

Chapter 4

Profit and Bayesian Optimality

In this chapter we consider the objective of *profit*. The objective of profit maximization adds significant new challenge over the previously considered objective of social surplus maximization. Fundamentally, where for social surplus there is always a single optimal mechanism (absent computational constraints), for profit there is no single optimal mechanism. For any mechanism, there is a setting 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. In this text we will consider two approaches. This chapter focuses on 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 their expected profit when preferences are indeed drawn from the distribution. In Chapter 6 we will consider the traditional in computer science approach. We will look for a single mechanism that is good in any setting. We know that it is not possible to be optimal, so instead the designer should look for the mechanism is the best approximation to the optimal one, for instance, to the Bayesian optimal one that follows from the aforementioned economics approach.

As in the setting of social surplus we will be considering general single-dimensional agent settings, i.e., each agent has a single value for service and the designer has a cost function over the sets of served agents (see Section 3.1). The profit of the mechanism with outcome \mathbf{x} and payments \mathbf{p} is $\text{Profit}(\mathbf{p}, \mathbf{x}) = \sum_i p_i - c(\mathbf{x})$ (Definition 3.5). Here the setting is given by $c(\cdot)$ and distribution assumption we may make on the values \mathbf{F} . For general feasibility settings, where the designer's cost for any feasible outcome is zero, we will refer to the profit as *revenue*.

A motivating example was given in Chapter 1 with a single-item auction with two agents whose values are drawn independently and identically from $U[0, 1]$. We calculated the expected revenue of the second-price auction and determined it was simply the expectation of the lower value, i.e, $1/3$. We then calculated the expected revenue of the second-price auction with reserve price $1/2$ and determined that it was $5/12$. We concluded that setting a reserve price can improve the revenue of an auction; however, we did not solve for the optimal auction. We show that the second-price auction with reserve $1/2$ is indeed optimal for

this two bidder example and furthermore we give a concise characterization of the optimal auction for any single-dimensional agent setting.

4.1 Revenue Curves

We start by removing all the complication of mechanisms for multiple agents and consider only a single agent, Alice, desiring a single item. Suppose Alice's value v is drawn from distribution F . How should we sell the item to Alice to maximize our profit?

A first step would be to consider offering Alice a take-it-or-leave-it price p . What does Alice do with such a price? If Alice's value $v \geq p$ then Alice buys the item and pays p , otherwise Alice does not buy the item and pays zero. Therefore, we can calculate the expected revenue as p times the probability that Alice's value exceeds p , which can be read directly from the density function as $1 - F(p)$. Our revenue as a function of p is $p \cdot (1 - F(p))$.

Definition 4.1 The *revenue curve* $B(p)$ is the revenue obtained from agent with value distribution F from a posted price p as a function of that price, i.e., $B(p) = p \cdot (1 - F(p))$.

We can clearly optimize this by taking the derivative and setting it equal to zero. If F is $U[0, 1]$ then $F(p) = p$ and the revenue is optimized at price $p = 1/2$ where an expected revenue of $1/4$ is achieved. The uniform distribution is well-behaved in the sense that the revenue, as a function of price, increases up to the optimal price of $1/2$ and then decreases. The importance of the derivative in solving for the optimal price can be noted by observing that the derivative is negative but increasing as p is lowered to $p = 1/2$, where it is zero, and then continues to be positive and increasing afterwards. This optimal pricing rule is allocating where the derivative of this revenue curve is negative!

4.2 Expected Revenue and Virtual Values

Suppose we are given the allocation rule of an agent (Alice) as $x(v)$. By the payment identity, the payment rule must be $p(v) = vx(v) - \int_0^1 x(z)dz$. Since v is drawn from F we can calculate Alice's ex ante expected payment.

$$\mathbf{E}_{v \sim F}[p(v)] = \int_0^\infty \left[vx(v) - \int_0^v x(z)dz \right] dv.$$

We could easily analyze this expression by noting the double integral, swapping order of integration, and simplifying. However, such a treatment, though formulaically correct, gives less intuition for the relevant quantities.

The following is a brief digression into integration by parts. Integrate the product rule for differentiation, e.g., with $g'(z)$ denoting the derivative of $g(z)$ with respect to z .

$$g(z)h(z)|_a^b = \int_a^b g'(z)h(z)dz + \int_a^b g(z)h'(z)dz.$$

Integration by parts is then usually formulated by rearranging slightly as follows:

$$g(z)h(z)|_a^b - \int_a^b g(z)h'(z)dz = \int_a^b g'(z)h(z)dz.$$

Notice that our payment identity looks like the left-hand side of this last formula with $g(z) = x(z)$ and $h(z) = z$. Therefore $p(v) = \int_0^v zx'(z)dz$. This latter formula is sometimes more convenient than the former.

We now make a concrete connection between the derivative of the revenue curve and the expected revenue of a mechanism.

Definition 4.2 The *virtual value* of an agent with value $v \sim F$ is:¹

$$\phi(v) = -\frac{B'(v)}{f(v)} = v - \frac{1-F(v)}{f(v)}.$$

Lemma 4.3 For allocation rule $x(\cdot)$ and $v \sim F$, $\mathbf{E}[p(v)] = \mathbf{E}[\phi(v)x(v)]$.

Proof: Consider the following analysis, where we use the definition of expectation, swap the order of integration, use the definition of the cumulative distribution function, use the definition of the revenue curve, integrate by parts, plug in $B(0) = B(\infty) = 0$, use the definition of expectation, and finally use the definition of virtual value.

$$\begin{aligned} \mathbf{E}_v[p(v)] &= \int_0^\infty \int_0^v zx'(z)dz f(v)dv \\ &= \int_0^\infty zx'(z) \int_z^\infty f(v)dv dz \\ &= \int_0^\infty zx'(z)(1 - F(z))dz \\ &= \int_0^\infty B(z)x'(z)dz \\ &= B(z)x(z)|_0^\infty - \int_0^\infty B'(z)x(z)dz \\ &= -\int_0^\infty \frac{B'(z)}{f(z)}x(z)f(z)dz \\ &= -\mathbf{E}_v\left[\frac{B'(v)}{f(v)}x(v)\right] \\ &= \mathbf{E}_v[\phi(v)x(z)]. \end{aligned}$$

■

It should be no surprise, as shown by the above lemma, that the expected revenue can be written entirely as a function of the allocation rule. What is perhaps surprising is that this expected revenue is just the expectation of the product of the allocation rule with the so-called virtual value of the agent.

¹As we will discuss later $B'(v)/f(v)$ is exactly the derivative of the revenue curve in “probability space”.

Definition 4.4 The *virtual surplus* of outcome \mathbf{x} and values \mathbf{v} is:

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

The following theorem follows from Lemma 4.3 and linearity of expectation.

Theorem 4.5 *The expected revenue of a mechanism is equal to its expected virtual surplus,*

$$\mathbf{E}_{\mathbf{v}} \left[\sum_i \phi_i(v_i)x_i(\mathbf{v}) - c(\mathbf{x}(\mathbf{v})) \right].$$

4.3 Optimal Mechanisms and Regular Distributions

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

The the non-game-theoretic optimization problem of maximizing virtual surplus is that of finding \mathbf{x} to maximize $\text{Surplus}(\phi(\mathbf{v}), \mathbf{x}) = \sum_i \phi_i(v_i)x_i - c(\mathbf{x})$.² Let OPT again be the surplus maximizing algorithm. We will care about both the allocation that $\text{OPT}(\phi(\mathbf{v}))$ selects, i.e., $\text{argmax}_{\mathbf{x}} \text{Surplus}(\phi(\mathbf{v}), \mathbf{x})$ and its virtual surplus $\text{max}_{\mathbf{x}} \text{Surplus}(\phi(\mathbf{v}), \mathbf{x})$. Where it is unambiguous we will use notation $\text{OPT}(\phi(\mathbf{v}))$ to denote either of these quantities. Recall that the formulation of OPT has no mention of the incentive constraints.

We know from the BIC characterization that the allocation rule of any BIC is monotone. Thus, the mechanism design problem of finding a BIC mechanism to maximize virtual surplus has an added monotonicity constraint. Even though we did not impose a monotonicity constraint on OPT, if the virtual valuation functions $\phi_i(\cdot)$ are monotone, $\text{OPT}(\phi(\mathbf{v}))$ is monotone.

Definition 4.6 Distribution F is *regular* if $\phi(\cdot)$ is non-decreasing.

Lemma 4.7 *For each agent i and any values of other agents \mathbf{v}_{-i} , if F_i is regular then i 's allocation rule from OPT on virtual values is monotone in i 's value v_i .*

Proof: Recall from Lemma 3.6 that maximizing surplus is monotone. Meaning, if we find \mathbf{x} to maximize $\text{Surplus}(\mathbf{v}, \mathbf{x})$ then $x_i(\mathbf{v}_{-i}, v_i)$ is monotone in v_i . Therefore $x_i(\phi_{-i}(\mathbf{v}_{-i}), \phi_i(v_i))$ is monotone in $\phi_i(v_i)$, i.e., increasing $\phi_i(v_i)$ does decrease x_i . By the regularity assumption on F_i , $\phi_i(v_i)$ is monotone in v_i . Therefore, increasing v_i cannot decrease $\phi_i(v_i)$ which cannot decrease x_i . ■

²Here the shorthand notation $\phi(\mathbf{v}) = (\phi_1(v_1), \dots, \phi_n(v_n))$.

Since OPT on virtual values is monotone for each agent and any values of other agents it satisfies our strongest incentive constraint. With the appropriate payments (i.e., the “critical values”) truth-telling is a dominant strategy equilibrium (recall Corollary 2.8). The resulting mechanism is due to Roger Myerson.

Mechanism 4.1 *The Myerson mechanism for regular distributions is:*

1. *Solicit and accept sealed bids \mathbf{b} ,*
2. *$(\mathbf{x}, \mathbf{p}') \leftarrow \text{VCG}(\phi(\mathbf{b}))$, and*
3. *for each i , $p_i \leftarrow \phi_i^{-1}(p'_i)$.*

Notice that the payments \mathbf{p} calculated can be viewed as the following. VCG on virtual values outputs virtual prices \mathbf{p}' . These correspond to the minimum virtual value an agent must have to win. The Myerson mechanism then applies the inverse virtual valuation function to determine the minimum value said agent must have to win.³

Theorem 4.8 *The Myerson mechanism for regular distributions is IC.*

Corollary 4.9 *The Myerson mechanism for regular distributions maximizes expected revenue.*

It is quite useful to view this result as a reduction from the problem of profit maximization to the problem of surplus maximization. As above, we used the VCG mechanism to construct Myerson’s mechanism. This general reduction applies to IC worst-case surplus approximation mechanisms as well.

Theorem 4.10 *For any IC mechanism \mathcal{M} that gives a β -approximation to the optimal social surplus, $\mathcal{M}(\phi(\cdot))$, with payments computed via inverse virtual valuations, is a β -approximation mechanism to the optimal expected profit.*

Notice that there is not a similar statement for BIC mechanisms and Bayesian approximation as BIC and Bayesian approximation requires the input to the mechanism be $\mathbf{v} \sim \mathbf{F}$. The distribution of virtual values is will not be the same as the distribution of values and therefore there is no implied guarantee of monotonicity nor expected revenue of the mechanism that results from the above construction.

4.4 Single-item Auctions

The above description of profit-optimal 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

³Assuming virtual valuations are strictly non-decreasing then the inverse virtual valuations are well defined. We defer discussion of the non-strict case to the subsequent section on irregular distributions.

auctions, i.e., settings where feasible outcomes serve at most one agent. So what is the mechanism that optimizes virtual surplus for single-item settings?

First notice that virtual values can be negative. Consider the uniform distribution $U[0, 1]$ where $F(z) = z$ and $f(z) = 1$. Here $\phi(v) = v - \frac{1-F(z)}{f(z)} = 2v - 1$. Thus, $\phi(0) = -1$. If our goal is to optimize virtual surplus we clearly do not want to allocate to any agent with negative virtual value. Recall that virtual values are the negative derivative of the revenue curve (normalized by the density function) and our analysis of the single-agent setting already suggested that we should not allocate to an agent for whom this quantity is negative.

Second notice that among the agents with positive virtual values the virtual surplus is maximized by allocating to the one with the highest virtual value. Conclude the following corollary.

Corollary 4.11 *For regular distributions, the optimal single-item auction allocates to the agent with the highest non-negative virtual valuation.*

As virtual valuations are the negative derivative of the revenue curve (normalized by the density function) this mechanism says to allocate to the agent whose revenue curve is the steepest at their value.

The case where the agents are independent and identically distributed is of special interest. Here the agent with the highest positive virtual value is also the one with the highest value (as the virtual valuation functions are identical). An agent's virtual value is non-negative when their value is at least $\phi^{-1}(0)$. What auction allocates to the agent with the highest value that is at least $\phi^{-1}(0)$? It is the second price auction with reserve $\phi^{-1}(0)$!

Corollary 4.12 *For regular i.i.d. distributions, F , the second-price auction with reserve $\phi^{-1}(0)$ is the single-item auction with the highest expected revenue.*

We conclude by returning to our two agent $U[0, 1]$ example. As we have calculated, $\phi(v) = 2v - 1$; therefore, $\phi^{-1}(0) = 1/2$. The second-price auction with reserve price $1/2$ has the optimal expected revenue. Our previous calculation showed that this revenue was $5/12$.

While this auction is optimal among BIC auctions, which is the class of mechanisms we restricted our attention to, the revelation principle implies that no auction has a BNE with higher expected revenue. Therefore, we conclude that in a very strong sense, that the second price auction with reserve price maximizes revenue.

4.5 Irregular Distributions and Ironed Virtual Values

We again turn our attention to the case where the non-game-theoretic optimization problem is not itself inherently monotone. An *irregular* distribution is one for which the virtual valuation functions are not monotone non-decreasing, i.e., where a higher value might result in a lower virtual value. Clearly $\text{OPT}(\phi(\cdot))$ is non-monotone for such settings.

4.5.1 Quantile Space

In Chapter 3 we saw that monotoning a non-monotone allocation rule was more intuitive in *quantile space* than value space. We will refer to the ex ante probability at which an agent with value $v \sim F$ accepts an offered price as the *quantile* of the price. As before, we will assume that the distribution function is continuous and therefore the distribution function $F(z)$ has a unique inverse, denoted $F^{-1}(z)$. The transformation from value space into *quantile space* is specified by the formula $q = 1 - F(v)$.

In the preceding section we wrote the expected payment as a function of the allocation rule and the revenue curve. Define $y(q) = x(F^{-1}(1 - q))$ as the allocation rule in quantile space. Define $R(q) = q \cdot (F^{-1}(1 - q)) = B(F^{-1}(1 - q))$ as the revenue curve in quantile space. Notice that increasing quantile corresponds to decreasing value therefore the derivatives in quantile and value space have opposite sign. The derivative of the revenue curve in quantile space is exactly the negative derivative of the revenue curve in value space normalized by the density function, i.e., $R'(q) = -B'(v)/f(v)$ for $q = 1 - F(v)$. Likewise derivative of the allocation rule in quantile space is exactly the negative derivative of the allocation rule in value space normalized by the density function, i.e., $y'(q) = -x'(v)/f(v)$ for $q = 1 - F(v)$.

Lemma 4.13 *For allocation rule $y(\cdot)$ and $v \sim F$,*

$$\mathbf{E}_v[p(v)] = \mathbf{E}_q[R'(q)y(q)] = -\mathbf{E}_q[R(q)y'(q)].$$

This lemma follows directly from the proof of Lemma 4.3 where the second equality follows from the definition of expectation and integration by parts. This alternative view point is useful as it immediately implies the following corollary.

Corollary 4.14 *For allocation rule $y(\cdot)$ and values $v_1 \sim F_1$ and $v_2 \sim F_2$, if $R_1(q) \geq R_2(q)$ for all q then $\mathbf{E}_{v_1}[p(v_1)] \geq \mathbf{E}_{v_2}[p(v_2)]$.*

Notice that virtual valuations are simply the derivative of the revenue curve in quantile space, i.e., $\phi(v) = R'(q)$ for $q = 1 - F(v)$ (Definition 4.2).

4.5.2 Ironed Revenue Curves

Recall the discussion of non-monotonicity in Chapter 3; if we treat Alice the same regardless of her value when her quantile is on some interval $[a, b]$ then we can replace her exact virtual valuation with her average virtual valuation on this interval. Figure 4.1(a) depicts a hypothetical non-concave revenue curve; the corresponding virtual value function, i.e., its derivative, is depicted in Figure 4.1(c). Figure 4.1(d) shows Alice's virtual value averaged on $[a, b]$. Finally, Figure 4.1(b) shows the resulting revenue curve. Notice that the constant virtual valuation over $[a, b]$ results in a linear revenue curve, specifically, the line segment connecting $(a, R(a))$ to $(b, R(b))$. This process of treating Alice the same on an interval to flatten the virtual valuation function is known as *ironing*.

It should be intuitively clear that if we restrict ourselves to allocation rules that treat Alice the same on appropriate subintervals of quantile space we can construct an effective

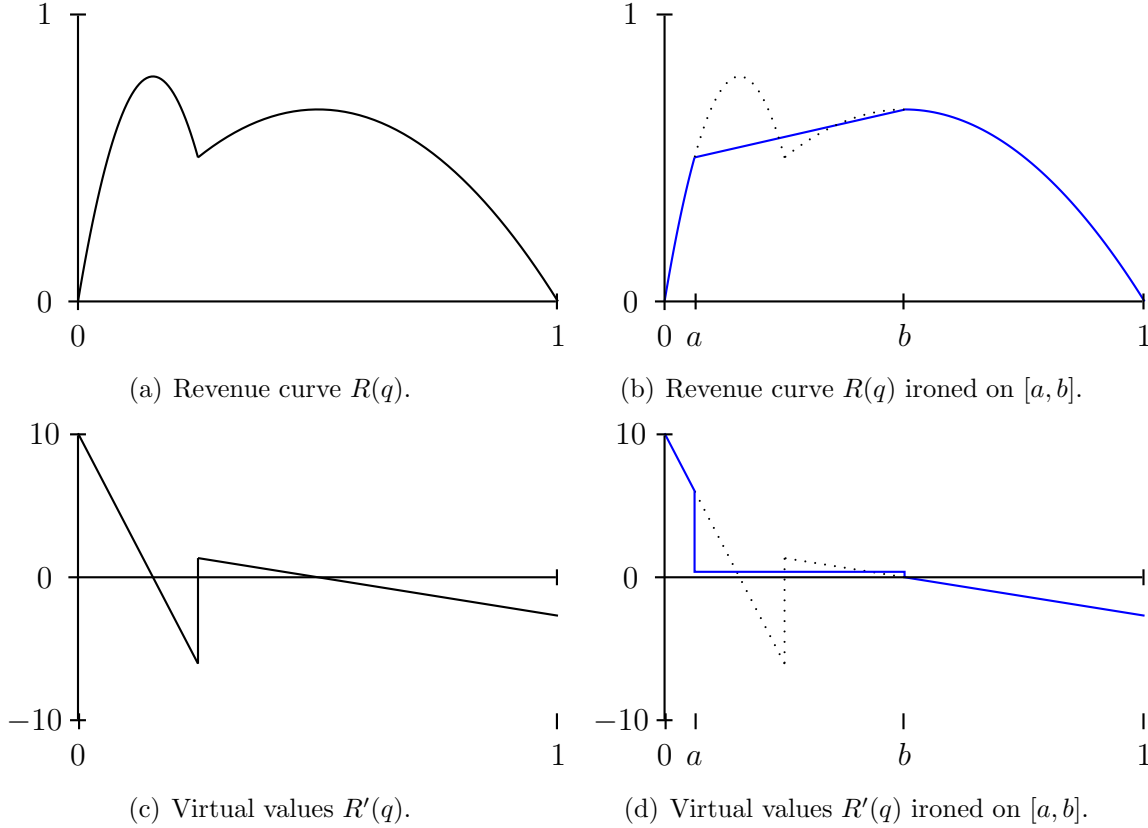


Figure 4.1: On the left is the revenue curve $R(q)$ and virtual valuations $R'(q)$ in quantile space. On the right is the effective revenue curve and virtual valuations when ironed on $[a, b]$. Though it is not necessary for understanding this example, this $R(\cdot)$ comes from bimodal distribution that is $U[0, 2]$ with probability $3/4$ and $U[2, 8]$ with probability $1/4$.

revenue curve $\bar{R}(\cdot)$ equal to the concave hull of the actual revenue curve $R(\cdot)$. This revenue curve is known as the *ironed revenue curve* and its derivative is the *ironed virtual valuation* function.

Definition 4.15 for $v \sim F$, the *ironed revenue curve*, $\bar{R}(\cdot)$, is the concave hull of $R(\cdot)$ and the *ironed virtual valuation* function is $\bar{\phi}(v) = \bar{R}'(q)$ for $q = 1 - F(v)$.

Notice the advantage of $\bar{R}(\cdot)$ over $R(\cdot)$ is two-fold. First, Lemma 4.14 suggests that we can get more revenue from $\bar{R}(\cdot)$ than from $R(\cdot)$. Second, $\bar{R}(\cdot)$ is concave by definition, so ironed virtual valuations are monotone, so ironed virtual surplus maximization results in a monotone allocation rule, so with the appropriate payment rule it is incentive compatible.

It may perhaps seem strange that for revenue curve $R(\cdot)$ somehow we are able to get more revenue than this revenue curve suggests, i.e., via the ironed revenue curve $\bar{R}(\cdot)$. In fact this is an artifact of our formulaic definition of the revenue curve as $R(q) = qF^{-1}(1 - q)$. Instead if look for optimal revenue curves see immediately that such curves are inherently concave.

A revenue curve should give the expected revenue as a function of a fixed probability of sale. In our definition of $R(\cdot)$ we assumed that this comes from posting the price that corresponds to this fixed probability of sale. There is another way we can sell with fixed probability \hat{q} : pick some interval $[a, b]$ with $a < \hat{q} < b$ and consider the allocation rule

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

Notice that when Alice's quantile q is realized (i.e., drawn from the uniform distribution) then the probability that Alice is served by $y^{\hat{q}}(\cdot)$ is $1 \times a + \frac{\hat{q}-a}{b-a} \times (b-a) = \hat{q}$. The revenue from such an allocation rule follows directly from Lemma 4.13. It is $R(a) + \frac{\hat{q}-a}{b-a}(R(b) - R(a))$. Notice that this is exactly the value at \hat{q} on the line segment connecting $(a, R(a))$ to $(b, R(b))$. Again this can be seen in Figure 4.1(b). Where $R(\cdot)$ is non-concave, this revenue can be higher than $R(\hat{q})$. Meaning, if we want to sell to Alice with ex ante probability \hat{q} there is a better way to do it than offer price $p = F^{-1}(1 - \hat{q})$; instead use allocation rule $y^{\hat{q}}(\cdot)$, above.

4.5.3 Optimal Mechanisms

So far we have said seen that the ironed revenue curve dominates the revenue curve. Furthermore, if our allocation rule treats agents the same in the ironed intervals then its virtual surplus equals its ironed virtual surplus. However, we still need to argue that allocation rules restricted in this manner are indeed optimal. Therefore, we need to relate the expected payment to $\bar{R}'(\cdot)$ for any (unrestricted) monotone allocation rule $y(\cdot)$.

Lemma 4.16 *For any monotone allocation rule $y(\cdot)$ and $v \sim F$,*

$$\mathbf{E}_v[p(v)] = \mathbf{E}_q[\bar{R}'(q)y(q)] + \mathbf{E}_q[(\bar{R}(q) - R(q)) \cdot y'(q)].$$

Proof:

$$\begin{aligned} \mathbf{E}_v[p(v)] &= \mathbf{E}_q[R'(q)y(q)] + \mathbf{E}_q[\bar{R}'(q)y(q)] - \mathbf{E}_q[\bar{R}'(q)y(q)] \\ &= \mathbf{E}_q[\bar{R}'(q)y(q)] - \mathbf{E}_q[(\bar{R}'(q) - R'(q)) \cdot y(q)] \\ &= \mathbf{E}_q[\bar{R}'(q)y(q)] + \mathbf{E}_q[(\bar{R}(q) - R(q)) \cdot y'(q)]. \end{aligned}$$

■

This lemma gives a concrete suggestion for how to proceed. Payment is equal to the expected ironed virtual value (first term) plus a second term. Inspecting the second term, $(\bar{R}(q) - R(q)) \cdot y'(q)$, more closely, notice that the difference in the revenue curves is non-negative, as $\bar{R}(\cdot)$ is the concave hull of $R(\cdot)$; and the derivative of the allocation rule is non-positive, as the allocation rule is monotone decreasing in quantile. Therefore, the second term is non-positive and maximizing it is equivalent to minimizing its magnitude. As we see in the lemma below, if we ignore the second term and maximize the first term then the second term will be zero, and therefore it is also maximized. We can conclude that maximizing ironed virtual surplus is optimal.

Lemma 4.17 For $v \sim F$ and monotone allocation rule $y(\cdot)$ satisfying $\bar{R}''(q) = 0 \Rightarrow y'(\cdot) = 0$ then,

$$\mathbf{E}_v[p(v)] = \mathbf{E}_q[\bar{R}'(q)y(q)].$$

Proof: We must argue that the assumption on the allocation rule implies that the second term of Lemma 4.16, $\mathbf{E}_q[[\bar{R}(q) - R(q)]y'(q)]$, is zero.

Notice that $\bar{R}(q) - R(q) > 0$ implies that $\bar{R}(\cdot)$ at q is on a line segment connecting points $a < q < b$ on $R(q)$. Therefore, $\bar{R}'(q)$ is constant, and $\bar{R}''(q)$, the second derivative of the ironed revenue curve, is zero. Our assumption then implies that $y'(q) = 0$. We conclude that if the first multiplicand of the expectation is non-zero, then the second is zero; therefore, the expectation is always identically zero. ■

Consider the optimization function given by $\text{OPT}(\bar{\phi}(\mathbf{v}))$, i.e., optimizing the *ironed virtual surplus*. Each $\bar{\phi}_i(\cdot)$ is monotone since it is the derivative of a concave function. This is good news because it means we can construct an incentive compatible mechanism from $\text{OPT}(\bar{\phi}(\cdot))$. Referring back to Lemma 4.16, by definition we have maximized the first term (i.e., the ironed virtual surplus). Coincidentally, we have also maximized the second term. To see this, notice that the input to the optimization is ironed virtual values; therefore, the resulting allocation rule must be constant on intervals of value (or quantile) where the ironed virtual values are constant. Therefore, the assumption of Lemma 4.17 is satisfied and the second term (of Lemma 4.16) is zero, i.e., at its maximum. We conclude that allocation rule of $\text{OPT}(\bar{\phi}(\cdot))$ maximizes expected profit. This argument proves the following lemma.

Lemma 4.18 For any distribution \mathbf{F} and any general single-dimensional agent setting, $\text{OPT}(\bar{\phi}(\cdot))$ maximizes virtual surplus subject to monotonicity.

Since OPT on ironed virtual values is monotone for each agent and all values of other agents, it satisfies our strongest incentive constraint. With the appropriate payments (i.e., the “critical values”) truth-telling is a dominant strategy equilibrium (recall Corollary 2.8).

Mechanism 4.2 The Myerson mechanism, $\text{Mye}_{\mathbf{F}}$, for product distribution \mathbf{F} is:

1. Solicit and accept sealed bids \mathbf{b} ,
2. $(\mathbf{x}, \mathbf{p}') \leftarrow \text{VCG}(\bar{\phi}(\mathbf{b}))$, and
3. calculate payments for each agent from the payment identity.

Unlike VCG and the Myerson mechanism for regular distributions (with strictly increasing virtual valuation functions) where the continuity assumption on the distribution implies that there is never a tie, the Myerson mechanism for irregular distributions may require a tie-breaking policy. Tie breaking can be done arbitrarily (as long as it is not a function of the agents’ values). Common tie-breaking rules are *lexicographical* and *random*. Lexicographical tie breaking will favor sets of agents with higher indices. Random tie breaking takes the lexicographical ordering on a random permutation of the agent indices. The randomized tie-breaking rule is nice because it is symmetric.

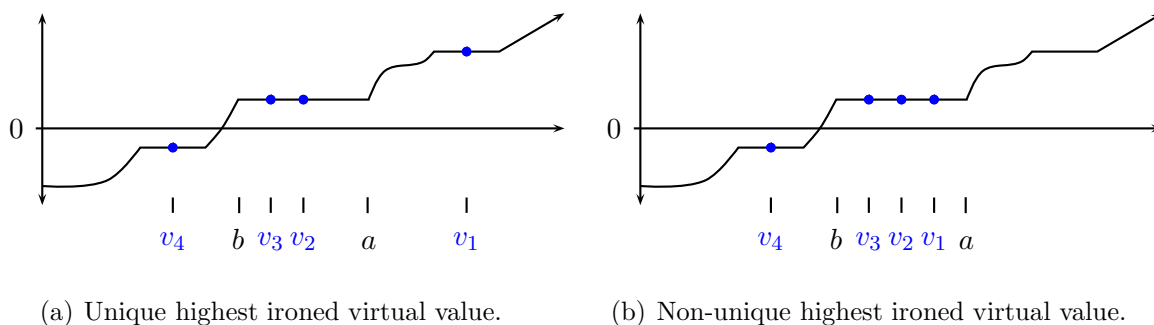


Figure 4.2: The ironed virtual valuation function $\bar{\phi}(v)$ under two realizations of agent values depicting both case that where the highest ironed virtual value is unique and the case where it is not unique.

Theorem 4.19 *The Myerson mechanism is IC.*

Corollary 4.20 *The Myerson mechanism maximizes expected revenue.*

Like in the regular case, it is quite useful to view this result as a reduction from the problem of profit maximization to the problem of surplus maximization. As above, we used the VCG mechanism to construct the Myerson mechanism. This general reduction applies to IC worst-case surplus approximation mechanisms as well.

Theorem 4.21 *For any IC mechanism \mathcal{M} that gives a β -approximation to the optimal social surplus, $\mathcal{M}(\bar{\phi}(\cdot))$, with appropriate payments, is a β -approximation mechanism to the optimal expected profit.*

4.5.4 Single-item Auctions

We consider the special case of single-item auctions to get a clearer picture of exactly what this optimal mechanism is in the case of i.i.d. irregular distributions. Figure 4.2 depicts hypothetical ironed virtual valuation function. Instantiating the agents' values corresponds to picking points on the x -axis. The agents' ironed virtual valuations can then be read off the plot. The Myerson auction assigns the item to the agent with the highest ironed virtual value. If there is a tie, it picks a random tied agent to win.

Figure 4.2(a) depicts a realization of values for $n = 4$ agents where the highest ironed virtual valuation is unique. What does Myerson do here? Myerson allocates the item to this agent, i.e., agent 1 in the figure. Figure 4.2(b) depicts a second realization of values where the highest ironed virtual valuation is not unique. In this setting Myerson, we will assume, breaks ties by picking a random tied agent as the winner, i.e., one of agents 1, 2, and 3 in the figure. In general when there is a k -agent tie for the highest ironed virtual valuation then each tied agent wins with probability $1/k$.

We now calculate the payments. Consider the case where there is a unique highest ironed virtual value. The agent with this ironed virtual value wins. To calculate their IC payment

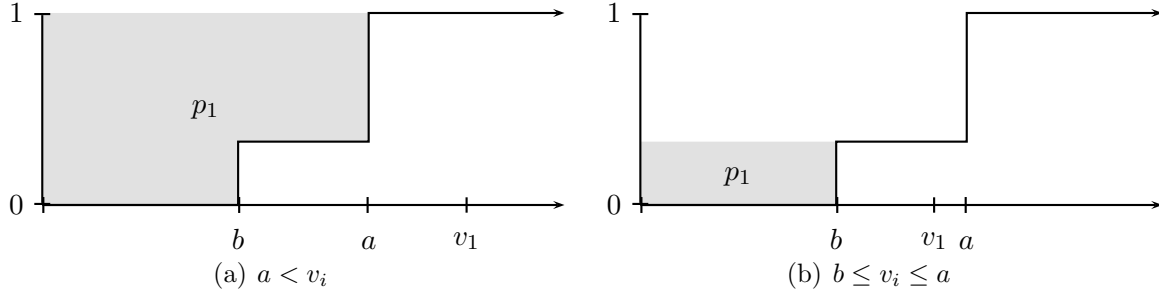


Figure 4.3: The allocation (black line) and payment rule (gray region) for agent 1 given \mathbf{v}_{-i} and the ironed virtual valuation function from Figure 4.2.

we need to consider agent i 's allocation rule for fixed values \mathbf{v}_{-i} of the other agents. Consider again the example in Figure 4.2(a) and imagine the probability we allocate to agent 1 as a function of v_1 . This is

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

when \mathbf{v}_{-i} has a $k - 1$ agents in $[b, a]$ tied for the highest ironed virtual valuation. The $1/k$ probability of winning for $z \in [b, a]$ arises from our analysis of what happens when in a k -agent tie. Figure 4.3(a) depicts the allocation and rule payment of this agent. When agent 1 has the unique highest ironed virtual value, i.e., $v_1 > a$ then clearly $p_1 = a - (a - b)/k$.

When agent 1 is tied for the highest ironed virtual value with $k - 1$ other agents, as depicted in Figure 4.3(b), their expected payment is $p_1 = b/k$. Of course, $x_1 = 1/k$ so such a payment can be implemented by charging b to the tied agent that wins.

4.6 Notes

The optimal single-item auction was derived by Myerson [18]. Its generalization to single-dimensional agent settings is an obvious extension. The relationship between optimal auctions, revenue curves, and *marginal revenue* (equivalent to virtual values) is due to Bulow and Roberts [8].