Chapter 3

Optimal Mechanisms

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

The objective of profit maximization, on the other hand, adds significant new challenge: for profit there is no single optimal mechanism. For any mechanism, there is a distributional 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. 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 we could 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, $\mathbf{E}\left[v_{(2)}\right] = 1/3$. 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 $1/3$). Perhaps surprisingly, a seller can make 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 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. Agents are single-dimensional, i.e., each has a single private value for receiving some abstract service and quasi-linear utility, i.e., the agent's utility is her value for the service less her payment. Recall that the outcome of a single-dimensional game is an allocation $\mathbf{x} = (x_1, \ldots, x_n)$, where x_i is an indicator for whether agent i is served, and payments $\mathbf{p} = (p_1, \ldots, p_n)$, where p_i is the payment made by agent i. Here we formalize the designer's constraints and objectives.

Definition 3.1. *A* general cost environment *is one where the designer must pay a service cost $c(\mathbf{x})$ for the allocation $\mathbf{x}$ produced.*

Definition 3.2. *A* general feasibility environment *is one where there is a feasibility constraint over the set of agents that can be simultaneously served.*

Definition 3.3. *A* downward-closed *feasibility constraint is one where subsets of feasible sets are feasible.*

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

The two most fundamental designer objectives are social surplus, a.k.a., social welfare,[1] and profit.

Definition 3.4. *The* social surplus *of an allocation is the cumulative value of agents served less the service cost:*

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

Definition 3.5. *The* profit *of allocation and payments is the cumulative payment of agents less the service cost:*

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

Implicit in the definition of social surplus is the fact that the payments from the agents are transferred to the service provider and therefore do not affect the objective.[2]

[1]A mechanism that optimizes social surplus is said to be *economically efficient*; though, we will not use this terminology because of possible confusion with *computational efficiency*.

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

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

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

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

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

We have yet to see any examples of general cost environments. One natural one is that of a *multicast auction*. The story for this problem comes from live video steaming. Suppose we wish to stream live video to viewers (agents) in a computer network. Because of the high-bandwidth nature of video streaming the content provider must lease the network links. Each link has a publicly known cost. To serve a set of agents, the designer must pay the cost of network links that connect each agent, located at different nodes in the network, to the "root", i.e., the origin of the multicast. The nature of multicast is that the messages need only be transmitted once on each edge to reach the agents. Therefore, the total cost to serve these agents is the minimum cost of the *multicast tree* that connects them.[3]

3.2 Social Surplus

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

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

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

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

Lemma 3.6. *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 OPT($\mathbf{v}$) or i is not served by OPT($\mathbf{v}$). We write out the surplus of OPT in both of these cases. Below, notation $(\mathbf{v}_{-i}, z)$ denotes the vector $\mathbf{v}$ with the ith coordinate replaced with z.

Case 1 ($i \in$ OPT):

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

Define $\mathrm{OPT}_{-i}(\mathbf{v})$ as the second term on the right hand side. Thus,

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

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

Case 2 ($i \notin$ OPT):

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

Define $\mathrm{OPT}(\mathbf{v}_{-i})$ as the term on the right hand side. Thus,

$$\mathrm{OPT}(\mathbf{v}) = \mathrm{OPT}(\mathbf{v}_{-i}).$$

Notice that $\mathrm{OPT}(\mathbf{v}_{-i})$ 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

$$v_i + \mathrm{OPT}_{-i}(\mathbf{v}) \geq \mathrm{OPT}(\mathbf{v}_{-i}).$$

Solving for v_i we conclude that OPT allocates to i whenever

$$v_i \geq \mathrm{OPT}(\mathbf{v}_{-i}) - \mathrm{OPT}_{-i}(\mathbf{v}).$$

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 $\tau_i = \mathrm{OPT}(\mathbf{v}_{-i}) - \mathrm{OPT}_{-i}(\mathbf{v})$. □

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 2.18). The resulting surplus maximization mechanism is often referred to as the *Vickrey-Clarke-Groves* (VCG) mechanism, named after William Vickrey, Edward Clarke, and Theodore Groves.

Mechanism 3.1. *The* surplus maximization (SM) *mechanism is:*

1. *Solicit and accept sealed bids* $\mathbf{b}$.

2. $\mathbf{x} \leftarrow \mathrm{OPT}(\mathbf{b})$*, and*

3. *for each* i*,* $p_i \leftarrow \mathrm{OPT}(\mathbf{b}_{-i}) - \mathrm{OPT}_{-i}(\mathbf{b})$.

An intuitive description of $\mathrm{OPT}(\mathbf{v}_{-i}) - \mathrm{OPT}_{-i}(\mathbf{v})$ is the *externality* that agent i imposes on the other agents by being served. I.e., because i is served the other agents obtain total surplus $\mathrm{OPT}_{-i}(\mathbf{v})$ instead of the surplus $\mathrm{OPT}(\mathbf{v}_{-i})$ that they would have received if i was not served. Hence, we can interpret the surplus maximization mechanism as serving agents to maximize the social surplus and charging each agent the externality imposed on the others.

By Corollary 2.18 and Lemma 3.6 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 3.7. *The surplus maximization mechanism is dominant strategy incentive compatible.*

Corollary 3.8. *The surplus maximization mechanism optimizes social surplus in dominant strategy equilibrium.*

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

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

Of course, 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., $p_i(\mathbf{v}_{-i}, 0)$ for agent i. For instance $p_i = \mathrm{OPT}_{-i}(\mathbf{v})$ is an DSIC payment rule as well with $p_i(\mathbf{v}_{-i}, 0) = \mathrm{OPT}(\mathbf{v}_{-i})$. This payment rule does not satisfy a natural *no-positive-transfers* condition which requires that agents not be paid to participate. It is also possible to design BNE mechanisms, e.g., with first-price semantics, that implement the same outcome in equilibrium as the surplus maximization mechanism (see Exercise 3.2), though unlike the surplus maximization mechanism given above, design of these BNE mechanisms often requires distributional knowledge.

3.3 Profit

Surplus maximization is singular among objectives in that there is a single mechanism that is optimal regardless of distributional assumptions. Essentially: the agents' incentives already aligned with the designer's objective and one only needs to derive the appropriate payments, i.e., the critical values. For general objectives we should not expect to be so lucky.

A non-game-theoretic optimization problem looks to maximize some objective subject to feasibility. Given the input, one can search over outcomes for the one with the highest objective value relative to this input. The outcome produced on one input need not bear any relation to the outcome produced on a (even slightly) different input. Mechanisms, on the other hand, additionally must address agent incentives which impose constraints on 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 the common prior, 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 Bayesian optimal mechanism for the objective of profit.

At various points in the remaining sections of this chapter it will be more convenient (and intuitive) to express certain functions in terms of the integral of their derivative. This notation is mathematically imprecise when the derivative is not defined, e.g., because the function is discontinuous. It can be made precise via the Dirac delta function which integrates to a step function; however, we will not describe these details formally. The reader is welcome to, instead, just assume the functions in question are continuous. An example of this is Theorem 3.10, below.

3.3.1 Quantile Space

In single-dimensional Bayesian mechanism design where an agent's value is distributed according to a continuous distribution F there is a one-to-one mapping between the agent's value and her strength relative to the distribution. For instance, an agent with value $v = 0.9$ drawn from $U[0, 1]$ is stronger than 90% of agents and weaker than 10% of agents with values drawn from the same distribution. We refer to indexing of agent from strong ($q = 0$) to weak ($q = 1$) as *quantile*. Importantly, the distribution of an agent's quantile is always $U[0, 1]$.

Definition 3.9. *The* quantile *q of an agent with value $v \sim F$ is the probability that the agent is weaker than a random draw from F. I.e., $q = 1 - F(v)$.*

It will be convenient to express an agent's value as a function of quantile as $v(q) = F^{-1}(1 - q)$. We will overload notation to define allocation and payment rules in quantile space as well. Specifically, "$x(q)$" and "$p(q)$" will be short-hand notation for $x(v(q))$ and $p(v(q))$, respectively. This convention will be extended to other functions as well: if the function is defined on values but applied to a quantile, then by this application we implicitly mean the function composed with the value function, $v(\cdot)$. We can rederive Theorem 2.7 in quantile space as follows.

Theorem 3.10. *Allocation and payment rules x and p are in BIC if and only if for all i,*

1. *(monotonicity) $x_i(q_i)$ is monotone non-increasing in q_i, and*
2. *(payment identity) $p_i(q_i) = -\int_{q_i}^{1} v_i(r)\, x_i'(r)\, dr + p_i(1)$,*

where $x_i'(q) = \frac{d}{dq}x_i(q)$ and $p_i(1)$ is the payment made by agent i when her value is at its lowest.

The payment identity of this theorem and Theorem 2.7 are related by a change of variables and integration by parts. Notice that as $x(q)$ is monotone non-increasing, its derivative $x'(q)$ is non-positive; hence, the negation of the integral guarantees a non-negative expected payment.

Proof. See Exercise 3.3. ☐

3.3.2 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?

Suppose we wish to sell to Alice with ex ante probability $\hat{q}$. The most direct way to do this would be to post a price $v(\hat{q})$ as this is the price at which $\mathbf{Pr}_{v\sim F}[v > v(\hat{q})] = \hat{q}$. The revenue obtained by posting such a price is exactly the price times the probability of sale, i.e., $v(\hat{q}) \cdot \hat{q}$.

Definition 3.11. *The* revenue curve $R(\cdot)$ *specifies the revenue as a function of ex ante probability of sale. I.e., $R(q) = v(q) \cdot q$. $R(1)$ and $R(0)$ are defined to be zero.*

We can clearly optimize revenue by taking the derivative of the revenue curve and setting it equal to zero. For example, if F is the uniform distribution $U[0,1]$ then $F(z) = z$, $v(q) = 1 - q$, $R(q) = q - q^2$, and $R'(q) = 1 - 2q$. The revenue is optimized by pricing at quantile $\hat{q} = 1/2$ (which corresponds to a price $v(1/2) = 1/2$). The uniform distribution is well-behaved in the sense that the revenue, as a function of quantile, increases up to quantile 1/2, which obtains a revenue of 1/4, and then decreases. The importance of the derivative in solving for the optimal price can be noted by observing that the derivative is positive but decreasing as r is increased to 1/2, where it is zero, and then continues to be negative and decreasing afterwards. This optimal revenue is obtained by allocating to Alice when the derivative of the revenue curve at her quantile, i.e., $R'(q)$, is non-negative.

3.3.3 Expected Revenue and Virtual Values

Suppose we are given the allocation rule (in quantile space) of an agent (Alice) as $x(q)$. By the payment identity (in quantile space), the payment rule must be $p(q) = -\int_q^1 v(r)\, x'(r)\, dr$. Since Alice's quantile q is drawn from $U[0, 1]$ we can calculate our expected revenue as follows.

$$\mathbf{E}_q[p(q)] = -\int_{q=0}^{1} \int_{r=q}^{1} v(r)\, x'(r)\, dr\, dq$$

This equation can be simplified by swapping the order of integration.

$$\begin{aligned}
\mathbf{E}_q[p(q)] &= -\int_{r=0}^{1} \int_{q=0}^{r} dq\, v(r)\, x'(r)\, dr \\
&= -\int_{r=0}^{1} r\, v(r)\, x'(r)\, dr \\
&= -\int_{q=0}^{1} R(q)\, x'(q)\, dq \qquad (3.1)
\end{aligned}$$

Equation (3.1) follows from substituting the definition of $R(\cdot)$ and making a change of variables. Denote $\frac{d}{dq}R(q)$ by $R'(q)$. If we integrate the above, by parts, we obtain:

$$\begin{aligned}
\mathbf{E}_q[p(q)] &= \int_{q=0}^{1} R'(q)\, x(q)\, dq - \Big[R(q)\, x(q)\Big]_{q=0}^{1} \\
&= \int_{q=0}^{1} R'(q)\, x(q)\, dq. \qquad (3.2)
\end{aligned}$$

Equation (3.2) follows from the definition of revenue curves which requires $R(0) = R(1) = 0$. We conclude this analysis by summarizing equations (3.1) and (3.2) as the following lemma.

Lemma 3.12. *An agent with revenue curve $R(\cdot)$ subject to allocation rule $x(\cdot)$ makes expected payment:*

$$\mathbf{E}_q[p(q)] = -\mathbf{E}_q[R(q)\, x'(q)] = \mathbf{E}_q[R'(q)\, x(q)]\,.$$

Both of the identities in Lemma 3.12 are useful for understanding the expected payments of agents in BNE. For instance, the former, from equation (3.1), implies that the same allocation rule (in quantile space) on a higher revenue curve gives more revenue.

Corollary 3.13. *If agents 1 and 2 with revenue curves satisfying $R_1(q) \geq R_2(q)$ for all q are subject to the same (in quantile space) allocation rule, i.e., satisfying $x_1(q) = x_2(q)$, then $\mathbf{E}_q[p_1(q)] \geq \mathbf{E}_q[p_2(q)]$.*

The latter identity from Lemma 3.12, from equation (3.2), gives an approach for optimizing revenue. It is instructive to view $R'(\cdot)$ as the marginal increase in revenue we get for selling to Alice with incrementally more probability. For our goal of optimizing expected profit, it suggests selecting x to maximize this *marginal revenue*. Informally: revenue is maximized by optimizing marginal revenue. This principle is standard in microeconomics.

The standard terminology in mechanism design for this marginal revenue is *virtual value*.

Definition 3.14. *The* virtual value *of an agent with quantile q and revenue curve $R(\cdot)$ is the marginal revenue at q:*

$$\phi(q) = R'(q).$$

The virtual surplus *of outcome $\mathbf{x}$ and profile of agent quantiles $\mathbf{q}$ is:*

$$\text{Surplus}(\boldsymbol{\phi}(\mathbf{q}), \mathbf{x}) = \sum\nolimits_i \phi_i(q_i)x_i - c(\mathbf{x}),$$

where $\boldsymbol{\phi}(\mathbf{q}) = (\phi_1(q_1), \ldots, \phi_n(q_n))$.

Often it is useful to write the virtual value in terms of the agent's value, v, and the distribution, F, from which the value is drawn. Evaluating, in value space, the derivative (with respect to quantile) of the revenue curve we obtain:

$$\phi(v) = v - \tfrac{1-F(v)}{f(v)}. \tag{3.3}$$

The following theorem is an immediate consequence of Lemma 3.12 and linearity of expectation.

Theorem 3.15. *A mechanism's expected revenue is equal to its expected virtual surplus, i.e., with allocation rule $\mathbf{x}(\cdot)$ on agents with virtual value functions $\boldsymbol{\phi}(\cdot)$ the expected revenue is:*

$$\mathbf{E}_{\mathbf{q}}\left[\sum\nolimits_i \phi_i(q_i)x_i(\mathbf{q}) - c(\mathbf{x}(\mathbf{q}))\right].$$

It should be noted that the distributional properties of an agent's value can be given equivalently by specifying the distribution F, the value function $v(\cdot)$, the revenue curve $R(\cdot)$, or the virtual value function $\phi(\cdot)$.

3.3.4 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 solve the remaining non-game-theoretic optimization problem. Since expected profit equals expected virtual surplus, this non-game-theoretic optimization problem is to optimize virtual surplus. We then verify that this solution does not violate the incentive constraints (under some conditions). We conclude that (under the same conditions) the resulting mechanism is optimal.

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

We know from the BIC characterization (Corollary 2.16) 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. Yet, even though we did not impose a monotonicity constraint on OPT, if the virtual valuation functions $\phi_i(\cdot)$ are monotone, $\text{OPT}(\boldsymbol{\phi}(\cdot))$ is monotone.

Definition 3.16. *Distribution F is* regular *if its associated revenue curve $R(q)$ is a concave function of q (equivalently: $\phi(\cdot)$ is monotone).*

Many distributions are regular, e.g., uniform, normal, exponential. On the other hand many relevant distributions are irregular, e.g., bimodal.

Lemma 3.17. *For each agent i and any values of other agents $\mathbf{v}_{-i}$, if F_i is regular then i's allocation rule from $\text{OPT}(\boldsymbol{\phi}(\cdot))$ 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 $x_i(\phi_i(v_i))$. ☐

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") truthtelling is a dominant strategy equilibrium (recall Corollary 2.18). One way to view the suggested virtual surplus maximization mechanism is as a reduction to surplus maximization, which is solved by the SM mechanism (Mechanism 3.1).

Mechanism 3.2. *The* virtual surplus maximization (VSM) *mechanism for regular distributions with virtual value functions $\boldsymbol{\phi}(\cdot)$ is:*

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

Notice that the payments $\mathbf{p}$ calculated can be viewed as follows. SM on virtual values outputs virtual prices $\mathbf{p}'$. These correspond to the minimum virtual value an agent must have to win. The price an agent pays is the minimum value it must have to win, this can be calculated from the virtual prices above via the inverse virtual valuation function.[4]

[4]Assuming 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.

Theorem 3.18. *For regular distributions, the virtual value maximization mechanism is dominant strategy incentive compatible.*

Corollary 3.19. *For regular distributions, the virtual surplus maximization mechanism optimizes expected profit in dominant strategy equilibrium.*

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

First notice that virtual values can be negative. Consider the uniform distribution $U[0,1]$ where $F(z) = z$ and $f(z) = 1$. From equation (3.3), $\phi(v) = v - \frac{1-F(z)}{f(z)} = 2v - 1$. 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 derivative of the revenue curve and our analysis of single-agent environments 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 3.20. *For regular, single-item environments, the auction that allocates to the agent with the highest non-negative virtual value optimizes expected revenue.*

As virtual valuations are the derivative of the revenue curve, the optimal auction allocate to the agent whose revenue curve is the steepest at her value.

The case where the agents are independent and identically distributed is of special interest. For i.i.d. and regular distributions, 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 her 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 $\phi^{-1}(0)$!

Definition 3.21 (Second-price Auction with reservation price r)**.** *The* second-price auction with reservation price r*, sells the item if any agent bids above* r*. The price the winning agent pays the maximum of the second highest bid and* r*.*

Corollary 3.22. *For identical, regular, single-item environments, the second-price auction with reserve* $\phi^{-1}(0)$ *optimizes expected revenue.*

Notice that the optimal reserve price is not a function of the number of agents. Furthermore, the result can easily be extended to single-item multi-unit auctions where the optimal reserve price is also not a function of the number of units that are for sale. As we will see from Theorem 4.21 in Chapter 4 the same result extends beyond single-item and multi-unit feasibility constraints to those that are downward-closed and satisfy a natural "augmentation" property that is related to substitutability, a.k.a., *matroids*.

While this auction is optimal among all BIC auctions, which is the class of mechanisms we restricted our attention to, (a) the revelation principle implies that no auction has a BNE with higher expected revenue, and (b) it actually satisfies the stronger dominant strategy incentive compatibility constraint. Therefore, we conclude that in a very strong sense, that the second-price auction with reserve price maximizes expected revenue.

We conclude by returning to the 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 calculation at the introduction of this chapter showed that this optimal revenue 5/12.

3.4 Irregular Distributions

We now 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 the revenue curve is non-concave (in quantile). The virtual valuation functions are non-monotone, therefore, a higher value might result in a lower virtual value. Clearly $\text{OPT}(\phi(\cdot))$ is non-monotone for such a distribution; therefore, there is no payment rule for which it is incentive compatible.

3.4.1 Ironed Revenue Curves

Consider again the problem of selling an item to Alice with ex ante probability $\hat{q}$. We could offer Alice price $v(\hat{q})$ to obtain revenue $R(\hat{q}) = \hat{q} \cdot v_i(\hat{q})$; however, when $R(\cdot)$ is not concave, this approach may not optimize expected revenue.

To see what is going wrong, notice that 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 3.1(a) depicts a hypothetical non-concave revenue curve; Figure 3.1(c) depicts the corresponding virtual value function. Figure 3.1(d) shows Alice's virtual value averaged on $[a,b]$. Finally, Figure 3.1(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 $(b, R(b))$. Since $R(\cdot)$ is non-concave this line segment at $\hat{q}$ can be strictly higher than $R(\hat{q})$, as pictured. This process of treating Alice the same on an interval to flatten the virtual valuation function is known as *ironing*.

To sell to Alice with ex ante probability $\hat{q}$ we can pick some interval $[a,b]$ with $a < \hat{q} < b$ and apply the allocation rule

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

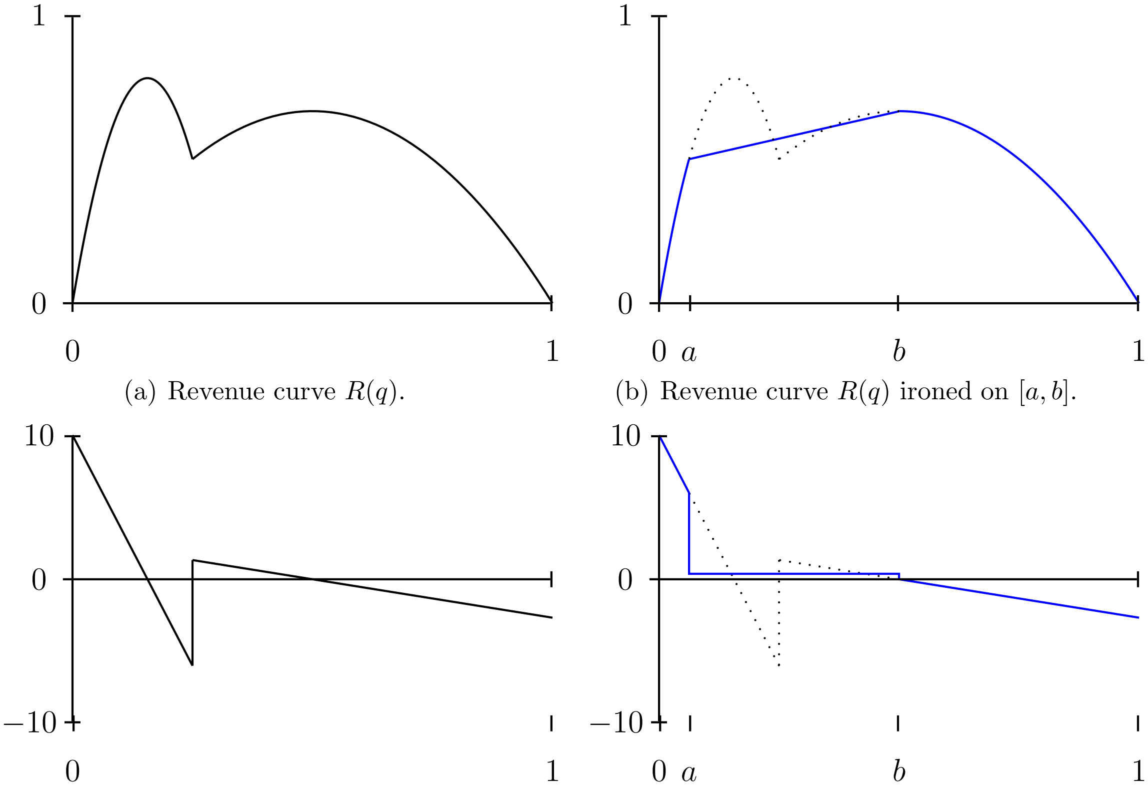

(a) Revenue curve $R(q)$.

(b) Revenue curve $R(q)$ ironed on $[a, b]$.

(c) Virtual values $R'(q)$.

(d) Virtual values $R'(q)$ ironed on $[a, b]$.

Figure 3.1: 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 $1/4$.

Notice that when Alice's quantile q is realized (i.e., drawn from the uniform distribution) then the probability that Alice is served by $x^{\hat{q}}(\cdot)$ 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 Theorem 3.10. It is $R(a) + \frac{\hat{q}-a}{b-a}(R(b) - R(a))$. Notice that this revenue is exactly the value at $\hat{q}$ on the line segment connecting $(a, R(a))$ to $(b, R(b))$, e.g., see Figure 3.1(b). Again, where $R(\cdot)$ is non-concave, the revenue obtained from this randomized rule can be higher than $R(\hat{q})$.

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 revenue curve $\bar{R}(\cdot)$ equal smallest concave function that upper-bounds 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 3.23. *For $v \sim F$, the* ironed revenue curve*, $\bar{R}(\cdot)$, is smallest concave function that upper-bounds $R(\cdot)$ and the* ironed virtual valuation *function is $\bar{\phi}(q) = \bar{R}'(q)$.*

Ironed intervals of the ironed revenue curve are those with $\bar{R}(q) > R(q)$. The usage of ironed virtual values in place of virtual values as a proxy for an agent's (Alice) expected payment is valid only for mechanisms treat her the same way regardless of where in the interval her quantile lies. Meaning: Alice with quantile $q \in [a, b]$ that is ironed will be served with the same probability as she would have been with any other quantile $q' \in [a, b]$. The following lemma formally states that ironed virtual surplus gives an upper bound on virtual surplus that is tight for mechanisms that respect the ironed intervals.

Lemma 3.24. *An agent's expected payment is upper-bounded by their expected ironed virtual surplus, i.e.,*

$$\mathbf{E}_v[p(v)] \leq \mathbf{E}_q\left[\bar{\phi}(q)x(q)\right].$$

Furthermore, this inequality holds with equality when if $\bar{R}(q) > R(q) \Rightarrow x'(q) = 0$.

Proof. We will start by showing a more precise statement.

$$\begin{aligned}
\mathbf{E}_q[p(q)] &= \mathbf{E}_q[R'(q) \cdot x(q)] + \mathbf{E}_q\left[\bar{R}'(q) \cdot x(q)\right] - \mathbf{E}_q\left[\bar{R}'(q) \cdot x(q)\right] \\
&= \mathbf{E}_q\left[\bar{R}'(q) \cdot x(q)\right] - \mathbf{E}_q\left[\left(\bar{R}'(q) - R'(q)\right) \cdot x(q)\right] \\
&= \mathbf{E}_q\left[\bar{R}'(q) \cdot x(q)\right] + \mathbf{E}_q\left[\left(\bar{R}(q) - R(q)\right) \cdot x'(q)\right]. \qquad (3.4)
\end{aligned}$$

The last line above follows from writing the expectation as an integral and integration by parts.

Inspecting the second term of equation (3.4) more closely, notice that the difference in the revenue curves is non-negative, as $\bar{R}(\cdot)$ is defined to be an upper-bound on $R(\cdot)$; and the derivative of the allocation rule is non-positive, as the allocation rule is monotone non-increasing in quantile. Therefore, the second term is non-positive and the inequality of the lemma is proven.

Of course, the assumption that $\bar{R}(q) > R(q) \Rightarrow x'(q) = 0$ implies that the second term of (3.4) is identically zero: whenever the first multiplicand is non-zero, the second multiplicand is zero. ∎

Notice the advantage of $\bar{R}(\cdot)$ over $R(\cdot)$ is two-fold. First, Corollary 3.13 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.

In retrospect it should be obvious that the optimal revenue as a function of ex ante sale probability is concave. Given any two IC mechanisms the convex combination of the two mechanisms is IC and its revenue is a convex combination of the two mechanisms revenue.

3.4.2 Optimal Mechanisms

We will now show that for any distribution, the mechanism that maximizes ironed virtual surplus obtains the optimal expected profit. Again we view this mechanism as a reduction to surplus maximization which is solved, e.g., by mechanism SM (Mechanism 3.1). The resulting mechanism is sometimes referred to as the Myerson auction (for single-item environments) or the Myerson mechanism (for general single-dimensional environments) after Roger Myerson.

Mechanism 3.3. *The* ironed virtual surplus maximization (IVSM) *mechanism for distributions with ironed virtual value functions* $\bar{\boldsymbol{\phi}}(\cdot)$ *is:*

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

By monotonicity of $\bar{\boldsymbol{\phi}}(\cdot)$ and $\mathrm{OPT}(\cdot)$, $\mathrm{OPT}(\bar{\boldsymbol{\phi}}(\cdot))$ is monotone for each agent and all values of other agents. Therefore, ironed virtual surplus maximization satisfies our strongest incentive constraint. With the appropriate payments (i.e., the "critical values") truthtelling is a dominant strategy equilibrium (recall Corollary 2.18).

Theorem 3.25. *The ironed virtual surplus maximization mechanism is dominant strategy incentive compatible.*

To show that ironed virtual surplus mechanism is optimal we need to argue that it respects the ironed intervals of the ironed revenue curve, i.e., any agent with value within an ironed interval receives the same outcome regardless of where in the interval her value lies.

Lemma 3.26. *For agents with revenue curves* $\mathbf{R}(\cdot)$*, the allocation rule* $\mathbf{x}(\cdot)$ *of the ironed virtual surplus maximization mechanism satisfies* $\bar{R}_i(q_i) > R_i(q_i) \Rightarrow x'_i(q_i) = 0$ *for all* i*.*

Proof. Observe that on ironed intervals, i.e., where $\bar{R}(q) > R(q)$, the ironed revenue curve, $\bar{R}(\cdot)$, is linear. This follows from the definition of the ironed revenue curve as the smallest concave function that upper-bounds the revenue curve. Since $\bar{R}(\cdot)$ is linear on this ironed interval, its derivative and, consequently, the ironed virtual function is constant on the interval. The allocation probability of the ironed virtual surplus maximization mechanism is determined by the optimization $\mathrm{OPT}(\bar{\boldsymbol{\phi}}(\cdot))$ which is a function only of the ironed virtual

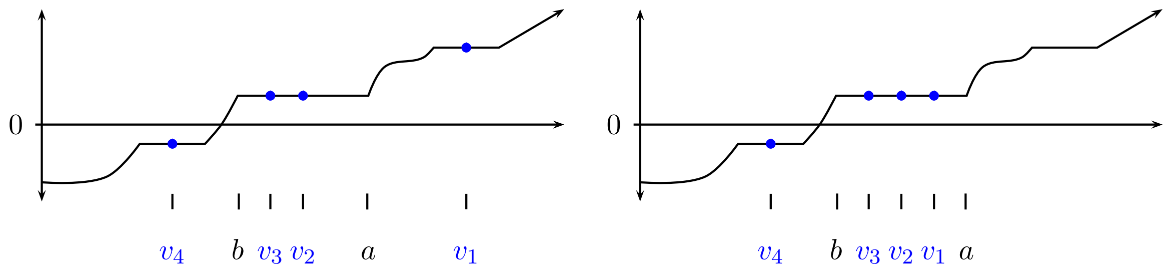

(a) Unique highest ironed virtual value. (b) Non-unique highest ironed virtual value.

Figure 3.2: The ironed virtual valuation function $\bar{\phi}(v)$ under two realizations of agent values depicting both the case where the highest ironed virtual value is unique and the case where it is not unique.

values. Since an agent with any quantile within an ironed interval has the same ironed virtual valuation, this optimization must produce an outcome that is constant on the interval. On ironed intervals, therefore, the derivative of the allocation rule is zero. □

Corollary 3.27. *The ironed virtual surplus maximization mechanism optimizes expected profit in dominant strategy equilibrium.*

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.

Note that unlike the surplus maximization mechanism and the virtual value maximization mechanism (for regular distributions) where the continuity assumption on the distribution implies that there is never a tie, the ironed virtual surplus maximization mechanism for irregular distributions may require a tie-breaking policy, for instance, when to agents with distinct values have the same ironed virtual value. Tie breaking can be implemented 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 often desired because it is symmetric.

3.4.3 Single-item Auctions

We consider the special case of single-item auctions to get a clearer picture of exactly what the optimal mechanism is in the case of i.i.d., irregular distributions. Figure 3.2 depicts hypothetical ironed virtual valuation function. Instantiating the agents' values corresponds to picking points on the horizontal axis. The agents' ironed virtual valuations can then be read off the plot. The optimal 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 3.2(a) depicts a realization of values for $n = 4$ agents where the highest ironed virtual valuation is unique. What does the ironed virtual surplus maximization do here? It

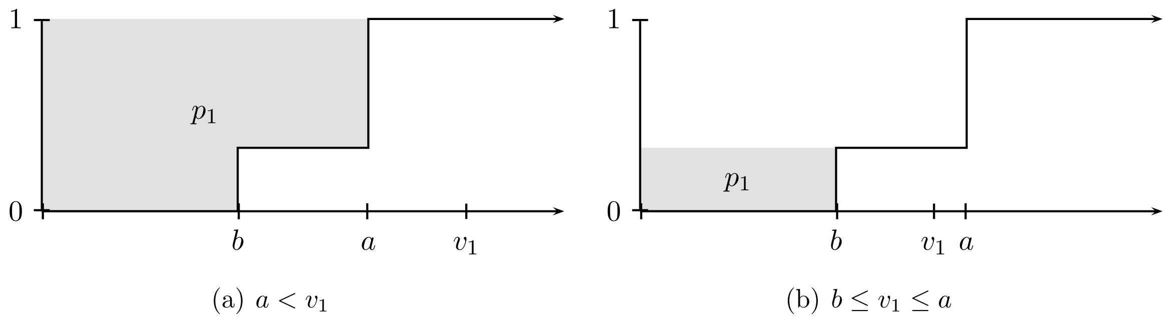

Figure 3.3: The allocation (black line) and payment rule (gray region) for agent 1 given $\mathbf{v}_{-i}$ and the ironed virtual valuation function from Figure 3.2.

allocates the item to this agent, i.e., agent 1 in the figure. Figure 3.2(b) depicts a second realization of values where the highest ironed virtual valuation is not unique. In this setting the mechanism, 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 $1/k$.

It is instructive to calculate the payment an agent must make in expectation over the random tie-breaking rule. Consider the case where there is a unique highest ironed virtual value. The agent with this ironed virtual value wins. To calculate their DSIC payment 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 3.2(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 3.3(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 $p_1 = a - (a - b)/k$.

When agent 1 is tied for the highest ironed virtual value with $k-1$ other agents, as depicted in Figure 3.3(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 and zero to the losers.

Exercises

3.1 In computer networks such as the Internet is is often not possible to use monetary payments to ensure 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 not direct value from and agent performing a computation. Define the *residual surplus* as be the social surplus less the payments, i.e., $\sum_i (v_i \cdot x_i - p_i) - c(\mathbf{x})$. (For more details, see the discussion of non-monetary payments in Chapter 1.)

Describe the mechanism that maximizes residual surplus when the distributions of on agent's 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 implications of that distribution under the monotone hazard rate assumption. Consider the following cases:

(a) a single-item auction with i.i.d. values.

(b) a single-item auction with non-identical values.

(c) an environment with general costs specified by $c(\cdot)$ and non-identical values.

3.2 Give a mechanism with first-price payment semantics that implements the social surplus maximizing outcome in equilibrium for any single-dimensional agent environment. Hint: your mechanism may be parameterized by the distribution.

3.3 Prove from first principles that BNE implies the payment identity of Theorem 3.10. You may assume that $\mathbf{x}(\cdot)$ and $\mathbf{p}(\cdot)$ are continuously differentiable with respect to quantile.

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 maximizing mechanism for general distributions.

(b) Describe the revenue maximizing 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 optimal public project mechanism when agents' values are i.i.d. from $U[0, 1]$.

Chapter Notes

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

The revenue-optimal single-item auction was derived by Roger Myerson (1981). Its generalization to single-dimensional agent environments is an obvious extension. The relationship between revenue-optimal auctions, revenue curves, and *marginal revenue* (equivalent to virtual values) is due to Bulow and Roberts (1989).