the envy-free optimal revenue is an economically meaningful benchmark by showing that it is pointwise normalized (which implies that it is normalized for any i.i.d. distribution).

Theorem 7.24. *For multi-unit environments and any virtual value*

function $\phi(\cdot)$, the envy-free revenue of virtual surplus maximization is at least its incentive-compatible revenue.

Proof. We show that the envy-free payment of agent i is at least her incentive-compatible payment. In particular if we let $x_i(\mathbf{v})$ be the allocation rule of the virtual surplus optimizer, then for $z \leq v_i$, $x_i(z, \mathbf{v}_{-i})$ (as a function of z) is at most the smallest $y(z)$ that satisfies the conditions of Theorem 7.21. Since the incentive-compatible and envy-free payments, respectively, correspond to the area “above the allocation curve” this inequality implies the desired payment inequality.

Since $x_i(z, \mathbf{v}_{-i})$ is monotone, we only evaluate it at $v_j \leq v_i$ and show that $x_i(v_j, \mathbf{v}_{-i}) \geq x_j(\mathbf{v})$. In particular,

$$x_i(v_j, \mathbf{v}_{-i}) = x_j(v_j, \mathbf{v}_{-i}) \geq x_j(\mathbf{v}).$$

The equality above comes from the symmetry of the environment and the fact that agent i and j have the same value in profile (v_j, \mathbf{v}_{-i}) . The inequality comes from greedy maximization with random tie breaking for multi-unit auctions: when agent i reduces her bid to tie agent j 's value v_j the probability that j receives a unit does not decrease as agent i is only less competitive. \square

We will see later that this theorem generalizes beyond multi-unit environments (see Section 7.5). In particular, the only properties of multi-unit environments that we employed in the proof were symmetry and that the greedy algorithm is optimal.

7.3.4 Permutation Environments

Envy freedom is less natural in asymmetric environments such as those given by matroid or downward-closed feasibility constraints. To extend the envy-free benchmark to asymmetric environments we assume a symmetry imposing prior-free analog of the (standard) Bayesian assumption that the agents' value distribution is independent and identically distributed. Specifically, the valuation profile can be arbitrary, but the roles the agents play with respect to the environment (e.g., feasibility constraint or cost function) are assigned by random permutation.

Definition 7.13. Given an environment, specified by cost function $c(\cdot)$, the *permutation environment* is the environment with the identities of the agents uniformly permuted. I.e., for permutation π drawn uniformly at random from all permutations, the permutation environment has cost function $c(\pi(\cdot))$.

Our goal is a prior-free analysis framework for which approximation implies prior-independent approximation in i.i.d. environments. Of course the expected revenue of the optimal auction in an i.i.d. environment is unaffected by a random permutation of the identities of the agents. Therefore, with respect to the goal of obtaining a prior-independent corollary from a prior-free analysis (by Proposition 7.1), it is without loss to assume a permutation environment. Importantly, while a matroid or downward-closed environment may be asymmetric, a matroid permutation or downward-closed permutation environment is inherently symmetric. This symmetry admits a meaningful study of envy freedom.

The environments considered heretofore have been given deterministically, e.g., by a cost function or set system (Chapter 3, Section 3.1). A generalization of this model would be to allow randomized environments. We view a randomized environment as a probability distribution over deterministic environments, i.e., as a convex combination. For the purpose of incentives and performance, we will view mechanism design in randomized environments as follows. First, the agents report their preferences; second, the designer's cost function (or feasibility constraint) is realized; and third, the mechanism for the realized cost function is run on the reported preferences. The performance in such probabilistic environment is measured in expectation over both the randomization in the mechanism and the environment. Agents act before the set system is realized and therefore from their perspective the game they are playing in is the composition of the randomized environment with the (potentially randomized) mechanism.

An example of such a probabilistic environment comes from *display advertising*. Banner advertisements on web pages are often sold by auction. Of course the number of visitors to the web page is not precisely known at the time the advertisers bid; instead, this number can be reasonably modeled as a random variable. Therefore, the environment is a convex combination of multi-unit auctions where the supply k is randomized.

7.3.5 The Envy-free Benchmark

We are now ready to formally define the envy-free benchmark. To do so we must address the potential asymmetry in the environment and the technicality that the envy-free revenue itself may have unbounded resolution (recall the discussion above Definition 7.5 on 208). Finally, we

must give economic justification to the benchmark by showing that it is normalized.

Definition 7.14. For any environment and valuation profile \mathbf{v} , the *envy-free benchmark*, denoted $\text{EFO}^{(2)}(\mathbf{v})$, is the optimal envy-free revenue in the permutation environment on the valuation profile $\mathbf{v}^{(2)}$ where the highest value $v_{(1)}$ is replaced with twice the second highest value $2v_{(2)}$, i.e., $\mathbf{v}^{(2)} = (2v_{(2)}, v_{(2)}, v_{(3)}, \dots, v_{(n)})$.

Theorem 7.25. *For i.i.d., regular, $n \geq 2$ agent, multi-unit environments, the envy-free benchmark is normalized.*

Proof. Recall that for i.i.d. regular distributions \mathbf{F} , the n -agent k -unit Bayesian optimal auction $\text{REF}_{\mathbf{F}}$ is the $k + 1$ st-price auction with the monopoly reserve \hat{v}^* for the distribution. We will show a stronger claim than the statement of the lemma. The *anonymous-reserve benchmark* APX for valuation profile \mathbf{v} is the revenue of the k -unit auction with the best reserve price for the valuation profile $\mathbf{v}^{(2)} = (2v_{(2)}, v_{(2)}, \dots, v_{(n)})$. For $k = 1$, $\text{APX}(\mathbf{v}) = 2v_{(2)}$ and in general $\text{APX}(\mathbf{v}) = \max_{2 \leq i \leq k} i v_{(i)}$. The outcome of the anonymous-reserve benchmark is envy-free for $\mathbf{v}^{(2)}$; therefore, it lower bounds the envy-free optimal revenue for $\mathbf{v}^{(2)}$; and therefore, the normalization of the anonymous-reserve benchmark implies normalization of the envy-free benchmark.

We first argue the $n = 2$ agent, $k = 2$ unit special case (a.k.a., the two-agent digital good environment). Fix an i.i.d. regular distribution \mathbf{F} over valuation profiles. We show that the expected anonymous-reserve benchmark (APX) is at least the performance of the Bayesian optimal mechanism ($\text{REF}_{\mathbf{F}}$) for the distribution, i.e., that $\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\text{APX}(\mathbf{v})] \geq \mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[\text{REF}_{\mathbf{F}}(\mathbf{v})]$.

Recall Theorem 5.1 (also Lemma 5.6) which states that for i.i.d. regular distributions that the revenue of the two-agent second-price auction exceeds that of the single-agent monopoly pricing. Thus, twice the second-price revenue exceeds twice the monopoly pricing revenue. For $n = k = 2$, the former is equal to the expected anonymous-reserve benchmark and the latter is equal to the expected Bayesian optimal revenue.

Under the regularity assumption, the normalization of the anonymous-reserve benchmark for the two-agent digital good environment implies its normalization for multi-unit environments with general $n \geq 2$ agents and k units. To show this extension, consider any $k \geq 2$ and any $n \geq 2$ and condition on the third-highest value $v_{(3)}$. The following argument

shows that $\mathbf{E}[\text{APX}(\mathbf{v}) \mid v_{(3)}] \geq \mathbf{E}[\text{REF}_{\mathbf{F}}(\mathbf{v}) \mid v_{(3)}]$ for $v_{(3)} < \hat{v}^*$ and $v_{(3)} \geq \hat{v}^*$ considered as separate cases.

When the third-highest value $v_{(3)}$ is less than the monopoly price \hat{v}^* , then all agents except for the top two are rejected. The conditional distribution on of the two highest valued agents is regular (the conditioning only truncates and scales the revenue curve; therefore, its convexity is preserved), moreover, the remaining feasibility constraint is that of a digital good. Hence, the normalization for the two-agent digital good environment implies normalization for this conditional environment.

When the third-highest value $v_{(3)}$ is more than the monopoly price \hat{v}^* , then the Bayesian optimal auction $\text{REF}_{\mathbf{F}}$ on \mathbf{v} sells at least two units at a uniform price and the empirical anonymous-reserve revenue from selling the same number of units is pointwise no smaller. Thus, the desired bound holds pointwise.

Now consider the final case of $k = 1$ unit, $n \geq 2$ agent environments. We will reduce normalization of the anonymous-reserve benchmark in this environment to that of the $k = 2$ unit environment. The benchmark in the two environments is the same: the one-unit benchmark is $2v_{(2)}$; the two-unit benchmark is $2v_{(2)}$. The Bayesian optimal revenue is only greater for two units than with one unit. Therefore, normalization for two units implies normalization for one unit. \square

It is evident from this proof that the anonymous-reserve benchmark is also normalized for multi-unit environments. We will prefer to use the envy-free benchmark for three reasons. First, the envy-free benchmark remains normalized for a larger class of distributions which admit a large degree of irregularity (though not arbitrary irregular distributions). Second, the envy-free benchmark is easier to work with as it is structurally a virtual surplus optimization. Third, for position environments discussed subsequently, the envy-free benchmark is linear in the position weights, while the anonymous reserve benchmark is not. This linearity will be important for our analysis.

7.4 Multi-unit Environments

In this section we will discuss two approaches for multi-unit environments. In the first, we will give an approximate reduction to digital good environments. This reduction will give a $\beta + 1$ approximation mechanism for multi-unit environments from any β approximation mechanism for

digital goods. Applied to the prior-free optimal digital good auction, a 2.42 approximation, this approach yields a multi-unit 3.42 approximation. The second approach will be to directly generalize the random sampling optimal price auction to multi-unit environments. This generalization randomly partitions the agents into a market and sample, calculates the empirical distribution of the sample, and then runs optimal multi-unit auction on the market according to the empirical distribution for the sample.

7.4.1 Reduction to Digital Goods

Our first approach is an approximate reduction. For i.i.d. irregular single-item environments, Corollary 4.16 shows that the second-price auction with anonymous reserve is a two approximation to the optimal auction. I.e., the loss in performance from not ironing when the distribution is irregular is at most a factor of two. In fact, this result extends to multi-unit environments (as does the prophet inequality from which it is proved) and the approximation factor only improves. Given the close connection between envy-free optimal outcomes and Bayesian optimal auctions, it should be unsurprising that this result translates between the two models.

Consider the revenue of the surplus maximization mechanism with the best (ex post) anonymous reserve price. For instance, for the k -unit environment and valuation profile \mathbf{v} , this revenue is $\max_{i \leq k} i v_{(i)}$. It is impossible to approximate this revenue with a prior-free mechanism so, as we did for the envy-free benchmark, we exclude the case that it sells to only the highest-valued agent at her value. Therefore, for k -unit environments the *anonymous-reserve benchmark* is $\max_{2 \leq i \leq k} i v_{(i)}$ for $k > 2$ (and $2v_{(2)}$ for $k = 1$), i.e., it is the optimal anonymous reserve revenue for the valuation profile $\mathbf{v}^{(2)} = (2v_{(2)}, v_{(2)}, \dots, v_{(n)})$. Notice that for digital goods, i.e., $k = n$, the anonymous-reserve benchmark is equal to the envy-free benchmark. Of course, an anonymous reserve is envy free so the envy-free benchmark is at least the anonymous-reserve benchmark.

We now give an approximate reduction from multi-unit environments to digital-good environments in two steps. We first show that the envy-free benchmark is at most twice the anonymous-reserve benchmark in multi-unit environments. We then show an approximation preserving reduction from multi-unit to digital-good environments with respect to the anonymous-reserve benchmark. In the next section a more sophisticated approach that attains a better bound is given.

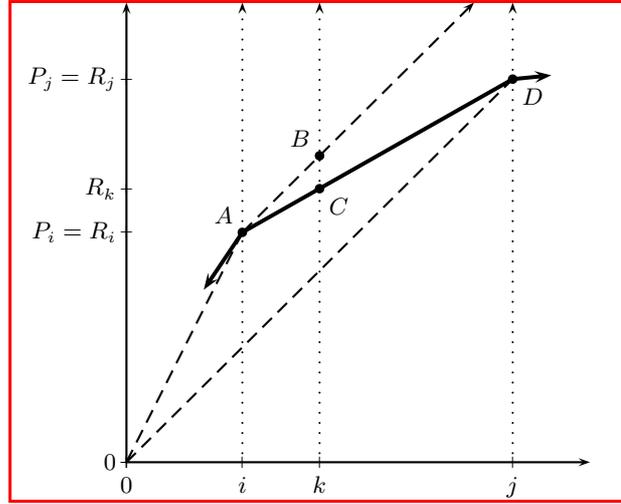


Figure 7.4. Depiction of ironed revenue curve \mathbf{R} for the geometric proof of Theorem 7.26. The solid piece-wise linear curve is \mathbf{R} , the convex hull of \mathbf{P} , and contains the line-segment connecting point $A = (i, P_i)$ and point $D = (j, P_j)$. The envy-free benchmark is achieved at point $C = (k, R_k)$. The parallel dashed lines have slope $v_{(j)}$, the other dashed line has slope $v_{(i)}$.

Theorem 7.26. *For any valuation profile, in multi-unit environments, the envy-free benchmark (resp. optimal revenue) is at most the sum of the anonymous-reserve benchmark (resp. optimal revenue) and $k + 1$ st-price auction revenue, which is at most twice the anonymous-reserve benchmark (resp. optimal revenue).*

Proof. We prove the statement with respect to the optimal revenues and any valuation profile \mathbf{v} and then apply the theorem to the valuation profile $\mathbf{v}^{(2)}$ to obtain the statement with respect to the benchmarks.

If the envy-free optimal revenue sells fewer than k units or the revenue curve is not ironed at k then the anonymous-reserve revenue equals the envy-free revenue and the theorem trivially holds. Otherwise, assume that the envy-free optimal revenue sells all k units and irons between index $i < k$ and $j > k$ (see Figure 7.4). In terms of empirical revenue curves (Definition 7.12), the envy-free optimal revenue for \mathbf{v} is REF = $R_k = C$. Note that the AC line has slope R'_k , i.e., $C = A + (k - i)R'_k$. The line from the origin to D has slope $v_{(j)}$. By geometry $v_{(j)} > R'_k$. Thus,

extending a line from $A = R_i$ with slope $v_{(j)}$ to point $B = A + (k - i)v_{(j)}$ satisfies $B > C$.

The anonymous-reserve revenue exceeds the $k + 1$ st-price auction revenue; thus, twice the anonymous-reserve revenue exceeds the sum of the anonymous-reserve revenue and the $k + 1$ st-price auction revenue. The anonymous-reserve revenue is at least R_i and the $k + 1$ st-price revenue is $k v_{(k+1)} \geq (k - i)v_{(j)}$ (as $j \geq k + 1$ and $i \geq 1$); thus the sum of their revenues exceeds $R_i + (k - i)v_{(j)} = B > C = \text{REF}$. The theorem follows. \square

Theorem 7.26 reduces the problem of approximating the envy-free benchmark to that of approximating the anonymous-reserve benchmark. There is a general construction for converting a digital good auction \mathcal{A} into a limited supply auction and if \mathcal{A} is a β approximation to the anonymous-reserve benchmark (which is identical to the envy-free benchmark for digital goods) then so is the resulting multi-unit auction.

Definition 7.15. The $k \geq 2$ unit restriction \mathcal{A}_k of digital good auction \mathcal{A} is the following:

- (i) Simulate the $k + 1$ st-price auction (i.e., the k highest valued agents win and pay $v_{(k+1)}$).
- (ii) Simulate \mathcal{A} on the k winners $v_{(1)}, \dots, v_{(k)}$.
- (iii) Serve the winners from the second simulation and charge them the higher of their prices in the two simulations.

The 1-unit restriction \mathcal{A}_1 is the second-price auction.

Implicit in this definition is a new notion of mechanism composition (cf. Chapter 5, Section 5.4.2). It is easy to see that this mechanism composition is dominant strategy incentive compatible. In general such a composition is DSIC whenever no winner of the first mechanism can manipulate her value to change the set of winners while simultaneously remaining a winner (see Exercise 7.3); mechanisms that satisfy this property are said to be *non-bossy*.

Theorem 7.27. *If \mathcal{A} is a β approximation to the envy-free benchmark in digital good environments then its multi-unit restriction \mathcal{A}_k is a 2β approximation in multi-unit environments.*

Proof. For 1-unit environments, the second-price auction (with revenue $v_{(2)}$) is a 2-approximation to the 1-unit envy-free benchmark $\text{EFO}^{(2)}(\mathbf{v}) =$

$2v_{(2)}$. For $k \geq 2$ unit environments, the k -unit restriction is a β approximation to the envy-free benchmark restricted to the k highest-valued agents. This benchmark is equal to the anonymous-reserve benchmark on the full set of agents. This benchmark, by Theorem 7.26, is at least half the envy-free benchmark on the full set of agents. Thus, the k -unit restriction is a 2β approximation to the envy-free benchmark. \square

This theorem can be applied to any digital good auction; for instance, from Theorem 7.7 we can conclude that there is a multi-unit auction that is a 4.84 approximation to the envy-free benchmark.

7.4.2 Combination of Benchmarks and Auctions

Theorem 7.26, which shows that the envy-free benchmark is bounded by the sum of the anonymous-reserve benchmark and the $k + 1$ st-price auction revenue, can be employed to construct a $\beta + 1$ approximation for multi-unit environments from a β approximation for digital goods. Applied to the prior-free optimal auction for digital goods, this yields an multi-unit 3.42 approximation. The approach is to view the envy-free benchmark as the sum of two benchmarks, design prior-free mechanisms for each benchmark, and then consider an appropriate convex combination of the two mechanisms to optimize the approximation with respect to the original benchmark. This approach provides two conclusions. First, it gives a modular approach to prior-free mechanism design. Second, it suggests that, even in pursuit of prior-free approximation with respect to the economically motivated envy-free benchmark, it may be useful to understand prior-free approximation for other benchmarks.

Definition 7.16. For benchmark $\mathcal{G}(\mathbf{v}) = \mathcal{G}_A(\mathbf{v}) + \mathcal{G}_B(\mathbf{v})$, mechanism \mathcal{M}_A giving a prior-free β_A approximation to benchmark \mathcal{G}_A , and mechanism \mathcal{M}_B giving a prior-free β_B approximation to benchmark \mathcal{G}_B ; the *prior-free combination* \mathcal{M} runs \mathcal{M}_A with probability $\beta_A/\beta_A + \beta_B$ and \mathcal{M}_B with probability $\beta_B/\beta_A + \beta_B$.

Theorem 7.28. *With respect to Definition 7.16, the prior-free combination \mathcal{M} is a prior-free $\beta = \beta_A + \beta_B$ approximation to benchmark $\mathcal{G}(\mathbf{v}) = \mathcal{G}_A(\mathbf{v}) + \mathcal{G}_B(\mathbf{v})$.*

Proof.

$$\begin{aligned} \mathcal{M}(\mathbf{v}) &= \frac{\beta_A}{\beta_A + \beta_B} \mathcal{M}_A(\mathbf{v}) + \frac{\beta_B}{\beta_A + \beta_B} \mathcal{M}_B(\mathbf{v}) \\ &\geq \frac{\beta_A}{\beta_A + \beta_B} \frac{\mathcal{G}_A(\mathbf{v})}{\beta_A} + \frac{\beta_B}{\beta_A + \beta_B} \frac{\mathcal{G}_B(\mathbf{v})}{\beta_B} = \frac{\mathcal{G}_A(\mathbf{v}) + \mathcal{G}_B(\mathbf{v})}{\beta_A + \beta_B}. \quad \square \end{aligned}$$

As described above, the k -unit restriction \mathcal{A}_k of β approximate digital good auction \mathcal{A} is an β approximation to the multi-unit anonymous reserve benchmark. The $k + 1$ st-price auction is, obviously, a one approximation to the $k + 1$ st-price auction revenue. Thus, we can invoke Theorem 7.26 and Theorem 7.28 to obtain the following corollary. The best multi-unit auction via this construction is attained by instantiating the reduction with the prior-free optimal digital good auction which is a 2.42 approximation (Theorem 7.7).

Corollary 7.29. *For prior-free β -approximate digital-good auction \mathcal{A} , the prior-free combination of the multi-unit restriction \mathcal{A}_k (with probability $\beta/\beta+1$) and the $k + 1$ st-price auction (with probability $1/\beta+1$) is a prior-free $\beta + 1$ approximation to the envy-free benchmark in multi-unit environments. For the prior-free optimal digital-good auction, this multi-unit auction is a prior-free 3.42 approximation.*

7.4.3 The Random Sampling Auction

An alternative approach to the multi-unit auction problem is to directly generalize the random sampling optimal price auction. Intuitively, the random sampling auction partitions the agents into a market and sample and then runs the optimal auction for the empirical distribution of the sample on the market. For digital goods the optimal auction for the empirical distribution sample is just the to post the empirical monopoly price. For multi-unit environments, the optimal auction irons when the empirical distribution of the sample is irregular.

Definition 7.17. *The random sampling (virtual surplus maximization) auction for the k -unit environment*

- (i) randomly partitions the agents into market M and sample S by assigning the highest-valued agent to M and flipping a fair coin for all other agents,
- (ii) computes virtual valuation function ϕ_S for the empirical distribution of S , and

- (iii) maximizes virtual surplus of selling at most k units to S with respect to ϕ_M .

The proof of the following theorem can be derived similarly to the proof of Lemma 7.13 (page 214); we omit the details.

Theorem 7.30. *For multi-unit environments and all valuation profiles, the random sampling auction is a constant approximation to the envy-free benchmark.*

The random sampling auction shares some good properties with optimal mechanisms. The first is that the mechanism on the market is a virtual-surplus optimization. I.e., it sorts the agents in the market by virtual value and allocates to the agents greedily in that order. This property is useful for two reasons. First, in environments where the supply k of units is unknown in advance, the mechanism can be implemented *incrementally*. Each unit of supply is allocated in to the agent remaining in the market with the highest virtual valuation. Second, as we will see in the next section, it can be applied without specialization to matroid permutation and position environments.

7.5 Matroid Permutation and Position Environments

Position environments are important as they model auctions for selling advertisements on Internet search engines such as Google and Microsoft's Bing. In these auctions agents bid for positions with higher positions being better. The feasibility constraint imposed by position auctions is a priori symmetric.

Definition 7.18. A *position environment* is one with n agents, n positions, each position j described by weight w_j . An auction assigns each position j to an agent i which corresponds to setting $x_i = w_j$. Positions are usually assumed to be ordered in non-increasing order, i.e., $w_j \geq w_{j+1}$. (Often w_1 is normalized to one.)

Position auctions correspond to advertising on Internet search engines as follows. Upon each search to the search engine, *organic search results* appear on the left-hand side and *sponsored search results*, a.k.a., advertisements, appear on the right-hand side of the search results page. Advertiser i receives a revenue of v_i in expectation each time her ad is

clicked (e.g., if the searcher buys the advertiser's product) and if her ad is shown in position j it receives *click-through rate* w_j , i.e., the probability that the searcher clicks on the ad is w_j . If the ad is not clicked on the advertiser receives no revenue. Searchers are more likely to click on the top slots than the bottom slots, hence $w_j \geq w_{j+1}$. An advertiser i shown in slot j receives value $v_i w_j$. Understandably, this model of Internet search advertising omits many details of the environment; nonetheless, it has proven to be quite relevant.

We now show that mechanism design for matroid permutation environments can be reduced to that for position environments which can be reduced to that for multi-unit environments. These reductions follow, essentially, because each of these environments are *ordinal*, i.e., because surplus is maximized by the greedy algorithm. The greedy algorithm does not compare magnitudes of the values of agents, it only considers their relative order. This intuition is summarized by the following definition.

Definition 7.19. The *characteristic weights* \mathbf{w} for a matroid are defined as follows: Set $v_i = n - i + 1$, for all i , and consider the surplus maximizing allocation when agents are assigned roles in the set system via random permutation and then the maximum feasible set is calculated, e.g., via the greedy algorithm. Let w_i be the probability of serving agent i , i.e., the i th highest-valued agent.

To see why the characteristic weights are important, notice that since the greedy algorithm is optimal for matroids, the cardinal values of the agents do not matter, just the sorted order. Therefore, e.g., when maximizing virtual value, w_i is the probability of serving the agent with the i th highest virtual value.

Theorem 7.31. *The problem of revenue maximization (or approximation) in matroid permutation environments reduces to the problem of revenue maximization (or approximation) in position environments.*

Proof. We show two things. First, we show that for any matroid permutation environment with characteristic weights \mathbf{w} , the position environment with weights \mathbf{w} has the same optimal expected revenue. Second, for any such environments any position auction can be converted into an auction for the matroid permutation environment that achieves the same expected revenue as the position auction in the position environment given by the characteristic weights of the matroid. These two results imply that any Bayesian, prior-independent, or prior-free approximation

results for position auctions extend to matroid permutation environments.

- (i) Revenue optimal auctions are virtual surplus optimizers. Let \mathbf{w} be the characteristic weights for the given matroid environment. By the definition of \mathbf{w} , the optimal auctions for both the matroid permutation and position environments serve the agent with the j th highest positive virtual value with probability w_j . (In both environments agents with negative virtual values are discarded.) Expected revenue equals expected virtual surplus; therefore, the optimal expected revenues in the two environments are the same.
- (ii) Consider the following matroid permutation mechanism which is based on the position auction with weights \mathbf{w} . The input is \mathbf{v} . First, simulate the position auction and let \mathbf{j} be the assignment where j_i is the position assigned to agent i , or $j_i = \perp$ if i is not assigned a slot. Reject all agents i with $j_i = \perp$. Now run the greedy matroid algorithm in the matroid permutation environment on input $v_i^\dagger = n - j_i + 1$ and output its outcome.

Notice that any agent i is allocated in the matroid permutation setting with probability equal to the expected weight of the position it is assigned in the position auction. Therefore the two mechanisms have the exact same allocation rule (and therefore, the exact same expected revenue). \square

We are now going to reduce the design of position auctions to that of multi-unit auctions. This reduction implies that the prior-free approximation factor for multi-unit environments extends to matroid permutation and position environments. Furthermore, the mechanism that gives this approximation can be derived from the multi-unit auction.

Theorem 7.32. *The problem of revenue maximization (or approximation) in position auctions reduces to the problem of revenue maximization (or approximation) in k -unit auctions.*

Proof. This proof follows the same high-level argument as the proof of Theorem 7.31.

Let $w'_j = w_j - w_{j+1}$ be the difference between successive position weights. Recall that without loss of generality $w_1 = 1$ so $\sum_j w'_j = 1$ and \mathbf{w}' can be interpreted as a probability measure over $[m]$.

- (i) The expected revenue of an optimal position auction is equal to the

expected revenue of the convex combination of optimal j -unit auctions under measure \mathbf{w}' . In the optimal position auction and the optimal auction for the above convex combination of multi-unit auctions the agent with the j th highest positive virtual value is served with probability w_j . (In both settings agents with negative virtual values are discarded.) Therefore, the expected revenues in the two environments are the same.

- (ii) Now consider the following position auction which is based on a multi-unit auction. Simulate a j -unit auction on the input \mathbf{v} for each $j \in [m]$ and let $x_i^{(j)}$ be the (potentially random) indicator for whether agent i is allocated in simulation j . Let $x_i = \sum_j x_i^{(j)} w'_j$ be the expected allocation to j in the convex combination of multi-unit auctions given by measure \mathbf{w}' . The vector of position weights \mathbf{w} majorizes the allocation vector \mathbf{x} in the sense that $\sum_i^k w_i \geq \sum_i^k x_i$ (and with equality for $k = m$). Therefore we can write $\mathbf{x} = S\mathbf{w}$ where S is a doubly stochastic matrix. Any doubly stochastic matrix is a convex combination of permutation matrices, so we can write $S = \sum_\ell \rho_\ell P_\ell$ where $\sum_\ell \rho_\ell = 1$ and each P_ℓ is a permutation matrix (Birkhoff–von Neumann Theorem). Finally, we pick an ℓ with probability ρ_ℓ and assign the agents to positions according to the permutation matrix P_ℓ . The resulting allocation is exactly the desired \mathbf{x} .

Let β be the worst case, over number of units k , approximation factor of the multi-unit auction in the Bayesian, prior-independent, or prior-free sense. The position auction constructed is at worst a β approximation in the same sense. \square

We conclude that matroid permutation auctions reduce to position auctions which reduce to multi-unit auctions. But multi-unit environments are the simplest of matroid permutation environments, i.e., the uniform matroid (Section 4.6.1, page 131), where even the fact that the agents are permuted is irrelevant because uniform matroids are inherently symmetric. Therefore, from the perspective of optimization and approximation all of these problems are equivalent.

It is important to note, however, that this reduction may not preserve non-objective aspects of the mechanism. For instance, we have discussed that anonymous reserve pricing is a two approximation to virtual surplus maximization in multi-unit environments (e.g., Corollary 4.16 and Theorem 7.26). The reduction from matroid permutation and position environments does not imply that surplus maximization with an anonymous reserve gives a two approximation in these more general envi-

ronments. This is because in the multi-unit two approximation via an anonymous reserve, the reserve is tailored to k , the number of units. Therefore, constructing a position auction or matroid mechanism would require simulating the multi-unit auction with potentially distinct reserve prices for each supply constraint; the resulting mechanism will not generally be an anonymous-reserve mechanism.

In fact, for i.i.d., irregular, position and matroid permutation environments the surplus maximization mechanism with anonymous reserve is not generally a constant approximation to the optimal mechanism. The approximation factor via the anonymous reserve in these environments is $\Omega(\log n / \log \log n)$, i.e., there exists a distribution and matroid permutation and position environments such that the anonymous-reserve mechanism has expected revenue that is a $\Theta(\log n / \log \log n)$ multiplicative factor from the optimal mechanism revenue (Exercise 7.4). The same inapproximation result holds with comparison between the anonymous-reserve and envy-free benchmarks.

Theorem 7.33. *There exists an i.i.d. distribution (resp. valuation profile), a matroid permutation environment, and position environment such that the (optimal) anonymous-reserve mechanism (resp. benchmark) is a $\Theta(\log n / \log \log n)$ approximation the Bayesian optimal mechanism (resp. envy-free benchmark).*

Implicit in the above discussion (and reductions) is the assumption that the characteristic weights for a matroid permutation setting can be calculated, or fundamentally, that the weights in the position auction are precisely known. Notice that in our application of position auctions to advertising on Internet search engines the position weights were the likelihood of a click for an advertisement in each position. These weights can be estimated but are not known exactly. The general reduction from matroid permutation and position auctions to multi-unit auctions requires foreknowledge of these weights.

Recall from the discussion of the multi-unit random sampling auction (Definition 7.17) that, as a virtual surplus maximizer, it does not require foreknowledge of the supply k of units. Closer inspection of the reductions of Theorem 7.32 reveals that if the given multi-unit auction is a virtual surplus maximizer then the weights do not need to be known to calculate the appropriate allocation. Simply maximize the virtual surplus for the realized environment.

In the definition of permutation environments, it is assumed that the agents are unaware of their roles in the set system, i.e., the agents' in-

centives are taken in expectation over the random permutation. A mechanism that is incentive compatible in this permutation model may not generally be incentive compatible if agents do know their roles. Therefore, matroid permutation auctions that result from the above reductions are not generally incentive compatible without the uniform random permutation. Of course the random sampling auction is a virtual surplus maximizer for the market and virtual surplus maximizers are dominant strategy incentive compatible (Theorem 3.14). Thus, the reduction applied to the random sampling auction is incentive compatible even if the permutation is known.

Corollary 7.34. *For any matroid environment and valuation profile, the random sampling auction is dominant strategy incentive compatible and when the values are randomly permuted, its expected revenue is a β approximation to the envy-free benchmark where β is its approximation factor for multi-unit environments.*

7.6 Downward-closed Permutation Environments

In this section we consider downward-closed permutation environments. In multi-unit, position, and matroid permutation environments, virtual surplus maximization is ordinal, i.e., it depends on the relative order of the virtual values and not their magnitudes. In contrast, the main difficulty of more general downward-closed environments is that virtual surplus maximization is not generally ordinal. Nonetheless, variants of the random sampling (virtual surplus maximization) and the random sampling profit extraction auctions give constant approximations to the envy-free benchmark in downward-closed environments. We will describe only the latter result, which can be viewed as transforming the non-ordinal environment into an ordinal one.

The first step in this construction is to generalize the notion of a profit extractor (from Section 7.2.4). Our approach to profit extraction in downward-closed permutation environments will be the following. The true (and unknown) valuation profile is \mathbf{v} . Suppose we knew a profile \mathbf{v}^\dagger that was a coordinate-wise lower bound on \mathbf{v} , i.e., $v_{(i)} \geq v_{(i)}^\dagger$ for all i (short-hand notation: $\mathbf{v} \geq \mathbf{v}^\dagger$). A natural goal with this side-knowledge would be to design an incentive compatible mechanism that obtains at least the envy-free optimal revenue for \mathbf{v}^\dagger . We refer to mechanism that obtains this revenue, in expectation over the random permutation and

whenever the coordinate-wise lower-bound assumption holds, as a profit extractor.

Definition 7.20. The *downward-closed profit extractor* for \mathbf{v}^\dagger is the following:

- (i) Sort \mathbf{v} and \mathbf{v}^\dagger in decreasing order.
- (ii) Reject all agents if there exists an i with $v_i < v_i^\dagger$.
- (iii) Calculate the empirical virtual values ϕ^\dagger for \mathbf{v}^\dagger .
- (iv) For all i , assign the i th highest-valued agent the i th highest virtual value ϕ_i^\dagger .
- (v) Serve the agents to maximize the virtual surplus.

Theorem 7.35. *For any downward-closed environment and valuation profiles \mathbf{v} and \mathbf{v}^\dagger , the downward-closed profit extractor for \mathbf{v}^\dagger is dominant strategy incentive compatible and if $\mathbf{v} \geq \mathbf{v}^\dagger$ then its expected revenue under a random permutation is at least the envy-free optimal revenue for \mathbf{v}^\dagger .*

Proof. See Exercise 7.5. □

To make use of this profit extractor we need to find a \mathbf{v}^\dagger that satisfies the assumption of the theorem and that is non-manipulable. The idea is to use biased random sampling. In particular, if the agents are partitioned into a sample with probability $p < 1/2$ and market with probability $1 - p$, then there is a high probability the valuation profile of the sample is a coordinate-wise lower bound on that of the market. Furthermore, we will show that even conditioned on this event, the expected optimal envy-free revenue of the sample approximates the envy-free benchmark. The approximate optimality of the mechanism follows.

Definition 7.21. The *biased (random) sampling profit extraction mechanism* for downward-closed environments (with parameters $p \in (0, 1/2)$ and $\ell \in \{0, 1, 2, \dots\}$) is:

- (i) Assign the top ℓ agents to the market M .
- (ii) Randomly partition the remaining agents into S (with probability p) and M (with probability $1 - p$).
- (iii) Reject agents in S .
- (iv) Run the downward-closed profit extractor for \mathbf{v}_S on M .³

³ The payments of the top ℓ agents are adjusted as follows. Flipping a biased coin for each such agent, but if she ends up in the sample (with probability p), she can buy her way into the market by agreeing to pay at least $v_{(\ell+1)}$. In such a

Lemma 7.36. *The biased sampling profit extraction mechanism is dominant strategy incentive compatible.*

Proof. Fix any outcome of the n coins. Each agent i faces a critical value. Pretend the agent is in the market, and simulate the rest of the auction. The profit extractor is deterministic and dominant strategy incentive compatible; thus by Theorem 7.6, it induces a critical value \hat{v}_i . Now consider i 's coin. If the coin puts i in the market then she is offered critical value \hat{v}_i ; if the coin puts i in the sample, then she is offered $\max(v_{(\ell+1)}, \hat{v}_i)$, i.e., she wins only if she is in the top ℓ and would win in the profit extractor. \square

The following lemma, which is key to the analysis, shows that the probability that $\mathbf{v}_M \geq \mathbf{v}_S$ in the biased sampling profit extraction auction is at least $1 - (p/1-p)^{\ell+1}$.

Lemma 7.37. *The probability of ruin of a biased random walk on the integers,⁴ that steps back with probability $p < 1/2$, steps forward with probability $1 - p$, and starts from position one; is exactly $p/1-p$. If it starts at position k the probability of ruin is $(p/1-p)^k$.*

Proof. The proof is similar to that of Lemma 7.14. See Exercise 7.6. \square

The remainder of this section follows the the approach of prior-free combination developed in Section 7.4.2. Lemma 7.39 will bound the envy-free benchmark by the sum of two benchmarks, the envy-free benchmark restricted to the two highest-valued agents and the envy-free optimal revenue excluding these two agents. Lemma 7.40 will show that the second-price auction (to serve at most one agent) is a two approximation to the first benchmark and Lemma 7.41 will show that a biased sampling profit extraction auction is a 4.51 approximation to the second benchmark. We will conclude by Theorem 7.28 that the prior-free combination (Definition 7.16) of the two auctions is a 6.51 approximation to the envy-free benchmark.

Theorem 7.38. *In downward-closed permutation environments, the prior-free combination of the second-price auction with a biased sampling profit extraction auction is a 6.51 approximation to the envy-free benchmark.*

case, her final payment is the maximum of her payment in the profit extraction mechanism and $v_{(\ell+1)}$.

⁴ Recall, the probability of ruin of a random walk is the probability that it ever reaches position zero

Lemma 7.39. *For any valuation profile \mathbf{v} , the envy-free optimal revenue for a subset S of agents is a subadditive function S . In particular, $\text{EFO}(\mathbf{v}) \leq \text{EFO}(v_1, v_2) + \text{EFO}(\mathbf{v}_{-1,2})$.*

Proof. Observe for disjoint sets A and B of agents,

$$\begin{aligned} \text{EFO}(\mathbf{v}_{A \cup B}) &= \text{EFO}_A(\mathbf{v}_{A \cup B}) + \text{EFO}_B(\mathbf{v}_{A \cup B}) \\ &\leq \text{EFO}(\mathbf{v}_A) + \text{EFO}(\mathbf{v}_B). \end{aligned}$$

The first line follows by definition where $\text{EFO}_A(\mathbf{v}_{A \cup B})$ denotes the contribution to the envy-free optimal revenue of $A \cup B$ from the agents in A , likewise for B . Of course, the envy-free optimal outcome for $A \cup B$ is envy free with respect to subset A . However, if we are only to consider envy-freeness constraints of A , then this outcome for $A \cup B$ is not necessarily optimal. Thus, $\text{EFO}_A(\mathbf{v}_{A \cup B}) \leq \text{EFO}(\mathbf{v}_A)$; likewise for B ; and the second line follows. The left- and right-hand side of this equation give the definition of subadditivity. \square

Lemma 7.40. *For any downward-closed environment, the second-price auction is a 2-approximation to the envy-free benchmark restricted to the two highest-valued agents.*

Proof. Assume all singleton sets are feasible with respect to the downward-closed environment and the two highest valued agents have values $v_1 \geq v_2$. The second-price auction, which always only serves a single agent, is feasible and its revenue is v_2 . For the valuation profile $\mathbf{v}^\dagger = (2v_2, v_2)$, the revenues are $\mathbf{R}^\dagger = (2v_2, 2v_2)$ and the marginal revenues are $(2v_2, 0)$. Thus, the envy-free optimal revenue is obtained by only serving the first agent at a price of $2v_2$. \square

Lemma 7.41. *For any downward-closed permutation environment and any valuation profile, the biased sampling profit extraction auction with $p = .29$ and $\ell = 2$ is a 4.51 approximation to the envy-free optimal revenue on the valuation profile without the two highest-valued agents.*

Proof. Index the two highest-valued agents by 1 and 2. Let $\text{REF}(\mathbf{v}) = \text{EFO}(\mathbf{v}_{-1,2})$ be the envy-free optimal revenue on the valuation profile without the two highest valued agents, and $\text{APX}(\mathbf{v})$ be the expected

revenue of the biased sampling profit extraction mechanism. We have,

$$\begin{aligned}
 \text{APX}(\mathbf{v}) &\geq \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathbf{v}_M \geq \mathbf{v}_S] \Pr[\mathbf{v}_M \geq \mathbf{v}_S] \\
 &= \mathbf{E}[\text{EFO}(\mathbf{v}_S)] - \mathbf{E}[\text{EFO}(\mathbf{v}_S) \mid \mathbf{v}_M \not\geq \mathbf{v}_S] \Pr[\mathbf{v}_M \not\geq \mathbf{v}_S]. \\
 &\geq p \text{EFO}(\mathbf{v}_{-1,2}) - \text{EFO}(\mathbf{v}_{-1,2}) \Pr[\mathbf{v}_M \not\geq \mathbf{v}_S] \\
 &\geq \left(p - \left(\frac{p}{1-p}\right)^3\right) \text{REF}(\mathbf{v}).
 \end{aligned}$$

The first line is by the definition of the mechanism and Theorem 7.35. The second line is by the definition of conditional expectation. The first and second part of the third line are by subadditivity (Lemma 7.39) and monotonicity of the envy-free optimal revenue, respectively. The last line is from $\Pr[\mathbf{v}_M \not\geq \mathbf{v}_S] \leq (p/1-p)^3$ as guaranteed by Lemma 7.37 for a random walk starting at position $\ell + 1 = 3$.

The expression $p - (p/1-p)^3$ is maximized at $p \approx 0.29$ giving an approximation of about 4.51 with respect to $\text{REF}(\mathbf{v})$. \square

Exercises

- 7.1 Complete the prior-free analysis framework for the objective of residual surplus in a two-agent single-item environment. The residual surplus is the sum of the values of the winners less any payments made.
- Identify a normalized benchmark.
 - Identify a distribution for which all auctions have the same residual surplus.
 - Give a lower bound on the resolution of your benchmark.
 - Give an upper bound on the prior-free optimal approximation with respect to your benchmark.

Ideally, your lower bound on resolution should match your upper bound on prior-free optimal approximation.

- 7.2 Consider the design of prior-free incentive-compatible mechanisms with revenue that approximates the (optimal) social-surplus benchmark, i.e., $\text{OPT}(\mathbf{v})$, when all values are known to be in a bounded interval $[1, h]$. For downward-closed environments, give a $\Theta(\log h)$ approximation mechanism.
- 7.3 Consider a generalization of the mechanism composition from the construction of the multi-unit variant of a digital good auction, i.e., where the $k + 1$ st-price auction and the given digital good

auction are composed (Definition 7.15). Two dominant strategy incentive compatible mechanisms A and B can be composed as follows: Simulate mechanism A ; run mechanism B on the winners of mechanism A ; and charge the winners of B the maximum of their critical values for A and B . A deterministic mechanism is *non-bossy* if there are no two values for any agent i such that the sets of winners of the mechanism are distinct but contain i .

- (a) Show that the composite mechanism is dominant strategy incentive compatible when mechanism A is non-bossy.
 - (b) Show that the surplus maximization mechanism in any single-dimensional environment is non-bossy.
- 7.4 Prove the envy-free variant of Theorem 7.33, i.e., that there exists a valuation profile and a position environment for which the anonymous-reserve benchmark is a $\Omega(\log n / \log \log n)$ approximation to the envy-free benchmark.
- 7.5 Show that for any downward-closed environment and valuation profiles \mathbf{v} and \mathbf{v}^\dagger , the downward-closed profit extractor for \mathbf{v}^\dagger is dominant strategy incentive compatible and if $\mathbf{v} \geq \mathbf{v}^\dagger$ then its expected revenue under random permutation is at least the envy-free optimal revenue for \mathbf{v}^\dagger . I.e., prove Theorem 7.35.
- 7.6 Prove Lemma 7.37: The probability of ruin of a biased random walk on the integers; that steps back with probability $p < 1/2$, steps forward with probability $1 - p$, and starts from position one; is exactly $p/1-p$. If it starts at position k the probability of ruin is $(p/1-p)^k$.

Chapter Notes

The prior-free auctions for digital good environments were first studied by Goldberg et al. (2001) where the deterministic impossibility theorem and the random sampling optimal price auction were given. The random sampling optimal price auction was shown to be a constant approximation by Goldberg et al. (2006). The proof that the random sampling auction is a prior-free 15 approximation is from Feige et al. (2005); the bound was improved to 4.68 by Alaei et al. (2009). The profit extraction mechanism and the random sampling profit extraction mechanism were given by Fiat et al. (2002). The extension of this auction to three partitions was studied by Hartline and McGrew (2005).

The lower-bound on the approximation factor of prior-free auctions for digital goods of 2.42 was given by Goldberg et al. (2004); this bound was proven to be tight by Chen et al. (2014b). For the special cases of $n = 2$ and $n = 3$ agents the form of the optimal auction is known. For $n = 2$, Fiat et al. (2002) showed that the second-price auction is optimal and its approximation ratio is $\beta^* = 2$. For $n = 3$, Hartline and McGrew (2005) identified the optimal three-agent auction and showed that its approximation ratio is $\beta^* = 13/6 \approx 2.17$.

The formal prior-free design and analysis framework for digital good auctions was given by Goldberg et al. (2006). This framework was refined for general symmetric auction problems and grounded in the theory of Bayesian optimal auctions by Hartline and Roughgarden (2008). The connection between prior-free mechanism design and envy-freeness was given by Devanur et al. (2015) (originally as Hartline and Yan, 2011).

The 2-approximate reduction from multi-unit to digital-good environments combines results from Fiat et al. (2002) and Devanur et al. (2015). The improved reduction via “prior-free combination” that gives a multi-unit $\beta + 1$ approximation from a digital-good β approximation is from Chen et al. (2014a).

Analysis of the random sampling auction for limited supply, position, matroid permutation, and downward-closed permutation environments was given by Devanur et al. (2015) (originally as Devanur and Hartline, 2009). For multi-unit auctions they prove the random sampling auction is a 9.6 approximation to the envy-free benchmark (i.e., Theorem 7.30) by extending the analysis of Alaei et al. (2009). They prove the equivalence between distributions over multi-unit environments, position environments, and matroid permutation environments which allows the 9.6 approximation bound for multi-unit environments to extend. For downward-closed permutation environments they give a variant of the random sampling auction that is a prior-free 189 approximation.

The downward-closed profit extractor is from Ha and Hartline (2011). Devanur et al. (2013) study the random sampling profit extraction auction, similar to the one described in this chapter, and show that it is a 7.5 approximation in downward-closed permutation environments. (They also give a variant of the auction for the case that the agents have a common budget.) The biased sampling profit extraction auction (Definition 7.21) and its analysis (Theorem 7.38) are from Chen et al. (2014a).

This chapter omitted discussion of a very useful technique for designing prior-free mechanisms using a “consensus mechanism” on statisti-

cally robust characteristics of the input. In this vein the consensus estimates profit extraction mechanism from Goldberg and Hartline (2003) obtains a 3.39 approximation for digital goods. This approach is also central in obtaining an asymmetric deterministic auction that gives a good approximation (Aggarwal et al., 2005). Ha and Hartline (2011) extend the consensus approach to downward-closed permutation environments.

This chapter omitted asymptotic analysis of the random sampling auction which is given Balcan et al. (2008). This analysis allows agents to be distinguished by publicly observable attributes and agents with distinct attributes may receive distinct prices.

8

Multi-dimensional and Non-linear Preferences

Chapters 2–7 focused on environments where the agent preferences are single-dimensional and linear, i.e., an agent’s utility for receiving a service at a given price is her value minus the price. In many settings of interest, however, agents’ preferences are multi-dimensional and non-linear. Common examples include (a) multi-item environments where an agent has different values for each item, (b) agents that are financially constrained, e.g., by a budget, where an agent’s utility is her value minus price as long as the price is at most her budget (if the budget is private knowledge of the agent then this agent is multi-dimensional and non-linear), or (c) agents who are risk averse; a common way to model risk averse preferences is to assume an agent’s utility is given by a concave function of her value minus price.

The challenge posed by multi-dimensional non-linear preferences is three-fold. First, multi-dimensional type spaces can be large, even optimizing single-agent problems (like those in Section 3.4 on page 79) may be analytically or computationally intractable. Second, we should not expect the revenue-linearity condition of Definition 3.16 on page 83 to hold when agents have non-linear preferences (in fact, it also does not generally hold for linear but multi-dimensional preferences). Third, often settings with multi-dimensional agents have multiple items and the externality an agent imposes on the other agents when she is served one of these items is, therefore, multi-dimensional as well.

Our approach to multi-dimensional and non-linear preferences will be to address the challenges above in the order given.

8.1 Optimal Single-agent Mechanisms

A general agent has a type t drawn from an abstract type space \mathcal{T} according to a distribution F . A mechanism can produce an outcome w from a general space of outcomes \mathcal{W} . Outcomes can be complicated objects but we will project them down onto our familiar notation for allocation and payment as follows. When the outcome includes a payment, it is denoted by p . In general an outcome can include multiple alternatives by which an agent is allocated or not allocated, $x = 1$ denotes the former and $x = 0$ denotes the latter. Both x and p are encoded by w , which may also encode other aspects of the outcome.

Recall any mechanism for a single agent, by the taxation principle, can be represented by a menu, i.e., a set of outcomes, where the agent picks her favorite outcome from the menu. Simple mechanisms can best be described by the set of outcomes they allow. For individually rational mechanisms it is without loss to assume that the outcome \emptyset , which does not allocate and requires no payment, is available in \mathcal{M} and that all types obtain zero utility for this outcome. When listing the outcomes of the mechanism we will omit \emptyset . When describing more complex mechanisms it will be convenient to index the outcomes by the types that prefer them, e.g., as $w(t)$ for $t \in \mathcal{T}$; thus,

$$\mathcal{M} = \{w(t) : t \in \mathcal{T}\}.$$

Indexing as such, incentive compatibility and individual rationality can be expressed as follows, respectively.

$$\begin{aligned} u(t, w(t)) &\geq u(t, w(s)), & \forall t, s \in \mathcal{T}, \\ u(t, w(t)) &\geq 0, & \forall t \in \mathcal{T}. \end{aligned}$$

The subsequent developments of this section and chapter will be illustrated two representative examples, (a) a single-dimensional agent with a public budget (a non-linear preference), and (b) a (multi-dimensional)

Chapter 8: Topics Covered.

- unit-demand and public budget preferences,
- revenue linearity (revisited from Section 3.4.4 on page 83),
- interim feasibility,
- implementation by stochastic weighted optimization, and
- multi-dimensional virtual values via amortized analysis.

unit-demand agent with linear utility given by her value for the alternative obtained minus her payment.

Definition 8.1. A *public budget* agent has a (single-dimensional) value t for service and a public budget B . Her utility is linear for outcomes with required payment that is within her budget, and infinitely negative for outcomes with payments that exceed her budget. For outcome $w = (x, p)$, where x denotes the probability she is allocated and p is her required payment, her utility is $u(t, w) = tx - p$ when $p \leq B$ (and $-\infty$, otherwise).

Definition 8.2. A *unit-demand* agent desires one of m alternatives. Her type $t = (\{t\}_1, \dots, \{t\}_m)$ is m -dimensional where $\{t\}_j$ is her value for alternative j . Her utility is linear; an outcome w is given by a payment and a probability measure over the m alternatives and nothing. For outcome $w = (\{x\}_1, \dots, \{x\}_m, p)$, where $\{x\}_j$ denotes the probability she obtains alternative j and p is her required payment, her utility is $u(t, w) = \sum_j \{t\}_j \{x\}_j - p$. The probability she receives any allocation is $x = \sum_j \{x\}_j$.

The single-agent problems we consider are (i) the unconstrained single-agent problem, (ii) the ex ante constrained single-agent problem, and (iii) the interim constrained single-agent problem. Similar to Section 3.4, we are looking to understand the single-agent mechanisms that correspond to $R(1)$, $R(\hat{q})$, and $\mathbf{Rev}[\hat{y}]$.

Definition 8.3. A *single-agent problem* is specified by the type space, outcome space, and distribution over types as $(\mathcal{T}, \mathcal{W}, F)$, and a feasibility constraint. The feasibility constraints of single-agent problems are:

- (i) *Unconstrained:* any mechanism is feasible. The optimal unconstrained mechanism's revenue is denoted $R(1)$. The *monopoly quantile* \hat{q}^* is defined to be its ex ante sale probability.
- (ii) *(Weak) ex-ante constrained:* for ex ante constraint \hat{q} , a mechanism is feasible if its ex ante allocation probability at most \hat{q} , i.e., $\mathbf{E}_{t \sim F}[x(t)] \leq \hat{q}$. The optimal \hat{q} ex ante mechanism's revenue is denoted $R(\hat{q})$.
- (iii) *(Weak) interim constrained:* for interim constraint \hat{y} (a monotone non-increasing function from $[0, 1]$ to $[0, 1]$), a mechanism \mathcal{M} is feasible if, for any subspace of types $S \subset \mathcal{T}$ with measure $\hat{q} = \mathbf{Pr}[t \in S]$ under distribution F , the probability that a type in the subset is allocated under \mathcal{M} is at most that of a quantile $q \in [0, \hat{q}]$ under constraint \hat{y} . In other words, the allocation constraint \hat{y} is satisfied if for type

distribution F , quantile $q \sim U[0, 1]$, and subspace of types S with $\Pr_{t \sim F}[t \in S] = \hat{q}$,

$$\mathbf{E}[x(t) \mid t \in S] \leq \mathbf{E}[\hat{y}(q) \mid q \leq \hat{q}].$$

The optimal \hat{y} interim mechanism's revenue is denoted $\mathbf{Rev}[\hat{y}]$.

The optimal revenues from these convex maximization problems satisfy the standard concavity properties. Thus,

- the revenue curve $R(\cdot)$ is concave, and
- the interim optimal revenue is concave, i.e., $\mathbf{Rev}[\hat{y}] \geq \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$ for $\hat{y} = \hat{y}^\dagger + \hat{y}^\ddagger$.

See the Technical Note on page 254 for further discussion.

8.1.1 Public Budget Preferences

In this section we will describe optimal mechanisms for the three single-agent problems and agents with a public budget. Formal derivations of these optimal mechanisms are deferred to Section 8.6.

Example 8.1. The exemplary *uniform public-budget agent* has (single-dimensional) private type t uniform on type space $\mathcal{T} = [0, 1]$ and public budget $B = 1/4$.

Technical Note. In Section 3.4 the ex ante constraints of $R(\cdot)$ and $\mathbf{Rev}[\cdot]$ were required to hold with equality, i.e., $\mathbf{E}[x(t)] = \hat{q}$ and $\mathbf{E}[x(t)] = \mathbf{E}[\hat{y}(q)]$, respectively. For multi-dimensional and non-linear preferences, single-agent mechanisms can be ill-behaved when required to serve types that the optimal unconstrained mechanism would reject. The “weak” definitions of Definition 8.3 avoid the resulting technicalities and are without loss for downward-closed environments. These weakened definitions of the single-agent problems, relative to those of Chapter 3, satisfy the following additional properties.

- The revenue curve $R(\cdot)$ is monotonically non-decreasing on $\hat{q} \in [0, \hat{q}^*]$ and constant on $\hat{q} \in [\hat{q}^*, 1]$.
- The unconstrained optimal revenue is given by $R(1) = R(\hat{q}^*)$.

The results of this chapter will be restricted to downward closed environments.

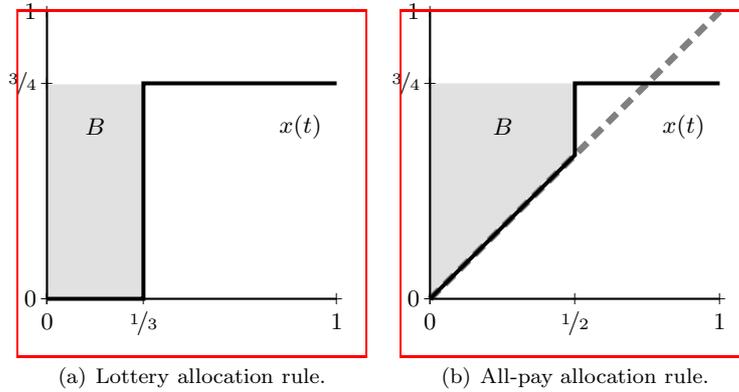


Figure 8.1. Depicted are the allocation rules for the $3/4$ -lottery and the two-agent all-pay auction (for public budget $B = 1/4$). In each case, the ex ante allocation probability is $\hat{q} = 1/2$ and the highest type $t = 1$ pays her full budget B . In subfigure (a), the $2/3$ measure of highest types buy the $3/4$ -lottery at price $B = 1/4$ for a total revenue of $1/6$. In subfigure (b), the allocation rule of the two-agent all-pay auction without the budget constraint is also depicted (think, gray, dashed line).

We begin by describing a few mechanisms for an agent with uniformly distributed private value and public budget $B = 1/4$ (Example 8.1). Recall any mechanism for a single agent, by the taxation principle, can be represented by a menu, where the agent picks her favorite outcome from the menu. An outcome is a pair $w = (x, p)$, and individual rationality requires that the outcome $\emptyset = (0, 0)$ is implicitly in the menu of any mechanism. The following are mechanisms that sell with ex ante probability $\hat{q} = 1/2$ and do not exceed the agent's budget.

- A $3/4$ lottery at price $1/4$ is $\mathcal{M} = \{(3/4, 1/4)\}$. The utility of an agent with type t for this lottery pricing is $3/4 t - 1/4$ and, thus, types $t \in [1/3, 1]$ buy. The ex ante sale probability is $3/4 \cdot 2/3 = \hat{q} = 1/2$ as desired. Its allocation rule is the following:

$$x(t) = \begin{cases} 0 & \text{if } t \leq 1/3, \text{ and} \\ 3/4 & \text{otherwise.} \end{cases}$$

- In a two-agent all-pay auction, types $t \in [1/2, 1]$ bid the budget $B = 1/4$, remaining types bid $1/2 t^2$ (the usual all-pay equilibrium, cf. Section 2.8 on page 38). The agents are symmetric; thus, each wins with ex ante

probability $\hat{q} = 1/2$. Each agent's allocation rule is the following:

$$x(t) = \begin{cases} t & \text{if } t \leq 1/2, \text{ and} \\ 3/4 & \text{otherwise.} \end{cases}$$

These allocation rules are depicted in Figure 8.1. In each the payment of an agent with the highest type $t = 1$ is exactly her budget $B = 1/4$. To understand where the second allocation rule comes from, consider running an all-pay auction for two agents with types uniformly distributed on $[0, 1]$. Absent a budget constraint, an agent with the highest type $t = 1$ would win with probability one and pay $1/2$. This exceeds the budget and this agent would prefer to lower her bid relative to the non-budgeted equilibrium. In fact, types in $[3/4, 1]$ prefer to lower their bids to B , this causes types in $(1/2, 3/4]$ to prefer to raise their bids to B , and leaves type $t = 1/2$ indifferent between bidding in the unconstrained all-pay equilibrium and bidding B (Exercise 8.1). The given allocation rule results. Naturally, there are many other possible mechanisms that satisfy the budget constraint and have ex ante sale probability of $\hat{q} = 1/2$; the optimal one, however, is the $3/4$ -lottery.

Section 8.6 derives a characterization of optimal mechanisms for the three single-agent optimization problems for an agent with a public budget and uniformly distributed type. For the example of budget $B = 1/4$ and uniformly distributed types, these optimal mechanisms are as follows.

- The unconstrained optimal mechanism posts price $B = 1/4$, sells to the $3/4$ measure of types $t \in [1/4, 1]$, and has expected revenue $3/16$. Its ex ante sale probability is $\hat{q}^* = 1 - B = 3/4$.
- The ex ante optimal mechanism for $\hat{q} \leq \hat{q}^*$ is the $\hat{q} + B$ lottery at price B , i.e., $\mathcal{M}^{\hat{q}} = \{(\hat{q} + B, B)\}$. The top $\hat{q}/\hat{q} + B$ measure of types choose to buy this lottery. See Figure 8.1(a) for the special case of $\hat{q} = 1/2$ and $B = 1/4$ where $R(1/2) = 1/6$. For $\hat{q} \geq \hat{q}^*$ the optimal mechanism with sale probability at most \hat{q} is the optimal unconstrained mechanism, above. The revenue curve is given by

$$R(\hat{q}) = \begin{cases} \hat{q} B / \hat{q} + B & \text{if } \hat{q} \leq \hat{q}^*, \text{ and} \\ 1 - B & \text{otherwise.} \end{cases}$$

- The interim optimal mechanism for allocation constraint \hat{y} is given by types $0 \leq \hat{t}^\dagger \leq \hat{t}^\ddagger \leq 1$ and has allocation rule given by reserve pricing the weak types $t \in [0, \hat{t}^\dagger)$, ironing the strong types $t \in (\hat{t}^\dagger, 1]$, and allocating maximally to intermediate types $t \in [\hat{t}^\ddagger, \hat{t}^\dagger)$. The reserve

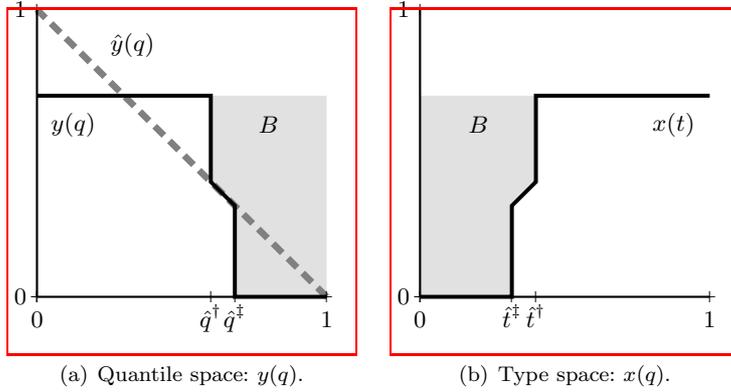


Figure 8.2. For an agent with uniform type on $[0, 1]$ and public budget $B = 1/4$, depicted is the allocation rule for the optimal mechanism that satisfies the interim allocation constraint $\hat{y}(q) = 1 - q$. The allocation rule in quantile space is depicted in subfigure (a); the allocation rule in type space is depicted in subfigure (b). For the uniform distribution, these are mirror images of each other.

price is set to optimize revenue; the ironed interval is set to meet the budget constraint with equality. The example of $\hat{y}(q) = 1 - q$ is given in Figure 8.2. It is most natural to describe the resulting allocation rule in quantile space. Recall that the quantile of a type is the measure of stronger types; for the uniform distribution the quantile of t is $q = 1 - t$. Thus, the allocation rule of this mechanism is $x(t) = y(1 - t)$ with $\hat{q}^\dagger = 1 - \hat{t}^\dagger$, $\hat{q}^\ddagger = 1 - \hat{t}^\ddagger$, and

$$y(q) = \begin{cases} 1/\hat{q}^\dagger \int_0^{\hat{q}^\dagger} \hat{y}(z) dz & \text{if } q \in [0, \hat{q}^\dagger), \\ \hat{y}(q) & \text{if } q \in [\hat{q}^\dagger, \hat{q}^\ddagger), \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

While in these single-dimensional public budget problems, there is a natural ordering on types by strength, it is not the case that optimal mechanisms always break ties in the same way. In particular the interval of ironing, i.e. $[0, \hat{q}^\dagger]$, in the interim mechanism design problem depends on the allocation constraint \hat{y} . Unlike the case of single-dimensional linear preferences described in Chapter 3, there is no fixed virtual value function for which optimization of virtual surplus gives the optimal mechanism. Instead, the appropriate virtual value function will

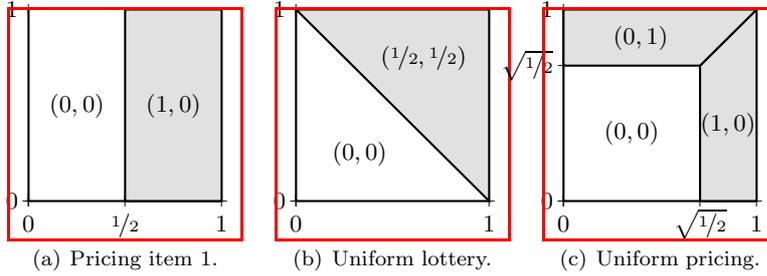


Figure 8.3. Depicted are the outcomes $x = (\{x\}_1, \{x\}_2)$ for the example mechanisms given in the text for the two-alternative uniform unit-demand agent (Example 8.2). The regions of allocation (gray) and non-allocation (white) have area $\hat{q} = 1 - \hat{q} = 1/2$.

depend on the environment and level of competition from other agents. This approach is further described subsequently in Section 8.6.

8.1.2 Unit-demand Preferences

In this section we will describe optimal mechanisms for the three single-agent problems and (multi-dimensional) agents with unit-demand. Formal derivations of these optimal mechanisms are deferred to Section 8.7.

Example 8.2. The exemplary *two-alternative uniform unit-demand agent* has type $t \in \mathcal{T} = [0, 1]^2$ uniformly distributed, i.e., $t \sim U[0, 1]^2$; equivalently her value for each alternative is i.i.d. and uniform on $[0, 1]$. Recall, $t = (\{t\}_1, \{t\}_2)$.

We begin by describing a few mechanisms for the two-alternative uniform unit-demand agent of Example 8.2. Recall any mechanism for a single agent, by the taxation principle, can be represented by a menu where the agent picks her favorite outcome from the menu. An outcome is a triple $w = (\{x\}_1, \{x\}_2, p)$ and individual rationality requires that the outcome $\emptyset = (0, 0, 0)$ is implicitly in the menu of any mechanism. The following are mechanisms that sell with ex ante probability $\hat{q} = 1/2$:

- Sell only item 1 for price $1/2$: $\mathcal{M} = \{(1, 0, 1/2)\}$.
- Sell the *uniform lottery* for price $1/2$: $\mathcal{M} = \{(1/2, 1/2, 1/2)\}$.
- Sell either item at a *uniform price* $\sqrt{1/2}$: $\mathcal{M} = \{(1, 0, \sqrt{1/2}), (0, 1, \sqrt{1/2})\}$.

The first two mechanisms obtain revenue $1/2$ when they sell and, thus, obtain an expected revenue of $1/4$. The final mechanism obtains revenue

$\sqrt{1/2} > 1/2$ when it sells. Of these three mechanisms, the latter has the highest revenue. Figure 8.3 depicts the outcomes of these mechanisms. Of course, these are just three of an infinite number of mechanisms that sell with ex ante probability $\hat{q} = 1/2$. As we will describe below, the ex ante optimal mechanism for $\hat{q} = 1/2$ is in fact the uniform pricing of both alternatives at $\sqrt{1/2}$.

Section 8.7 derives a characterization of optimal mechanisms for the three single-agent optimization problems (Definition 8.3) for a unit-demand agent with a uniformly distributed type. For $m = 2$ alternatives, these optimal mechanisms are as follows.

- The unconstrained optimal mechanism is the uniform pricing at $\sqrt{1/3}$, i.e., $\mathcal{M}^{\hat{q}} = \{(1, 0, \sqrt{1/3}), (0, 1, \sqrt{1/3})\}$. This mechanism serves with ex ante probability $\hat{q}^* = 2/3$ and has revenue $R(1) = \sqrt{4/27} \approx 0.38$.
- The ex ante optimal mechanism for $\hat{q} \leq 2/3$ is the uniform pricing at $\sqrt{1 - \hat{q}}$. For $\hat{q} > 2/3$ it is the $2/3$ ex ante optimal mechanism, i.e., the optimal unconstrained mechanism, above. The revenue curve is given by

$$R(\hat{q}) = \begin{cases} \hat{q} \sqrt{1 - \hat{q}} & \text{if } \hat{q} \leq 2/3, \text{ and} \\ \sqrt{4/27} \approx 0.38 & \text{otherwise.} \end{cases}$$

- The interim optimal mechanism for allocation constraint \hat{y} is has allocation rule (cf. Section 3.4.2 on page 81 and see Figure 8.4).

$$y(q) = \begin{cases} \hat{y}(q) & \text{if } q \leq 2/3, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Figure 8.4(c) is the example of $\hat{y}(q) = 1 - q$; the menu of the optimal mechanism for this interim constraint is

$$\mathcal{M} = \left\{ (x, 0, \int_0^x \sqrt{\max(z, 1/3)} dz) : x \in [1/3, 1] \right\} \cup \left\{ (0, x, \int_0^x \sqrt{\max(z, 1/3)} dz) : x \in [1/3, 1] \right\}.$$

Unlike the example of a single dimensional agent with a public budget (Example 8.1), there is no implicit ordering on types by strength. Which is stronger type $t^\dagger = (.8, .1)$ or type $t^\ddagger = (.6, .6)$? In fact, which of these types is stronger generally depends on the mechanism. The uniform pricing of $\sqrt{1/2} \approx .71$ (Figure 8.3(c)) serves type t^\dagger and rejects type t^\ddagger while the uniform lottery at price $1/2$ (Figure 8.3(b)) serves type t^\ddagger and rejects type t^\dagger . If we consider optimal mechanisms for any strictly decreasing interim constraint \hat{y} , however, it is clear that types are ranked as stronger based on their maximum coordinate $\max(\{t\}_1, \{t\}_2)$. Moreover,

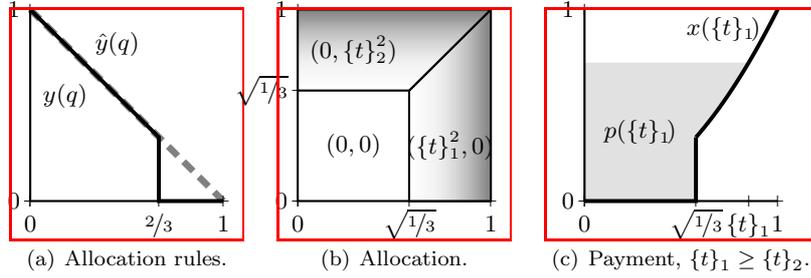


Figure 8.4. For a unit-demand agent with uniformly distributed type for $m = 2$ alternatives (Example 8.2), the allocation constraint $\hat{y}(q) = 1 - q$, its optimal allocation rule $y(\cdot)$, the allocation in the two-dimensional type space $\mathcal{T} = [0, 1]^2$, and the payment as a function of $\{t\}_1 > \{t\}_2$ are depicted. In subfigure (b), the degree of shading represents the probability with which a type $t = (\{t\}_1, \{t\}_2)$ receives her preferred alternative. When alternative 1 is preferred (i.e., $\{t\}_1 > \{t\}_2$) and at least the reservation value (i.e., $\{t\}_1 \geq \sqrt{1/3}$), the probability of receiving alternative 1 is $\{t\}_1^2$. For a type $t = (\{t\}_1, \{t\}_2)$ with $\{t\}_1 > \{t\}_2$, subfigure (c) depicts the allocation rule $x(\{t\}_1)$ as a function of the agent's value for alternative 1. The agent's payment for her preferred outcome with such a type is calculated as the area above this allocation rule $x(\cdot)$, by integrating $x^{-1}(\cdot)$ vertically, as $p(\{t\}_1) = \int_0^{\{t\}_1^2} \sqrt{\max(z, 1/3)} dz$ for $\{t\}_1 \geq \sqrt{1/3}$.

the optimal mechanism for an interim constraint can be viewed as a convex combination of optimal mechanisms for ex ante constraints. Though these simplifying properties of optimal single-agent mechanisms do not hold in general for unit-demand agents, in Section 8.7 we will describe sufficient conditions, beyond the uniform distribution, under which they extend. As we already observed in the preceding study of public budgets these properties do not hold for non-linear preferences.

Important differences in families of single-agent problems have been exhibited above in our study unit-demand and public budget preferences under the uniform distribution. In particular, unit-demand preferences drawn from the uniform distribution (Example 8.2) behave similarly to the single-dimensional linear preferences of Chapter 3, whereas public budget preferences do not (Example 8.1). In the subsequent sections as we describe optimal multi-agent mechanisms both for families of preferences that behave similarly to the uniform unit-demand example and families of preferences that behave similarly to the uniform public budget example. This latter class of mechanisms will be completely general.

8.2 Service Constrained Environments

We will consider environments, like the single-dimensional environments of Section 3.1 on page 54, where agents only impose a single-dimensional externality on each other. In these environments each agent would like to receive an abstract service and there is a feasibility constraint over the set of agents who can be simultaneously served. Agents may also have preferences over unconstrained attributes that may accompany service. Payments are one such attribute; for example, a seller of a car can only sell one car, but she can assign arbitrary payments to the agents (subject to the agents' incentives, of course). Likewise, the seller of the car could paint the car one of several colors as it is sold and the agents may have multi-dimensional preferences over colors. Of course, if the car is sold to one agent then it cannot be sold to other agents so, while the color plays an important role in an agents' multi-dimensional incentive constraints, it plays no role in the feasibility constraints. We refer to environments with single-dimensional externalities as service constrained environments. The more general case of environments that exhibit multi-dimensional externalities is deferred to Section 8.5.

Definition 8.4. A *service constrained environment* is one where a feasibility constraint restricts the set of agents who can be simultaneously served, but imposes no restriction on how they are served. Subsets of the n agents N that can be simultaneously served are given by $\mathcal{X} \subset 2^N$.

In the subsequent sections we will reduce the problem of multi-agent mechanism design to a collection of single-agent mechanism design problems. These sections will not further address the details of how to solve these single-agent problems, instead they will focus on how the multi-agent mechanism is constructed from the single-agent components. We begin in Section 8.3 where the simplifying assumption of revenue linearity enables optimal multi-agent mechanisms to be constructed from the single-agent ex ante optimal mechanisms. In Section 8.4 we consider the more general case where revenue linearity does not hold. In this case we describe how to construct multi-agent mechanisms from the solution to the single-agent interim optimal mechanism design problems.

The definition of service constrained environments, above, corresponds to the general feasibility environments of Section 3.1. The framework can be easily extended to incorporate service costs that are a function of the set of agents served (see Exercise 8.8).

8.3 The Ex Ante Reduction

In this section we construct optimal multi-agent mechanisms for agents whose single-agent problems behave similarly to the single-dimensional linear agents of Chapter 3. We will use, as a running example of such agents, the uniform unit-demand agent (Example 8.2 on page 258). Our approach follows and extends that of Section 3.4 on page 79. In this approach the multi-agent mechanism design problem is reduced to the single-agent ex ante optimal mechanism design problem. This single-agent optimization gives rise to a revenue curve $R(\cdot)$. The revenue linearity property (Definition 3.16), specifically that $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$ for $\hat{y} = \hat{y}^\dagger + \hat{y}^\ddagger$, implies that any interim optimal mechanism can be expressed in terms of marginal revenue $\mathbf{Rev}[\hat{y}] = \mathbf{E}_q[R'(q) \hat{y}(q)]$ (Proposition 3.17).¹ The optimal mechanism is, thus, the one that maximizes marginal revenue (cf. Theorem 3.20).

Recall from Definition 8.3, the ex ante optimal mechanism design problem is given an upper bound \hat{q} on the ex ante probability of serving the agent, over randomization in the agent's type and the mechanism, and determines the outcome rule $w^{\hat{q}}$ which maps types to outcomes. Encoded in an outcome $w^{\hat{q}}(t)$ for type t is a probability of service, denoted $x^{\hat{q}}(t)$, and a payment, denoted $p^{\hat{q}}(t)$. This mechanism can alternatively be thought of as the menu $\mathcal{M}^{\hat{q}} = \{w^{\hat{q}}(t) : t \in \mathcal{T}\}$. The revenue of the ex ante optimal mechanism for every $\hat{q} \in [0, 1]$ defines the revenue curve $R(\hat{q})$.

We will use the uniform unit-demand agent of Example 8.2, which is revenue linear, to illustrate this construction and then give the formal definition, derivation, and proof of correctness. For this example, recall that the \hat{q} ex ante optimal mechanism posts the uniform price $\sqrt{1 - \hat{q}}$ for each of the two alternatives (see Figure 8.3(c)).

¹ For review: View allocation constraint \hat{y} , a monotone non-increasing function from $[0, 1]$ to $[0, 1]$, as a convex combination of reverse step functions each of which steps from 1 to 0 at some \hat{q} . In this convex combination \hat{q} is drawn with cumulative distribution function $G^{\hat{y}}(q) = 1 - \hat{y}(q)$ and density function $g^{\hat{y}}(q) = -\hat{y}'(q)$. The optimal revenue for each \hat{q} , of the \hat{q} ex ante optimal mechanism, defines $R(\hat{q})$, the revenue of the convex combination is thus $\mathbf{E}[(-\hat{y}'(q) R(q))]$. Integration by parts with $R(1) = R(0) = 0$ gives the marginal revenue $\mathbf{MargRev}[\hat{y}] = \mathbf{E}[R'(q) \hat{y}(q)]$.

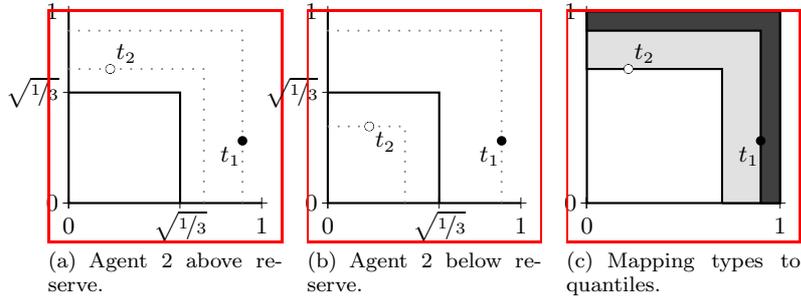


Figure 8.5. The optimal auction for the red-blue car environment of Example 8.3 is illustrated. In subfigure (a), the value agent 2 has for her preferred color $\{t_2\}_2$ exceeds the reserve price of $\sqrt{1/3}$, agent 1 wins her preferred alternative 1 and pays $\{t_2\}_2$. In subfigure (b), the value agent 2 has for her preferred item is below the reserve, agent 1 wins alternative 1 and pays the reserve price. In subfigure (c), the types stronger than agent 1 (dark gray region) and agent 2 (light and dark gray regions) are depicted. As the distribution on types is uniform, the quantile q_1 of agent 1 with type t_1 and q_2 of agent 2 with t_2 can be calculated as the areas of these regions, respectively.

8.3.1 Example: Uniform Unit-demand Preferences

This section illustrates the general construction of optimal mechanisms for revenue-linear agents. The first example considers two (identically distributed, i.e., symmetric) uniform unit-demand agents (Example 8.2) in a single-item environment, and the second considers (non-identically distributed, i.e., asymmetric) agents.

Example 8.3. There are two agents, the seller has one car that he can paint red or blue on its sale. The agents' types, i.e., values for each color, are independently, identically, and uniformly distributed on $[0, 1]$. The second-price auction for each agent's preferred color with a reserve of $\sqrt{1/3}$ is revenue optimal.

To explain Example 8.3, we will follow the construction of optimal mechanisms for single-dimensional linear agents as described in Chapter 3. In particular, we map types to quantiles by their relative strength. We calculate marginal revenue of a type t with quantile q as $R'(q) = \frac{d}{dq}R(q)$. We allocate the car to the type with the highest positive marginal revenue. To determine the payment and what color to paint the car, we look at the weakest quantile, given the quantile(s) of the other agent(s), at which the winner still wins and allocate according to ex ante mech-

anism for this critical quantile. These four steps are described in detail below; the mechanism is illustrated in Figure 8.5.

Observe that agent 2 may win the car and this imposes an interim constraint on agent 1. As we have observed previously, the optimal single-agent mechanism for agent 1 orders her types by her value for her preferred alternative. This ordering on types allows types to be mapped to quantiles. Recall, the quantile of a type designates its strength relative to the distribution of types and is defined as the measure of stronger types. For the example type $t = (\{t\}_1, \{t\}_2)$ is weaker than all types s with $\max(\{s\}_1, \{s\}_2) > \max(\{t\}_1, \{t\}_2)$ and stronger than all types s with $\max(\{s\}_1, \{s\}_2) < \max(\{t\}_1, \{t\}_2)$. As the distribution F is uniform on $[0, 1]^2$, the quantile of a type t is $q = 1 - [\max(\{t\}_1, \{t\}_2)]^2$; see Figure 8.5(c).

The revenue curves $R(\cdot)$ is defined from the solution to the ex ante optimal mechanism for each $\hat{q} \in [0, 1]$. For the two-alternative uniformly distributed types of the example and $\hat{q} \leq 2/3$, \hat{q} ex ante optimal mechanism posts price $\sqrt{1 - \hat{q}}$ and obtains revenue $R(\hat{q}) = \hat{q} \sqrt{1 - \hat{q}}$. For two symmetric agents, as in our example, the details of the revenue curve before its maximum, $\hat{q}^* = 2/3$ for the example, are irrelevant as long as it is strictly concave (unlike the asymmetric example given subsequently). The agent with the stronger quantile wins, as long as that quantile is at least the quantile reserve, which is given by the unconstrained optimal mechanism and is $\hat{q}^* = 2/3$.

Agent 1 will win the auction whenever her quantile is less than agent 2's quantile and the quantile reserve. Agent 1's critical quantile is thus $\hat{q}_1 = \min(q_2, \hat{q}^*)$. Fixing agent 2's type and quantile, agent 1 faces the \hat{q}_1 ex ante optimal mechanism. For the example, this mechanism is given by the menu $\{(1, 0, \sqrt{1 - \hat{q}_1}), (0, 1, \sqrt{1 - \hat{q}_1})\}$, i.e., it is a uniform pricing of $\sqrt{1 - \hat{q}_1}$. When this critical quantile comes from agent 2, the uniform price is exactly the value agent 2 has for her preferred alternative. When this critical quantile comes from the reserve, then the uniform price is from the optimal unconstrained mechanism, i.e., it is $\sqrt{1/3}$. Thus, agent 1 is offered a uniform price that is the higher of the reserve and agent 2's value for her preferred alternative. The optimal mechanism is the second-price auction for the agents' preferred alternative with a uniform reserve of $\sqrt{1/3}$.

The next example environment shows that the approach taken above can treat asymmetric agent preferences and that the dimensionality of the preferences need not be the same. Due to the asymmetry, however,

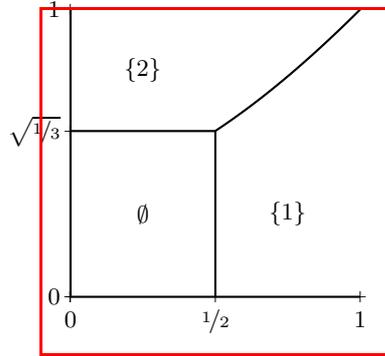


Figure 8.6. The allocation of the optimal auction for the red-blue-green car environment of Example 8.4 is depicted. The horizontal axis is the value of agent 1 for the green car t_1 ; the vertical axis is the value of agent 2 for her preferred color of red and blue $\{t_2\}_{\max} = \max_j \{t_2\}_j$.

the resulting optimal mechanism is more complex and will depend more finely on the details of the agents' revenue curves.

Example 8.4. There are two agents. A seller has one car that he can paint red, blue, or green on its sale. Agent 1 is single-dimensional and has uniformly distributed value for a green car (and no value for a red or blue car); agent 2 has independent and uniformly distributed values for red or blue cars (and no value for a green car). The winner of the optimal mechanism is depicted in Figure 8.6, the winner always receives her preferred color car.

The optimal mechanism of the red-blue-green car example (Example 8.4) is constructed as follows. We map agent 1's single-dimensional type to a virtual value via the standard transformation from value to quantile to marginal revenue as per the construction of Section 3.4. With type t_1 , agent 1's virtual value is $\phi_1(t_1) = 2t_1 - 1$. We map agent 2's multi-dimensional type to a (single-dimensional) virtual value similarly:

- (i) $q_2 = 1 - [\max(\{t_2\}_1, \{t_2\}_2)]^2$.
- (ii) $R'_2(q) = \frac{d}{dq} [q \sqrt{1-q}] = \frac{1-3/2q}{\sqrt{1-q}} = [1 - 3/2q] [1 - q]^{-1/2}$.
- (iii) $\phi_2(t_2) = 3/2 \{t_2\}_{\max} - 1/2 \{t_2\}_{\max}^{-1}$ where $\{t_2\}_{\max} = \max_j \{t_2\}_j$.

The winner is the agent i with the highest virtual value. She receives her preferred alternative at a price of $\phi_i^{-1}(\max\{\phi_{3-i}(t_{3-i}), 0\})$.

The remainder of this section will formalize the construction of the

marginal revenue mechanism for revenue-linear agents in general and prove its optimality among all mechanisms.

8.3.2 Orderability

Fundamental to the examples above is the identification of an ordering on types from which types can be mapped to quantiles (and then to marginal revenues). In this section we show that the existence of such an ordering is a consequence of the revenue linearity property (restating Definition 3.16):

$$\mathbf{Rev}[\hat{y}^\dagger + \hat{y}^\ddagger] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger].$$

The uniform unit-demand agent (Example 8.2) is revenue linear. This observation follows from the fact that the interim optimal mechanism for constraint \hat{y} is a convex combination of ex ante optimal mechanisms. Revenue linearity immediately implies that the surplus of marginal revenue (Definition 3.15) is equal to the optimal revenue (Proposition 3.17):

$$\mathbf{MargRev}[\hat{y}] = \mathbf{E}[R'(q) \hat{y}(q)] = \mathbf{Rev}[\hat{y}].$$

However, it does not immediately suggest how to implement surplus of marginal revenue maximization as the mapping from type space to quantile space is not explicit in a multi-dimensional type space (as it is in a single-dimensional type space).

Definition 8.5. A single-agent problem given by a type space, outcome space, and distribution over types is *orderable* if there is an equivalence relation on types and an ordering over equivalence classes, such that for every allocation constraint \hat{y} , an optimal mechanism for \hat{y} , i.e., solving $\mathbf{Rev}[\hat{y}]$, induces an allocation rule that is greedy by the given ordering with ties broken uniformly at random and with types in a special lowest equivalence class \perp (if any) rejected.

Notice that the single-dimensional budgeted agent (Example 8.1) is not orderable by the above definition. Though, the agent's value for service gives a natural ordering on types, the optimal mechanism for \hat{y} irons the strongest quantiles so as to meet the budget constraint with equality and this ironed interval depend on the allocation constraint \hat{y} (see Figure 8.2 on page 257).

Theorem 8.1. *For any single-agent problem, revenue linearity implies orderability.*

The main intuition behind Theorem 8.1 comes from the observation that revenue linearity implies that the allocation rule y that is obtained from optimization subject to interim constraint \hat{y} must be equal to \hat{y} at all quantiles where the revenue curve is strictly concave; the equivalence classes in the theorem statement then correspond to types with equal marginal revenue (which have non-zero measure only on intervals of quantile space where the marginal revenue is constant). The proof illustrates how revenue linearity enables interim optimal mechanisms to be understood in terms of ex ante optimal mechanisms.

Recall the definition of the cumulative allocation rule as $Y(\hat{q}) = \int_0^{\hat{q}} y(q) dq$ and, by integration by parts, we can express the marginal revenue of any allocation rule y as (recall that $Y(0) = 0$):

$$\begin{aligned} \mathbf{MargRev}[y] &= \left[R'(q) Y(q) \right]_0^1 - \int_0^1 R''(q) Y(q) dq \\ &= R'(1) Y(1) - \mathbf{E}[R''(q) Y(q)]. \end{aligned} \quad (8.1)$$

We will prove Theorem 8.1 by combining the following two lemmas. The first lemma shows that for any quantile \hat{q} where marginal revenue is strictly decreasing, i.e., where $R''(\hat{q}) < 0$, the interim optimal allocation rule y for interim constraint \hat{y} allocates to the maximum extent possible, i.e., the interim constraint of Definition 8.3 is tight at \hat{q} , i.e., $Y(\hat{q}) = \hat{Y}(\hat{q})$.

Lemma 8.2. *For a revenue-linear agent, allocation rule y that is optimal for allocation constraint \hat{y} , and any ex ante probability \hat{q} with $R''(\hat{q}) \neq 0$, the cumulative allocation rule and constraint satisfy $Y(\hat{q}) = \hat{Y}(\hat{q})$.*

Proof. If we optimize revenue for allocation constraint \hat{y} and obtain a mechanism with allocation rule y , then $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[y]$. Revenue linearity implies that optimal revenues are equal to marginal revenues, i.e., $\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}]$ and $\mathbf{Rev}[y] = \mathbf{MargRev}[y]$, respectively. Writing the difference between these marginal revenues and employing equation (8.1), we have:

$$\begin{aligned} 0 &= \mathbf{MargRev}[\hat{y}] - \mathbf{MargRev}[y] \\ &= R'(1) [\hat{Y}(1) - Y(1)] + \mathbf{E}[-R''(q) [\hat{Y}(q) - Y(q)]]. \end{aligned} \quad (8.2)$$

By interim feasibility of y for \hat{y} , $[\hat{Y}(q) - Y(q)] \geq 0$. By concavity of revenue curves $[-R''(q)] \geq 0$. By monotonicity of revenue curves $R'(1) \geq 0$. Thus, every term in equation (8.2) is non-negative. The only way it

can be identically zero if $[-R''(q)] > 0$ implies that $[\hat{Y}(q) - Y(q)] = 0$ as required by the lemma. (Also observe, though unnecessary for the lemma, that $R'(1) > 0$ implies that $[\hat{Y}(1) - Y(1)] = 0$.) \square

An immediate corollary of Lemma 8.2 is that the ex ante optimal mechanism for any \hat{q} with $R''(\hat{q}) < 0$ deterministically serves or rejects each type, i.e., the allocation rule in type space is $x^{\hat{q}}(t) \in \{0, 1\}$. Contrast this corollary to the uniform public-budget example where the optimal mechanism for ex ante constraint $\hat{q} = 1/2$ is pricing the $3/4$ lottery while $R(\hat{q}) = \hat{q}^{B/\hat{q}+B}$ is strictly convex at $\hat{q} = 1/2$ (Section 8.1.1).

Corollary 8.3. *For a revenue-linear agent and any ex ante probability \hat{q} with $R''(\hat{q}) \neq 0$, the \hat{q} ex ante optimal mechanism deterministically serves or rejects each type $t \in \mathcal{T}$.*

Proof. Notice that an ex ante constraint of \hat{q} is equivalent to the interim constraint given by the reverse-step function that steps from 1 to 0 at \hat{q} , denoted $\hat{y}^{\hat{q}}$. By Lemma 8.2, the optimal allocation rule for this constraint is the reverse-step function itself. A mechanisms whose allocation rule is a reverse-step function deterministically allocates or rejects each type. \square

For \hat{q} where the revenue curve is strictly concave, Corollary 8.3 implies the type space \mathcal{T} is partitioned into types that are served and those that are rejected. Denote the allocation rule of the \hat{q} ex ante optimal mechanism in type space by $x^{\hat{q}}(\cdot)$ and the subset of types it serves by

$$S^{\hat{q}} = \{t \in \mathcal{T} : x^{\hat{q}}(t) = 1\}. \quad (8.3)$$

The following lemma shows that the sets of types served by the ex ante optimal mechanisms are nested.

Lemma 8.4. *For a revenue-linear agent and any ex ante probabilities $\hat{q}^\dagger < \hat{q}^\ddagger$ with $R''(\hat{q}^\dagger) \neq 0$ and $R''(\hat{q}^\ddagger) \neq 0$, then $S^{\hat{q}^\dagger} \subset S^{\hat{q}^\ddagger}$.*

Proof. This proof is illustrated in Figure 8.7. Let \hat{y}^\dagger and \hat{y}^\ddagger denote the interim allocation constraint corresponding to the ex ante constraints \hat{q}^\dagger and \hat{q}^\ddagger , respectively. Consider the interim constraint $\hat{y} = 1/2 \hat{y}^\dagger + 1/2 \hat{y}^\ddagger$. The constraint \hat{y} is a reverse stair function that steps from 1 to $1/2$ at \hat{q}^\dagger and from $1/2$ to 0 at \hat{q}^\ddagger . Suppose for a contradiction that $S^{\hat{q}^\dagger}$ contains a measurable (with respect to the distribution F) set of types that is not also contained in $S^{\hat{q}^\ddagger}$. By revenue linearity the optimal mechanism for \hat{y} is the convex combination of the optimal mechanisms for \hat{y}^\dagger and \hat{y}^\ddagger ; denote

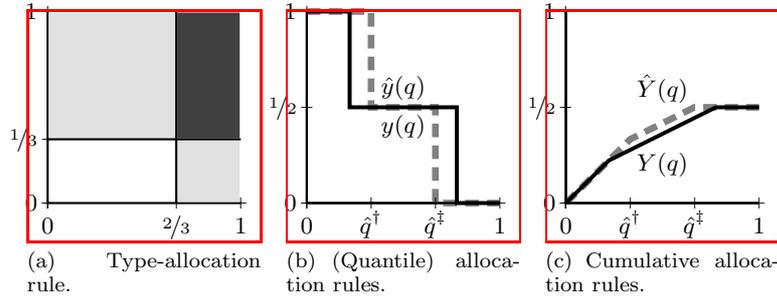


Figure 8.7. For a two-alternative unit-demand agent, the proof of Lemma 8.4 is illustrated. In this hypothetical situation the \hat{q}^\dagger (resp. \hat{q}^\ddagger) optimal mechanism posts a price for selling alternative 1 (resp. 2) only. In subfigure (a) the two-dimensional type space $\mathcal{T} = [0, 1]$ and the allocation of the convex combination of the \hat{q}^\dagger and \hat{q}^\ddagger optimal mechanism are depicted. The types $t \in S^{\hat{q}^\dagger} \cap S^{\hat{q}^\ddagger}$ (dark gray region) have allocation probability $x(t) = 1$, types t in the symmetric difference of $S^{\hat{q}^\dagger}$ and $S^{\hat{q}^\ddagger}$ (light gray region) have allocation probability $x(t) = 1/2$, and the remaining types have allocation probability $x(t) = 0$. In subfigure (b) the allocation constraint \hat{y} and allocation rule y are depicted. In subfigure (c) the cumulative allocation constraint \hat{Y} and cumulative allocation rule Y are depicted. Allocation constraints are depicted with thick, dashed, gray lines and allocation rules are depicted with thin, solid, black lines.

its allocation rules in quantile and type space by y and x , respectively. The following table tallies the measure and service probability of each relevant subset of types (with ρ defined by the probability of the second line):

S	$\Pr[t \in S]$	$x(t)$
$S^{\hat{q}^\dagger} \cap S^{\hat{q}^\ddagger}$	$\hat{q}^\dagger - \rho$	1
$S^{\hat{q}^\dagger} \setminus S^{\hat{q}^\ddagger}$	ρ	1/2
$S^{\hat{q}^\ddagger} \setminus S^{\hat{q}^\dagger}$	$\hat{q}^\ddagger - \hat{q}^\dagger + \rho$	1/2
$\mathcal{T} \setminus S^{\hat{q}^\dagger} \setminus S^{\hat{q}^\ddagger}$	$1 - \hat{q}^\ddagger + \rho$	0

The allocation rule y is a reverse-stair function that steps from 1 to 1/2 at $\hat{q}^\dagger - \rho$ and from 1/2 to 0 at $\hat{q}^\ddagger + \rho$. Inspection reveals, for a contradiction, that Lemma 8.2 is violated at \hat{q}^\dagger and \hat{q}^\ddagger . For example, $\hat{Y}(\hat{q}^\dagger) = \hat{q}^\dagger > Y(\hat{q}^\dagger) = \hat{q}^\dagger - 1/2\rho$. \square

Proof of Theorem 8.1. Define type subspace $S^{\hat{q}}$ as in equation (8.3). Define the marginal revenue of a type t as $\inf\{R'(\hat{q}) : S^{\hat{q}} \ni t\}$. The equiv-

alence classes of Definition 8.5 are sets of types with the same marginal revenue; types with marginal revenue zero are in the equivalence class \perp . By Lemma 8.4, these definitions are well defined.

Consider the \hat{q} optimal mechanism. If $R''(\hat{q}) \neq 0$ then it is optimal to serve types $S^{\hat{q}}$. Greedy by marginal revenue (as defined above) serves these types. If $R''(\hat{q}) = 0$, then let $(\hat{q}^\dagger, \hat{q}^\ddagger)$ be the interval on which $R'(\hat{q})$ is constant. An optimal mechanism randomizes between the \hat{q}^\dagger ex ante optimal and \hat{q}^\ddagger ex ante optimal mechanisms so that the total sale probability is \hat{q} . By Lemma 8.4, types in $S^{\hat{q}^\dagger}$, with marginal revenue strictly greater than $R'(\hat{q})$, are served with certainty and types in $S^{\hat{q}^\ddagger} \setminus S^{\hat{q}^\dagger}$, with marginal revenue equal to $R'(\hat{q})$ are served with probability $\hat{q} - \hat{q}^\dagger / \hat{q}^\ddagger - \hat{q}^\dagger$. One way to achieve these service probabilities is to randomly order the types by marginal revenue with ties broken randomly and to greedily serve the first \hat{q} measure of types. Thus, all ex ante optimal mechanisms order types by marginal revenue and serve them greedily.

By revenue linearity the optimal mechanism for allocation constraint \hat{y} is a convex combination of ex ante optimal mechanisms. As these ex ante optimal mechanisms all order the types greedily by marginal revenue with ties broken randomly, so does the optimal mechanism for \hat{y} . \square

Theorem 8.1 says that while there is not an inherent ordering on type space that is respected by all mechanisms, there is one that, for all interim allocation constraints, is consistent with an optimal mechanism for the constraint.

Definition 8.6. For an orderable agent (Definition 8.5) and an implicit arbitrary total order on types that is consistent with the partial order on types, the *quantile q of a type t* is the probability that a random type $s \in \mathcal{T}$ from the distribution F precedes type t in the total order.

8.3.3 The Marginal Revenue Mechanism

We now define the marginal revenue mechanism for orderable agents.

Definition 8.7. The *marginal revenue mechanism for orderable agents* works as follows:

- (i) Map the profile of agents' types \mathbf{t} to a profile of quantiles \mathbf{q} via Definition 8.6.
- (ii) Calculate the profile of marginal revenues for the profile of quantiles.

- (iii) Calculate a feasible allocation to optimize the surplus of marginal revenue, i.e., $\mathbf{x} = \operatorname{argmax}_{\mathbf{x}^\dagger} \sum_i x_i^\dagger R_i'(q_i) - c(\mathbf{x}^\dagger)$. For each agent i , calculate the supremum quantile \hat{q}_i she could possess for which she would be allocated in the above calculation of \mathbf{x} .
- (iv) Offer each agent i the \hat{q}_i optimal single-agent mechanism.

Theorem 8.5. *The marginal revenue mechanism for revenue-linear agents is (a) dominant strategy incentive compatible, (b) feasible, (c) revenue-optimal, and (d) deterministically selects the set of winners.*

Proof. Consider any agent i . Analogously to the proof of Theorem 3.5, monotonicity of marginal revenue curves and Lemma 3.1 implies that, for every profile of reports of the other agents, there is a critical quantile \hat{q}_i for agent i . The \hat{q}_i ex ante mechanism is dominant strategy incentive compatible. Thus, the composition is dominant strategy incentive compatible.

The critical quantile \hat{q}_i for agent i is at the boundary of service and non-service thus $R_i''(\hat{q}_i) \neq 0$ which implies that the \hat{q}_i optimal mechanism, by Corollary 8.3, deterministically serves or rejects the agent. The set of agents served are exactly those served by \mathbf{x} which is feasible.

The mechanism is revenue optimal because (a) its revenue is equal to its surplus of marginal revenue, (b) its surplus of marginal revenue is pointwise at least that of any other feasible mechanism, and (c), by revenue linearity, the revenue of any mechanism is at most its marginal revenue. \square

While it is often assumed that the optimality of the marginal revenue mechanism is special to single-dimensional linear agents (as in Chapter 3), we have seen here that the condition required is revenue linearity not single dimensionality. It is instructive to contrast the simple optimal mechanism for revenue-linear agents to the complex optimal mechanism for non-revenue-linear agents that is derived in the next section.

8.4 The Interim Reduction

Without the revenue linearity property, which was assumed in the preceding section, single-agent interim optimal mechanisms can not be described solely in terms of the single-agent ex ante optimal mechanisms. For this reason, more sophisticated single-agent mechanisms are needed to enable the optimization of general multi-agent mechanisms.

In this section we characterize optimal multi-agent mechanisms for service constrained environments (Definition 8.4) in terms of the solution to single-agent interim optimal mechanism design problems (and without the simplifying revenue-linearity property). It is useful to contrast the complexity of the optimal mechanism for general preferences with that of the optimal mechanism with the revenue linearity assumption of Section 8.3.

The (quantile) allocation rule of a mechanism can be determined as follows. Recall that a single-agent mechanism \mathcal{M} is given by an outcome rule $w(\cdot)$ which maps types to outcomes. The mechanism can alternatively be thought of as the menu $\{w(t) : t \in \mathcal{T}\}$. Encoded in an outcome $w(t)$ for type t is a probability of service denoted $x(t)$ and a payment denoted $p(t)$. For finite type spaces where $f(\cdot)$ denotes the probability mass function of the distribution F , the (quantile) allocation rule $y(\cdot)$ of the mechanism can be found by making a rectangle for each type t with height $x(t)$ and width $f(t)$ and sorting these rectangles in decreasing order of height. Equivalently and generally for continuous type spaces, consider the function defined as the measure of types that are served with at least a given service probability, the (quantile) allocation rule is the inverse of this function, i.e., $y(q) = \sup\{x^\dagger : \Pr_{t \sim F}[x(t) \geq x^\dagger] \geq q\}$.

Recall from Definition 8.3 that the single-agent interim optimal mechanism for allocation constraint \hat{y} has allocation rule y that is no stronger than \hat{y} , i.e., the cumulative allocation rules satisfy $Y(\hat{q}) \leq \hat{Y}(\hat{q})$. Moreover, it optimizes revenue, denoted $\mathbf{Rev}[\hat{y}]$, over all such mechanism.

8.4.1 Symmetric Single-item Environments

We will illustrate the approach of this section with an example environment where symmetry enables the optimal mechanism to be easily identified.

Example 8.5. There are two agents competing for a single item each with private value t independently, identically, and uniformly distributed on $[0, 1]$ and a public budget of $B = 1/4$ (as in Example 8.1). The revenue-optimal mechanism fixes $(\hat{t}^\ddagger, \hat{t}^\dagger) \approx (0.32, 0.40)$, rejects agents with values less than \hat{t}^\ddagger , and allocates to the item to the remaining agent i for which $\min(t_i, \hat{t}^\dagger)$ is highest, breaking ties randomly. Each agent makes deterministic payment according to the interim allocation rule $x(\cdot)$. In particular, types $t_i \geq \hat{t}^\dagger$ pay B . See Figure 8.2 or Figure 8.8.

The main observation that enables the identification of an optimal

mechanism in symmetric environments, i.e., with identically distributed agent types and symmetric feasibility constraint, is that convexity of the mechanism design problem, i.e., that convex combinations of mechanisms are valid mechanisms, implies that there is always an optimal mechanism that is symmetric. The search for the optimal mechanism is then facilitated by symmetry.

Proposition 8.6. *In any symmetric environment there is an optimal mechanism that is symmetric.*

Proof. Consider any optimal incentive compatible mechanism that is asymmetric. Symmetry of the environment implies that permuting the identities of the agents gives a (potentially distinct) incentive compatible mechanism that is also optimal. The convex combination of, i.e., randomization over, incentive compatible mechanisms is incentive compatible. In particular, the convex combination of mechanisms for the uniform distribution over all permutations of the identities of agents is optimal, incentive compatible, and symmetric. \square

We will separate the process of finding the (symmetric) optimal mechanism into two parts. The first part will be to identify a symmetric profile of allocation constraints $\hat{\mathbf{y}} = (\hat{y}, \dots, \hat{y})$ that is feasible for a mechanism. The second part will be to find the single-agent mechanism, with allocation rule y , that is optimal for the identified constraint \hat{y} . The optimal multi-agent mechanism is then found by optimizing over the first part and combining with the second part. In the subsequent discussion, this approach is illustrated for the two-agent public-budget environment of Example 8.5.

The following theorem resolves the first part by identifying a single allocation constraint that is feasible and stronger than all other symmetric interim feasible allocation rules. In particular, the optimization over solutions to the first part is given by this allocation constraint.

Theorem 8.7. *Let $\hat{\mathbf{y}} = (\hat{y}, \dots, \hat{y})$ be the n -agent allocation constraints induced by the k strongest-agents-win mechanism and $\mathbf{y} = (y, \dots, y)$ the allocation rules induced by any symmetric k -unit mechanism for n i.i.d. agents, then y is feasible for \hat{y} .*

Proof. We prove the $n = 2$ agent $k = 1$ unit special case; the general result is left for Exercise 8.2. The two agent strongest-agent-wins mechanism induces allocation constraint $\hat{y}(q) = 1 - q$.

We argue, as follows, that \hat{y} is the strongest symmetric allocation rule.

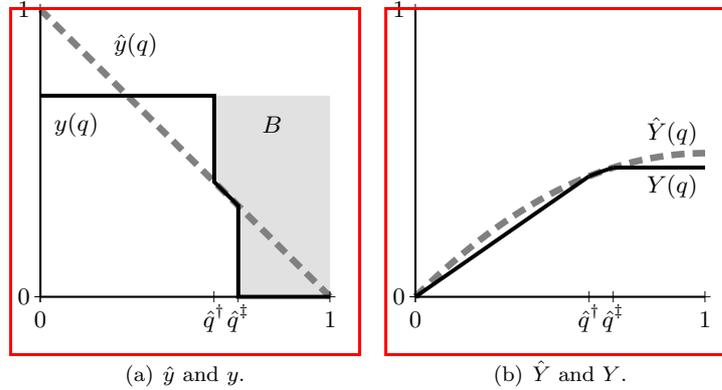


Figure 8.8. Depicted in (a) are the allocation constraint (thick, gray, dashed line) and allocation rule (thin, black, solid line) for the optimal mechanism of the two-agent public budget environment of Example 8.5. Similarly, in (b) are the cumulative allocation rule and constraint.

For one of the agents, the probability she has quantile stronger than \hat{q} is \hat{q} while the probability she has quantile stronger than \hat{q} and is served by the mechanism is $\hat{Y}(\hat{q}) = \int_0^{\hat{q}} \hat{y}(q) dq = \hat{q} - 1/2 \hat{q}^2$. For both agents, the probability at least one has quantile stronger than \hat{q} is

$$1 - \Pr[q \geq \hat{q}]^2 = 1 - (1 - \hat{q})^2 = 2\hat{q} - \hat{q}^2. \quad (8.4)$$

Meanwhile and by linearity of expectation, the expected number of agents with stronger quantile than \hat{q} that are served is

$$2\hat{Y}(\hat{q}) = 2\hat{q} - \hat{q}^2. \quad (8.5)$$

Importantly the quantity of (8.4) is equal to the quantity of (8.5).

Consider the allocation rule y induced by any symmetric mechanism. For a contradiction, assume that y is stronger than \hat{y} at \hat{q} , i.e., with $Y(\hat{q}) > \hat{Y}(\hat{q})$. This allocation rule cannot result from any ex post feasible mechanism as the analogous quantity to (8.5) for y would exceed quantity (8.4). Because there is only one unit available, it is impossible for the expected number of units allocated to agents with quantile in $[0, \hat{q}]$ to be greater than the probability that there is at least one such agent. \square

For the second part, recall the interim optimal mechanism for the uniform public-budget agent (Example 8.1) and allocation constraint $\hat{y}(q) = 1 - q$ described in Section 8.1.1. It has allocation rule $y(\cdot)$ that

is equal to $\hat{y}(\cdot)$ except that it irons quantiles on interval $[0, \hat{q}^\dagger \approx 0.60]$ and rejects those below quantile reserve $\hat{q}^\ddagger \approx 0.68$. See Figure 8.8. For the uniform public-budget agent, these quantiles map back to type space via $t = 1 - q$ as $\hat{t}^\dagger = 0.40$ and $\hat{t}^\ddagger = 0.32$.

These two parts can be easily combined to observe that the mechanism described in Example 8.5 is revenue optimal. Shortly, in Section 8.4.4 we will want to generalize this construction, so it is instructive to see how we might construct the optimal mechanism knowing nothing about \hat{y} except that there is an ex post mechanism, i.e., a mapping from quantile profiles to an ex post allocation that is feasible, that induces the interim allocation \hat{y} , and that we want to iron the strongest $[0, \hat{q}^\dagger]$ quantiles and reject the weakest $(\hat{q}^\ddagger, 1]$ quantiles. Denote the ex post allocation for quantile profile \mathbf{q} by $\hat{\mathbf{y}}^{EP}(\mathbf{q})$; e.g., the two-agent strongest-quantile-wins ex post allocation is

$$\hat{y}_i^{EP}(\mathbf{q}) = \begin{cases} 1 & \text{if } q_i < q_{3-i}, \\ 0 & \text{otherwise.} \end{cases}$$

The following steps suffice to convert this ex post allocation $\hat{\mathbf{y}}^{EP}(\mathbf{q})$ that implements $\hat{\mathbf{y}} = (\hat{y}, \hat{y})$ to an ex post allocation rule $\mathbf{y}^{EP}(\mathbf{q})$ that implements the desired $\mathbf{y} = (y, y)$.

- (i) Calculate \mathbf{q}^\dagger by ironing on $[0, \hat{q}^\dagger]$ as

$$q_i^\dagger = \begin{cases} U[0, \hat{q}^\dagger] & \text{if } q_i \in [0, \hat{q}^\dagger], \\ q_i & \text{otherwise.} \end{cases}$$

- (ii) Calculate \mathbf{y} with quantile reserve \hat{q}^\ddagger as

$$y_i = \begin{cases} \hat{y}_i^{EP}(\mathbf{q}^\dagger) & \text{if } q_i \in [0, \hat{q}^\ddagger], \\ 0 & \text{otherwise.} \end{cases}$$

These operations are easy to interpret on the cumulative allocation rule (see Figure 8.8). We are given the cumulative allocation constraint \hat{Y} and we wish to implement cumulative allocation rule Y that satisfies $Y(q) \leq \hat{Y}(q)$ for all $q \in [0, 1]$; both the cumulative constraint and rule are concave. Ironing by resampling quantiles on an interval replaces the original curve with a line segment. A quantile reserve replaces the original curve with a horizontal line from the quantile reserve and over weaker quantiles. Combinations of these operations can produce any such Y from any such \hat{Y} .

The following proposition summarizes and generalizes the discussion

of optimal mechanism for symmetric single-item environments. In the next section, these methods are further generalized to asymmetric environments.

Proposition 8.8. *For symmetric n -agent single-item environments, the optimal mechanism has expected revenue $n \mathbf{Rev}[\hat{y}]$ with allocation constraint $\hat{y}(q) = (1 - q)^{n-1}$. For uniform public budget preferences, there exists quantiles $0 \leq \hat{q}^\dagger \leq \hat{q}^\ddagger \leq 1$ such that the optimal mechanism irons the strongest $[0, \hat{q}^\dagger]$ and reserve prices the weakest $(\hat{q}^\ddagger, 1]$ quantiles.*

8.4.2 Interim Feasibility

Consider any multi-agent mechanism \mathcal{M} . When the agent types \mathbf{t} are drawn from the product distribution \mathbf{F} , each agent i has an induced *interim mechanism* \mathcal{M}_i . This interim mechanism maps the agent's type t_i to a distribution over outcomes (including the service received and non-service-constrained attributes such as payments). Any single-agent mechanism \mathcal{M}_i induces an allocation rule y_i as described at the onset of Section 8.4; since this is an interim mechanism we refer to this allocation rule as the interim allocation rule. Repeating this construction for each of the n agents we obtain a profile of interim allocation rules \mathbf{y} that is feasible in the sense that there exists an ex post feasible mechanism (in particular, \mathcal{M}) that induces it.² Note that interim feasibility is unrelated to incentives, any function \mathcal{M} that maps profiles of types \mathbf{t} to feasible allocations induces interim feasible allocation rules.

Definition 8.8. A profile of allocation rules \mathbf{y} is *interim feasible* if it is induced by some ex post feasible mechanism \mathcal{M} and type distribution \mathbf{F} .

Example 8.6. Consider selling a single item to one of two agents each with one of two interim allocation rules:

$$y^\dagger(q) = 1/2, \quad y^\ddagger(q) = \begin{cases} 1 & \text{if } q \in [0, 1/2], \\ 0 & \text{otherwise.} \end{cases} \quad (8.6)$$

Notice that both allocation rules have an ex ante probability of $1/2$ of allocating (as agent quantiles are always drawn from the uniform distribution). Consider the profile of allocation rules $\mathbf{y} = (y^\dagger, y^\ddagger)$, i.e., where

² Such a profile of interim allocation rules is sometimes also called the *reduced form* of the mechanism.

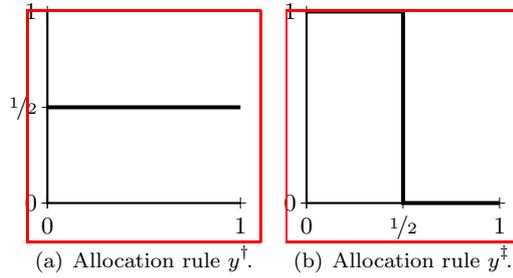


Figure 8.9. The allocation rules y^\dagger and y^\ddagger demonstrating interim feasibility. For single-item environments (y^\dagger, y^\dagger) and (y^\dagger, y^\ddagger) are feasible while (y^\ddagger, y^\ddagger) is infeasible.

both agents have interim allocation rule y^\dagger . This profile is interim feasible as it is the outcome of the *fair-coin-flip* mechanism. Similarly the profile $\mathbf{y} = (y^\ddagger, y^\dagger)$ is interim feasible, it is induced by the *serial-dictator* mechanism that serves agent 1 if she has a high type (i.e., $q_1 \in [0, 1/2]$) and agent 2 otherwise. The *serve-high-types* profile of interim allocation rules $\mathbf{y} = (y^\ddagger, y^\ddagger)$, on the other hand, is not interim feasible. If both agents have high types, which happens with probability $1/4$, the interim allocation rules require that both agents be served, but doing so would not be ex post feasible as there is only one item.

Our goal is to maximize expected revenue over (Bayesian incentive compatible and interim individually rational) mechanisms subject to ex post feasibility. Decomposing this goal into optimization of single-agent revenue subject to interim feasibility, we obtain the following program.

$$\begin{aligned} \max_{\hat{\mathbf{y}}} \quad & \sum_i \mathbf{Rev}[\hat{g}_i] & (8.7) \\ \text{s.t.} \quad & \text{“}\hat{\mathbf{y}} \text{ is interim feasible.”} \end{aligned}$$

Recall that the optimal revenue for a single agent as solved by $\mathbf{Rev}[\cdot]$ is a convex optimization problem and thus $\mathbf{Rev}[\cdot]$ is concave, i.e., $\mathbf{Rev}[\hat{g}^\dagger + \hat{g}^\ddagger] \geq \mathbf{Rev}[\hat{g}^\dagger] + \mathbf{Rev}[\hat{g}^\ddagger]$. Observe that while the constraint of interim feasibility on $\hat{\mathbf{y}}$ is somewhat opaque at this point, it is nonetheless a convex constraint. Simply, the convex combination of two interim feasible mechanisms is interim feasible. The ex post mechanism that implements the convex combination is exactly the convex combination of the ex post mechanisms that implement the two original mechanisms. It will be the

task of the remainder of this section to further elucidate the constraint imposed by interim feasibility.

The following proposition shows that the revenue optimal mechanism can be found by optimizing expected revenue over profiles of allocation constraints subject to interim feasibility.

Proposition 8.9. *The optimal multi-agent revenue is given by optimizing single-agent revenue subject to interim feasibility, i.e., solving program (8.7).*

Proof. We will argue the two directions of this proof separately. First note that the optimal revenue from the program (8.7) is at least the revenue of the optimal mechanism. To see this, observe that any mechanism induces a profile of interim allocation rules \mathbf{y} . The ex post feasibility of this mechanism implies that this profile of allocation rules is interim feasible. The revenue from each agent i in this mechanism is at most the revenue of the interim optimal mechanism subject to allocation rule y_i as a constraint, i.e., at most $\mathbf{Rev}[y_i]$. Thus, the program upper bounds the optimal revenue.

For the other direction we will construct, from any ex post mechanism \mathcal{M} that induces the profile $\hat{\mathbf{y}}$ of interim allocation rules that attains the maximum of the program and each agent i 's \hat{y}_i -optimal mechanism \mathcal{M}_i , a Bayesian incentive compatible mechanism with revenue equal to the revenue of the program, i.e., $\sum_i \mathbf{Rev}[\hat{y}_i]$. The remainder of this proof is deferred to Section 8.4.4 where the construction is generalized by Definition 8.13 and shown to be correct by Theorem 8.18. \square

Optimization of mathematical program (8.7) in asymmetric environments relies on better understanding the constraint posed by interim feasibility. Consider first interim feasibility in single-item environments. Take any profile of ex ante constraints $\hat{\mathbf{q}} = (\hat{q}_1, \dots, \hat{q}_n)$ and consider the ex ante probability by which each agent i with quantile at most \hat{q}_i is served, i.e., $Y_1(\hat{q}_1), \dots, Y_n(\hat{q}_n)$. The expected number of these agents served is thus $\sum_i Y_i(\hat{q}_i)$. Of course the probability that one or more agents agent i with quantile bounded by \hat{q}_i are realized is $1 - \prod_i (1 - \hat{q}_i)$.³ Given the single-item ex post feasibility constraint that allows only one such agent to be served at once, the expected number served must be at least the probability that at least one is realized. In fact this necessary condition is also sufficient, as we will see by the max-flow-min-cut

³ The probability that one or more such agents show up is one minus the probability that none show up.

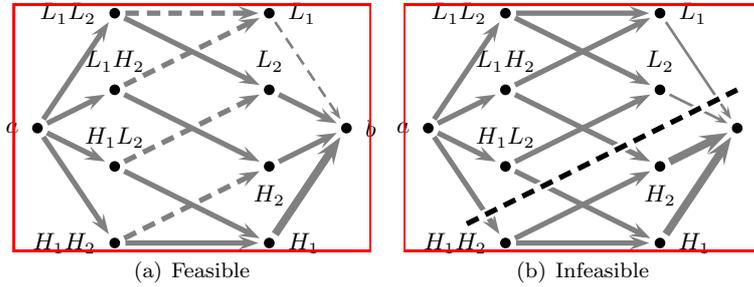


Figure 8.10. The flow constructions are illustrated for (a) the feasibility of the serial-dictator mechanism and (b) the infeasibility of the serve-high-types profile of interim allocation rules from Example 8.6. In both flow graphs the edges depicted in heavy, medium, and light weight correspond to capacities $1/2$, $1/4$, and zero, respectively. To translate between the quantile-space allocation rules of Example 8.6 and the type-space allocation rules in the proof of Theorem 8.10, type H will correspond to strong quantiles $[0, 1/2]$ and type L corresponds to weak quantiles $(1/2, 1]$. Subfigure (a) depicts the flow graph that corresponds to the profile of interim allocation rules $\mathbf{y} = (y^\dagger, y^\ddagger)$ and flow (solid gray edges) that corresponds to the ex post allocation rule of the serial-dictator mechanism. Recall that this serial-dictator mechanism allocates to agent 1 if she has a high type H_1 and to agent 2 otherwise. This ex post allocation rule can be determined by inspecting the out-going flow from vertices corresponding to type profiles, i.e., the left-side column. Subfigure (b) depicts the flow graph that corresponds to the profile of interim allocation rules $\mathbf{y} = (y^\dagger, y^\ddagger)$ which require that an agent is served if and only if she has a high type. This profile is infeasible, which can be seen as the minimum a - b cut, depicted with the black dashed line, has cost $3/4$ (it cuts three edges with capacity $1/4$ each and two edges with capacity zero) while the total capacity of edges incident on sink b is one. Inequality (8.9) is violated for subsets of types \mathbf{S}^* with $S_i^* = \{H_i\}$ for each i , i.e., corresponding to the vertices in the right-side column that are on the sink b side of the cut.

style argument of the proof below. The following theorem is often called Border's Theorem in recognition of Kim Border's pioneering study of interim feasibility.

Theorem 8.10. *For single-item environments, a profile of allocation rules \mathbf{y} (with cumulative allocation profile \mathbf{Y}) is interim feasible if and only if,*

$$\sum_i Y_i(\hat{q}_i) \leq 1 - \prod_i (1 - \hat{q}_i), \quad \forall \hat{\mathbf{q}} \in [0, 1]^n. \quad (8.8)$$

Proof. This proof is most instructive to see in type space. For finite type spaces, the inequality (8.8) of the theorem is equivalent to the following

(see Exercise 8.7). For any subsets of agents' types $S_i \subset \mathcal{T}_i$ for all i ,

$$\sum_i \sum_{t_i \in S_i} x_i(t_i) f_i(t_i) \leq 1 - \prod_i (1 - f_i(S_i)) \quad (8.9)$$

with $f_i(t_i)$ denoting the probability that agent i has type t_i and $f_i(S_i) = \sum_{t_i \in S_i} f_i(t_i)$ denoting the probability that i has type $t_i \in S_i$. Notice that the left-hand side of the equation is the expected number of items allocated to agents i with types $t_i \in S_i$. Notice that the right-hand side of the equation is simply the probability that one or more agents i have realized type $t_i \in S_i$. As describe above, the necessity of the condition for interim feasibility is straightforward.

The following argument shows sufficiency, specifically, that if a profile of allocation rules is infeasible that there exists subsets of agents' types $S_1^* \subset \mathcal{T}_1, \dots, S_n^* \subset \mathcal{T}_n$ for which inequality (8.9) is violated. The approach of this proof is (a) to show that interim feasibility is equivalent to whether a specific cut in a network flow graph is the minimum cut, and (b) to use the minimum cut corresponding to interim infeasible allocation rules to identify the subsets of agents' types that violate inequality (8.9).

Consider the following network flow problem, equivalently a weighted directed graph where the weights are referred to as capacities; see Figure 8.10. This graph is defined on the following vertices (left to right in Figure 8.10):

- a source vertex a ,
- a vertex \mathbf{t} for each type profile $\mathbf{t} \in \mathcal{T}$,
- a vertex t_i for each type $t_i \in \mathcal{T}_i$ of each agent i , and
- a sink vertex b .

Directed weighted edges connect these vertices as follows (left to right in Figure 8.10):

- source a is connected to each vertex \mathbf{t} with capacity $f(\mathbf{t}) = \prod_i f_i(t_i)$, i.e., the probability that type profile \mathbf{t} is realized;
- each vertex \mathbf{t} is connected to vertex t_i for each i with capacity $f(\mathbf{t})$; and
- each vertex t_i is connected to sink b with capacity $x_i(t_i) f_i(t_i)$, i.e., the probability that agent i has type t_i and is allocated by the interim allocation rule x_i .

A profile of interim allocation rules \mathbf{x} is feasible if and only if there is a flow in the flow graph constructed above that saturates all edges

incident on the sink b ; see Figure 8.10(a). For the “only if” direction, consider any ex post feasible mechanism that induces interim allocation rules \mathbf{x} and construct a flow as follows. Flow from source a to vertex \mathbf{t} represents the probability that type profile \mathbf{t} is realized. The flow from vertex \mathbf{t} to vertices t_i for each i represents the probability that \mathbf{t} is realized and agent i is served by the ex post mechanism. Since the total flow into vertex \mathbf{t} is $f(\mathbf{t})$ the cumulative flow out can be at most $f(\mathbf{t})$ which satisfies the ex post feasibility constraint that at most one of the agents is served. The flow on the edge from vertex \mathbf{t} to vertex t_i is $x_i(\mathbf{t}) f(\mathbf{t})$ as follows. Vertex t_i aggregates flow from each type profile \mathbf{t} containing t_i and thus the flow that can go from vertex t_i to sink b is the cumulative probability that t_i is realized and is served, i.e., $x_i(t_i) f_i(t_i)$. Thus, the edges incident on sink b are saturated. For the “if” direction, given any flow that saturates all the edges incident on the sink b , an ex post mechanism can be inferred. The ex post allocation on type profile \mathbf{t} picks an agent with probability equal to the flow from \mathbf{t} to t_i divided by $f(\mathbf{t})$.

Non-existence of a flow that saturates the edges incident on sink b implies that the profile of allocation rules \mathbf{x} is infeasible. We will now show that this non-existence of a flow will enable us to identify subsets of types $\mathbf{S}^* = (S_1^*, \dots, S_n^*)$ that violate inequality (8.9) and thus the inequality is sufficient for the interim feasibility of \mathbf{x} ; see Figure 8.10(b). An a - b cut in a directed graph is partitioning of the vertices into two sets $\{a\} \cup A$ and $\{b\} \cup B$. The capacity of the cut is the sum of the capacities of edges that cross from $\{a\} \cup A$ to $\{b\} \cup B$. The proof will show that the inequality (8.9) is satisfied for interim allocation rules \mathbf{x} only if $B = \emptyset$ is minimum capacity a - b cut, i.e., the capacity of the minimum cut is equal to the sum of the capacities of edges from vertices $t_i \in \mathcal{T}_i$ and all i to vertex b , specifically, $\sum_i \sum_{t_i \in \mathcal{T}_i} x_i(t_i) f_i(t_i)$.

Observe that the value of the maximum a - b flow in the graph is upper bounded by the capacity of any a - b cut in the graph. Simply, there is no way to get more flow across this cut than the total capacity of the cut. More precisely, the well known max-flow min-cut theorem states that the value of the maximum a - b flow in a flow graph is equal to the capacity of its minimum a - b cut. We now show that there is a flow that saturates the edges incident on sink b , equivalently, that $B = \emptyset$ is a minimum cut, if and only if there is no profile of subsets of type space (S_1, \dots, S_n) for which inequality (8.9) is violated.

We will calculate the difference between the capacity of the $B = \emptyset$ cut, i.e., the capacity edges incident on sink b , and the minimum cut.

When this difference is strictly positive we will identify a violation of inequality (8.9). Denote by (A^*, B^*) the minimum capacity a - b cut. The subsets of each agent's type space that are candidates for violation of inequality (8.9) are $S_i^* = B^* \cap \mathcal{T}_i$. The difference between the capacities of these two cuts is calculated as follows.

- For edges incident on sink b : The capacity of the cut (A, B) (with $B = \emptyset$) is equal to the sum of the capacities of edges incident on sink b . Subtracting from this the sum of capacities of edges crossing cut (A^*, B^*) , the difference is the sum of capacities of edges that are not cut by (A^*, B^*) . These uncut edges are the ones from vertices in B^* to sink b which correspond to types $t_i \in S_i^*$ for all i . The total contribution from these edges to the difference is thus, $\sum_i \sum_{t_i \in S_i^*} x_i(t_i) f_i(t_i)$, i.e., the left-hand side of inequality (8.9).
- For edges incident on vertices $\mathbf{t} \in \mathcal{T}$: Vertices corresponding to type profiles, e.g., \mathbf{t} , are either in B^* , in which case we have cut the edge from source a and must subtract $f(\mathbf{t})$, or in A^* , in which case we have cut edges with capacity $f(\mathbf{t})$ for each i with $t_i \in S_i^*$ (i.e., with $t_i \in B^*$) and must subtract $f(\mathbf{t}) \cdot |\{i : t_i \in S_i^*\}|$. Since (A^*, B^*) is a minimum cut, we must have chosen the smaller of these two quantities, i.e., $f(\mathbf{t}) \cdot \min(1, |\{i : t_i \in S_i^*\}|)$. Summing this quantity to be subtracted over all type profiles \mathbf{t} equates to the probability that one or more types $t_i \in S_i^*$ are realized, i.e., the right-hand side of inequality (8.9) of $1 - \prod_i (1 - f_i(S_i^*))$.

Combine these two contributions to the difference, and observe that when the difference is strictly positive then inequality (8.9) is violated for the subsets of types S_1^*, \dots, S_n^* . \square

This characterization of interim feasibility extends naturally to matroid environments where the right-hand side becomes the *expected rank*, with short-hand notation $\text{rank}(\hat{\mathbf{q}})$ representing $\mathbf{E}_S[\text{rank}(S)]$ where each i is in S independently with probability \hat{q}_i (cf. Section 4.3).

Theorem 8.11. *For matroid environments, a profile of allocation rules \mathbf{y} is interim feasible if and only if,*

$$\sum_i Y_i(\hat{q}_i) \leq \text{rank}(\hat{\mathbf{q}}), \quad \forall \hat{\mathbf{q}} \in [0, 1]^n.$$

One way this characterization of interim feasibility is helpful is as follows. It can be shown that the matroid rank function $\text{rank}(\cdot)$ is submodular. This submodularity implies that the interim feasibility constraint has a polymatroidal structure which, in turn, implies that the vertices

corresponding to the feasible region can be implemented by greedily ordering types and serving each type to the maximum extent possible. Instead of introducing this polymatroidal theory of optimization, we will give an alternative first-principles proof of this result in the next section (see Corollary 8.15).

For feasibility constraints beyond matroid, we will not get a succinct formula like inequality (8.8) in Theorem 8.10 that characterizes interim feasibility. Nonetheless, in the next section we will describe a simple family of ex post mechanisms from which any profile of interim feasible allocation rules can be derived.

8.4.3 Interim Feasibility by Stochastic Weighted Optimization

In this section we show that, for any service constrained environment, any interim feasible profile of allocation rules is implementable as a stochastic weighted optimization. This characterization is derived by observing that:

- (i) there is an isomorphism between profiles of interim allocation rules to points in a high dimensional Euclidean space,
- (ii) the set of interim feasible points by this isomorphism is convex and, specifically, a polytope,
- (iii) any point in the interior of this polytope can be given as a convex combination of points on the exterior of the set, specifically, vertices of the polytope, and
- (iv) these vertices can be implemented by weighted optimization, i.e., a mapping of each type of each agent in the profile to a weight and selection of the ex post feasible set of agents with highest cumulative weight.

We relax two constraints from the Bayesian mechanism design problem. Relaxing incentive compatibility, allocation rules need not be monotone and we will, thus, work in type space rather than quantile space.⁴ Relaxing the independence of the distribution of types across agents, we

⁴ The approach we take is similar to that of the characterization of interim feasibility for single-item and matroid environments (Theorem 8.10 and Theorem 8.11) which was also described in type space. At the end of this section we will describe how to modify the construction for quantile space and when this approach is helpful.

draw type profile \mathbf{t} from joint distribution \mathbf{F} and denote by $f_i(t)$ the marginal probability that agent i has type t , i.e., $\Pr_{\mathbf{t} \sim \mathbf{F}}[t = t_i]$.⁵

For Step (i), consider a finite space of type profiles $\mathcal{T} = \mathcal{T}_1 \times \cdots \times \mathcal{T}_n$ with size $\ell = \sum_i |\mathcal{T}_i|$ and map profiles of interim allocation rules \mathbf{x} , with $x_i : \mathcal{T}_i \rightarrow [0, 1]$, to points $\mathbf{z} \in [0, 1]^\ell$, in a high dimensional Euclidean space. In this mapping the vector \mathbf{z} will be indexed by agent-type pairs it as z_{it} .

Definition 8.9. For joint type space \mathcal{T} with size $\ell = \sum_i |\mathcal{T}_i|$ and joint distribution \mathbf{F} , the *flattened ex post allocation rule* $\mathbf{z}^{EP} : \mathcal{T} \rightarrow [0, 1]^\ell$ and *flattened interim allocation* $\mathbf{z} \in [0, 1]^\ell$ are induced by ex post and interim allocation rules \mathbf{x}^{EP} and \mathbf{x} and indexed by it for agent i and type $t \in \mathcal{T}_i$ as:

$$z_{it}^{EP}(\mathbf{t}) = \begin{cases} x_i^{EP}(\mathbf{t}) & \text{if } t = t_i \\ 0 & \text{otherwise,} \end{cases} \quad z_{it} = x_i(t) f_i(t).$$

Notice that, by the above definition, the flattened interim allocation is in fact specifying the ex ante probability that each type of each agent is served. The normalization by the density function in the definition of the flattened interim allocation serves a similar purpose in the geometry of interim feasibility as the mapping of types to quantiles. Definition 8.9 is useful as it immediately gives the following propositions; the second of which concludes Step (ii).

Proposition 8.12. *The flattened interim allocation is the expectation of the flattened ex post allocation rule:*

$$\mathbf{z} = \mathbf{E}_{\mathbf{t} \sim \mathbf{F}}[\mathbf{z}^{EP}(\mathbf{t})]. \quad (8.10)$$

Proposition 8.13. *For joint type space \mathcal{T} with size $\ell = \sum_i |\mathcal{T}_i|$ and distribution \mathbf{F} , the space $\mathcal{Z} \subset [0, 1]^\ell$ of feasible flattened interim allocations \mathbf{z} is convex.*

Proof. Randomized mechanisms are feasible, and flattened interim allocations are linear with respect to convex combinations. \square

The flattened ex post allocation rule is a redundant representation of the ex post allocation rule, it specifies allocation probabilities for all types an agent might possess. For a given type profile, all of these probability must be zero except for the ones that correspond to types in the

⁵ The relaxation to (possibly) correlated distributions over type profiles will allow the results of this section to generalize beyond service constrained environments as in Section 8.5.

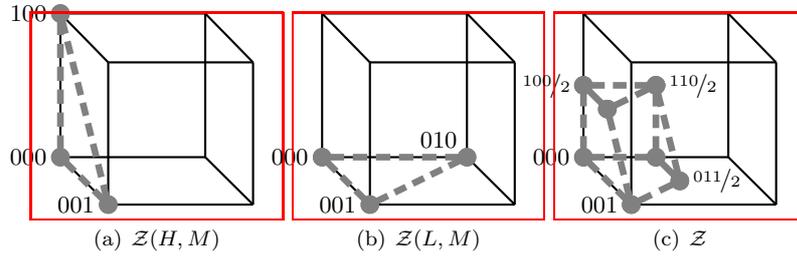


Figure 8.11. The polytopes that define the feasible flattened allocations of Example 8.7 are depicted as convex subsets of the unit cube. Selected vertices of these polytopes are labeled in short-hand as bit vectors and correspond to $\mathbf{z} = (z_{1H}, z_{1L}, z_{2M})$. The vertical axis is z_{1H} ; the horizontal axis is z_{1L} ; and the outward axis is z_{2M} . The ex post feasible flattened allocations for type profiles $\mathbf{t} = (H, M)$ and $\mathbf{t} = (L, M)$ are depicted in figures (a) and (b), respectively. The interim feasible flattened allocations when t_1 is uniform on $\{L, H\}$ and $t_2 = M$ (deterministically) are depicted in figure (c).

type profile. Specifically, consider an agent-type pair it with $t \in \mathcal{T}_i$ and type profile \mathbf{t} , if $t_i \neq t$ in type profile \mathbf{t} then $z_{it}^{EP}(\mathbf{t}) = 0$ as a mechanism cannot serve a type that “does not show up.” The types served must also satisfy the feasibility constraint of the service constrained environment.

Denote the feasibility constraint imposed by the service constrained environment by \mathcal{X} . Randomized mechanisms are allowed, thus \mathcal{X} is convex. For example in a single-item environment deterministic outcomes correspond to allocating to a single agent i or not allocating, the convex closure of these outcomes gives $\mathcal{X} = \{\mathbf{x} \in [0, 1]^n : \sum_i x_i \leq 1\}$. Ex post feasibility for flattened allocation rules $z^{EP}(\mathbf{t})$ is the projection of the service constrained feasibility constraint onto the indices $\{it_i\}_{i \in [n]}$ of the types in the given type profile \mathbf{t} . Specifically, for any profile \mathbf{t} , $\mathbf{z}^{EP}(\mathbf{t}) \in \mathcal{Z}(\mathbf{t})$ is ex post feasible, where

$$\mathcal{Z}(\mathbf{t}) = \{\mathbf{z} \in [0, 1]^\ell : \times_i z_{it_i} \in \mathcal{X} \wedge z_{\{it : t \neq t_i\}} = \mathbf{0}\}.$$

This projection is depicted in Figure 8.11 (for Example 8.7, below).

Example 8.7. Consider two agents and a single-item environment. Agent 1 has type t_1 uniformly drawn from $\mathcal{T}_1 = \{L, H\}$; agent 2 has type deterministically $t_2 = M$ (i.e., with $\mathcal{T}_2 = \{M\}$). A flattened allocation is $\mathbf{z} = (z_{1H}, z_{1L}, z_{2M})$. Ex post feasible allocations for type profile $\mathbf{t} = (H, M)$ are convex combinations of $\{(0, 0, 0), (1, 0, 0), (0, 0, 1)\}$; ex post feasible allocations for type profile $\mathbf{t} = (L, M)$ are convex combinations

of $\{(0, 0, 0), (0, 1, 0), (0, 0, 1)\}$. Interim feasible flattened allocations are convex combinations of $\{(0, 0, 0), (1/2, 0, 0), (0, 1/2, 0), (0, 0, 1), (1/2, 1/2, 0), (1/2, 0, 1/2), (0, 1/2, 1/2)\}$. The vertices of the interim feasible polytope are given by the ordinal mechanisms, i.e., where there is an ordered subset of types and after the types are realized the first one in order receives the item (or none if no type in the order is realized); see Corollary 8.15. The ordered subsets that correspond to the vertices above are $\{\emptyset, (1H), (1L), (1M), (1H, 1L), (1H, 2M), (1L, 2M)\}$. See Figure 8.11.

We have seen that interim feasibility is convex (Proposition 8.13); as previously observed, the single-agent optimal revenues given by $\mathbf{Rev}[\cdot]$ are concave. Suppose that, instead of the concave objective given by the sum of the single agent revenues, the objective was linear and given by weights $\mathbf{w} \in \mathbb{R}^\ell$ indexed in the flattened space by agent-type pair it . The optimization of the expected *surplus of weights*, i.e., $\sum_{it} z_{it} w_{it}$, subject to interim feasibility, is achieved by optimizing the surplus of weights pointwise for each profile of types \mathbf{t} subject to ex post feasibility. Moreover, for any such weights (which we view as a direction in the flattened space) the corresponding interim feasible allocation vector, denoted $\mathbf{z}^{\mathbf{w}}$, is given by Proposition 8.12.

A *vertex* of a convex subset \mathcal{Z} of ℓ -dimensional Euclidean space is a point that is uniquely optimal for some direction. Vertices can be specified equivalently as the direction, e.g., \mathbf{w} , or the point, e.g., $\mathbf{z}^{\mathbf{w}}$. Any other point in \mathcal{Z} can be represented as a convex combination of $\ell + 1$ vertices; therefore, we can implement any interim feasible allocation rule by sampling a direction from a distribution over $\ell + 1$ vectors of weights, and then for the type profile realized, optimizing the weights given by the direction subject to ex post feasibility.

A weights \mathbf{w} in the space of flattened allocation rules correspond, in the original space of allocation rules, to a profile of functions that map each type to a weight. The following theorem summarizes the construction above in the original space of allocation rules.

Definition 8.10. A *stochastic weighted optimizer* is given by a joint distribution over profiles of weight functions \mathbf{w} with $w_i : \mathcal{T}_i \rightarrow \mathbb{R}$ as follows for type profile \mathbf{t} :

- (i) Draw weight functions \mathbf{w} from the distribution.
- (ii) Output allocation $\mathbf{x} = \operatorname{argmax}_{\mathbf{x}^\dagger \in \mathcal{X}} \sum_i w_i(t_i) x_i^\dagger$.

Theorem 8.14. For any joint distribution on type profiles and service

constrained environment, any interim feasible allocation profile can be ex post implemented by a stochastic weighted optimization.

In the special case that the service constrained environment is ordinal, e.g., multi-unit environments, matroid environments, and position environments, the surplus of weights is optimized by the greedy algorithm (See Section 4.6 on page 129). The greedy algorithm, by definition, considers only the order of weights of each type and not magnitudes of the weights. The following corollary refines Theorem 8.14 for ordinal environments.

Definition 8.11. A *stochastic ordered-subset algorithm* is given by a joint distribution over ordered subsets of joint type space $\bigcup_i \mathcal{T}_i$ as follows for type profile \mathbf{t} :

- (i) Draw an ordered subset from the distribution.
- (ii) Output allocation \mathbf{x} obtained by the greedy algorithm on agents ordered by the rank of their types in the ordering; agents whose types are not present in the subset are discarded.

Corollary 8.15. *For any joint distribution on type profiles and service constrained environment that is given by the independent sets of a matroid, any interim feasible allocation profile can be ex post implemented by a stochastic ordered-subset algorithm.*

This characterization of interim feasible allocation rules as a convex subset of high-dimensional Euclidean space is central to the design of computationally-efficient revenue-optimal mechanisms. The main challenges to be resolved is in quickly finding the distribution over weights Theorem 8.14. Discussion of the computational issues involved are deferred to ??.

Though this section approached the characterization of interim feasibility in type space, it can be equivalently characterized in quantile space as well. For example, discretize quantile space into intervals and apply the construction in this section with each discrete interval for each agent as a type. This approach is advantageous when the agents' type spaces are very large, or high dimensional as quantiles are always single dimensional. Again, further discussion is deferred to ??.

8.4.4 Combining Ex post Feasibility and Bayesian Incentive Compatibility

This section formalizes the general constructions for ex post implementation of interim mechanisms. Let $\hat{\mathcal{M}}$ denote an ex post feasible mechanism (that is not necessarily incentive compatible or possessing revenue guarantees). Its ex post allocation rule maps quantile profiles to distributions over ex post feasible allocations via $\hat{\mathbf{y}}^{EP} : [0, 1]^n \rightarrow \Delta(\mathcal{X})$. Recall from Section 2.4 that the induced interim allocation rule $\hat{y}_i : [0, 1] \rightarrow [0, 1]$ for agent i is defined as $\hat{y}_i(q_i) = \mathbf{E}_q[\hat{y}_i^{EP}(\mathbf{q}) \mid q_i]$. Let \mathcal{M} denote a Bayesian incentive compatible mechanism (potentially with good revenue properties, but that is not necessarily ex post feasible). Its interim allocation rules are denoted by \mathbf{y} with $y_i : [0, 1] \rightarrow [0, 1]$ for each agent i . We can compose these two mechanisms to obtain a mechanism with the ex post feasibility of $\hat{\mathcal{M}}$ and the Bayesian incentive compatibility (and revenue properties) of \mathcal{M} if and only if the allocation rules \mathbf{y} are feasible for allocation constraints $\hat{\mathbf{y}}$.

This construction can be instantiated with the revenue-optimal mechanisms of the previous section. In such an instantiation, $\hat{\mathcal{M}}$ is the ex post mechanism that induces the interim allocation rules $\hat{\mathbf{y}}$ that optimize $\sum_i \mathbf{Rev}[\hat{y}_i]$ subject to interim feasibility, and \mathcal{M} is the profile of interim mechanism that optimize revenue subject to the allocation constraints $\hat{\mathbf{y}}$, i.e., with \mathcal{M}_i as the \hat{y}_i interim optimal mechanism.

First, given any single-agent allocation constraint $\hat{y} : [0, 1] \rightarrow [0, 1]$ and single-agent mechanism \mathcal{M} with allocation rule y that satisfies constraint \hat{y} , we give an ex post implementation of y from \hat{y} . Second, given any ex post implementation $\hat{\mathbf{y}}^{EP}$ that induces interim constraints $\hat{\mathbf{y}}$ and a profile of single-agent mechanisms $(\mathcal{M}_1, \dots, \mathcal{M}_n)$, where y_i satisfies \hat{y}_i for each agent i , we give an ex post implementation of a mechanism with allocation rules \mathbf{y} .

Recall that a single-agent mechanism \mathcal{M} is equivalently a menu of outcomes $\{w(t) : t \in \mathcal{T}\}$. A deterministic outcome is either a service outcome or a non-service outcome. Outcomes are closed under convex combination, i.e., they may be randomized. The (type) allocation rule $x : \mathcal{T} \rightarrow [0, 1]$ gives a probability of service for each type $t \in \mathcal{T}$. For $x(t) \in [0, 1]$, the outcome distribution $w(t)$ is a distribution over service and non-service outcomes. Denote by $w^x(t)$ the distribution of outcomes conditioned on the allocation $x \in \{0, 1\}$.

Any single-agent mechanism \mathcal{M} induces an ordering on types which in turn induces a mapping from types to quantiles. Correctness requires

that the distribution of quantiles from this mapping be uniform on the $[0, 1]$ interval.

Definition 8.12. The *quantile mapping* for mechanism \mathcal{M} with (type) allocation rule $x(\cdot)$ is $Q(\cdot)$ defined as follows. For any type t , calculate interval $[\hat{q}^\dagger, \hat{q}^\ddagger]$ as

$$\hat{q}^\dagger = \Pr_{t^\dagger \sim F} [x(t^\dagger) > x(t)], \quad \hat{q}^\ddagger = \Pr_{t^\ddagger \sim F} [x(t) < x(t^\ddagger)].$$

The stochastic mapping from types to quantiles is Q defined as:

$$Q(t) \sim U[\hat{q}^\dagger, \hat{q}^\ddagger].$$

Lemma 8.16. For types $t \in \mathcal{T}$ from distribution F , the quantile distribution of the quantile mapping $Q(t)$ (of Definition 8.12) is uniform on $[0, 1]$.

From the above induced mapping from types to quantiles and a procedure for the allocation rule $y(\cdot)$ of mechanism \mathcal{M} , the mechanism can be implemented with $y(\cdot)$ as follows:

- (i) Calculate the agent's quantile as $q = Q(t)$.
- (ii) Calculate the agent's service as

$$x = \begin{cases} 1 & \text{w.p. } y(q) \\ 0 & \text{otherwise.} \end{cases}$$

- (iii) Calculate the agent's outcome as $w = w^x(t)$.

To generalize this construction to procedures for allocation constraints $\hat{y}(\cdot)$ that allocation rule $y(\cdot)$ satisfies, we need to convert the procedure for \hat{y} to a procedure for y .

Suppose we have an interim allocation constraint \hat{y} and a mechanism \mathcal{M} with allocation rule y that satisfies the constraint, i.e., $Y(q) \leq \hat{Y}(q)$ for all q . The following lemma shows that we can implement y from \hat{y} ; thus, by the above construction, we can implement \mathcal{M} .

Lemma 8.17. Any allocation rule y that satisfies allocation constraint \hat{y} can be implemented by a quantile reserve pricing \hat{q} and a stationary quantile resampling transformation σ .

Proof. This proof will be by construction; see Figure 8.12. First, we will construct \hat{y}^\dagger from \hat{y} with a quantile reserve so as to equate the ex ante service probabilities $\hat{Y}^\dagger(1) = Y(1)$ while preserving feasibility of

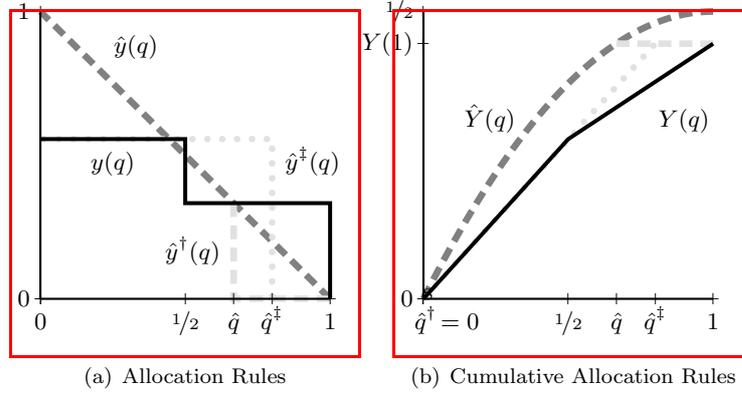


Figure 8.12. The reserve price and resampling transformation construction of Lemma 8.17 is depicted for a piecewise constant allocation rule y with $\ell = 2$ pieces. The allocation constraint \hat{y} is thick, dashed, and dark gray, the allocation rule y is thin, solid, and black. The allocation constraint \hat{y}^\dagger (think, dashed, and light gray) is constructed from \hat{y} by reserve pricing at quantile \hat{q} . The allocation constraint \hat{y}^{\ddagger} (think, dotted, and light gray) is constructed from \hat{y}^\dagger by ironing on interval $[\hat{q}^\dagger, \hat{q}^\ddagger]$.

y for \hat{y}^\dagger , i.e., so $Y(q) \leq \hat{Y}^\dagger(q)$ for all q . Second, we will give a stationary quantile resampling transformation σ , i.e., with $\sigma(q)$ uniformly distributed on $[0, 1]$ if q is uniform on $[0, 1]$, that transforms \hat{y}^\dagger to y , i.e., $y(q) = \mathbf{E}_\sigma[\hat{y}^\dagger(\sigma(q))]$.

The allocation constraint \hat{y}^\dagger is obtained from allocation constraint \hat{y} by quantile reserve pricing at $\hat{q} = \hat{Y}^{-1}(Y(1))$. Recall, quantile reserve pricing has the effect of replacing the cumulative allocation rule with a constant function after the quantile reserve; thus $\hat{Y}^\dagger(1) = Y(1)$.

Assume that y is piece-wise constant, equivalently that Y is piece-wise linear, with ℓ equal-width pieces. This assumption can be removed by considering y in the limit, as ℓ goes to infinity, of such a piece-wise constant allocation rule. The following inductive procedure gives a stationary resampling transformation σ that constructs y from \hat{y}^\dagger . Let \hat{q}^\dagger be the lower end point of the first piece on which $\hat{Y}^\dagger(\cdot)$ and $Y(\cdot)$ are distinct, equivalently, where the upper end point q of the piece satisfies $\hat{Y}^\dagger(q) > Y(q)$. Note that the right slope of $Y(\cdot)$ at \hat{q}^\dagger is equal to the right value of $y(\cdot)$ at \hat{q}^\dagger . Set $\hat{q}^\ddagger > \hat{q}^\dagger$ to the quantile at which the line through point $(\hat{q}^\dagger, Y(\hat{q}^\dagger))$ with slope $y(\hat{q}^\dagger)$ next intersects the cumulative allocation constraint $\hat{Y}^\dagger(\cdot)$. Define σ^\dagger to be the interval resampling

transformation that irons on quantile interval $[\hat{q}^\dagger, \hat{q}^\ddagger]$, i.e.,

$$\sigma^\dagger(q) = \begin{cases} q^\dagger \sim U[\hat{q}^\dagger, \hat{q}^\ddagger] & \text{if } q \in [\hat{q}^\dagger, \hat{q}^\ddagger], \\ q & \text{otherwise.} \end{cases}$$

By the line-segment interpretation of ironing on the cumulative allocation rule, this resampling transformation gives an allocation constraint $\hat{y}^\dagger(q) = \mathbf{E}_{\sigma^\dagger}[\hat{y}^\dagger(\sigma^\dagger(q))]$ with $Y^\dagger(q) = Y(q)$ for q in the piece.

By this construction \hat{y}^\dagger differs from y on (at least) one fewer piece than \hat{y}^\ddagger . By induction we can construct a sequence of interval resampling transformations that, when composed, transform \hat{y}^\dagger to y . The transformation σ in the statement of the lemma is this composition of interval resampling transformations. The stationarity of each interval resampling transformation implies that the transformation σ is stationary. \square

Definition 8.13. For distribution \mathbf{F} and mechanisms $\hat{\mathcal{M}}$ and \mathcal{M} with allocation rules \hat{y} and y satisfying $y \preceq \hat{y}$, the *interim composite mechanism* is:

- (i) For each agent i , map type to quantile from $q_i = Q_i(t_i)$ according to \mathcal{M}_i .
- (ii) For each agent i , calculate the quantile reserve \hat{q}_i and resampling transformation σ_i by which y_i can be constructed from \hat{y}_i . Set $q_i^\dagger = \sigma_i(q_i)$.
- (iii) Run $\hat{\mathcal{M}}$ on q^\dagger to get $x^\dagger = \hat{y}^{EP}(q^\dagger)$. Set x to incorporate the reserves \hat{q} as

$$x_i = \begin{cases} x_i^\dagger & \text{if } q_i^\dagger \leq \hat{q}_i, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

- (iv) For each agent i , select the outcome distribution of \mathcal{M}_i conditioned on x_i , i.e., $w_i^{x_i}(t_i)$.

Theorem 8.18. For any type distribution \mathbf{F} and mechanisms $\hat{\mathcal{M}}$ and \mathcal{M} with allocation rules \hat{y} and y satisfying $y \preceq \hat{y}$, the composite mechanism (Definition 8.13) induces a distribution over allocations that is in the downward closure of the distribution of allocations of $\hat{\mathcal{M}}$ and the same interim mechanisms as \mathcal{M} .⁶

Proof. See Exercise 8.9. \square

⁶ One distribution of allocations is in the downward closure of a second distribution of allocations if there is a coupling of the distributions so that the set of agents served by the first is a subset of those served by the second.

Corollary 8.19. *For any downward-closed service constrained environment, type distribution \mathbf{F} , ex post feasible mechanism $\hat{\mathcal{M}}$, and Bayesian incentive compatible \mathcal{M} with allocation rules $\hat{\mathbf{y}}$ and \mathbf{y} satisfying $\mathbf{y} \preceq \hat{\mathbf{y}}$, the composite mechanism (Definition 8.13) is ex post feasible, Bayesian incentive compatible, and has the same expected revenue as \mathcal{M} .*

8.5 Multi-dimensional Externalities

This section considers optimal mechanisms for agents with multi-dimensional preferences where the way an agent is served imposes a multi-dimensional externality on the other agents via the feasibility constraint of the environment. For example, in the *multi-dimensional matching environment*, there are n unit-demand agents and m unit-supply items. The environment exhibits a multi-dimensional externality because when an item j is assigned to agent i then it cannot be assigned to another agent $i^\dagger \neq i$ but other items $j^\dagger \neq j$ can be so assigned.

Definition 8.14. In a *multi-service service-constrained environment* there are n agents N and m services M . The subset of agent-service pairs that can be simultaneously assigned is given by $\mathcal{X} \subset \{0, 1\}^{N \times M}$.

The multi- to single-agent reduction that was described in the previous section separates the problem of producing an outcome that is ex post feasible from the problem of ensuring that the mechanism is incentive compatible for each agent. This section takes the same approach; it applies generally to environments where each agent's utility linearly separates across distinct services in which she is interested. For simplicity we will state all results for the special case of additive agents.

Definition 8.15. An *additive* agent desires subsets of m services. Her type $t = (\{t\}_1, \dots, \{t\}_m)$ is m -dimensional where $\{t\}_j$ is her value for alternative j . Her utility is linear; an outcome w is given by a payment and a marginal probability for each of the m services. For outcome $w = (\{x\}_1, \dots, \{x\}_m, p)$, where $\{x\}_j$ denotes the marginal probability with which she obtains alternative j and p is her required payment, her utility is $u(t, w) = \sum_j \{t\}_j \{x\}_j - p$.

Note that it is possible within an additive multi-service environment to model more complex preferences. As a first example, unit-demand preferences (Definition 8.2) can be incorporated into the model by modifying the feasibility constraints \mathcal{X} of Definition 8.14 so that it is infea-

sible to serve an agent more than one unit. It is also possible to model any general utility function over bundles of services as a unit-demand utility function over the power set of services, i.e., $\{0, 1\}^M$. We will see shortly, however, that the complexity of the construction depends on the number m of services, and thus moving to the power-set representation comes at an exponential blowup in complexity.

The assumption that the agents are additive implies, as is stated in the definition, the expected utility of an agent is determined by the marginal probability by which she is allocated each service. This property does not hold for general multi-dimensional utility functions. For example, a problematic case is when the agent views the services as complementary, e.g., she has high value to receive two services together but low value to receive either of the services individually. A mechanism that serves her both services or neither service with equal probability has the same marginal probabilities of allocating each service as the mechanism that serves her one or the other with equal probability. The agent has a higher utility for the former outcome distribution than the latter; thus, marginal probabilities are insufficient for determining such an agent's utility. In fact, any non-linearity of utility renders marginal probabilities similarly insufficient.

It is possible to decompose this mechanism design problem into a collection of single-agent problems that can be combined into a multi-agent mechanism, as we did in the previous section. In such a decomposition the single-agent problem is specified by m allocation constraints, one for each service. The difference between mechanisms for additive multi-service service constrained environments and the (single-service) service constrained environments is that the incentive compatibility constraints of the agents bind across multiple services not a single service. We will not formalize this approach here, instead we show that for additive multi-service service constrained environments, the optimal mechanism is a stochastic weighted optimizer (cf. Definition 8.10).

Our objective is to optimize expected revenue subject to Bayesian incentive compatibility and ex post feasibility. As we did in previous sections, we will replace the ex post feasibility constraint with an equivalent interim feasibility constraint. Ex post feasibility of a multi-service service constrained environment is equivalent to ex post feasibility of the following representative environment which is service constrained as per Definition 8.4. The intuition behind this representative environment is that we replace each multi-service agent, i.e., who desires subsets of the m services, with m single-service agents.

Definition 8.16. The *representative environment* for an n -agent m -service multi-service service-constrained environment is given by $n^\dagger = nm$ agents $N^\dagger = N \times M$ and service-constrained feasibility constraint $\mathcal{X}^\dagger = \mathcal{X} \subset 2^{N^\dagger}$. The type profile \mathbf{t} for the original environment is extended to representative environment by duplicating each agent i 's type across her m representatives, i.e., $t_{ij} = t_i$ for all i and j .

Importantly, ex post feasibility of the representative environment and the original multi-service environment is the same. Consequently, our discussion of interim feasibility extends directly. Recall that our discussion of interim feasibility relaxed the requirement that the agent types be independently distributed. This relaxation is important as in the representative environment for the multi-service environment all the representatives ij for $j \in M$ have the same type t_i , i.e., they are perfectly correlated. Thus, the characterization of interim feasibility and ex post implementations (Theorem 8.14) of the previous section hold for the representative environment.

Corollary 8.20. *For any joint distribution on type profiles and multi-service service-constrained environment, any interim feasible allocation profile can be ex post implemented by a stochastic weighted optimizer with weights that correspond to each type-service pair of each agent, i.e., \mathbf{w} with $w_i : \mathcal{T}_i \times M \rightarrow \mathbb{R}$ for each i .*

Recall, that a weighted optimizer for the representative environment assigns a weight to each type t_{ij} of each representative ij (see Definition 8.10); in the original environment such an assignment of weights corresponds to a weight for each type t_i , agent i , and service j (though weights for services that are infeasible for agent i can be omitted).

Theorem 8.21. *For any additive multi-service service-constrained environment, there is a stochastic weighted optimizer, with weights that map each feasible type-service pair, that is Bayesian incentive compatible and revenue-optimal among all Bayesian incentive compatible mechanisms.*

Proof. Consider any optimal mechanism \mathcal{M}^* . The optimal mechanism must produce interim feasible allocation rules (specifying the marginal probability by which an agent of a given type receives each service). By Corollary 8.20 any profile of interim feasible allocation rules can be implemented as a stochastic weighted optimizer. Consider the mechanism given by this stochastic weighted optimizer and the same payment rule of \mathcal{M}^* .

By the definition of additive utility agents, each agent's utility for a randomized outcome depends only on the marginal probability that she receives each service (and expected payment). These marginal probabilities and expected payments are the same for both mechanisms; thus, incentive compatibility of \mathcal{M}^* implies incentive compatibility of the stochastic weighted optimizer. Since both mechanisms have the same payment rule, they both have the same revenue; the optimality of \mathcal{M}^* implies the optimality of the stochastic weighted optimizer. \square

8.6 Public Budget Preferences

In the sections below we will prove the optimality of the single-agent mechanisms described in Section 8.1 for agents with public budgets. In particular, we will show that for a large family of well-behaved distributions the revenue-optimal single-agent mechanism will have an all-pay payment rule and will reserve-price the low valued agents and iron the top valued agents.

Recall the public budget preference where the agent has a single dimensional value t drawn from distribution F and public budget B . The agent's utility for allocation x and payment p is $tx - p$ when $p \leq B$ and negative infinity of $p > B$. We will assume that the distribution F is continuous and supported on types $\mathcal{T} = [0, h]$.

We begin by observing that, for an agent with a public budget, the optimal mechanism, satisfying the usual Bayesian incentive compatibility (BIC) and interim individual rationality (IIR) constraints, is an all-pay mechanism. In other words, the agent makes a bid and pays this bid always, though she may only win some of time. All-pay mechanisms may seem unnatural as they are not ex post individually rational, i.e., an agent will sometimes have negative utility. Notice, however, that in most economic interactions there are upsides and downsides that strategic agents must trade off; ex post individual rationality is the exception rather than the rule. Moreover, as non-all-pay mechanisms will generally be suboptimal, ex post individual rationality comes at a loss in performance, in this case revenue, relative to the optimal all-pay format.

Proposition 8.22. *The revenue-optimal Bayesian-incentive-compatible and interim-individually-rational mechanism for single-dimensional agents with public budgets is an all-pay mechanism.*

Proof. An agent with public budget is quasi-linear except for her budget

constraint. Therefore, unless the budget constraint is violated, revenue equivalence of Section 2.7 on page 37 implies two mechanisms with the same allocation rule in equilibrium have the payment rule (in the interim stage of the mechanism).

Consider any mechanism where, in equilibrium, the agent's budget constraint is not violated. Recall that the payment rule $p(t)$ is defined as the expected payment of the agent. For a given valuation t the payment of an agent is a random variable, potentially a function of randomization in the mechanism and randomization in the types of other agents. By the definition of expectation, the maximum payment in the support of the distribution of payments is at least the expected payment. As the budget is not violated for this maximum payment, it is not violated for the expected payment, i.e., $p(t) \leq B$. Thus, the all-pay mechanism that requires deterministic payment $p(t)$ does not violate the budget either. Therefore, it is incentive compatible and obtains the same revenue. \square

We now proceed to characterize the optimal single-agent all-pay mechanism subject to an interim feasibility constraint $\hat{y}(\cdot)$. This optimization problem is similar to that for the single-dimensional linear agent that was previously solved in Section 3.3; however, the solution to the optimization must additionally satisfy the (all-pay) budget constraint that $p(t) \leq B$ for all $t \in \mathcal{T}$. As payments are non-decreasing in the agent's type, the budget constraint for all types $t \in \mathcal{T} = [0, h]$ is implied by the budget constraint for the highest type h . In other words, the revenue-optimization problem has only the additional constraint $p(h) \leq B$.

Our approach will be to write a mathematical program for the revenue maximization problem where the budget appears as a constraint. We will then use Lagrangian relaxation to move the budget constraint into the objective.⁷ We will proceed by optimizing this Lagrangian objective in

⁷ For maximization problems, Lagrangian relaxation of a constraint (a) rewrites it as a quantity that is at least zero and (b) and moves the terms of the constraint, scaled by a Lagrangian parameter λ , to the objective scales. Thus, satisfying the constraint is consistent with the objective, i.e., this term is larger when the constraint is satisfied. The Lagrangian parameter λ allows the emphasis of the Lagrangian objective to be traded off between the original objective and satisfaction of the constraint. At $\lambda = 0$ the constraint is ignored and is only satisfied if it was not binding in the first place. At $\lambda = \infty$, the objective is ignored and the constraint is satisfied with slack (if it is satisfiable by any assignment of the variables of the program). The original program with the constraint is optimized by finding the λ where the constraint is met with equality; at such a point the Lagrangian program trades off emphasis on the objective and the constraint perfectly. Notice that when the constraint is met with equality, the contribution to the objective is zero and the objective of the Lagrangian program is the optimal value of the original program.

the same manner as our revenue optimization for single-dimensional linear agents, cf. Section 3.3.4. We will rewrite the objective in terms of the allocation rule and Lagrangian revenue curves. For a given Lagrangian parameter λ , and these revenue curves, we will be able to identify the optimal allocation rule for any allocation constraint, cf. Section 3.4.5. Finally, we choose the Lagrangian parameter so that the budget is met with equality.

We begin by writing a mathematical program for the interim revenue maximization problem and using Lagrangian relaxation to move the budget constraint into the objective. The other constraints of the problem will not play a roll in most of our discussion so we will not write them formally. The original and relaxed programs are as follows; recall the value of the highest type is denoted h .

$$\begin{array}{ll} \sup_{(x,p)} \mathbf{E}_{t \sim F}[p(t)] & \sup_{(x,p)} \mathbf{E}_{t \sim F}[p(t)] + \lambda B - \lambda p(h) \\ \text{s.t. } (x, p) \text{ are BIC, IIR,} & \text{s.t. } (x, p) \text{ are BIC, IIR,} \\ \text{and feasible for } \hat{y}; & \text{and feasible for } \hat{y}. \\ p(h) \leq B. & \end{array}$$

We will fix the Lagrangian parameter λ and characterize the optimizer of the Lagrangian objective. Notice that this Lagrangian objective is linear and therefore can be treated with the methods of Chapter 3. With such a characterization, the Lagrangian parameter λ can be chosen to be zero if the budget is not binding or so that the budget constraint is met with equality if it is binding.

With Lagrangian λ fixed, the λB term in the objective is a constant and does not affect optimization. The optimization is to find allocation and payment rules (x, p) to maximize $\mathbf{E}[p(t)] - \lambda p(h)$. Our approach to this optimization problem will mirror our approach to revenue optimization without budgets, cf. Section 3.3.4. We will define Lagrangian revenue curves, we will write the Lagrangian objective of any allocation rule in terms of these revenue curves, and then we will directly interpret the form of the Lagrangian optimizer.

Recall that price-posting revenue curves are defined by considering the ex ante constraint \hat{q} and the mechanism $(x^{\hat{q}}, p^{\hat{q}})$ that posts the price $V(\hat{q})$ that is accepted with probability \hat{q} . Consider the Lagrangian objective $\mathbf{E}_t[p^{\hat{q}}(t)] - \lambda p^{\hat{q}}(h)$ for the mechanism that posts price $V(\hat{q})$. The revenue from such a price is $\mathbf{E}_t[p^{\hat{q}}(t)] = P(\hat{q}) = \hat{q} V(\hat{q})$ where, recall, $P(\hat{q})$ denotes

the price-posting revenue curve of the single-dimensional linear utility agent.

For $\hat{q} > 0$ (strictly positive), the price $V(\hat{q})$ is strictly less than the value of the highest type $V(0) = h$, so $p^{\hat{q}}(h) = V(\hat{q})$. Thus, the Lagrangian objective for $\hat{q} \in (0, 1]$ is $P_{\lambda}(\hat{q}) = P(\hat{q}) - \lambda V(\hat{q})$. For $\hat{q} = 0$ the highest type is indifferent between buying and not buying. This indifference does not matter as highest type (quantile $q = 0$) is realized with measure zero (i.e., never) and so this type cannot affect the optimization. It will be technically convenient and consistent with the highest type “not mattering” to assume, with respect to the budget constraint, that indifference of the highest type to the price h is resolved in favor of rejecting the price. For posting price $V(\hat{q} = 0) = h$, the expected revenue is $\mathbf{E}[p^0(t)] = 0$ and, by this indifference-resolution assumption, the payment of the highest type is $p^0(h) = 0$; therefore, the expected Lagrangian objective from posting price $V(0)$ is identically zero. Having identified the expected revenue and payment of the highest type for every \hat{q} price posting we have identified the Lagrangian price-posting revenue curve; see Figure 8.13(a). Notice that this revenue curve is discontinuous at $\hat{q} = 0$ (unless $\lambda = 0$, i.e., when the budget constraint is not binding).

Proposition 8.23. *The Lagrangian price-posting revenue curve for an agent with public budget satisfies*

$$P_{\lambda}(\hat{q}) = \begin{cases} 0 & \text{if } \hat{q} = 0, \text{ and} \\ P(\hat{q}) - \lambda V(\hat{q}) & \text{otherwise.} \end{cases}$$

Notice that on $q \in (0, 1]$ this Lagrangian price-posting revenue curve is the difference between the original revenue curve and the scaled value function. If the original revenue curve is concave and the value function $V(q) = F^{-1}(1 - q)$ is convex (equivalently, the cumulative distribution function $F(\cdot)$ is convex; equivalently, the density function $f(\cdot)$ is monotone non-decreasing), then this Lagrangian price-posting revenue curve is concave (on $q \in (0, 1]$).

Definition 8.17. A single-dimensional public budget agent is *regular* if for all Lagrangian parameters $\lambda \geq 0$ the Lagrangian price-posting revenue curve is concave on interval $(0, 1]$. The value distribution F of such a regular public-budget agent is *public-budget regular*.

Proposition 8.24. *A single-dimensional public budget agent is regular*

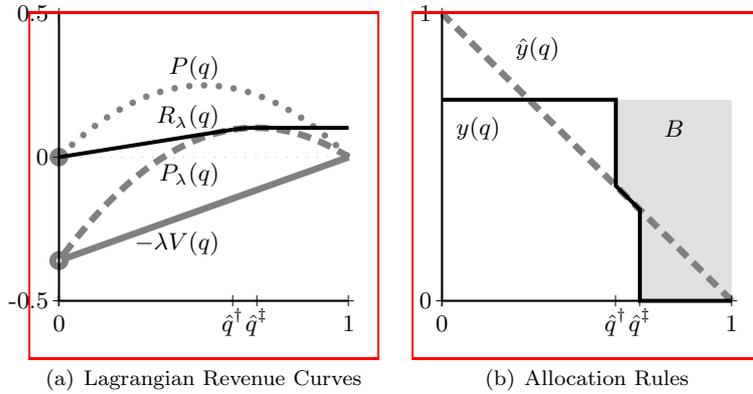


Figure 8.13. Depicted in (a) are the Lagrangian revenue curves corresponding to an agent with type distributed uniformly on $[0, 1]$ and Lagrangian parameter $\lambda > 0$. Note that the Lagrangian price-posting revenue curve $P_\lambda(\cdot)$ is discontinuous at $q = 0$ with $P_\lambda(0) = 0$ and $\lim_{q \rightarrow 0} P_\lambda(q) = -\lambda V(0)$. The Lagrangian price-posting revenue curve (thick, gray, dashed line) on $q \in (0, 1]$ is the sum of the single-dimensional linear price-posting revenue curve $P(q)$ (thick, gray, dotted line) and the relaxed budget constraint $-\lambda V(q)$ (thick, gray, solid line). Depicted in (b) are the allocation rules corresponding to optimization of the Lagrangian objective for two i.i.d. agents with budget $B = 1/4$. The allocation rule y (thin, black, solid line) is derived from the allocation constraint \hat{y} (thick, gray, dashed line) by averaging on $[0, \hat{q}^\dagger]$ and zeroing on $(\hat{q}^\dagger, 1]$. The Lagrangian parameter was selected to meet the budget constraint with equality for this example.

if (a) her type distribution F is regular (for single-dimensional linear agents) and (b) its cumulative distribution function $F(\cdot)$ is convex.

Due to the discontinuity at $q = 0$, public-budget regularity does not imply that the entire Lagrangian price-posting revenue curve is concave. Recall that when a price-posting revenue curves is not concave, as in the case of an irregular distribution with a single-dimensional linear agent, optimization subject to incentive compatibility (i.e., monotonicity of the allocation rule) is simplified by ironing (cf. Section 3.3.5 on page 75). With respect to the Lagrangian price-posting revenue curve, ironing is equivalent to taking the concave hull, i.e., the smallest concave upper bound. Geometrically, it is easy to see that this ironing replaces the revenue curve with a line segment from the origin to the point where it is tangent to the original curve. This point is uniquely identified by \hat{q}^\dagger satisfying $P_\lambda(\hat{q}^\dagger) = \hat{q}^\dagger P'_\lambda(\hat{q}^\dagger)$. The resulting revenue curve is continuous and concave (Proposition 8.25). The *Lagrangian revenue curve* denotes,

for every ex ante constraint \hat{q} , the optimal Lagrangian objective value from a mechanism with ex ante sale probability \hat{q} .

Proposition 8.25. *The Lagrangian revenue curve $R_\lambda(\cdot)$ for an agent with public budget and value drawn from a public-budget regular distribution satisfies*

$$R_\lambda(q) = \begin{cases} q P'_\lambda(\hat{q}^\dagger) & \text{if } q \in [0, \hat{q}^\dagger], \\ P_\lambda(q) & \text{if } q \in [\hat{q}^\dagger, \hat{q}^\ddagger], \text{ and} \\ P_\lambda(\hat{q}^\ddagger) & \text{if } q \in [\hat{q}^\ddagger, 1], \end{cases}$$

with \hat{q}^\dagger and \hat{q}^\ddagger set to satisfy $P_\lambda(\hat{q}^\dagger) = \hat{q}^\dagger P'_\lambda(\hat{q}^\dagger)$ and $P'_\lambda(\hat{q}^\ddagger) = 0$, respectively.

The revenue curves $P_\lambda(\cdot)$ and $R_\lambda(\cdot)$ that correspond to price posting and ex ante optimization can be extended to describe the Lagrangian objective, respectively, for any allocation rule $y(\cdot)$ and for optimization with respect to constraint $\hat{y}(\cdot)$. These extensions follow from reinterpreting an allocation rule or constraint as a distribution over ex ante constraints. For example, the a mechanism with allocation rule $y(\cdot)$ can be obtained by drawing a random quantile \hat{q} from distribution G^y with cumulative distribution function $G^y(z) = 1 - y(z)$ and offering the agent the price that corresponds to this quantile i.e., $V(\hat{q})$ (and then applying the revenue equivalence to convert this mechanism to its all-pay equivalent). Thus, the expected Lagrangian objective for allocation rule $y(\cdot)$ is $\mathbf{E}_{\hat{q} \sim G^y} [P_\lambda(\hat{q})] = P_\lambda(1) y(1) + \mathbf{E}_{q \sim U[0,1]} [P_\lambda(q) [-y'(q)]]$. Similarly, the optimal Lagrangian objective for allocation constraint \hat{y} is $\mathbf{Rev}[\hat{y}] = \mathbf{E}[R_\lambda(q) [-\hat{y}'(q)]]$.

The usual integration by parts approach, with the fact that the Lagrangian revenue curve satisfies $R_\lambda(0) = 0$, implies that the optimal Lagrangian objective can be rewritten in terms of the Lagrangian marginal revenue curve $R'_\lambda(\cdot)$ as $\mathbf{Rev}[\hat{y}] = \mathbf{E}[R'_\lambda(q) \hat{y}(q)]$. Monotonicity of this marginal revenue curve, the theory of Lagrangian relaxation, and Corollary 3.6 on page 65 gives the following theorem.

Theorem 8.26. *For a public-budget agent, the revenue-optimal mechanism is given by optimizing the surplus of Lagrangian marginal revenue with Lagrangian parameter $\lambda > 0$ when the budget constraint met with equality for the highest type, or with $\lambda = 0$ when the budget constraint is not binding.*

As in the linear utility case, the optimal mechanism for the Lagrangian

objective can be interpreted from the Lagrangian revenue curves. In particular, we get the optimal allocation rule y subject to constraint \hat{y} from ironing on the intervals where the Lagrangian price-posting revenue curve is ironed, and reserve pricing at its peak (cf. Section 3.4.5 on page 85). Thus, it is optimal to iron the high-valued types, i.e., quantiles $q \in [0, \hat{q}^\dagger)$ with \hat{q}^\dagger as described above, and reserve price the low-valued types, i.e., quantiles $q \in (\hat{q}^*, 1]$ with \hat{q}^* defined as the maximizer of $R_\lambda(\cdot)$. The remaining types, which correspond to quantiles $q \in [\hat{q}^*, \hat{q}^\dagger]$, are served greedily according to the allocation constraint $\hat{y}(\cdot)$. The resulting allocation rule $y(\cdot)$ can be interpreted as averaging $\hat{y}(\cdot)$ on $[0, \hat{q}^\dagger)$ and setting it to zero on $(\hat{q}^*, 1]$; see Figure 8.13(b).

Corollary 8.27. *For a regular public-budget agent and interim allocation constraint $\hat{y}(\cdot)$, the optimal single-agent mechanism allocates as by $\hat{y}(\cdot)$ except that types with quantiles in $[0, \hat{q}^\dagger)$ are ironed, and types with quantiles in $(\hat{q}^*, 1]$ are reserve priced.*

8.7 Unit-demand Preferences

In the sections below we will prove the optimality of the single-agent mechanisms described in Section 8.1 for agents with unit-demand preferences. More generally, we will show that for a large family of well-behaved and alternative-symmetric distributions, the optimal mechanism is given by a uniform reserve price (i.e., the same across all alternative) and sells the agent her favorite alternative.

This section gives a generalization to multi-dimensional preferences of the framework of virtual values (cf. Section 3.3.1). Recall that virtual values are an amortization of revenue in the sense that they can be evaluated pointwise, but equate to revenue in expectation. The pointwise optimization of virtual surplus, then, gives a revenue-optimal incentive compatible mechanism. The approach is to (a) cover type space with paths, (b) solve the problem restricted to the path, and then (c) find sufficient conditions on the distribution of over types that implies that the optimal mechanisms on the path are consistent. This approach will generally fail unless the right paths are identified.

We will use this approach to solve the single-agent problems corresponding to an unit-demand agent with uniformly distributed types on the unit square (Example 8.2). For such an agent the optimal mechanism projects the multi-dimensional agent type onto a single dimension

that corresponds to the agent's value for her favorite alternative. In addition to proving this result, we will give sufficient conditions on the distribution, beyond uniform, under which this projection continues to be optimal.

This single-dimensional projection result gives insight on the role of second-degree price discrimination, i.e., whether a seller can make more money with a differentiated product. For example, a seller might introduce a high-quality and low-quality product to segment the market between high-valued consumers (to buy the high-quality product at a premium) and low-valued consumers (to buy the low-quality product at a discount). Intuitively, this approach can be profitable if high-valued consumers are more sensitive to quality than low-valued consumers. This section develops a proof of the inverse, that if high-valued consumers are less sensitive to quality than low-valued consumers, then there is no benefit to quality-based second-degree price discrimination. For example, movie tickets are predominantly sold with a uniform price. Such a mechanism is suggested by the results of this section under the assumption that film buffs tend not to have a higher willingness to pay than the general public.

We begin the section with a simple warmup exercise that single-agent problems for the two-alternative uniform unit-demand agent of Example 8.2. The approach is to solve the mechanism design problem independently on rays from the origin and relies solely on the single-dimensional theory of mechanism design from Chapter 3. To solve more complex multi-dimensional problems we generalize the characterization of incentive compatible mechanisms to multi-dimensional agents. We then solve the multi-dimensional mechanism design problem for more general families of paths. For the right choice of paths the approach of the warmup can be generalized. To identify the right paths, we develop a multi-dimensional framework of virtual values.

8.7.1 Warmup: The Uniform Distribution

As a warmup, consider selling one of two alternatives to a unit-demand agent with type drawn from the uniform distribution over type space $\mathcal{T} = [0, 1]^2$ (Example 8.2). We claimed without proof in Section 8.1.2 that the optimal (unconstrained) single-agent mechanism is to post a uniform price of $\sqrt{1/3}$. A simple argument for this result is as follows.

First, restrict the problem to the alternative-1 preferred subspace of types, i.e., where $\{t\}_1 > \{t\}_2$ (the solution for the other part will be

symmetric). A uniform pricing always sells the agent her favorite alternative, so with this restriction, the uniform pricing sells alternative 1 only. The conditional distribution on $\{t\}_1$ is the distribution of the maximum of two i.i.d. uniform random variables and has cumulative distribution function $F_{\max}(z) = \Pr[\{t\}_1 \leq z \wedge \{t\}_2 \leq z] = z^2$, density function $f_{\max}(z) = 2z$, single-dimensional virtual value $\phi_{\max}(z) = z - 1 - z^2/2z$, and monopoly price $\hat{v}_{\max}^* = \sqrt{1/3}$.

Now consider restricting the type space to paths that coincide with rays from the origin. Parameterizing such a path by its slope θ , a type on this path can be expressed in terms of $\{t\}_1$ as $t = (\{t\}_1, \theta \{t\}_1)$. Notice that all types $t \in \mathcal{T}_\theta = \{(v, \theta v) : v \in [0, 1]\}$ have the same value for receiving alternative 1 with probability θ or alternative 2 with certainty. Thus, restricting the type space to the path \mathcal{T}_θ , the problem of selling the agent alternative 1 or 2 is equivalent to that of selling the agent alternative 1 with probability one or alternative 1 with probability θ . Recall from Section 3.3 that the optimal single-dimensional mechanism, which is allowed to probabilistically allocate, is always deterministic. It sells to the agent with probability one if she has a non-negative virtual value and with probability zero otherwise. In other words, it posts the monopoly price. Thus, the optimal mechanism for \mathcal{T}_θ posts a price for alternative 1.

This restriction on type space to a path can be equivalently viewed as giving the mechanism designer extra knowledge, specifically, the knowledge of θ . With this extra knowledge, the conditional distribution on $\{t\}_1$ is F_{\max} , thus, the designer with this knowledge would post a price of $\hat{v}_{\max}^* = \sqrt{1/3}$ for alternative 1 (and not sell alternative 2). This solution is independent of θ , and the designer can do as well without knowledge of θ as with it. Thus, there is no loss with respect to the optimal mechanism from relaxing the incentive constraints between types that are not on the same path. The optimal mechanism for a unit-demand agent with types uniformly drawn from the full type space $\mathcal{T} = [0, 1]^2$ is the uniform price of $\sqrt{1/3}$.

In the remainder of the section this approach is generalized to a richer family of distributions. In particular, the same uniform-pricing result holds for any distribution where the conditional distributions of $\theta = \{t\}_2/\{t\}_1$ with respect to $\{t\}_1$ is ordered according to $\{t\}_1$ by first-order stochastic dominance. In other words, $\Pr[\{t\}_2/\{t\}_1 \leq \theta \mid \{t\}_1]$, for all fixed θ , is monotone in $\{t\}_1$.

8.7.2 Multi-dimensional Characterization of Incentive Compatibility

Chapter 2 characterized incentive compatible mechanisms for single-dimensional linear agents in Theorem 2.2 and Corollary 2.12. These results concluded that a mechanism with allocation and payment rules (x, p) is incentive compatible if and only if

- the allocation rule $x(\cdot)$ is monotone non-decreasing, and
- the payment rule satisfies the payment identity:

$$p(v) = v x(v) - \int_0^v x(z) dz.$$

Recall that the first term in the payment identity is the surplus and the second term is, thus, the agent's utility. We can reinterpret this characterization in terms of utility as follows. The utility function $u(\cdot)$ corresponds to an incentive compatible mechanism with allocation rule x if and only if

- it is convex, and
- related to the allocation rule by the *utility derivative identity*:

$$x(v) = \frac{d}{dv} u(v).$$

Moreover, under our usual interpretation of the allocation rule $x(v)$ as denoting the probability that the agent with value v is served, the utility derivative identity combined with $x(v) \in [0, 1]$ imply that the utility function is non-decreasing and has derivative at most one.

The multi-dimensional characterization of incentive compatibility generalizes this reinterpretation.

Theorem 8.28. *For an agent with linear utility, allocation rule and utility functions (x, u) correspond to an incentive compatible mechanism if and only if*

- (i) (convexity) $u(\cdot)$ is convex, and
- (ii) (utility gradient identity) $x(t) = \nabla u(t)$.⁸

⁸ Technically, the gradient ∇u of a convex function is only guaranteed to exist almost everywhere (and not everywhere). For types t where the gradient does not exist, the allocation $x(t)$ can be any *subgradient*, i.e., the gradient of any plane through point $(t, u(t))$ that lower bounds the utility function $u(\cdot)$; convexity of the utility function guarantees that such a plane exists.

Proof. Incentive compatibility is equivalent to the following inequality holding for all pairs of types (t, t^\dagger) :

$$u(t) \geq u(t^\dagger) + (t - t^\dagger) \cdot x(t^\dagger). \quad (8.11)$$

The right-hand side of equation (8.11) is the utility that t obtains for the outcome of t^\dagger . The only difference between the utility of t for an outcome and the utility of t^\dagger for an outcome is the surplus from the allocation. Thus, t 's utility for the outcome of t^\dagger is equal to the utility t^\dagger for this outcome plus the difference in surplus for t and t^\dagger for the outcome.

Like the proof of Theorem 2.2, this proof is broken into three parts.

- (i) The allocation rule and utility function (x, u) correspond to an incentive compatible mechanisms if convexity and the utility gradient identity hold.

Consider any type t^\dagger and the plane orthogonal to the surface of $u(\cdot)$ at t^\dagger . By convexity, this plane is lower-bounds the utility at any other type t . In other words,

$$u(t) \geq u(t^\dagger) + (t - t^\dagger) \cdot \nabla u(t^\dagger).$$

In this equation, the right-hand side is the point on this plane at t . By the gradient utility identity, we can substitute the $x(t^\dagger)$ for $\nabla u(t^\dagger)$ in the right-hand side to obtain the defining inequality (8.11) of incentive compatibility.

- (ii) The allocation rule and utility function (x, u) correspond to an incentive compatible mechanisms only if convexity holds.

Consider types t^\dagger, t^\ddagger , and convex combination $t = \gamma t^\dagger + (1 - \gamma) t^\ddagger$. Incentive compatibility requires equation (8.11) hold for all pairs of types. Restating the equation for type pairs $(t^\dagger, t), (t^\ddagger, t)$ we have:

$$\begin{aligned} u(t^\dagger) &\geq u(t) + (t^\dagger - t) \cdot x(t). \\ u(t^\ddagger) &\geq u(t) + (t^\ddagger - t) \cdot x(t). \end{aligned}$$

A convex combination of these equations gives:

$$\begin{aligned} \gamma u(t^\dagger) + (1 - \gamma) u(t^\ddagger) &\geq u(t) + (\gamma t^\dagger + (1 - \gamma) t^\ddagger - t) \cdot x(t) \\ &= u(t) \end{aligned} \quad (8.12)$$

The final equation above comes from the definition of t as the convex combination of t^\dagger and t^\ddagger . Inequality (8.12) implies convexity of utility as desired.

- (iii) The allocation rule and utility function (x, u) correspond to an incentive compatible mechanisms only if the utility gradient equality holds.

Let e_j be the unit vector corresponding to allocation of alternative j , i.e., $\{e_j\}_j = 1$ and $\{e_j\}_{j^\dagger} = 0$ for $j \neq j^\dagger$. For small constant ϵ , apply equation (8.11) to type pairs $(t + \epsilon e_j, t)$ and $(t - \epsilon e_j, t)$ to conclude:

$$u(t + \epsilon e_j) - u(t) \geq \epsilon \{x(t)\}_j, \text{ and}$$

$$u(t - \epsilon e_j) - u(t) \geq -\epsilon \{x(t)\}_j.$$

Combine these equations to obtain upper and lower bounds on $\{x(t)\}_j$ as:

$$1/\epsilon [u(t + \epsilon e_j) - u(t)] \geq \{x(t)\}_j \geq 1/\epsilon [u(t - \epsilon e_j) - u(t)]$$

Assuming the partial derivative of $u(\cdot)$ with respect to $\{t\}_j$ is defined at t , the limit as ϵ goes to zero is defined and both the upper and lower bound, above, are equal to the partial derivative of u with respect to $\{t\}_j$ at t which is the j th coordinate of the gradient $\{\nabla u(t)\}_j$. If the partial derivative is not defined, then the same limit argument implies that $x(t)$ is a subgradient of the utility function at type t . \square

The subsequent developments of this section will rely heavily on Theorem 8.28.

8.7.3 Optimal Mechanisms for Paths

A seller who segments the market by offering a differentiated product line is engaging in what is called *second-degree price discrimination*. One way to offer a differentiated product is to offer lotteries for the same product. For example, a seller could offer (1) the good at a high price or (2) the same good with probability $1/2$ (and nothing otherwise) at a low price. Our analysis of single-dimensional agents of Chapter 3 has shown pricing these lotteries is never beneficial.

In the example above, a buyer who has a value v for the allocation of (1) and will have value $v/2$ for the allocation of (2). Viewing these allocations as two alternatives, the buyer's type can be mapped into the two dimensional space corresponding to her value for each alternative. The buyer's type space is degenerate and lies on the line with slope $1/2$. Of course, the optimal mechanism when the buyer's value is drawn from a distribution is to post the monopoly price for the distribution for alternative 1 (and never sells alternative 2). In this section,

this monopoly-pricing result is generalized to type spaces given by more general families of paths.

Definition 8.18. A *path-based agent* is specified by a path $C : [0, 1] \rightarrow [0, 1]$ with $C(v) \leq v$ (thus, $C(0) = 0$) and distribution F_{\max} where the agent's type is given by $t^v = (v, C(v))$ with v drawn from F_{\max} . The type space is $\mathcal{T} = \{(v, C(v)) : v \in [0, 1]\}$.

Now we follow the same approach as Section 3.3 and convert the problem of optimizing revenue in expectation to the problem of optimizing a virtual surplus pointwise. The approach is the following. Expected profit is equal to expected surplus minus expected utility. We will use integration by parts on the path to write the expected utility as the integral of the gradient of the utility. By Theorem 8.28, the gradient of utility is equal to the allocation. The two terms for expected surplus and utility can then be combined to give a virtual surplus.

Lemma 8.29. For a path-based agent with path C and distribution F_{\max} , and any incentive compatible mechanism with allocation rule x , the agent's expected utility is

$$\mathbf{E}[u(t^v)] = u(t^0) + \mathbf{E}[x(t^v) \cdot (1, C'(v)) (1 - F_{\max}(v))^{1/f_{\max}(v)}].$$

Proof. An explanation of the following calculus is given below.

$$\begin{aligned} \mathbf{E}[u(t^v)] &= \int_0^1 u(t^v) f_{\max}(v) \, dv \\ &= - \int_0^1 u(t^v) \frac{d}{dv} [1 - F_{\max}(v)] \, dv \\ &= - \left[u(t^v) [1 - F_{\max}(v)] \right]_0^1 + \int_0^1 \frac{d}{dv} [u(t^v)] [1 - F_{\max}(v)] \, dv \\ &= u(t^0) + \int_0^1 \nabla u(t^v) \cdot (1, \frac{d}{dv} C(v)) [1 - F_{\max}(v)] \, dv \\ &= u(t^0) + \mathbf{E}[x(t^v) \cdot (1, C'(v)) (1 - F_{\max}(v))^{1/f_{\max}(v)}]. \end{aligned}$$

The first line is from the definition of expectation. The second line is from the definition of the density function f_{\max} as the derivative of the cumulative distribution function F_{\max} . The third line is by integration by parts. The first part of the third line simplifies by substituting $1 - F_{\max}(0) = 1$ and $1 - F_{\max}(1) = 0$; the second part of the third line simplifies by taking the derivative of the utility; and the fourth line results. The final line is from the definition of expectation. \square

Theorem 8.30. For a path-based agent with path C and distribution F_{\max} , and any incentive compatible mechanism (x, p) , the expected revenue is

$$\mathbf{E}[p(t^v)] = p(t^0) + \mathbf{E}[x(t^v) \cdot [t^v - (1, C'(v)) [1 - F_{\max}(v)]^{1/f_{\max}(v)}]] .$$

Proof. Expected revenue is equal to expected surplus minus expected utility, i.e.

$$\mathbf{E}[p(t^v)] = \mathbf{E}[x(t^v) \cdot t^v] - \mathbf{E}[u(t^v)] .$$

Lemma 8.29 allows the expected utility to be rewritten in terms of the allocation rule and the utility of type $t^0 = (0, 0)$. The agent with type t^0 has no surplus, so her only utility is from the negation of her payment, i.e., $p(t^0) = -u(t^0)$. Combining these two equations, we have the theorem. \square

Theorem 8.30 shows that the vector field $\phi(t^v) = t^v - (1, C'(v)) [1 - F_{\max}(v)]^{1/f_{\max}(v)}$ gives a virtual value function for revenue. The remaining question for revenue maximization is to choose the allocation rule to optimize virtual surplus with respect to ϕ subject to incentive compatibility. As in Section 3.3, we relax the incentive compatibility constraint, and choose allocation x to optimize the virtual surplus $\phi(t^v) \cdot x$ pointwise for each $t^v \in \mathcal{T}$. We then check for conditions on the environment, in this case the path C and distribution F_{\max} , that imply that the resulting allocation rule is incentive compatible.

Our goal in this section is not to identify the optimal incentive compatible mechanism. Instead we are looking for sufficient conditions on the distribution to imply that the optimal mechanism posts a price for alternative 1 only. Notice that the first coordinate of the virtual value function is exactly the single-dimensional virtual value corresponding to distribution F_{\max} , i.e., $\{\phi(t^v)\}_1 = v - \frac{1 - F_{\max}(v)}{f_{\max}(v)}$. If the distribution is regular (i.e., this function is monotone non-decreasing; Definition 3.4) then selling alternative 1 to maximize virtual surplus will post the monopoly price \hat{v}_{\max}^* that solves $\hat{v}_{\max}^* - \frac{1 - F_{\max}(\hat{v}_{\max}^*)}{f_{\max}(\hat{v}_{\max}^*)} = 0$. Pointwise optimization of $\phi(t^v) \cdot x$ serves the agent the alternative with the highest positive virtual value. Thus, for virtual surplus maximization to be equivalent to posting a price for alternative 1 only, it better be that when $\{\phi\}_1 > 0$ that $\{\phi\}_1 \geq \{\phi\}_2$ and when $\{\phi\}_1 \leq 0$ that $\{\phi\}_2 \leq 0$.

Definition 8.19. A path $C(\cdot)$ is *ratio monotone* if the ratio $C(v)/v$ is monotone non-decreasing in v ; see Figure 8.14.

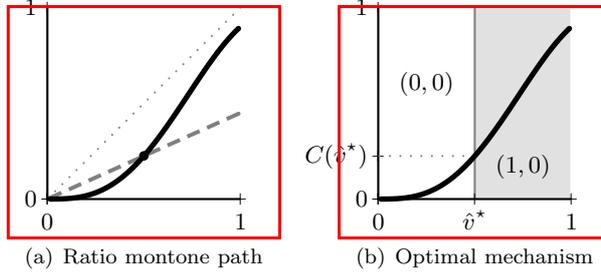


Figure 8.14. In subfigure (a), a ratio-monotone path (solid, thick, black line) is depicted. Ratio monotonicity of the path implies that the slope of the path at a type t is greater than that of the line through the type and the origin (depicted for type $t^{1/2}$ ($1/2, C(1/2)$); gray, dashed line). Subfigure (b) depicts the optimal mechanism, when value v is drawn uniformly from $[0, 1]$. This mechanism post price $\hat{v}^* = 1/2$ for alternative 1.

Theorem 8.31. *For a path-based agent with ratio-monotone path C and regular distribution F_{\max} , the optimal mechanism is to post a price for alternative 1.*

Proof. For the allocation that optimizes virtual surplus $\phi(t^v) \cdot x$ pointwise to never sell alternative 2, it better be that when $\{\phi\}_1 > 0$ that $\{\phi\}_1 \geq \{\phi\}_2$ and when $\{\phi\}_1 \leq 0$ that $\{\phi\}_2 \leq 0$. A sufficient condition is, for all v ,

$$\frac{C(v)}{v} \{\phi(t^v)\}_1 \geq \{\phi(t^v)\}_2. \tag{8.13}$$

Ratio-monotonicity is equivalent to the property that rays from the origin only cross the path from above to below, i.e., at the point of intersection, the slope of the ray is at most the slope of the path, i.e.,

$$C'(v) \geq C(v)/v.$$

The sufficient condition of equation (8.13) can be derived from ratio monotonicity as follows,

$$\begin{aligned} \frac{C(v)}{v} \{\phi(t^v)\}_1 &= \frac{C(v)}{v} \left[v - \frac{1-F_{\max}(v)}{f_{\max}(v)} \right] \\ &\geq C(v) - C'(v) \frac{1-F_{\max}(v)}{f_{\max}(v)} \\ &= \{\phi(t^v)\}_2. \end{aligned}$$

Thus, pointwise virtual surplus maximization never sells alternative 2. For virtual surplus maximization to additionally correspond to posting a price alternative 1, regularity of the distribution F_{\max} , as assumed in the statement of the theorem, is sufficient. \square

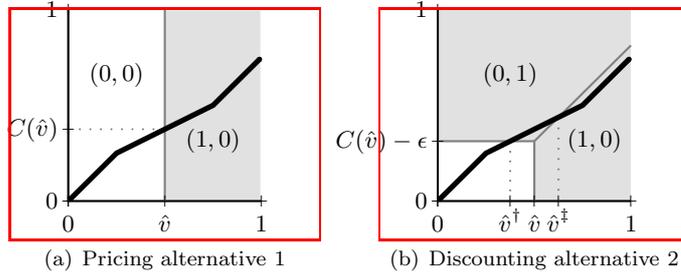


Figure 8.15. Depicted is a path (thick, solid, black line) that is not ratio monotone at \hat{v} . In subfigure (a), the allocation from posting price \hat{v} for alternative 1 is depicted by the shaded regions. In subfigure (b), the allocations from posting prices \hat{v} for alternative 1 and $C(\hat{v}) - \epsilon$ for alternative 2 are depicted by the shaded region. On the path, types with value $\{t\}_1$ for alternative 1 in interval $[\hat{v}^\dagger, \hat{v}^\ddagger]$ will buy alternative 2. Relative to simply posting a price for alternative 1, offering alternative 2 at a discount adds revenue from types with $\{t\}_1 \in [\hat{v}^\dagger, \hat{v}]$ and loses revenue from types with $\{t\}_1 \in [\hat{v}, \hat{v}^\ddagger]$. With constant density f_{\max} , types that add revenue have measure $(\hat{v} - \hat{v}^\dagger) f_{\max} = \epsilon f_{\max}/C'(\hat{v})$; types that lose revenue have measure $(\hat{v}^\ddagger - \hat{v}) f_{\max} = \epsilon f_{\max}/1 - C'(\hat{v})$.

One way to view the result above is that, fixing a ratio-monotone path C , regularity of the distribution F_{\max} of value for alternative 1 implies that the optimal mechanism is to post a price for alternative 1. We now show that regularity implies that price posting is optimal only if the path is ratio monotone.

Theorem 8.32. *For any non-ratio-monotone path C there exists a regular distribution F_{\max} such that posting a price for alternative 1 is not optimal for the path-based agent defined by C and F_{\max} .*

Proof. This proof is by counter example. Suppose that the path C is not monotone at some value $\hat{v} \in (0, 1)$, i.e.,

$$C'(\hat{v}) < C(\hat{v})/\hat{v}. \quad (8.14)$$

Consider the uniform distribution on $[0, 2\hat{v}]$ truncated at 1 with a point-mass. By construction the monopoly price is \hat{v} (thus, \hat{v} is the optimal price to post for alternative 1) and the density function is constant at $f_{\max} = 1/2\hat{v}$.

For the remainder of the proof, assume that $C(\cdot)$ is locally linear at \hat{v} (the general proof is deferred to Exercise 8.13). Consider adding the option of buying alternative 2 at price $C(\hat{v}) - \epsilon$. There is a gain from types who were not buying before who now buy and a loss from types

who were buying alternative 1 before but now switch to the lower cost alternative 2; see Figure 8.15. These are:

$$\begin{aligned}\text{Gain}(\epsilon) &= (C(\hat{v}) - \epsilon) \frac{\epsilon f_{\max}}{C'(\hat{v})}, \\ \text{Loss}(\epsilon) &= (\hat{v} - C(\hat{v}) + \epsilon) \frac{\epsilon f_{\max}}{1 - C'(\hat{v})}.\end{aligned}$$

The first term in each expression above is the gain or loss from each type; the second term is the measure of such types.

To see that the gain is more than the loss in the limit as ϵ goes to zero, we can divide each by ϵf_{\max} and take their limits.

$$\begin{aligned}\lim_{\epsilon \rightarrow 0} [\text{Gain}(\epsilon) / \epsilon f_{\max}] &= C(\hat{v}) / C'(\hat{v}) > \hat{v}, \\ \lim_{\epsilon \rightarrow 0} [\text{Loss}(\epsilon) / \epsilon f_{\max}] &= \hat{v} - C(\hat{v}) / 1 - C'(\hat{v}) < \hat{v}.\end{aligned}$$

The inequalities of both lines follow from equation (8.14), by rearranging as $C(\hat{v}) / C'(\hat{v}) > \hat{v}$ for the first line and because it implies $\hat{v} - C(\hat{v}) < \hat{v} [1 - C'(\hat{v})]$ for the second line. Thus, the gain is strictly more than the loss and posting price \hat{v} for alternative 1, which is optimal among such price postings, is not optimal among all mechanisms. \square

8.7.4 Uniform Pricing for Ratio-monotone Distributions

This section duplicates the analysis from the warmup (Section 8.7.1) for more general distributions F on the alternative-1 preferred type space, i.e., $\mathcal{T} = \{t \in [0, 1]^2 : \{t\}_2 \leq \{t\}_1\}$. In the previous analysis the mechanism design problem was decomposed into a collection of paths, solved on each path in the collection, and then it was argued that these solutions are consistent with a single mechanism. Critical in this analysis is the choice of the paths. In The results of this section are based on a natural guess at the right paths; a principled method for determining the right paths will be described in subsequent sections.

The following discussion gives a natural guess as to the right paths on which to decompose the multi-dimensional mechanism design problem. We would like to solve the mechanism design problem independently on these paths and for the optimal mechanism on each path to post the same price for alternative 1 (and never sell alternative 2). One sufficient condition to guarantee that the optimal mechanisms for selling alternative 1 on each path posts the same price is to require the distribution of the agent's value for alternative 1, conditioned on the path on which the agent's type lies, be the same for all paths.

Definition 8.20. The *equiquantile path* C_θ solves

$$\Pr_{t \sim F}[\{t\}_2 \leq C_\theta(v) \mid \{t\}_1 = v] = \theta.$$

The *equiquantile type subspace* is $\mathcal{T}_\theta = \{(v, C_\theta(v)) : v \in [0, 1]\}$.

Lemma 8.33. For any quantile θ , the conditional distribution of $\{t\}_1$ given $t \in \mathcal{T}_\theta$ for $t \sim F$ is equal to the unconditional distribution of $\{t\}_1$, i.e., for all $z \in [0, 1]$,

$$\Pr_{t \sim F}[\{t\}_1 \leq z \mid t \in \mathcal{T}_\theta] = \Pr_{t \sim F}[\{t\}_1 \leq z].$$

Proof. By definition, given $\{t\}_1$ the quantile θ corresponding to the path type space \mathcal{T}_θ that contains type t is uniformly distributed. Therefore, θ is independent of $\{t\}_1$; equivalently, $\{t\}_1$ is independent of θ . Thus, the conditional distribution of $\{t\}_1$ given θ is the same as its unconditional distribution. \square

We are now ready to complete the construction. The set of equiquantile paths $\{C_\theta : \theta \in [0, 1]\}$ partition type space. Supposing the designer knew the path C_θ on which the type was drawn, then the designer would employ the optimal mechanism for that path. If the distribution of $\{t\}_1$ is regular and the equiquantile curves are ratio monotone, then by Theorem 8.31 the optimal mechanism for each path is to post the monopoly price \hat{v}_{\max}^* for the distribution F_{\max} of $\{t\}_1$. This mechanism is the same regardless of the path; thus, posting price \hat{v}_{\max}^* for alternative 1 is revenue optimal. From this argument, we can conclude the following theorem.

Theorem 8.34. For distribution F on the alternative-1 preferred type space $\mathcal{T} = \{t \in [0, 1]^2 : \{t\}_2 \leq \{t\}_1\}$ satisfying (a) the distribution of $\{t\}_1$ is regular and (b) the equiquantile paths are ratio monotone, the revenue optimal mechanism is to post the monopoly price for alternative 1.

This theorem can easily be extended to distributions on the unit square for which the distribution of the agent's value for her preferred alternative is independent of which alternative is preferred. Under this assumption, the problem can be independently solved under each conditioning, and the mechanism that posts a uniform price for each alternative is optimal.

This theorem can also be easily generalized to selling a single item in one of two configurations (which we will continue to refer to as alternatives) to several agents. Notice that if there was a cost c for serving

the agent, then the optimal mechanism of Theorem 8.34 would still be a posted price for alternative 1. The price \hat{v} however would be increased to solve $\hat{v} - \frac{1 - F_{\max}(\hat{v})}{f_{\max}(\hat{v})} = c$. We conclude that the optimal mechanism for several multi-dimensional agents (that each satisfy the assumptions of the theorem) is simply the optimal mechanism that projects each agent into a single dimension according her value for her favorite alternative.

The next section will formalize the method of virtual values employed above to prove the optimality of posting the monopoly price for alternative 1. The section following will use this formulation to give a general method for solving for the appropriate paths on which to solve the mechanism design problem.

8.7.5 Multi-dimensional Virtual Values

In this section we generalize the virtual-value-based approach to mechanism design from the single-dimensional agents of Section 3.3.2 on page 64 to multi-dimensional agents. Assume that the seller has a cost function for producing a given allocation of alternatives x that is specified by $c(x)$. The definitions below are given for a single agent and the objective of profit, i.e., expected payment minus expected cost, but they could be equally well defined for multiple agents and any objective.

Definition 8.21. A *virtual value function* ϕ is a vector field that satisfies three properties:

- (i) *Amortization of revenue:* For any incentive compatible mechanism (x, p) , the agent's expected virtual surplus is an upper bound on expected revenue, i.e., $\mathbf{E}[\phi(t) \cdot x(t)] \geq \mathbf{E}[p(t)]$.
- (ii) *Incentive compatibility:* A point-wise virtual surplus maximizer $x^*(t) \in \operatorname{argmax}_x \phi(t) \cdot x - c(x)$ is incentive compatible, i.e., there exists a payment rule p^* such that mechanism (x^*, p^*) is incentive compatible.
- (iii) *Tightness:* For this point-wise virtual surplus maximizer (x^*, p^*) , the agent's expected virtual surplus is equal to the expected revenue, i.e., $\mathbf{E}[\phi(t) \cdot x^*(t)] = \mathbf{E}[p^*(t)]$.

Definition 8.21 makes a distinction between the agent's virtual surplus $\phi(t) \cdot x$ and the virtual surplus of the mechanism $\phi(t) \cdot x - c(x)$ which includes the seller's cost. A special case of interest is a *uniform cost* c where virtual surplus is $\phi(t) \cdot x - c \sum_j \{x\}_j$. This uniform cost could represent the *opportunity cost* the seller faces for serving this agent. For example, with two agents with virtual value functions and a single-item

environment, the opportunity cost of serving one agent is the maximum of the virtual value of the other agent and zero.

Proposition 8.35. *For any mechanism design problem that admits a virtual value function, a virtual surplus maximizer is the optimal mechanism.*

Proof. Denote the incentive compatible virtual surplus maximizer (guaranteed to exist by Definition 8.21) by allocation and payment rules (x^*, p^*) ; denote any other incentive compatible mechanism by allocation and payment rules (x, p) ; then,

$$\begin{aligned} \mathbf{E}_t[p^*(t) - c(x^*(t))] &= \mathbf{E}[\phi(t) \cdot x^*(t) - c(x^*(t))] \\ &\geq \mathbf{E}[\phi(t) \cdot x(t) - c(x(t))] \geq \mathbf{E}[p(t) - c(x^*(t))]. \end{aligned}$$

The first equality is by tightness (the expected cost term $\mathbf{E}[c(x^*(t))]$ is the same on both sides of the equality), the second inequality is by the fact that x^* is a virtual surplus maximizer, the third inequality is because ϕ is an amortization of revenue (again, the expected cost term $\mathbf{E}[c(x(t))]$ is the same on both sides of the inequality). \square

In Section 8.7.3, we derived a virtual value function for types on a ratio-monotone path and with regularly distributed value for the preferred alternative. In Section 8.7.4, we guessed a set of paths, solved the mechanism design problem on each path, and argued that under some distributional assumptions these optimal mechanisms are consistent with one mechanism. In this section we will develop a general framework for deriving virtual value functions absent a good guess of the decomposition to paths. The main idea is to leave the paths as variables, that can then be solved for later. To do this we will employ a multi-dimensional integration by parts (which is defined with respect to any vector field; see the Mathematical Note on page 315), with the constraint that this vector field corresponds to paths.

Our goal is to use integration by parts to rewrite the expected utility $\mathbf{E}[u(t)] = \int_{t \in \mathcal{T}} u(t) f(t) dt$ in terms of the gradient of the utility ∇u , which by Theorem 8.28 is equal to the allocation x , and a boundary integral. Thus, we need to identify a vector field α with divergence equal to the (negated) density. The term corresponding to the boundary integral, we would prefer to be zero, but for an upper bound it would be sufficient for it to be negative. Thus, we seek a vector field α to satisfy properties of the following definition.

Definition 8.22. For distribution F and type space \mathcal{T} , a vector field $\alpha : \mathbb{R}^m \rightarrow \mathbb{R}^m$ satisfies

- the *divergence density equality* if $\nabla \cdot \alpha(t) = -f(t)$ for all $t \in \mathcal{T}$, and
- *boundary influx* if $\alpha(t) \cdot \eta(t) \leq 0$ for all $t \in \partial\mathcal{T}$.

Theorem 8.36. If $\alpha : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a vector field satisfying the divergence density equality and boundary influx on type space \mathcal{T} then vector field $\phi(\mathcal{T}) = t - \alpha(t)/f(t)$ is an amortization of revenue. Moreover, the amortization ϕ is tight for incentive compatible mechanisms that have binding individual rationality constraint $u(t) = 0$ on all boundary types t with non-trivial flux, i.e., $\alpha(t) \cdot \eta(t) \neq 0$.

Mathematical Note. Multi-dimensional integration by parts is defined for function $u : \mathbb{R}^m \rightarrow \mathbb{R}$ and vector field $\alpha : \mathbb{R}^m \rightarrow \mathbb{R}^m$ on region \mathcal{T} with boundary $\partial\mathcal{T}$ as follows:

$$\int_{t \in \mathcal{T}} \nabla u(t) \cdot \alpha(t) dt = \int_{t \in \partial\mathcal{T}} u(t) (\alpha(t) \cdot \eta(t)) dt - \int_{t \in \mathcal{T}} u(t) (\nabla \cdot \alpha(t)) dt.$$

In the formula above, $\nabla \cdot \alpha(t)$ is the divergence of α at point $t \in \mathcal{T}$ and is defined as $\nabla \cdot \alpha(t) = \sum_j \{\nabla \alpha(t)\}_j$; and $\eta(t)$ is a unit-length normal vector to the boundary at point $t \in \partial\mathcal{T}$.

The *divergence theorem* is the application of multi-dimensional integration by parts to the vector field α and the function $u(\cdot) = 1$ (which has trivial gradient $\nabla u(\cdot) = (0, \dots, 0)$). Viewing the vector field as a flow, the divergence theorem shows that the divergence of a flow in a region \mathcal{T} is equal to magnitude of the flux out of the region.

$$\int_{t \in \mathcal{T}} \nabla \cdot \alpha(t) dt = \int_{t \in \partial\mathcal{T}} \alpha(t) \cdot \eta(t) dt.$$

Proof. Rewrite expected utility as

$$\begin{aligned}
\mathbf{E}[u(t)] &= \int_{t \in \mathcal{T}} u(t) f(t) dt \\
&= - \int_{t \in \mathcal{T}} u(t) (\nabla \cdot \alpha(t)) dt \\
&= - \int_{t \in \partial \mathcal{T}} u(t) (\alpha(t) \cdot \eta(t)) dt + \int_{t \in \mathcal{T}} \nabla u(t) \cdot \alpha(t) dt \\
&\geq \int_{t \in \mathcal{T}} \nabla u(t) \cdot \alpha(t) dt \\
&= \int_{t \in \mathcal{T}} x(t) \cdot \alpha(t) dt \\
&= \mathbf{E}[x(t) \cdot \alpha(t)/f(t)].
\end{aligned}$$

The first line is the definition of expectation, the second line applies the divergence density equality, the third line is integration by parts, the fourth line follows from individual rationality, i.e., $u(t) \geq 0$ for all $t \in \mathcal{T}$, and boundary influx (implying that the first term on the third line is non-negative), the fifth line is from Theorem 8.28, and the sixth line is the definition of expectation.

A type t with binding participation constraint has zero utility, i.e., $u(t) = 0$. If all boundary types with non-trivial boundary influx, i.e., with $\alpha(t) \cdot \eta(t) \neq 0$, have binding participation constraint $u(t) = 0$ then $u(t) (\alpha(t) \cdot \eta(t)) = 0$ at all boundary types $t \in \partial \mathcal{T}$. In this case, the first term on the third line is identically zero and the whole sequence of inequalities is tight.

Expected revenue is equal to the expected surplus less the agent's expected utility, i.e.,

$$\begin{aligned}
\mathbf{E}[p(t)] &= \mathbf{E}[t \cdot x(t)] - \mathbf{E}[u(t)] \\
&\leq \mathbf{E}[t \cdot x(t)] - \mathbf{E}[\alpha(t)/f(t) \cdot x(t)] \\
&= \mathbf{E}[\phi(t) \cdot x(t)].
\end{aligned}$$

The second line is from the previous derivation, and the third line is from the definition of vector field ϕ . This sequence of inequalities is tight when the previous sequence of inequalities is tight. \square

While Theorem 8.36 does not explicitly mention paths, paths are implicit in the choice of vector fields α that give tight amortizations. With two-dimensional type space oriented as for considering the optimality of pricing only alternative 1 with weaker types towards the left and

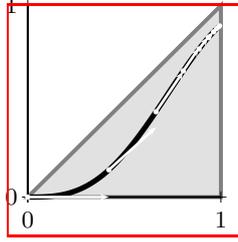


Figure 8.16. Depicted is the type space $\mathcal{T} = \{t \in [0, 1]^2 : \{t\}_2 \leq \{t\}_1\}$ (light gray region), its boundary $\partial\mathcal{T}$ (dark gray boarder), and a path (thick, black line) through type space. The vector field α (white arrows) that corresponds to this path is shown. The lengths of the arrows are proportional to the magnitude of vectors in the field (the first coordinate of which is the remaining probability density f left to distribute on the path). All such paths originate at type $(0, 0)$ (the left boundary) and terminate on the right boundary, i.e., with types $\{t \in \mathcal{T} : \{t\}_1 = 1\}$.

stronger types towards the right, the boundary will have four regions; see Figure 8.16. Paths will originate on the left boundary with an influx of flow. On this path the direction of vector field α is the direction of the path; the magnitude of the first coordinate $\{\alpha(t)\}_1$ of this flow is cumulative the density on the remainder of the path. The path terminates on the right boundary where vector field α is the zero vector (there is no remaining density according to f). The top and bottom boundary regions are parallel to paths and thus the dot-product of vector field α with the normal to the boundary is zero. By these interpretation, the influx on the boundary is trivial (equal to zero) on all types except those on the left where the paths originate. If these types t on the left are chosen so that the mechanisms under consideration have binding individual rationality constraint, i.e., $u(t) = 0$, then vector field ϕ constructed in Theorem 8.36 is tight, as desired.

Definition 8.23. An amortization of revenue ϕ is *canonical* if is derived as $\phi(t) = t - \alpha(t)/f(t)$ from vector field α that satisfies the divergence density equality and boundary influx.

Notice that a vector field α that satisfies the divergence density equality and boundary influx will have divergence -1 on type space \mathcal{T} . Consequentially, by the divergence theorem, the outflux on the boundary must also be -1 . For mechanism where this non-trivial outflux is concentrated on boundary types that have zero utility, then the amortization of revenue ϕ defined from α is tight.

Example 8.8. Consider a single-dimensional agent with type t uniformly distributed on type space $\mathcal{T} = [1, 2]$. The density function is $f(t) = 1$. There is only one path and it goes from type 1 to type 2. The (single-dimensional) vector field α at t is the remaining cumulative density on $[t, 2]$, i.e., $\alpha(t) = 2 - t$. Notice that there is boundary influx only at type $t = 1$. Any mechanism where type 1 has zero utility, e.g., posting a price at 1 or higher, will have a binding participation constraint for type 1. Thus, $u(1)\alpha(1) \cdot \eta(1) = 0$ where the (single-dimensional) normal vector at 1 is $\eta(1) = -1$. The resulting amortization of revenue is vector field $\phi(t) = t - \alpha(t)/f(t) = 2t - 2$. Notice that $\alpha(t)$ is equal to $1 - F(t)$ (for cumulative distribution function $F(t) = t - 1$), so this formula is identical to the single-dimensional virtual value derived in Chapter 3. Notice that for mechanisms that post prices less than one, say, at $1/2$, the amortization of revenue is not tight. The expected virtual surplus of this mechanism is one, while its revenue is $1/2$, the virtual surplus less the utility of the weakest type, i.e., type 1.

Unfortunately, except in single dimensional environments, canonical amortizations of revenue are not unique. Any covering of type space by paths will give a canonical amortization. Generally, at most one of these canonical amortizations can be a virtual value function. In the next section we will develop a systematic method for identifying a virtual value function, or equivalently, the right set of paths.

We conclude this section by observing that the existence of a virtual value function for the family of single-agent environments with uniform costs implies revenue linearity (Definition 3.16), i.e., $\mathbf{Rev}[y] = \mathbf{Rev}[y^\dagger] + \mathbf{Rev}[y^\ddagger]$ for $y = y^\dagger + y^\ddagger$. Essentially, virtual surplus is a linear objective. Thus, as described in Section 8.3, multi-agent service constrained mechanism design problems reduce to single-agent ex ante problems.

Theorem 8.37. *Consider a unit-demand agent (given by type space and distribution), if vector field ϕ is a virtual value function for the single-agent environment with any non-negative uniform cost c then the agent is revenue linear.*

Proof. Sort the types t in decreasing order of the virtual value of the alternative with the highest virtual value, i.e. $\max_j \{\phi(t)\}_j$. Let \hat{q}^* be the measure of types where this highest virtual value is non-negative. The \hat{q} ex ante optimal mechanism serves the first $\min(\hat{q}, \hat{q}^*)$ measure of types in this order. The \hat{q} interim optimal mechanism serves the first

\hat{q}^* measure of types greedily by this order (and discards the remaining types). By the linearity of virtual surplus, the expected virtual surplus of the latter is the appropriate convex combination of the expected virtual surplus of the former. Since expected virtual surplus equals expected revenue, the agent is revenue linear. \square

8.7.6 Reverse Solving for Virtual Values

In multi-dimensional environments, because there are multiple ways to cover type space by paths, there are multiple canonical amortizations of revenue. If we can find an amortization of revenue that is incentive compatible, i.e., for which its pointwise optimization gives an incentive compatible mechanism (Definition 8.21), then Proposition 8.35 implies that this mechanism is optimal. Except in edge cases, optimal mechanisms for a single agent are unique. Thus, we are searching among these canonical amortizations for the one, if any, that is incentive compatible. In Section 8.7.4, we guessed the right paths, in this section we give a principled approach for identifying them.

Consider an agent with the two-dimensional alternative-1 preferred type space, i.e., $\mathcal{T} = \{t \in [0, 1]^2 : \{t\}_1 \geq \{t\}_2\}$. The goal of this setting is to describe sufficient conditions on the distribution F (as specified by density function f) so that posting a price for alternative 1 (only) is revenue optimal. As in Section 8.7.4 the solution to this problem will generalize to the full type space $[0, 1]^2$; moreover, it will also generalize to $m \geq 2$ alternatives.

Definition 8.24. The *single-dimensional favorite-alternative projection* is given by value $v = \{t\}_1$, distribution function F_{\max} , density function f_{\max} , amortization of revenue $\phi_{\max}(v) = v - \frac{1 - F_{\max}(v)}{f_{\max}(v)}$, and monopoly price \hat{v}_{\max}^* that solves $\phi_{\max}(\hat{v}_{\max}^*) = 0$.

Proposition 8.38. For non-negative uniform costs c , a vector field ϕ is a virtual value that proves the the optimality of the favorite-alternative single-dimensional projection, i.e., the mechanism that projects the agent's type to her value for alternative-1 and is optimal for this projection, if (a) the alternative-1 virtual value $\{\phi(\cdot)\}_1$ is a virtual value for the single-dimensional projection, and (b) the alternative 2 virtual value never maximizes virtual surplus, i.e., $\{\phi(t)\}_1 \geq 0$ implies $\{\phi(t)\}_2 \leq \{\phi(t)\}_1$ and $\{\phi(t)\}_1 \leq 0$ implies $\{\phi(t)\}_2 \leq 0$ for all types t .

Proof. By property (b), virtual surplus maximization only serves al-

ternative 1 (or nothing if $\{\phi(t)\}_1 < c$);⁹ by property (a) and Proposition 8.35, virtual surplus maximization is optimal among all mechanisms that only sell alternative alternative 1. \square

The goal of this section is to identify a vector field ϕ that satisfies the conditions of Proposition 8.38. The approach is to use property (a) of the proposition to reduce a degree of freedom in defining a canonical amortization, and then to identify conditions on the distribution that are sufficient to imply property (b). Specifically, set the first coordinate of the virtual value function, denoted $\{\phi(t)\}_1$, to the virtual value of the single-dimensional projection, denoted $\phi_{\max}(\{t\}_1)$. The definition of the canonical amortization ϕ for vector field α (Definition 8.23) gives $\{\alpha(t)\}_1$ from $\{\phi(t)\}_1$; the divergence density equality gives $\{\alpha(t)\}_2$ from $\{\alpha(t)\}_1$ (and identifies the right paths); and Definition 8.23, again, gives $\{\phi(t)\}_2$ from $\{\alpha(t)\}_2$. With amortization $\phi(\cdot)$ fully defined, sufficient conditions on the distribution to imply property (b) can be identified.

Definition 8.25. The two-dimensional extension of the favorite-alternative projection defines vector fields ϕ and α as follows:

- (i) $\{\phi(t)\}_1 = \phi_{\max}(\{t\}_1)$,
- (ii) $\{\alpha(t)\}_1 = [\{t\}_1 - \{\phi(t)\}_1] f(t) = \frac{1 - F_{\max}(\{t\}_1)}{f_{\max}(\{t\}_1)} f(t)$,
- (iii) $\{\alpha(t)\}_2 = - \int_0^{\{t\}_2} [f(\{t\}_1, z) + d/d\{t\}_1 \{\alpha(\{t\}_1, z)\}_1] dz$.
- (iv) $\{\phi(t)\}_2 = \{t\}_2 - \{\alpha(t)\}_2 / f(t)$.

Lemma 8.39. The two-dimensional extension of the favorite-item projection defines vector field α that satisfies the divergence density equality and boundary inflow, and vector field ϕ is a canonical amortization that is tight for mechanisms for which individual rationality binds on type $(0, 0)$.

Proof. By Theorem 8.36, it suffices to show that vector field α satisfies the divergence density equality and trivial boundary influx at $t \in \partial\mathcal{T} \setminus \{(0, 0)\}$.

The divergence density equality is satisfied by definition; to see this, differentiate both sides of the definition of $\{\alpha(t)\}_2$ with respect to $\{t\}_2$.

⁹ Notice that the assumptions of the proposition are insufficient if the uniform service cost c is negative. With a negative service cost the agent may receive a non-trivial alternative even when her virtual value for this alternative is negative. This restriction limits the applicability of the proposition to single-agent problems where the ex ante constraint holds as an inequality, and to multi-agent service-constrained environments that are downward closed.

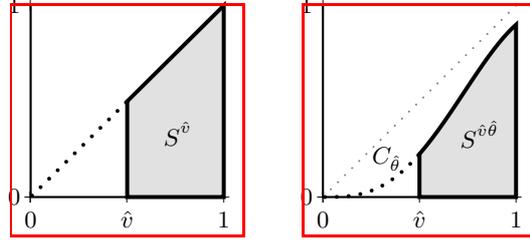


Figure 8.17. Depicted are the regions to which the divergence theory is applied in Lemma 8.39 and Lemma 8.40. The flux on the right and bottom boundaries are zero by definition. In subfigure (a), the divergence of the subspace of types $S^{\hat{v}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v}\}$ (gray region) and the outflux of the left boundary are both equal to $-(1 - F_{\max}(\hat{v}))$. In subfigure (b), the divergence of the subspace of types $S^{\hat{v}\hat{\theta}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v} \wedge \{t\}_2 \leq C_{\hat{\theta}}(\{t\}_1)\}$ (gray region) and the outflux of the left boundary are both equal to $-\hat{\theta}(1 - F_{\max}(\hat{v}))$. For both, the remaining flux out the top boundary must be zero.

We now analyze the flux $\alpha(t) \cdot \eta(t)$ for types t on the right, bottom, and top boundaries.

- Right boundary, i.e., t with $\{t\}_1 = 1$: Normal $\eta(t) = (1, 0)$ and $\{\alpha(t)\}_1 = 0$ as $1 - F_{\max}(1) = 0$, so $\alpha(t) \cdot \eta(t) = 0$.
- Bottom boundary, i.e., t with $\{t\}_2 = 0$: Normal $\eta(t) = (0, -1)$ and $\{\alpha(t)\}_2 = 0$ as the integral from 0 to 0 of any function is zero, so $\alpha(t) \cdot \eta(t) = 0$.
- Top boundary, i.e., t with $\{t\}_1 = \{t\}_2$: Apply the divergence theorem to vector field α on the type subspace with value at least \hat{v} for alternative 1, i.e., $S^{\hat{v}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v}\}$; see Figure 8.17. The divergence theorem requires that the divergence of α on this subspace is equal the outflux. By the divergence density equality, the divergence of subspace $S^{\hat{v}}$ is $-\int_{t \in S^{\hat{v}}} f(t) dt = -(1 - F_{\max}(\hat{v}))$. The outflux on the left-boundary is $-\int_0^1 \{\alpha(\hat{v}, z)\}_1 dz = -\int_0^1 \frac{1 - F_{\max}(\hat{v})}{f_{\max}(\hat{v})} f(\hat{v}, z) dz$. Note that the density of the favorite-alternative projection is $f_{\max}(\hat{v}) = \int_0^1 f(\hat{v}, z) dz$ and, thus, this outflux is $-(1 - F_{\max}(\hat{v}))$. The remaining total outflux for the top boundary is zero. This equality holds for all values \hat{v} , thus the outflux at each type t on the top boundary is identically zero.

Though unnecessary for the proof, observe that the above calculation of the outflux on the left-boundary applied to full type space \mathcal{T} (equal to

the subspace $S^{\hat{v}}$ when $\hat{v} = 0$) implies that the outflux at type $(0, 0)$ in the direction of $(-1, 0)$ is -1 . \square

The remaining task is to identify sufficient conditions on the distribution F so that uniform pricing optimizes virtual surplus with respect to the amortization defined by the two-dimensional extension of the favorite-alternative projection. Recall that the monopoly price \hat{v}_{\max}^* for the distribution F_{\max} is the optimal price to post for alternative 1 when the agent's value for the alternative is drawn from distribution F_{\max} . A sufficient condition on the virtual value function ϕ is that (a) for types t with $\{t\}_1 \geq \hat{v}_{\max}^*$ that $\{\phi(t)\}_2 \leq \{\phi(t)\}_1$; and (b) for types t with $\{t\}_1 < \hat{v}_{\max}^*$, that $\{\phi(t)\}_2 \leq 0$.

Our approach will be to show that the paths defined by vector field α , i.e., the direction of α , are the equiquantile paths (Definition 8.20) and that ratio-monotonicity of these paths (Definition 8.19) implies both conditions (a) and (b) when the distribution of the agents value $\{t\}_1$ for her favorite alternative is regular.

Lemma 8.40. *The vector field α of the two-dimensional extension of the favorite-item projection corresponds to the equiquantile paths.*

Proof. Consider subspace of types who value alternative 1 more than \hat{v} but lie below the $\hat{\theta}$ -equiquantile path $C_{\hat{\theta}}$, i.e., $S^{\hat{v}\hat{\theta}} = \{t \in \mathcal{T} : \{t\}_1 \geq \hat{v} \wedge \{t\}_2 \leq C_{\hat{\theta}}(\{t\}_1)\}$; see Figure 8.17. We will show that the outflux of α on the top boundary, namely the types on the path $C_{\hat{\theta}}$, is zero by applying the divergence theorem to subspace $S^{\hat{v}\hat{\theta}}$.

By the divergence density equality, the divergence of α on subspace $S^{\hat{v}\hat{\theta}}$ can be evaluated as:

$$\begin{aligned} \int_{t \in S^{\hat{v}\hat{\theta}}} \nabla \cdot \alpha(t) dt &= - \int_{t \in S^{\hat{v}\hat{\theta}}} f(t) dt \\ &= -\mathbf{Pr}[t \in S^{\hat{v}\hat{\theta}}] \\ &= -\hat{\theta} [1 - F_{\max}(\hat{v})]. \end{aligned}$$

The first line is by the divergence density, and the second line is by the definition of probability. The third line follows because the probabilities that $\{t\}_1 \geq \hat{v}$ and $\{t\}_2 \leq C_{\hat{\theta}}(\{t\}_1)$ are independent (Lemma 8.33); equal to $[1 - F_{\max}(\hat{v})]$ and $\hat{\theta}$, respectively; and the probability of the intersection of two independent events is the product of their probabilities.

As before, the outflux of α on the right and bottom boundary is zero. The outflux of α on the left boundary can be calculated by integrat-

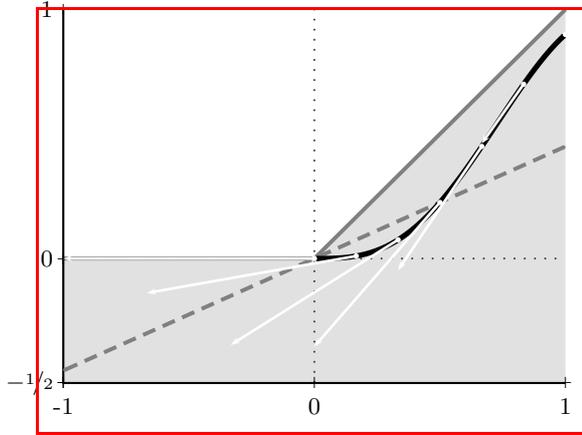


Figure 8.18. An equiquantile path C_θ (solid, thick, black line) is depicted. For the case where the distribution of $\{t\}_1$ is uniform on $[0, 1]$ the virtual values for six types on this path are depicted. By the definition of canonical amortization, $\phi(t) = t - \alpha(t)/f(t)$; thus, the virtual values can be calculated by starting at the type t (white bullets) and moving backwards (white arrows) a magnitude of $|\alpha(t)/f(t)|$ in the direction of $-\alpha(t)$. The gray shaded region represents the allowable space $\{\phi \in [-1, 1]^2 : \{\phi\}_2 \leq \{\phi\}_1 \vee \{\phi\}_2 \leq 0\}$ of virtual values for which selling only alternative 1 is optimal. Ratio monotonicity of the path implies that the slope of the path at a type t is greater than that of the line through the type and the origin (depicted for type $(1/2, C_\theta(1/2))$; gray, dashed line). Since the direction of α is tangent to the path, the virtual value for the type is below the line and, consequentially, within the allowable space.

ing the formula $\{\alpha(t)\}_1 = \frac{1 - F_{\max}(\{t\}_1)}{f_{\max}(\{t\}_1)} f(t)$ with respect to $\{t\}_2$ when $\{t\}_2 = \hat{v}$. The first term in this integral is independent of $\{t\}_2$ and can be factored out. Integrating second-term, i.e., the density, on the left boundary up to quantile $\hat{\theta}$ gives $\hat{\theta} f_{\max}(\hat{v})$ as the total density of types with $\{t\}_1 = \hat{v}$ is $f_{\max}(\hat{v})$. The product of these quantities gives influx

$$\frac{1 - F_{\max}(\hat{v})}{f_{\max}(\hat{v})} \times \hat{\theta} f_{\max}(\hat{v}) = \hat{\theta} (1 - F_{\max}(\hat{v})),$$

and its negation is the outflux. Thus, the remaining outflux for the top boundary of type subspace $S^{\hat{\theta}}$ is zero. This equality holds for all values \hat{v} , thus the outflux at each type t on the top boundary is identically zero. In other words, α is tangent to the equiquantile paths. \square

Theorem 8.41. For distribution F on type space $\mathcal{T} = \{t \in [0, 1]^2 : \{t\}_2 \leq \{t\}_1\}$ satisfying (c) the distribution of $\{t\}_1$ is regular and (d) the

equiquantile paths are ratio monotone, the two-dimensional extension of the favorite-alternative projection ϕ satisfies properties (a) and (b) of Proposition 8.38, i.e., the alternative 1 coordinate of ϕ is a virtual value for the single dimensional projection and alternative 2 coordinate of ϕ never maximizes virtual surplus.

Proof. Property (c), the regularity of the distribution of $\{t\}_1$, implies that the amortization of revenue for the single-dimensional projection is a virtual value function. Hence, property (a) holds.

We argue that property (d) implies property (b) as follows. Recall from Definition 8.23 that $\phi(t) = t - \alpha(t)/f(t)$. Consider this vector addition geometrically as the vector from the origin to t plus the vector back in the direction of $-\alpha(t)$ (the magnitude will not be important). Since the paths defined by vector field α are ratio monotone, this vector back lies below the line that connects the type t to the origin; see Figure 8.18. In other words,

$$\{\phi(t)\}_2 \leq \frac{\{t\}_2}{\{t\}_1} \{\phi(t)\}_1. \quad (8.15)$$

Thus, when $\{\phi\}_1 > 0$ then $\{\phi\}_2 \leq \{\phi\}_1$, and when $\{\phi\}_1 \leq 0$ then $\{\phi\}_2 \leq 0$. \square

We wrap this section up by applying this virtual value theory to the two-alternative uniform unit-demand agent of Example 8.2 to solve the three single-agent problems of Section 8.1.

Example 8.9. For the two-alternative uniform unit-demand agent of Example 8.2 the two-dimensional extension of the single-dimensional projection satisfies the assumptions of Proposition 8.38. The virtual value function (on $t \in \{t \in \mathcal{T} : \{t\}_2 \leq \{t\}_1\}$):

$$\phi(t) = (1, \{t\}_2/\{t\}_1) [\{t\}_1 - 1 - F_{\max(\{t\}_1)/f_{\max(\{t\}_1)}}].$$

The three single-agent problems are solved by optimizing revenue under the single-dimensional favorite-alternative projection. The optimal unconstrained mechanism posts a uniform price of $\sqrt{1/3}$, the \hat{q} ex ante optimal mechanism for post a uniform price that sells with ex ante probability \hat{q} (for $\hat{q} \leq \hat{q}^* = 2/3$), the \hat{y} interim optimal mechanism is the distribution over uniform prices that give allocation rule $y(q) = \hat{y}(q)$ for $q \leq \hat{q}^*$ (and $y(q) = 0$, otherwise).

Exercises

- 8.1 Consider two agents with independent, identical, and uniformly distributed values on $[0, 1]$ and budget $B = 1/4$. Solve for the equilibrium of the highest-bid-wins all-pay auction by identifying the critical type that is indifferent between following the budget unconstrained equilibrium of bidding $s(t) = t^2/2$ (see Section 2.8) and bidding the budget B assuming all higher types bid the budget. Is this equilibrium unique?
- 8.2 Prove Theorem 8.7: Let $\hat{\mathbf{y}} = (\hat{y}, \dots, \hat{y})$ be the n -agent allocation constraints induced by the k strongest-agents-win mechanism and $\mathbf{y} = (y, \dots, y)$ the allocation rules induced by any symmetric k -unit mechanism for n i.i.d. agents, then y is feasible for \hat{y} .
- 8.3 Let $\hat{\mathbf{y}} = (\hat{y}, \dots, \hat{y})$ be the n -agent allocation constraints induced by the assortative matching of agents to n positions with weights $\mathbf{w} = (w_1, \dots, w_n)$, i.e., stronger agents are matched to positions with larger weights. The position weights correspond to a stochastic probability of service. Let $\mathbf{y} = (y, \dots, y)$ be the allocation rules induced by any symmetric mechanism for position weights \mathbf{w} and n i.i.d. agents. Prove that y is feasible for \hat{y} . (Hint: Use Theorem 8.7.)
- 8.4 Consider a three-agent position environment with position weights 1, $1/2$, and 0 for the first, second, and third positions; respectively. Recall that in a position environment (Definition 7.18), an agent assigned to the k th position is served with probability given by the k th position weight. Consider three agents with values drawn independently, identically, and uniformly from the interval $[0, 1]$ and each with a public budget constraint of $B = 1/4$ (Example 8.1). Derive the revenue optimal auction.
- 8.5 Consider the all-or-none set system that corresponds to a public project. Consider two agents with types drawn uniformly from type space $\mathcal{T} = \{L, H\}$ and characterize the class of symmetric interim feasible allocation rules. Specifically, what are the pairs of allocation probabilities $(x(L), x(H))$ that are induced by a symmetric ex post feasible mechanism?
- 8.6 Consider the all-or-none set system that corresponds to a public project. Consider two agents with quantiles drawn independently, identically, and uniformly from $[0, 1]$ and characterize the class of interim feasible allocation rules $y : [0, 1] \rightarrow [0, 1]$ that are induced by a symmetric ex post feasible mechanism. (Recall, y is monotone non-increasing by definition.)

- 8.7 Prove that the quantile space and type space inequalities (8.8) and (8.9) that characterize interim feasibility are equivalent (Theorem 8.10).
- 8.8 Extend Theorem 8.14 from general feasibility environments, i.e., where only $\mathbf{x} \in \mathcal{X}$ are feasible, to general cost environments, i.e., where $\mathbf{x} \in \{0, 1\}^n$ has service cost $c(\mathbf{x})$. Prove that if there exists a mechanism that induces the profile of allocation rules \mathbf{y} with expected cost C , then there is a stochastic weighted optimizer that induces \mathbf{y} with expected cost C .
- 8.9 Prove Theorem 8.18: *For any type distribution \mathbf{F} and mechanisms $\hat{\mathcal{M}}$ and \mathcal{M} with allocation rules $\hat{\mathbf{y}}$ and \mathbf{y} satisfying $\mathbf{y} \preceq \hat{\mathbf{y}}$, the composite mechanism (Definition 8.13) induces a distribution over allocations that is in the downward closure of the distribution of allocations of $\hat{\mathcal{M}}$ and the same interim mechanisms as \mathcal{M} .*¹⁰
- 8.10 Use the analysis of the public budget agent from Section 8.6 with the uniform public-budget agent (Example 8.1; budget $B = 1/4$) to solve the single-agent problems as described in Section 8.1.1. In particular derive the unconstrained optimal mechanism, the $1/2$ ex ante optimal mechanism, and the \hat{y} interim optimal mechanism for allocation constraint $\hat{y}(q) = 1 - q$. In your derivation of each of these mechanisms explicitly identify the correct Lagrangian parameter λ that gives the right Lagrangian revenue curves.
- 8.11 Consider the objective of optimizing welfare for an agent with a public budget. Adapt a version of Corollary 8.27 for the welfare objective. Specifically:
- Identify the optimal auction for any interim allocation constraint \hat{y} under sufficiently general assumptions on the distribution of types, e.g., regularity.
 - Clearly state the necessary assumptions on the distribution.
 - Identify the optimal auction for two agents with uniformly distributed types on $[0, 1]$ and public budget $B = 1/4$.
- 8.12 Consider a single-item, i.i.d., public-budget regular environment. Prove that the expected revenue of the highest-bid-wins all-pay auction (with no reserve or explicit intervals on which types are ironed) on $n + 1$ agents obtains at least optimal revenue for n agents. In other words, generalize Theorem 5.1 to public budgets.

¹⁰ One distribution of allocations is *in the downward closure* of a second distribution of allocations if there is a coupling of the distributions so that the set of agents served by the first is a subset of those served by the second.

- 8.13 Extend the proof Theorem 8.32 given in the text to relax the assumption of locally linearity of the path $C(\cdot)$ at the price \hat{v} where it is non-ratio-monotone. Hint: Instead of deriving limit equations for $\lim_{\epsilon \rightarrow 0} [\text{Gain}(\epsilon)/\epsilon f_{\max}]$ give a general expression for $\text{Gain}(\epsilon)$, take its derivative, and evaluate at $\epsilon = 0$; likewise for $\text{Loss}(\epsilon)$.

Chapter Notes

Exact reductions from multi-agent to single-agent mechanism design problems for multi-dimensional and non-linear agents and the objective of revenue were considered by Alaei et al. (2012) and Alaei et al. (2013). The former considered the general case of non-revenue-linear agents; the latter defined revenue linearity as a property of interest and generalized the optimality of the marginal revenue mechanism of Myerson (1981) and Bulow and Roberts (1989).

For single-item environments the necessary and sufficient conditions for ex post implementation of an interim mechanism were developed by Maskin and Riley (1984), Matthews (1984), and Border (1991). For symmetric single-item environments, the latter gave a characterization of interim feasible mechanisms that is similar to the one presented here (characterizing them as stochastic weighted optimizers). These results were generalized to asymmetric single-item environments by Border (2007) and Mierendorff (2011) and to matroid environments by Alaei et al. (2012) and Che et al. (2013). In Cai et al. (2012a,b) and Alaei (2012) these results were generalized beyond matroids to show that any interim feasible allocation for a general feasibility environment could be implemented as a stochastic weighted optimizer optimization. Cai et al. (2012a,b) additionally address environments with multi-dimensional externalities and the results presented here for multi-service service constrained environments, such as n -agent m -item matching environments, are an adaptation of their results.

Optimization of revenue and welfare for single-dimensional agents with public budgets was considered by Laffont and Robert (1996) and Maskin (2000), respectively. The derivation in this text is a simplification of the one from Laffont and Robert (1996) that is enabled by the marginal revenue framework of Bulow and Roberts (1989) and can be found in Devanur et al. (2013). For the analogous, and more challenging, optimization problem where the budget of the agent is private see Pai and Vohra (2014) and Alaei et al. (2012).

The multi-dimensional characterization of Bayesian incentive compatibility is due to Rochet (1985). Though it was not discussed in this text, there is generalization of this characterization to agents with non-linear utility by McAfee and McMillan (1988). The canonical amortizations which underlie the theory of multi-dimensional virtual values were characterized by Rochet and Choné (1998) and further refined in Rochet and Stole, 2003. Armstrong (1996) developed the approach of integration by parts on rays from the origin which, as we described in the text, can be used to prove the optimality of uniform pricing for an agent with multi-dimensional type drawn from the uniform distribution. The methods given in the text for solving for optimal mechanisms on paths and for reverse-solving for the right paths are from Haghpanah and Hartline (2015).

Haghpanah and Hartline (2015) apply their framework for multi-dimensional virtual values to prove the optimality of uniform pricing, broadly. This result can be viewed as a “no haggling” result for substitutes. Single-dimensional no-haggling theorems come from Stokey (1979) and Riley and Zeckhauser (1983). The no-haggling characterization of Haghpanah and Hartline for multi-dimensional types on paths shows that these single-dimensional no-haggling result are on the boundary between haggling and no haggling. The part of this characterization that shows when haggling can be expected is a simplification of an example from Thanassoulis (2004).

A similar characterization of expected revenue (to the canonical amortizations presented in the text) was given by Daskalakis et al. (2015) where, instead of integration by parts to rewrite expected revenue in terms of the allocation rule, they use integration by parts to rewrite expected revenue in terms of expected utility. This alternative approach is also useful in identifying optimal mechanisms, e.g., see Giannakopoulos and Koutsoupias (2014).

Appendix

Mathematical Reference

Contained herein is reference to mathematical notations and conventions used throughout the text.

A.1 Big-oh Notation

We give asymptotic bounds using big-oh notation. Upper bounds are given with O , strict upper bounds are given with o , lower bounds are given with Ω , strict lower bounds are given with ω , and exact bounds are given with Θ . Formal definitions are given as follows:

Definition A.1. Function $f(n)$ is $O(g(n))$ if there exists a $c > 0$ and $n_0 > 0$ such that

$$\forall n > n_0, f(n) \leq c g(n).$$

Definition A.2. Function $f(n)$ is $\Omega(g(n))$ if there exists a $c > 0$ and $n_0 > 0$ such that

$$\forall n > n_0, f(n) \geq c g(n).$$

Definition A.3. Function $f(n)$ is $\Theta(g(n))$ if it is $O(g(n))$ and $\Omega(g(n))$.

Definition A.4. Function $f(n)$ is $o(g(n))$ if it is $O(g(n))$ but not $\Theta(g(n))$.

Definition A.5. Function $f(n)$ is $\omega(g(n))$ if it is $\Omega(g(n))$ but not $\Theta(g(n))$.

A.2 Common Probability Distributions

Common continuous probability distributions are *uniform* and *exponential*. Continuous distributions can be specified by their *cumulative distribution function*, denoted by F , or its derivative $f = F'$, the *probability density function*.

Definition A.6. The *uniform distribution* on support $[a, b]$, denoted $U[a, b]$, is defined as having a constant density function $f(z) = 1/(b - a)$ over $[a, b]$.

For example, the distribution $U[0, 1]$ has distribution $F(z) = z$ and density $f(z) = 1$. The expectation of the uniform distribution on $[a, b]$ is $\frac{a+b}{2}$. The monopoly price for the uniform distribution is $\max(b/2, a)$ (see Definition 3.7).

Definition A.7. The *exponential distribution* with rate λ has distribution $F(z) = 1 - e^{-\lambda z}$ and density $f(z) = \lambda e^{-\lambda z}$. The support of the exponential distribution is $[0, \infty)$.

The exponential distribution with rate λ has expectation $1/\lambda$ and monopoly price $1/\lambda$. The exponential distribution has constant *hazard rate* λ .

A.3 Expectation and Order Statistics

The *expectation* of a random variable $v \sim F$ is its “probability weighted average.” For continuous random variables this expectation can be calculated as

$$\mathbf{E}[v] = \int_{-\infty}^{\infty} z f(z) dz. \quad (\text{A.1})$$

For continuous, non-negative random variables this expectation can be reformulated as

$$\mathbf{E}[v] = \int_0^{\infty} (1 - F(z)) dz \quad (\text{A.2})$$

which follows from (A.1) and integration by parts.

An *order statistic* of a set of random variables is the value of the variable that is at a particular rank in the sorted order of the variables. For instance, when a valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ is drawn from a distribution then the i th largest value, which we have denoted $v_{(i)}$, is an

order statistic. A fact that is useful for working out examples with the uniform distribution.

Fact A.1. *In expectation, i.i.d. random variables chosen uniformly from a given interval will evenly divide the interval.*

A.4 Integration by Parts

Integration by parts is the integration analog of the product rule for differentiation. We will denote the derivative of a function $\frac{d}{dz}g(z)$ by $g'(z)$. The product rule for differentiation is:

$$[g(z)h(z)]' = g'(z)h(z) + g(z)h'(z). \quad (\text{A.3})$$

The formula for integration by parts can be derived by integrating both sides of the equation and rearranging.

$$\int g'(z)h(z) dz = g(z)h(z) - \int g(z)h'(z) dz. \quad (\text{A.4})$$

As an example we will derive (A.2) from (A.1). Plug $g(z) = 1 - F(z)$ and $h(z) = z$ into equation A.4.

$$\begin{aligned} \mathbf{E}[v] &= \int_0^\infty z f(z) dz \\ &= - \int_0^\infty h(z) g'(z) dz \\ &= - \left[h(z) g(z) \right]_0^\infty + \int_0^\infty h'(z) g(z) dz \\ &= - \left[z(1 - F(z)) \right]_0^\infty + \int_0^\infty 1(1 - F(z)) dz \\ &= \int_0^\infty (1 - F(z)) dz. \end{aligned}$$

The last equality follows because $z(1 - F(z))$ is zero at both zero and ∞ .

A.5 Hazard Rates

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

A distribution is completely specified by its hazard rate via the following formula.

$$F(z) = 1 - e^{-\int_{-\infty}^z h(z) dz}.$$

References

- Aggarwal, G., Feder, T., Motwani, R., and Zhu, A. 2004. Algorithms for Multi-Product Pricing. In: *Proc. of ICALP*.
- Aggarwal, G., Fiat, A., Goldberg, A., Hartline, J., Immorlica, N., and Sudan, M. 2005. Derandomization of Auctions. In: *Proc. 37th ACM Symposium on the Theory of Computing*. ACM Press, New York.
- Agrawal, S., Ding, Y., Saberi, A., and Ye, Y. 2010. Correlation robust stochastic optimization. Pages 1087–1096 of: *Proc. 21st ACM Symp. on Discrete Algorithms*. Society for Industrial and Applied Mathematics.
- Alaei, S. 2011. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In: *Proc. 52th IEEE Symp. on Foundations of Computer Science*.
- Alaei, S., Malekian, A., and Srinivasan, A. 2009. On random sampling auctions for digital goods. Pages 187–196 of: *ACM Conference on Electronic Commerce*.
- Alaei, S., Fu, H., Haghpanah, N., and Hartline, J. 2013. The Simple Economics of Approximately Optimal Auctions. In: *Proc. 54th IEEE Symp. on Foundations of Computer Science*.
- Alaei, Saeed. 2012. *Mechanism design with general utilities*. Ph.D. thesis, University of Maryland.
- Alaei, Saeed, Fu, Hu, Haghpanah, Nima, Hartline, Jason, and Malekian, Azarakhsh. 2012. Bayesian optimal auctions via multi-to single-agent reduction. Page 17 of: *Proceedings of the 13th ACM Conference on Electronic Commerce*.
- Amann, E., and Leininger, W. 1996. Asymmetric All-Pay Auctions with Incomplete Information: The Two-Player Case. *Games and Economic Behavior*, **14**(1), 1 – 18.
- Archer, A., and Tardos, E. 2001. Truthful mechanisms for one-parameter agents. In: *Proc. 42nd IEEE Symp. on Foundations of Computer Science*.
- Archer, A., Papadimitriou, C., Talwar, K., and Tardos, E. 2003. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In: *Proc. 14th ACM Symp. on Discrete Algorithms*. ACM/SIAM.

- Armstrong, Mark. 1996. Multiproduct Nonlinear Pricing. *Econometrica*, **64**(1), 51–75.
- Ausubel, L., and Milgrom, P. 2006. The lovely but lonely Vickrey auction. Chap. 1, pages 17–40 of: Cramton, P., Shoham, Y., and Steinberg, R. (eds), *Combinatorial Auctions*. MIT Press.
- Babaioff, M., Blumrosen, L., Naor, M., and Schapira, M. 2008. Informational overhead of incentive compatibility. Pages 88–97 of: *Proc. 9th ACM Conf. on Electronic Commerce*.
- Babaioff, M., Sharma, Y., and Slivkins, A. 2009a. Characterizing truthful multi-armed bandit mechanisms: extended abstract. Pages 79–88 of: *Proc. 10th ACM Conf. on Electronic Commerce*.
- Babaioff, M., Lavi, R., and Pavlov, E. 2009b. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *Journal of the ACM*.
- Babaioff, M., Kleinberg, R., and Slivkins, A. 2010. Truthful Mechanisms with Implicit Payment Computation. In: *Proc. 11th ACM Conf. on Electronic Commerce*.
- Bajari, P. 2001. Comparing competition and collusion: a numerical approach. *Economic Theory*, **18**, 187–205.
- Balcan, M.-F., Blum, A., Hartline, J., and Mansour, Y. 2008. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, **74**(8), 1245–1270.
- Bei, X., and Huang, Z. 2011. Bayesian Incentive Compatibility via Fractional Assignments. In: *Proc. 22nd ACM Symp. on Discrete Algorithms*.
- Bhattacharya, S., Goel, G., Gollapudi, S., and Munagala, K. 2010. Budget Constrained Auctions with Heterogeneous Items. In: *Proc. 41st ACM Symp. on Theory of Computing*.
- Border, Kim C. 1991. Implementation of reduced form auctions: A geometric approach. *Econometrica: Journal of the Econometric Society*, 1175–1187.
- Border, Kim C. 2007. Reduced form auctions revisited. *Economic Theory*, **31**(1), 167–181.
- Briest, P. 2008. Uniform Budgets and the Envy-Free Pricing Problem. Pages 808–819 of: *ICALP*.
- Briest, P., Krysta, P., and Vöcking, B. 2005. Approximation techniques for utilitarian mechanism design. Pages 39–48 of: *Proc. 36th ACM Symp. on Theory of Computing*.
- Briest, P., Chawla, S., Kleinberg, R., and Weinberg, S. 2010. Pricing Randomized Allocations. Pages 585–597 of: *Proc. 21st ACM Symp. on Discrete Algorithms*.
- Bulow, J., and Klemperer, P. 1996. Auctions versus negotiations. *American Economic Review*, **86**, 180–194.
- Bulow, J., and Roberts, J. 1989. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, **97**, 1060–90.
- Cai, Y., and Daskalakis, C. 2011. Extreme-Value Theorems for Optimal Multidimensional Pricing. In: *Proc. 52th IEEE Symp. on Foundations of Computer Science*.

- Cai, Yang, Daskalakis, Constantinos, and Weinberg, S Matthew. 2012a. An algorithmic characterization of multi-dimensional mechanisms. Pages 459–478 of: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM.
- Cai, Yang, Daskalakis, Constantinos, and Weinberg, S Matthew. 2012b. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. Pages 130–139 of: *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*. IEEE.
- Chawla, S., and Hartline, J. 2013. Auctions with Unique Equilibria. Pages 181–196 of: *Proc. 14th ACM Conf. on Electronic Commerce*.
- Chawla, S., Hartline, J., and Kleinberg, R. 2007. Algorithmic Pricing via Virtual Valuations. In: *Proc. 8th ACM Conf. on Electronic Commerce*.
- Chawla, S., Malec, D., and Sivan, B. 2010a. The power of randomness in bayesian optimal mechanism design. Pages 149–158 of: *ACM Conference on Electronic Commerce*.
- Chawla, S., Hartline, J., Malec, D., and Sivan, B. 2010b. Sequential Posted Pricing and Multi-parameter Mechanism Design. In: *Proc. 41st ACM Symp. on Theory of Computing*.
- Che, Yeon-Koo, Kim, Jinwoo, and Mierendorff, Konrad. 2013. Generalized Reduced-Form Auctions: A Network-Flow Approach. *Econometrica*, **81**(6), 2487–2520.
- Chen, Ning, Gravin, Nick, and Lu, Pinyan. 2014a. *Competitive Auction Design: How to Play with Benchmarks*. Manuscript.
- Chen, Ning, Gravin, Nick, and Lu, Pinyan. 2014b. Optimal Competitive Auctions. Pages 253–262 of: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. New York, NY, USA: ACM.
- Clarke, E. H. 1971. Multipart Pricing of Public Goods. *Public Choice*, **11**, 17–33.
- Daskalakis, Constantinos, Deckelbaum, Alan, and Tzamos, Christos. 2015. Strong Duality for a Multiple-Good Monopolist. Pages 449–450 of: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM.
- Day, R., and Milgrom, P. 2007. Core selecting package auctions. *International Journal of Game Theory*, **36**, 393–407.
- Demers, A., Keshav, S., and Shenker, S. 1989. Analysis and simulation of a fair queueing algorithm. Pages 1–12 of: *Symposium proceedings on Communications architectures & protocols*.
- Devanur, N., and Hartline, J. 2009. Limited and online supply and the bayesian foundations of prior-free mechanism design. Pages 41–50 of: *Proc. 10th ACM Conf. on Electronic Commerce*.
- Devanur, N., and Kakade, S. 2009. The price of truthfulness for pay-per-click auctions. Pages 99–106 of: *Proc. 10th ACM Conf. on Electronic Commerce*.
- Devanur, N., Ha, B., and Hartline, J. 2013. Prior-free Auctions for Budgeted Agents. Pages 287–304 of: *Proc. 14th ACM Conf. on Electronic Commerce*.

- Devanur, Nikhil R, Hartline, Jason D, and Yan, Qiqi. 2015. Envy freedom and prior-free mechanism design. *Journal of Economic Theory*, **156**, 103–143.
- Dhangwatnotai, P., Roughgarden, T., and Yan, Qiqi. 2010. Revenue Maximization with a Single Sample. In: *ECOM10*.
- Diakonikolas, Ilias, Papadimitriou, Christos, Pierrakos, George, and Singer, Yaron. 2012. Efficiency-revenue trade-offs in auctions. Pages 488–499 of: *Automata, Languages, and Programming*. Springer.
- Dughmi, S., and Roughgarden, T. 2010. Black-Box Randomized Reductions in Algorithmic Mechanism Design. Pages 775–784 of: *Proc. 51th IEEE Symp. on Foundations of Computer Science*.
- Dughmi, S., Roughgarden, T., and Sundararajan, M. 2009. Revenue Submodularity. In: *Proc. 10th ACM Conf. on Electronic Commerce*.
- Dwork, C., and Naor, M. 1992. Pricing via Processing or Combatting Junk Mail. Pages 139–147 of: *12th Annual International Cryptology Conference*.
- Edelman, B., Ostrovsky, M., and Schwarz, M. 2007. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *The American Economic Review*, **97**(1), 242–259.
- Even, S, Itai, A, and Shamir, A. 1976. On the Complexity of Timetable and Multicommodity Flow Problems. *SIAM Journal on Computing*, **5**(4), 691–703.
- Feige, U., Flaxman, A., Hartline, J., and Kleinberg, R. 2005. On the Competitive Ratio of the Random Sampling Auction. Pages 878–886 of: *Proc. 1st Workshop on Internet and Network Economics*.
- Fiat, A., Goldberg, A., Hartline, J., and Karlin, A. 2002. Generalized Competitive Auctions. Pages 72–81 of: *Proc. 34th ACM Symp. on Theory of Computing*. ACM Press.
- Giannakopoulos, Yiannis, and Koutsoupias, Elias. 2014. Duality and optimality of auctions for uniform distributions. Pages 259–276 of: *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM.
- Goldberg, A., Hartline, J., Karlin, A., and Saks, M. 2004. A Lower Bound on the Competitive Ratio of Truthful Auctions. In: *Proc. 21st Annual Symp. on Theoretical Aspects of Computer Science*. Springer.
- Goldberg, A. V., and Hartline, J. D. 2003. Competitiveness via Consensus. In: *Proc. 14th ACM Symp. on Discrete Algorithms*. ACM/SIAM.
- Goldberg, A. V., Hartline, J. D., and Wright, A. 2001. Competitive Auctions and Digital Goods. Pages 735–744 of: *Proc. 12th ACM Symp. on Discrete Algorithms*. ACM/SIAM.
- Goldberg, A. V., Hartline, J. D., Karlin, A., Saks, M., and Wright, A. 2006. Competitive Auctions. *Games and Economic Behavior*, **55**, 242–269.
- Griesmer, J., Levitan, R., and Shubik, M. 1967. Towards a study of bidding processes, part four: games with unknown costs. *Naval Research Quarterly*, **14**, 415–443.
- Griffin, T., Shepherd, B., and Wilfong, G. 2002. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, **10**(April), 232–243.
- Groves, T. 1973. Incentives in Teams. *Econometrica*, **41**, 617–631.

- Guruswami, V., Hartline, J., Karlin, A., Kempe, D., Kenyon, C., and McSherry, F. 2005. On Profit-Maximizing Envy-Free Pricing. In: *Proc. 16th ACM Symp. on Discrete Algorithms*.
- Ha, B., and Hartline, J. 2011. Mechanism Design via Consensus Estimates, Cross Checking, and Profit Extraction. *Transactions on Economics and Computation*, **1**. to appear.
- Haghpanah, Nima, and Hartline, Jason. 2015. Reverse Mechanism Design. Pages 757–758 of: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM.
- Harsanyi, John C. 1967. Games with Incomplete Information Played by “Bayesian” Players; Part I. The Basic Model. *Management Science*, **14**(3), 159–182.
- Hartline, J., and Lucier, B. 2010. Bayesian Algorithmic Mechanism Design. In: *Proc. 41st ACM Symp. on Theory of Computing*.
- Hartline, J., and McGrew, R. 2005. From Optimal Limited to Unlimited Supply Auctions. Pages 175–182 of: *Proceedings of the 7th ACM Conference on Electronic Commerce*.
- Hartline, J., and Roughgarden, T. 2008. Optimal Mechanism Design and Money Burning. In: *Proc. 39th ACM Symp. on Theory of Computing*.
- Hartline, J., and Roughgarden, T. 2009. Simple versus Optimal Mechanisms. In: *Proc. 10th ACM Conf. on Electronic Commerce*.
- Hartline, J., and Yan, Q. 2011. Envy, truth, and profit. Pages 243–252 of: *Proceedings of the 12th ACM conference on Electronic commerce*.
- Hartline, J., Mirrokni, V., and Sundararajan, M. 2008. Optimal marketing strategies over social networks. Pages 189–198 of: *Proceedings of the 17th international conference on World Wide Web*. ACM.
- Hartline, J., Kleinberg, R., and Malekian, A. 2011. Bayesian Incentive Compatibility via Matchings. In: *Proc. 22nd ACM Symp. on Discrete Algorithms*.
- Hartline, Jason, Hoy, Darrell, and Taggart, Sam. 2014. Price of anarchy for auction revenue. Pages 693–710 of: *Proceedings of the fifteenth ACM conference on Economics and computation*.
- Jackson, M., and Kremer, I. 2007. Envy-freeness and implementation in large economies. *Review of Economic Design*, **11**, 185–198.
- Jansen, B., and Mullen, T. 2008. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, **6**, 114–131.
- Kaplan, Todd R, and Zamir, Shmuel. 2012. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, **50**(2), 269–302.
- Karp, R. 1972. Reducibility Among Combinatorial Problems. Pages 85–103 of: Miller, R., and Thatcher, J. (eds), *Complexity of Computer Computations*. Plenum Press.
- Kirkegaard, René. 2006. A short proof of the Bulow-Klemperer auctions vs. negotiations result. *Economic Theory*, **28**, 449–452.
- Kleinberg, R., and Weinberg, S. 2012. Matroid prophet inequalities. Pages 123–136 of: *Proc. 43rd ACM Symp. on Theory of Computing*. ACM.

- Koutsoupias, Elias, and Papadimitriou, Christos. 1999. Worst-case equilibria. Pages 404–413 of: *STACS 99*. Springer.
- Kruskal, J. 1956. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. Pages 48–50 of: *Proceedings of the American Mathematical Society*, vol. 7.
- Laffont, Jean-Jacques, and Robert, Jacques. 1996. Optimal auction with financially constrained buyers. *Economics Letters*, **52**(2), 181–186.
- Lebrun, B. 2006. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, **55**(1), 131 – 151.
- Lehmann, D., O’Callaghan, L., and Shoham, Y. 2002. Truth revelation in rapid, approximately efficient combinatorial auctions. *Journal of the ACM*, **49**.
- Leiner, B.M., Cerf, V.G., Clark, D.D., Kahn, R.E., Kleinrock, L., Lynch, D.C., Postel, J., Roberts, L.G., and Wolff, S.S. 1997. The past and future history of the Internet. *Communications of the ACM*, **40**(2), 102–108.
- Lucier, Brendan, and Borodin, Allan. 2010. Price of anarchy for greedy auctions. Pages 537–553 of: *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*.
- Maskin, E., and Riley, J. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, **45**(2), 395 – 409.
- Maskin, Eric. 2000. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, **44**(4–6), 667–681.
- Maskin, Eric, and Riley, John. 1984. Optimal auctions with risk averse buyers. *Econometrica*, 1473–1518.
- Matthews, Steven A. 1984. On the implementability of reduced form auctions. *Econometrica*, 1519–1522.
- McAfee, R Preston, and McMillan, John. 1988. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, **46**(2), 335–354.
- Mierendorff, Konrad. 2011. Asymmetric reduced form auctions. *Economics Letters*, **110**(1), 41–44.
- Myerson, R. 1981. Optimal Auction Design. *Mathematics of Operations Research*, **6**, 58–73.
- Myerson, R.B., and Satterthwaite, M.A. 1983. Efficient Mechanisms for Bilateral Trade. *Journal of Economic Theory*, **29**, 265–281.
- Nisan, N., and Ronen, A. 2001. Algorithmic Mechanism Design. *Games and Economic Behavior*, **35**, 166–196.
- Oxley, J. 2006. *Matroid Theory*. Oxford graduate texts in mathematics. Oxford University Press.
- Pai, Mallesh M, and Vohra, Rakesh. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, **150**, 383–425.
- Reynolds, S., and Wooders, J. 2009. Auctions with a buy price. *Economic Theory*, **38**, 9–39.
- Riley, John, and Zeckhauser, Richard. 1983. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 267–289.

- Rochet, Jean-Charles. 1985. The taxation principle and multi-time Hamilton-Jacobi equations. *Journal of Mathematical Economics*, **14**(2), 113–128.
- Rochet, Jean-Charles, and Choné, Philippe. 1998. Ironing, sweeping, and multidimensional screening. *Econometrica*, 783–826.
- Rochet, Jean-Charles, and Stole, Lars A. 2003. The economics of multidimensional screening. *Econometric Society Monographs*, **35**, 150–197.
- Roughgarden, Tim. 2012a. Intrinsic Robustness of the Price of Anarchy. *Communications of the ACM*, **55**(7), 116–123.
- Roughgarden, Tim. 2012b. The price of anarchy in games of incomplete information. Pages 862–879 of: *Proceedings of the 13th ACM Conference on Electronic Commerce*.
- Roughgarden, Tim, and Tardos, Éva. 2002. How bad is selfish routing? *Journal of the ACM (JACM)*, **49**(2), 236–259.
- Samuel-Cahn, E. 1984. Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *The Annals of Probability*, **12**(4), 1213–1216.
- Stokey, Nancy L. 1979. Intertemporal price discrimination. *The Quarterly Journal of Economics*, 355–371.
- Syrkanis, Vasilis, and Tardos, Eva. 2013. Composable and efficient mechanisms. Pages 211–220 of: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*.
- Taghi-Hajiaghayi, M., Kleinberg, R., and Sandholm, T. 2007. Automated Online Mechanism Design and Prophet Inequalities. Pages 58–65 of: *AAAI*.
- Talwar, K. 2003. The Price of Truth: Frugality in Truthful Mechanisms. Pages 608–619 of: *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science*.
- Thanassoulis, J. 2004. Hagglng over Substitutes. *J. of Economic Theory*, **117**, 217–245.
- Varian, H. 2007. Position auctions. *International Journal of Industrial Organization*, **25**(6), 1163–1178.
- Vickrey, W. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, **16**, 8–37.
- Vincent, D., and Manelli, A. 2007. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, **137**(1), 153–185.
- Wang, X., Montgomery, A., and Srinivasan, K. 2008. When auction meets fixed price: a theoretical and empirical examination of buy-it-now auctions. *Quantitative Marketing and Economics*, **6**, 339–370.
- Welsh, D. 2010. *Matroid Theory*. Dover Books on Mathematics. Dover Publications.
- Wilson, R. 1987. Game Theoretic Analysis of Trading Processes. *Advances in Economic Theory*.
- Wilson, R. 1997. *Nonlinear Pricing*. Oxford University Press.
- Yan, Q. 2011. Mechanism design via correlation gap. Pages 710–719 of: *Proc. 22nd ACM Symp. on Discrete Algorithms*. SIAM.