

3

Optimal Mechanisms

In this chapter we discuss the objectives of surplus and profit. As we will see, the economics of designing mechanisms to maximize 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. 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 2.4.1). 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 3.1.1. *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

(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 $5/12$.

Chapter 3: Topics Covered.

- single-dimensional environments;
- surplus maximization and externality pricing;
- profit maximization;
- virtual values and ironing;
- revenue curves and marginal revenue;
- welfare maximization with a balanced budget.

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 where 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.

The two most fundamental economics objectives are surplus, a.k.a., social welfare² and profit. Implicit in the definition of surplus is the fact that the payments from the agents are transferred to the service provider and therefore do not affect the objective³.

Definition 3.1.2. *The surplus of an allocation is the cumulative value of the agents served less the service cost:*

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

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

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

A key additional constraint that comes into play in mechanism design problems for environments with non-trivial designer costs $c(\cdot)$ or non-downward-closure and mixed-sign agent values, e.g., bilateral exchange, is budget balance. Budget balance requires that the mechanism does not sometimes subsidize the transaction.

Definition 3.1.3. *An outcome (\mathbf{x}, \mathbf{p}) is budget balanced if its profit is non-negative, i.e.,*

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

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

² A mechanism that optimizes 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 [Chapter 10](#)).

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 *consumer surplus* was discussed in detail in [Chapter 1](#) mechanisms for optimizing consumer surplus are the subject of [Exercise 3.1](#).

environments have costs:

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

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

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

The *bilateral exchange* (or *double auction*) environment is one where the mechanism is mediating the exchange of an item between two agents, a buyer and a seller. Agent 1 is a buyer and has a positive value for trading; agent 2 is a seller and has a negative value, i.e., a cost, for trading (but is willing to do so if paid). For example, surplus is maximized by trading if the buyer's value is greater than the seller's cost. The environment is not downward closed in the sense that either both agents trade or neither trade. Bilateral exchanges have costs:

$$c(\mathbf{x}) = \begin{cases} 0 & \text{if } x_1 = x_2, \text{ and} \\ \infty & \text{otherwise.} \end{cases}$$

The final environment we will discuss in this section is that of a *non-excludable public project*. In this environment 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. For example, once the bridge is built, it is not possible to exclude some agents, i.e., feasibility is *all or none*. In this sense, public project environments generalize the bilateral exchange environment to more than two agents and to non-zero costs. Excludable public projects have costs:

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

Non-excludable public project environments are difficult ones for mechanism design as incentives tend to be weak. When the cost C and agent values are positive (e.g., agents are buyers), there is a free-rider problem where an agent can bid zero and hope that the other agents support the project. In practice this free-rider problem is often solved by taxation where agents are taxed proportionally to wealth.

The non-excludable public project environment also has practical application where the cost C and values are negative (e.g., agents are sellers). For example, the government desires to construct a new highway and must buy the properties in the path of the highway from their individual owners. In this case, the weak incentives manifest in a hold-out problem where each seller would like to hold out and be paid the full benefit for the project $|C|$ (a negative cost C is a benefit). In practice, mechanisms like *eminent domain*, where a third-party assesses a fair market value to be paid to each agent, are common.

We can see from the examples above that the single-dimensional environments admit a wide variety of practical applications. Much of the text will focus on these environments. This chapter will further consider surplus maximization (Section 3.2), profit maximization (Section 3.3) and Section 3.4), and surplus maximization subject to budget balance (Section 3.5).

3.2 Surplus

We now derive the optimal mechanism for 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 equilibrium or 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 Bayesian incentive compatible 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 3.2.1. *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 $(\mathbf{z}, \mathbf{v}_{-i})$ denotes the vector \mathbf{v} with the i th coordinate replaced with \mathbf{z} .

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 2.10.4). 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. The price p_i that should be charged to agent i via Corollary 2.10.4 is \hat{v}_i if i is served (with $v_i \geq \hat{v}_i$) and zero otherwise.

The calculation of payments can be simplified to $p_i = \text{OPT}(0, \mathbf{v}_{-i}) - \text{OPT}_{-i}(\mathbf{v})$ 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 resulting mechanism as serving agents to maximize the surplus and charging each agent the externality she imposes on the others. This mechanism is often referred to as the *Vickrey-Clarke-Groves* (VCG) mechanism, named after William Vickrey, Edward Clarke, and Theodore Groves.

Definition 3.2.1. *The externality pricing (EP) mechanism for surplus maximization with single-dimensional agents is:*

- (i) Solicit and accept sealed bids \mathbf{b} with $b_i \in \mathbb{R}_+$.
- (ii) Find the optimal outcome $\mathbf{x} \leftarrow \text{OPT}(\mathbf{b})$.
- (iii) Set prices \mathbf{p} with $p_i \leftarrow \text{OPT}(0, \mathbf{b}_{-i}) - \text{OPT}_{-i}(\mathbf{b})$ for each agent i .

By Corollary 2.10.4 and Lemma 3.2.1 we have the following theorem, and by the optimality of OPT and the assumption that agents follow the dominant truthtelling strategy we have the following corollary.

Theorem 3.2.2. *The externality pricing mechanism is dominant strategy incentive compatible.*

Corollary 3.2.3. *The externality pricing mechanism optimizes surplus in dominant strategy equilibrium.*

Example 3.2.1. *The second-price routing auction from Chapter 1 is an instantiation of the externality pricing mechanism where feasible outcomes are subsets of agents whose messages can be simultaneously routed.*

It is useful to view the externality pricing 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 externality pricing 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, the optimal mechanism is distribution dependent.

There are other mechanisms, besides the externality pricing mechanism of [Definition 3.2.1](#) that maximize surplus in equilibrium. 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 $p_i(\mathbf{v}) = \text{OPT}_{-i}(\mathbf{v})$ 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., in the common winner-pays-bid format, that maximize surplus in equilibrium. We already saw that for i.i.d. distributions, the first-price auction and the second-price auction are equivalent ([Section 2.7](#)). In particular, for single-item environments, both maximize surplus. Moreover, for asymmetric distributions and general environments, it is possible to use distributional knowledge to carefully construct a winner-pays-bid surplus-maximizing mechanism (see [Exercise 3.2](#)).

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

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 3.3.6](#) on page [76](#). 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:

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

prior (Definition 2.3.2) 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 surplus and payments can be similarly considered (e.g., consumer surplus, see Exercise 3.1).

The following value distributions will serve as examples in the subsequent exposition.

Example 3.3.1. A uniform agent has single-dimensional linear utility with value v drawn uniformly from $[0, 1]$, i.e., cumulative distribution function $F(z) = z$ and density function $f(z) = 1$.

Example 3.3.2. A (exemplary) 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 73).

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). Via the payment identity of Theorem 2.5.1, namely

$$p(v) = v x(v) - \int_0^v x(v^\dagger) dv^\dagger, \quad (3.3.1)$$

the payment of an agent depends on her allocation rule, in particular, on $x(v^\dagger)$ for $v^\dagger \leq v$ when the agent's value is 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 surplus maximization, cf. Section 3.2). Consider selling an item to a single agent with value v drawn uniformly from $[0, 1]$ (Example 3.3.1). If her value is 0.2, then it is pointwise optimal to offer her the item at price $\hat{v} = 0.2$. 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

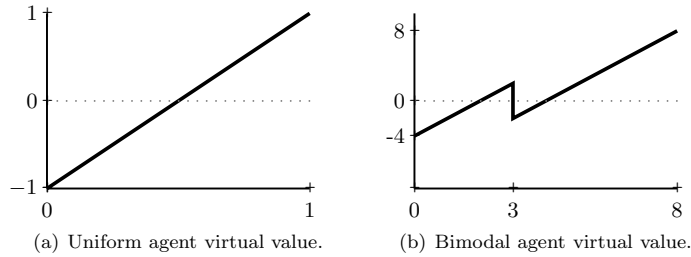


Figure 3.1. Depicted are virtual value functions $\varphi(v) = v - \frac{1-F(v)}{f(v)}$ for the uniform and bimodal agent examples [Example 3.3.1](#) and [Example 3.3.2](#). Notice that the virtual value function in the uniform example is monotone non-decreasing in value while in the bimodal example it is not. For reference, the line $v_2 = v_1$ is depicted (grey dotted line).

$\hat{v} = 0.7$. 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 $\hat{v}(1 - F(\hat{v}))$. For the uniform agent where $F(z) = z$, this optimal price is $\hat{v}^* = 1/2$ ⁴.

The payment identity [\(3.3.1\)](#) gives a formula for the expected payment 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 sub-

⁴ Set $\frac{d}{d\hat{v}} [\hat{v}(1 - \hat{v})] = 1 - 2\hat{v} = 0$ and solve for \hat{v} to get the optimal price to post of $\hat{v}^* = 1/2$.

sequently in [Section 3.3.4](#) will give virtual value function

$$\varphi(v) = v - \frac{1 - F(v)}{f(v)}. \quad (3.3.2)$$

In equation [\(3.3.2\)](#), 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

$$\mathbf{E}_{v \sim F}[p(v)] = \mathbf{E}_{v \sim F}[\varphi(v) x(v)] \quad (3.3.3)$$

for any allocation and payment rules (x, p) that satisfy the Bayes-Nash equilibrium characterization ([Theorem 2.5.1](#)) i.e., monotonicity of x and the payment identity [\(3.3.1\)](#). Equation [\(3.3.3\)](#) 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 [\(3.3.2\)](#) the virtual value function for the uniform agent example is $\varphi(v) = 2v - 1$; for the bimodal agent example it is depicted in [Figure 3.1](#). Notice that $\varphi(0) < 0$ 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 $\varphi(1) = 1$ and the bimodal agent with values on interval $[0, 8]$ has $\varphi(8) = 8$.

The importance of equation [\(3.3.3\)](#) 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 $\text{Surplus}(\varphi(\mathbf{v}), \mathbf{x}) = \sum_i \varphi_i(v_i) x_i - c(\mathbf{x})$ ⁵. Let OPT again be the surplus maximizing algorithm. We will care about both the allocation that $\text{OPT}(\varphi(\mathbf{v}))$ selects, i.e., $\text{argmax}_{\mathbf{x}} \text{Surplus}(\varphi(\mathbf{v}), \mathbf{x})$ and its virtual surplus $\text{max}_{\mathbf{x}} \text{Surplus}(\varphi(\mathbf{v}), \mathbf{x})$. Where it is unambiguous we will use notation $\text{OPT}(\varphi(\mathbf{v}))$ to denote either of these quantities. Note that this formulation of $\text{OPT}(\varphi(\cdot))$ 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

⁵ Here, $\varphi(\mathbf{v})$ denotes the profile of virtual values $(\varphi_1(v_1), \dots, \varphi_n(v_n))$.

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 Bayesian incentive compatibility characterization (Corollary 2.10.2) 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 $\text{OPT}(\varphi(\cdot))$, if the virtual value functions $\varphi_i(\cdot)$ for each agent i are monotone then $\text{OPT}(\varphi(\cdot))$ is monotone.

Lemma 3.3.1. *For any profile of virtual value functions φ , monotonicity of $\varphi_i(\cdot)$ implies the monotonicity of the allocation to agent i of $\text{OPT}(\varphi(\mathbf{z}, \mathbf{v}_{-i}))$ with respect to \mathbf{z} .*

Proof. Let $\mathbf{x}(\cdot)$ be the allocation rule of OPT , i.e., $\mathbf{x}(\mathbf{v}) = \text{argmax}_{\mathbf{x}} \text{Surplus}(\mathbf{v}, \mathbf{x})$. Recall from Lemma 3.2.1 that maximizing surplus is monotone in that $\mathbf{x}_i(\mathbf{z}, \mathbf{v}_{-i})$ is monotone in \mathbf{z} . Therefore $\mathbf{x}_i(\varphi_i(\mathbf{z}), \varphi_{-i}(\mathbf{v}_{-i}))$ is monotone in $\varphi_i(\mathbf{z})$, i.e., increasing $\varphi_i(\mathbf{z})$ does not decrease \mathbf{x}_i . By assumption $\varphi_i(\mathbf{z})$ is monotone in \mathbf{z} ; therefore, increasing \mathbf{z} cannot decrease $\varphi_i(\mathbf{z})$ which cannot decrease $\mathbf{x}_i(\varphi_i(\mathbf{z}), \varphi_{-i}(\mathbf{v}_{-i}))$. \square

For many distributions the virtual value function $\mathbf{v} - \frac{1-F(\mathbf{v})}{f(\mathbf{v})}$ of equation (3.3.2) is monotone, e.g., uniform (Example 3.3.1), 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 3.3.1. *A distribution F is regular if $\mathbf{v} - \frac{1-F(\mathbf{v})}{f(\mathbf{v})}$ is monotone non-decreasing.*

On the other hand, many relevant distributions are irregular, e.g., bimodal (Example 3.3.2; 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 mechanisms are virtual surplus maximizers, we now give a generic and formal description of this sort of mechanism. For monotone virtual value functions, [Lemma 3.3.1](#) 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”) [truthtelling is a dominant strategy equilibrium](#) (recall [Corollary 2.10.4](#)). One way to view the suggested virtual surplus maximization mechanism is as a reduction to surplus maximization, which is solved by the externality pricing mechanism (EP, see [Definition 3.2.1](#)).

Definition 3.3.2. *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*

$$(\mathbf{x}, \mathbf{p}^\dagger) \leftarrow \text{EP}(\varphi(\mathbf{b})),$$

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

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

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

Notice that the payments \mathbf{p} calculated by VSM can be viewed as follows. EP 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 $\varphi(\cdot)$ that are discontinuous or not strictly increasing this inverse virtual value function is defined as $\varphi^{-1}(z) = \inf\{v^\dagger : \varphi(v^\dagger) \geq z\}$.)

Theorem 3.3.2. *For monotone virtual value functions $\varphi = (\varphi_1, \dots, \varphi_n)$, the virtual surplus maximization mechanism VSM is dominant strategy incentive compatible.*

Proof. The theorem follows from [Lemma 3.3.1](#) applied to each agent, the definition of VSM, and [Corollary 2.10.4](#). \square

Corollary 3.3.3. *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 3.3.3](#) 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 3.3.3. *A virtual value function $\varphi(\cdot)$ 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 $\varphi(\cdot)$, there is some value $\hat{v}^* = \varphi^{-1}(0)$ 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 3.3.4. *For an agent with value v drawn from distribution F and virtual value function φ , the monopoly price $\hat{v}^* = \varphi^{-1}(0)$ 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, $\varphi_i(v_i) \geq \max(\varphi_j(v_j), 0)$ for all $j \neq i$, so i 's critical value is

$$\hat{v}_i = \max(\varphi_i^{-1}(\varphi_j(v_j)), \varphi_i^{-1}(0)) \quad (3.3.4)$$

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 3.3.3](#) below). Notice that the second term in the maximization is the monopoly price $\hat{v}_i^* = \varphi_i^{-1}(0)$. If the other agents are not competitive, i.e., all agents j have $\varphi_j(v_j) < 0$, then the optimization problem reduces to the single-agent case and agent i should see a reserve price of \hat{v}_i^* .

Corollary 3.3.4. *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 3.3.3. *Consider a two-agent single-item environment with agent 1's (Alice) value from $U[0, 1]$ (as in [Example 3.3.1](#)) 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 [\(3.3.2\)](#) are $\varphi_1(v_1) = 2v_1 - 1$ and $\varphi_2(v_2) = 2v_2 - 2$. The virtual surplus maximization mechanism serves Alice whenever $\varphi_1(v_1) > \max(\varphi_2(v_2), 0)$, 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 . In such a scenario, $\varphi_i^{-1}(\varphi_j(v_j)) = \varphi^{-1}(\varphi(v_j)) = v_j$, and equation (3.3.4) 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 the monopoly price $\hat{v}^* = \varphi^{-1}(0)$. What auction has this equilibrium outcome? The second-price auction with monopoly reserve \hat{v}^* .

Definition 3.3.5. *The second-price auction with reservation price \hat{v} , sells the item to the highest bidder 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 3.3.5. *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 $\hat{v}^* = \varphi^{-1}(0)$.*

Example 3.3.4. *Consider a two-agent single-item environment with i.i.d. uniform agents (as in Example 3.3.1). As we have calculated, $\varphi(v) = 2v - 1$ is monotone and strictly increasing, the monopoly price is $\hat{v}^* = \varphi^{-1}(0) = 1/2$, 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 environment'.*

Notice that the optimal reserve price depends on the distribution but 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. As we have seen, this single-agent tradeoff is optimized by setting a reserve equal to the monopoly price. The invariance of the optimal reserve price to details of the market besides the agents' distribution holds more broadly as well, see Proposition 4.6.1 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 , $\Pr[\varphi_i(v_i) = \varphi_j(v_j)] = 0$. On the other hand, when virtual value functions are constant on an interval $[a, b]$

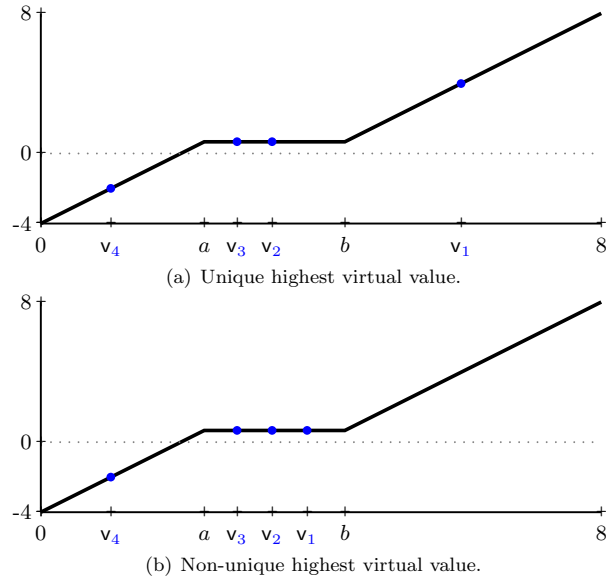


Figure 3.2. The weakly monotone virtual valuation function $\varphi(v)$ under two realizations of four agent values depicting both the case where the highest virtual value is unique and the case where it is not unique.

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 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 maxi-

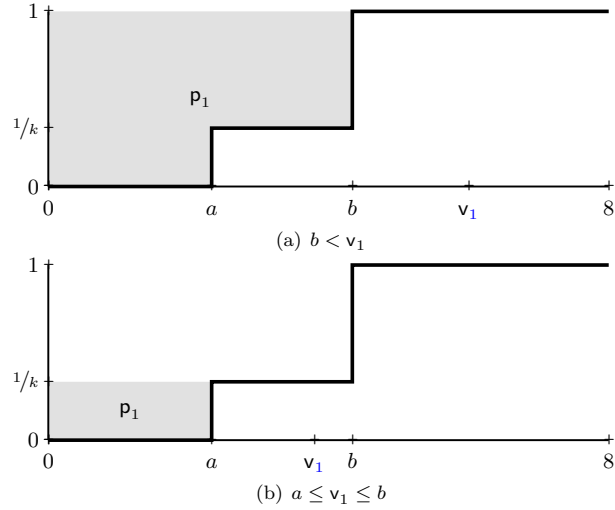


Figure 3.3. The allocation (black line) and payment rule (gray region) for agent 1 given fixed \mathbf{v}_{-1} with $k - 1$ of the other agents tied for having the highest virtual value, i.e., with values in $[a, b]$ (e.g., from virtual valuation function of [Figure 3.2](#)). For $v_1 \in [a, b]$, agent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

mization 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 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

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

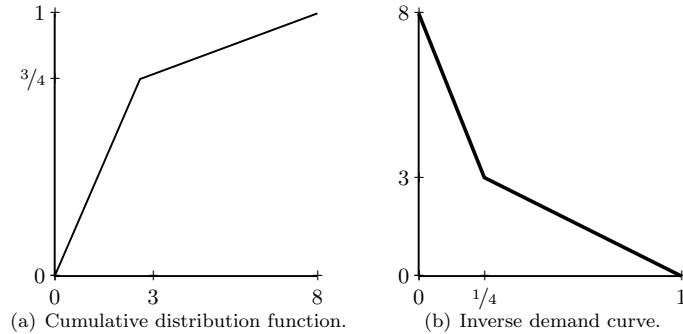


Figure 3.4. Depicted are the cumulative distribution function $F(v)$ and inverse demand curve $v(q)$ corresponding to the bimodal agent of Example 3.3.2. The inverse demand curve is obtained from the cumulative distribution function by rotating it 90 degrees counterclockwise.

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 $\mathbf{p}_1 = b - b^{-a/k}$, 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 $\mathbf{p}_1 = a/k$. 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.

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 2.4.1). 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 3.3.6. *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 3.3.5. *For the example of a uniform agent (Example 3.3.1) where $F(z) = z$, the inverse demand curve is $v(q) = 1 - q$; for the example of a bimodal agent (Example 3.3.2), 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 by y and defined as $y(q) = x(v(q))$. 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, steps from one to zero at \hat{q} , and is zero thereafter. 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 3.3.7. *The price-posting revenue curve of a single-dimensional linear agent specified by inverse demand curve $v(\cdot)$ is $P(\hat{q}) = \hat{q} v(\hat{q})$ for any $\hat{q} \in [0, 1]$.*

Notice that with this definition of price-posting revenue, $P(0) = 0$ for finite support distributions and $P(1) = 0$ for value distributions supported with lower endpoint 0. Simply, the revenues from never selling and always selling are zero.

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

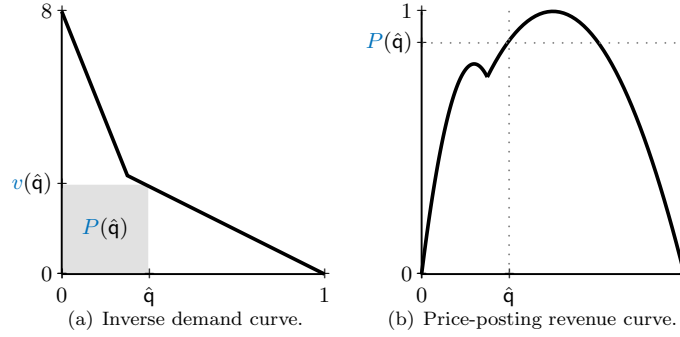


Figure 3.5. Depicted are the inverse demand curve $v(q)$ and price-posting revenue curve $P(q)$ corresponding to the bimodal agent of [Example 3.3.2](#). The price-posting revenue curve is given by $P(\hat{q}) = \hat{q} v(\hat{q})$, i.e., the area of the rectangle of width \hat{q} and height $v(\hat{q})$ that fits under the inverse demand curve.

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 with cumulative distribution function $G^y(z) = 1 - y(z)$ 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 can be expressed as follows.

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

The first line follows from a change of variables from $\hat{q} \sim G^y$ to $q \sim U[0, 1]$ according to G^y 's density function $g^y(z) = \frac{d}{dz}G^y(z) = \frac{d}{dz}(1 - y(z)) = -y'(z)$ and the pointmass of $y(1)$ at $\hat{q} = 1$. The second line follows from integration by parts 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 $P(\cdot)$.

Finally, we apply revenue equivalence ([Theorem 2.5.1](#)) to observe that any mechanism with the same allocation rule y generates the same revenue as calculated above. In the resulting proposition the additive term $P(0)y(0)$ is often zero as $P(0) = 0$ for distributions with finite support.

Proposition 3.3.6. *A single-agent mechanism with allocation rule y has expected revenue equal to the surplus of marginal price-posting revenue $\mathbf{E}_q[P'(q)y(q)] + P(0)y(0)$.*

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 reamortization of revenues across the different agent quantiles.

The marginal price-posting revenues are exactly the virtual values described previously by equation ([3.3.2](#)).

$$P'(q) = \frac{d}{dq}(q v(q)) = v(q) + q v'(q) = v - \frac{1-F(v)}{f(v)}, \quad (3.3.5)$$

where the first equality follows from the definition of price-posting revenue ([Definition 3.3.7](#)) 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 $\frac{1}{f(v)} = -\frac{d}{dq}v(q) = -v'(q)$. 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 P is concave.

Proposition 3.3.7. *A distribution F is regular if and only if its corresponding price-posting revenue curve is concave.*

[Proposition 3.3.6](#) 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 3.3.2) is dominant strategy incentive compatible and profit optimal (Corollary 3.3.3).

Theorem 3.3.8. *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} , which by Proposition 3.3.6 can be expressed as

$$\begin{aligned} \max_y \mathbf{E}[P'(q)y(q)] + y(0)P(0) \\ \text{s.t. } \mathbf{E}[y(q)] = \hat{q}. \end{aligned}$$

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 $\text{OPT}(\varphi(\cdot))$ 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 2.10.2). 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 the direct approach to an amortized analysis, e.g., of [Proposition 3.3.6](#) gives a non-monotone virtual value function, the following generic *ironing procedure* gives an ironed virtual value function which is monotone and for which point-wise optimization is equivalent to the optimization of expected virtual surplus subject to monotonicity of the allocation rule.

There are three key ideas to this ironing procedure. First, when maximizing a non-monotone virtual value function subject to monotonicity, the monotonicity constraint will bind on intervals of non-monotone virtual value and the resulting allocation rule on the interval will be constant, i.e., quantiles on that interval will be allocated with the same probability. Second, 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 third part implies the assumption of the second 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 $\varphi(\cdot)$ in quantile space. We denote the cumulative virtual value for quantiles at most \hat{q} as $\Phi(\hat{q}) = \int_0^{\hat{q}} \varphi(q) dq$. 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., $\varphi(q) = P'(q)$ and $\Phi(q) = P(q)$. The ironing procedure we will describe, however, can be applied to non-monotone virtual value functions corresponding to any objective.

The goal of ironing is arrive at a monotone (ironed) virtual value function, equivalently, a concave cumulative virtual value function, for which virtual surplus maximization has no loss when constrained to monotone allocation rules. We now investigate the consequences of the ironing pro-

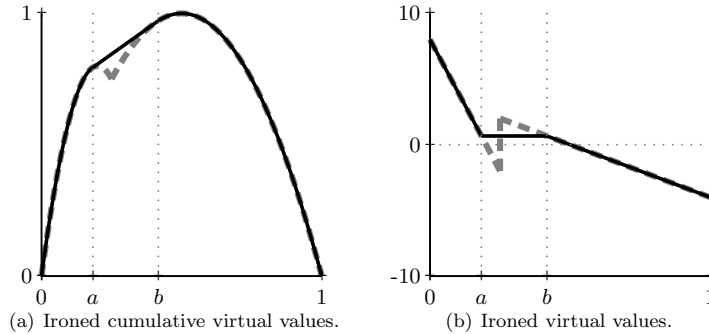


Figure 3.6. Consider the bimodal agent of [Example 3.3.2](#) and virtual value function equal to the marginal price-posting revenue curve. The cumulative virtual value and virtual value functions in quantile space are depicted (thick, gray, dashed lines) in the left and right diagram, respectively. After ironing on interval $[a, b]$, the resulting cumulative virtual value and virtual value functions are depicted (thin, black, solid lines). Subsequent discussion will justify this choice of interval $[a, b]$.

cedure 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 3.3.2](#) these quantities are depicted in [Figure 3.6](#), with a choice of interval $[a, b]$ that will be justified subsequently.

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 3.3.9](#) 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 3.3.8. *The ironing procedure for (non-monotone) virtual value function φ (in quantile space)⁶ is:*

- (i) Define the cumulative virtual value function as $\Phi(\hat{q}) = \int_0^{\hat{q}} \varphi(\mathbf{q}) \, d\mathbf{q}$.
- (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 $\bar{\varphi}(\mathbf{q}) = \frac{d}{d\mathbf{q}} \bar{\Phi}(\mathbf{q}) = \bar{\Phi}'(\mathbf{q})$.

Theorem 3.3.9. *For any monotone allocation rule $y(\cdot)$ and any virtual value function $\varphi(\cdot)$, the expected virtual surplus of an agent is upper-bounded by her expected ironed virtual surplus, i.e.,*

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

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

Proof. By integration by parts for any virtual value function $\varphi^\dagger(\cdot)$ and monotone allocation rule $y(\cdot)$ ([Mathematical Note](#) on page [62](#)),

$$\mathbf{E}[\varphi^\dagger(\mathbf{q}) y(\mathbf{q})] = \mathbf{E}[-y'(\mathbf{q}) \Phi^\dagger(\mathbf{q})]. \quad (3.3.6)$$

Notice that the (non-increasing) monotonicity of the allocation rule $y(\cdot)$ implies the non-negativity of $-y'(\mathbf{q})$. With the left-hand side of equation [\(3.3.6\)](#) 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 $\Phi(\cdot)$, $\bar{\Phi}(\mathbf{q}) \geq \Phi(\mathbf{q})$ and, therefore, for any monotone allocation rule, in expectation, the ironed virtual surplus is at least the virtual surplus. I.e., $\mathbf{E}[-y(\mathbf{q}) \bar{\Phi}(\mathbf{q})] \geq \mathbf{E}[-y(\mathbf{q}) \Phi(\mathbf{q})]$.

To see the equality under the assumption that $y'(\mathbf{q}) = 0$ for all \mathbf{q} where $\bar{\Phi}(\mathbf{q}) > \Phi(\mathbf{q})$, rewrite the difference between the ironed virtual surplus and the virtual surplus via equation [\(3.3.6\)](#) as,

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

The assumption on y' implies the term inside the expectation on the right-hand side is zero for all \mathbf{q} . \square

⁶ The ironing procedure can also be expressed in value space by first mapping 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.

Corollary 3.3.10. *For any virtual value function $\varphi(\cdot)$ with ironed virtual value $\bar{\varphi}(\cdot)$ from the ironing procedure [\(Definition 3.3.8\)](#), 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 $\bar{\varphi}$. This profile $\bar{\varphi}$ 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 3.3.11. *For any (non-monotone) virtual value functions φ , the ironed virtual surplus maximization mechanism maximizes expected virtual surplus in dominant strategy equilibrium.*

Corollary 3.3.12. *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 [Chapter 8](#) 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 2.10.1). 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 are 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, p) 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 according to the inverse demand curve v as $v(q)$. 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 3.4.1. *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 $R(\hat{q})$.

For the single-dimensional linear agents discussed in this chapter, where revenue equivalence holds, the revenue curve is calculated by the following mathematical program.

$$\begin{aligned} R(\hat{q}) &= \max_y \mathbf{E}[P'(q)y(q)] + y(0)P(0) \\ \text{s.t. } &\mathbf{E}[y(q)] = \hat{q}. \end{aligned}$$

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.⁷ 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 following proposition which, with [Proposition 3.3.6](#) (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 3.4.1. *The optimal ex ante pricing problems induce a concave revenue curve and, for any ex ante service probability, the optimal*

⁷ 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.

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 3.4.2. *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 3.4.3. *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} .*

For the single-dimensional linear agents discussed in this chapter, where revenue equivalence holds, the optimal revenue for an allocation constraint \hat{y} is calculated by the following mathematical program.

$$\begin{aligned} \mathbf{Rev}[\hat{y}] &= \max_y \mathbf{E}[P'(\mathbf{q}) y(\mathbf{q})] + y(0) P(0) \\ \text{s.t. } Y(\mathbf{q}) &\leq \hat{Y}(\mathbf{q}), \quad \forall \mathbf{q} \in [0, 1]. \end{aligned}$$

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 3.3.6](#) using optimal lottery pricings improves on price postings in that for each $\hat{\mathbf{q}}$ the optimal ex ante revenue $R(\hat{\mathbf{q}})$ may exceed the price-posting revenue $P(\hat{\mathbf{q}})$. Consider the mechanism that draws $\hat{\mathbf{q}}$ from the distribution $G^{\hat{y}}(z) = 1 - \hat{y}(z)$ and offers Alice the optimal ex ante pricing for $\hat{\mathbf{q}}$. The optimal revenue for allocation constraint \hat{y} must be at least the revenue of this mechanism. By the [Mathematical Note](#) on page [62](#), we have:

$$\begin{aligned} \mathbf{Rev}[\hat{y}] &\geq \mathbf{E}_{\mathbf{q} \sim G^{\hat{y}}}[R(\hat{\mathbf{q}})] \\ &= \mathbf{E}_{\mathbf{q}}[-\hat{y}'(\mathbf{q}) R(\mathbf{q})] + \hat{y}(1) R(1) \end{aligned} \quad (3.4.1)$$

$$= \mathbf{E}_{\mathbf{q}}[R'(\mathbf{q}) \hat{y}(\mathbf{q})] + R(0) \hat{y}(0), \quad (3.4.2)$$

where $R'(\mathbf{q}) = \frac{d}{d\mathbf{q}} R(\mathbf{q})$ is the *marginal revenue* at \mathbf{q} . Notice that for valuation distributions supported on $[0, h]$, the lower bound of the support of 0 gives $R(1) = 0$, the upper bound of the support of h gives $R(0) = 0$, and the additive terms in equations [\(3.4.1\)](#) and [\(3.4.2\)](#) are both zero.

Definition 3.4.4. *The surplus of marginal revenue of an allocation constraint \hat{y} is $\mathbf{MargRev}[\hat{y}] = \mathbf{E}_{\mathbf{q}}[R'(\mathbf{q}) \hat{y}(\mathbf{q})] + R(0) \hat{y}(0)$.*

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 $Y(1) = \hat{Y}(1)$.

Lemma 3.4.2. *Consider the alternative ex ante and interim pricing definitions.*

- (i) *For downward-closed ex ante pricing the marginal revenue satisfies $R'(q) \geq 0$ for all quantiles q .*
- (ii) *For non-downward-closed interim pricings y for \hat{y} the cumulative allocation rules satisfy $Y(1) = \hat{Y}(1)$.*

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 derivation of [Section 3.4.2](#) shows that 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 that for any allocation constraint the optimal revenue and surplus of marginal revenue are equal.

Definition 3.4.5. *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 $\mathbf{Rev}[\hat{y}] = \mathbf{Rev}[\hat{y}^\dagger] + \mathbf{Rev}[\hat{y}^\ddagger]$ ⁸

Proposition 3.4.3. *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. 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, obtain more revenue with a weaker allocation rule than the price postings, which define price-posting revenue curves (Proposition 3.4.1); and that revenue equivalence allows revenue to be expressed in terms of price-posting revenue curves (Proposition 3.3.6). Optimal revenue equaling surplus of marginal revenue for single-dimensional linear agents, then, is an immediate corollary of this revenue linearity and Proposition 3.4.3

Theorem 3.4.4. *A single-dimensional linear agent is revenue linear.*

Proof. In this proof we will assume that the distribution is bounded with finite upper endpoint $v(0)$ and lower endpoint $v(1) = 0$; consequently, the revenue curves are zero at their endpoints, i.e., $R(0) = R(1) = P(0) = P(1) = 0$; thus the additive terms in the expression of surplus of marginal revenue are zero. Recall as well that the cumulative allocation rules always satisfy $\hat{Y}(0) = Y(0) = 0$. The theorem statement under this assumption implies the general statement via a limit argument.

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 62). 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

⁸ It is assumed that the ex ante and interim problem are consistent with respect to downward closure, see Section 3.4.3

no lower a revenue (as $-y'(\cdot)$ is non-negative). In particular, the allocated marginal revenue exceeds the surplus of marginal price-posting revenue as $R(\mathbf{q}) \geq P(\mathbf{q})$ for all \mathbf{q} (by [Proposition 3.4.1](#)).

- (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} ⁹.

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 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 3.3.6](#)), the revenue of any allocation rule is equal to its surplus of marginal price-posting revenue, so

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

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

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

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

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

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

$$\mathbf{Rev}[\hat{y}] = \mathbf{MargRev}[\hat{y}] = \mathbf{E}[R'(\mathbf{q}) \hat{y}(\mathbf{q})] + R(0) \hat{y}(0).$$

⁹ Consistency with respect to downward-closure ([Lemma 3.4.2](#)) 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 $R'(1) \hat{Y}(1) = R'(1) Y(1)$.

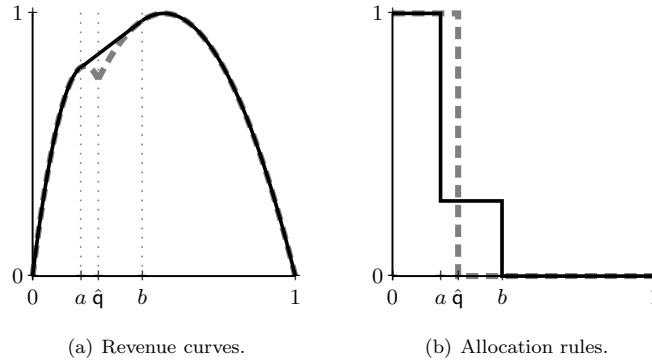


Figure 3.7. Depicted are the revenue curve $R(q)$, price-posting revenue curve $P(q)$, and the allocation rules corresponding to ex ante allocation constraint \hat{q} for the bimodal agent of [Example 3.3.2](#). For this agent the revenue curve $R(\cdot)$ (thin, black, solid line) is obtained from the 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 $\hat{y}(\cdot)$ 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} .

Observe that [Corollary 3.4.5](#) 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 3.4.6. *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 3.4.1](#)). 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 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 3.3.6](#) in terms of the revenue curve for the regular case.

Corollary 3.4.7. *For regular single-agent environments, allocation rule y has expected revenue equal to the surplus of marginal revenue $\mathbf{E}_q[R'(\mathbf{q})y(\mathbf{q})]$.*

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,

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

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 3.4.1](#). Optimal two-price lotteries (a) induce a concave revenue curve, (b) have at least the revenue of price posting, and (c) have allocation rules is no stronger than those of price posting. Consequently, via the exact same proof as [Theorem 3.4.4](#) (and [Corollary 3.4.5](#)) 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.

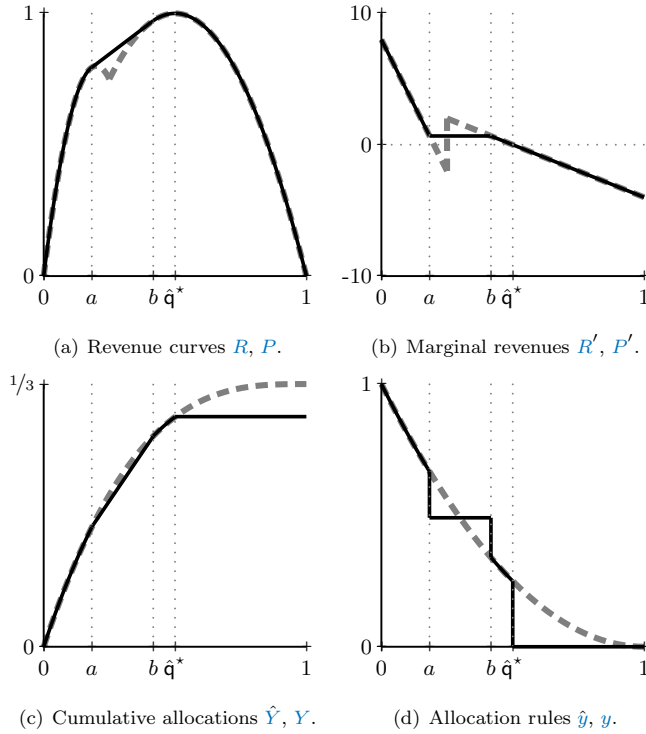


Figure 3.8. The optimal single-item auction is depicted for three bimodal agents [Example 3.3.2](#). The price-posting revenue curve $P(\cdot)$ is depicted by a thick, grey, dashed line in [Figure 3.8\(a\)](#). The revenue curve (thin, black, solid line) is its concave hull. The ironed interval (a, b) 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.

Theorem 3.4.8. *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., $[a, b]$ where $\mathbf{q} \in [a, b]$ satisfies $R''(\mathbf{q}) = 0$. 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 \mathbf{q} in some ironed interval $[a, b]$, resamples the quantile from this interval, i.e., as $y(\mathbf{q}) = \mathbf{E}_\sigma[\hat{y}(\sigma(\mathbf{q}))]$. 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 $\hat{\mathbf{q}}^*$ of the monopoly price; equivalently Y is flat after $\hat{\mathbf{q}}^*$. 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., $\hat{Y}(1) = Y(1)$ as discussed in [Section 3.4.3](#).

3.5 Surplus with a Balanced Budget

In this section we explore the role that a designer's budget constraint plays on mechanism design for the objective of surplus. Assume that the mechanism designer would like to maximize 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 in-

stead can be instead represented as a constraint on total payments which, because revenue is linear (Theorem 3.4.4), is a linear constraint.

Recall that with outcome (\mathbf{x}, \mathbf{p}) the 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 3.2.1) 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 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

$$\begin{aligned} \max_{\mathbf{x}, \mathcal{P}} \mathbf{E}_{\mathbf{v}} \left[\sum_i v_i x_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v})) \right] & \quad (3.5.1) \\ \text{s.t. mechanism } (\mathbf{x}, \mathcal{P}) \text{ is IC and IR} & \\ \mathbf{E}_{\mathbf{v}} \left[\sum_i p_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v})) \right] \geq 0 & \end{aligned}$$

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 (3.5.1) so as to be able to describe its solution. First, we will employ Proposition 3.3.6 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 surplus and profit (see Figure 3.9(a)). From this Pareto frontier we can see how to optimize surplus subject to a constraint on profit (such as budget balance) or optimize profit subject to a constraint on 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.

$$\begin{aligned} \max_{\mathbf{y}} \mathbf{E}_{\mathbf{q}} \left[\sum_i v_i(\mathbf{q}_i) \mathbf{y}_i(\mathbf{q}) - c(\mathbf{y}(\mathbf{q})) \right] \\ + \lambda \mathbf{E}_{\mathbf{q}} \left[\sum_i P'(\mathbf{q}_i) \mathbf{y}_i(\mathbf{q}) - c(\mathbf{y}(\mathbf{q})) \right] \\ \text{s.t. } \mathbf{y}(\cdot) \text{ is monotone.} \end{aligned} \quad (3.5.2)$$

Simplifying the objective with the identity (3.3.5) of $P'(\mathbf{q}) = \frac{d}{d\mathbf{q}}(\mathbf{q} v(\mathbf{q})) = v(\mathbf{q}) - \mathbf{q} v'(\mathbf{q})$, we have

$$\mathbf{E}_{\mathbf{q}} \left[\sum_i [(1 + \lambda) v_i(\mathbf{q}_i) + \lambda \mathbf{q} v'_i(\mathbf{q}_i)] \mathbf{y}_i(\mathbf{q}) - (1 + \lambda) c(\mathbf{y}(\mathbf{q})) \right].$$

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

$$\varphi_i^\lambda(\mathbf{q}) = (1 + \lambda) v_i(\mathbf{q}_i) + \lambda \mathbf{q} v'_i(\mathbf{q}_i). \quad (3.5.3)$$

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 $\varphi_i^\lambda(\cdot)$ is monotone non-increasing in \mathbf{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 3.3.3).

Lemma 3.5.1. *For a regular distribution (Definition 3.3.1 on page 66) given by inverse demand function $v(\cdot)$ and any non-negative Lagrangian*

parameter λ , the Lagrangian virtual value function $\varphi^\lambda(\mathbf{q}) = (1+\lambda)v(\mathbf{q}) + \lambda \mathbf{q} v'(\mathbf{q})$ is monotonically decreasing.

Proof. The Lagrangian virtual value function of equation (3.5.3) is a convex combination of the inverse demand curve $v(\cdot)$ and the marginal price-posting revenue curve $P'(\mathbf{q}) = v(\mathbf{q}) - \mathbf{q} v'(\mathbf{q})$, i.e., virtual values for revenue. The inverse demand curve is strictly decreasing by definition (Definition 3.3.6) and the marginal price-posting revenue curve is non-increasing by the regularity assumption (Proposition 3.3.7). 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 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 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 3.5.2. *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 3.5.3. *For regular general-costs environments, an Lagrangian virtual values from equation (3.5.3), there exists a Lagrangian parameter for which the virtual surplus maximization mechanism has zero expected*

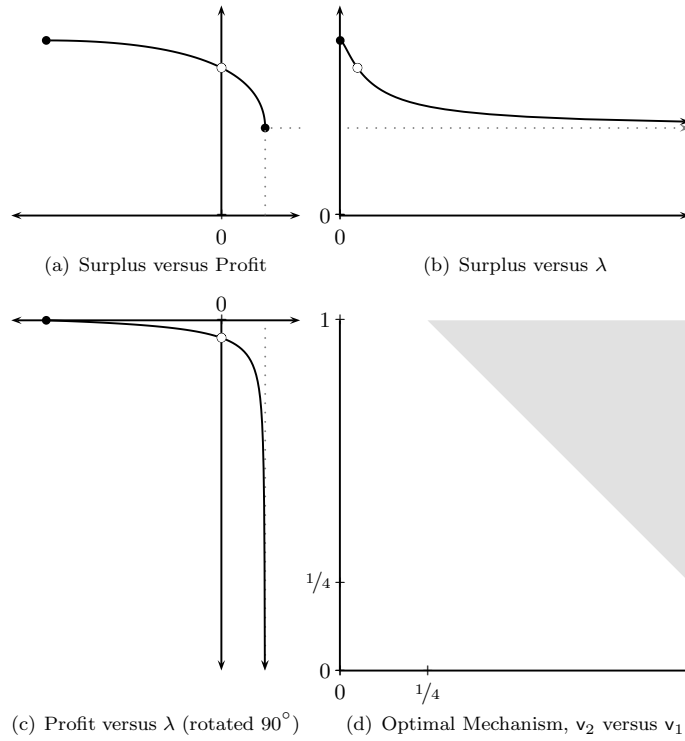


Figure 3.9. Depiction of the Pareto frontier for surplus (vertical axis) and profit (horizontal axis). On the Pareto frontier, the surplus maximizing point is profit minimizing (with negative profit) and the profit maximizing point is surplus minimizing. The surplus optimal point subject to budget balance 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.

profit and with this parameter the mechanism maximizes expected surplus subject to budget balance in dominant strategy equilibrium.

Example 3.5.1. *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 $\varphi(\mathbf{v}) = (2\lambda + 1)\mathbf{v} - \lambda$. 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

$$v_1 + v_2 \geq \frac{3\lambda + 1}{2\lambda + 1}. \quad (3.5.4)$$

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 (3.5.4) 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 $5/4$, the optimal Lagrangian parameter is $\lambda^* = 1/2$, and the 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 3.3.8](#)) 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 $\varphi(\cdot)$ and ironed virtual value function $\bar{\varphi}(\cdot)$ constructed for $\varphi(\cdot)$ for distribution F via the ironing procedure (Definition 3.3.8). 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 $\bar{\Phi}(\mathbf{q}) > \Phi(\mathbf{q})$ for $\mathbf{q} \in (a, b)$ and $\bar{\Phi}(\mathbf{q}) = \Phi(\mathbf{q})$ for $\mathbf{q} \in \{a, b\}$. 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 3.3.9). 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 3.5.4. *For general-cost environments, and Lagrangian virtual values from equation (3.5.3), 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 *consumer surplus* as the surplus less the payments, i.e., $\sum_i (v_i x_i - p_i) - c(\mathbf{x})$. (For more details, see the discussion of non-monetary payments in [Chapter 1](#))

Describe the mechanism that maximizes consumer 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 surplus maximizing outcome in equilibrium for any single-dimensional agent environment. Hint: Your mechanism may be parameterized by the distribution.
- 3.3 Derive equation [\(3.3.3\)](#),

$$\mathbf{E}_{v \sim F}[p(v)] = \mathbf{E}_{v \sim F}[\varphi(v) x(v)] \quad (3.3.3)$$

by taking expectation of the payment identity [\(3.3.1\)](#),

$$p(v) = v x(v) - \int_0^v x(z) dz, \quad (3.3.1)$$

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 [71](#)). 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 agents 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 3.5.1 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 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 surplus subject to budget balance is much more complex as the space of deterministic mechanisms is not convex ([Diakonikolas et al., 2012](#)).