2
Equilibrium

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

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

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

2.1 Complete Information Games

In games of compete information all players are assumed to know precisely the payoff structure of all other players for all possible outcomes of the game. A classic example of such a game is the *prisoner's dilemma*, the story for which is as follows.

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

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

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

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

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

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

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

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

A reasonable bimatrix game depicting this story is the following.

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

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

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

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

2.2 Incomplete Information Games

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

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

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

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

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

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

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

2.3 Bayes-Nash Equilibrium

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

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

The distribution $\boldsymbol{F}$ over $\boldsymbol{t}$ may generally be correlated. Which means that an agent with knowledge of her own type must do *Bayesian updating* to determine the distribution over the types of the remaining bidders. We denote this conditional distribution as $\boldsymbol{F}_{-i}\big|_{t_i}$. Of course, when the distribution of types is independent, i.e., $\boldsymbol{F}$ is the *product distribution* $F_1 \times \cdots \times F_n$, then $\boldsymbol{F}_{-i}\big|_{t_i} = \boldsymbol{F}_{-i}$.

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

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

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

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

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

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

Thus,

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

Third, we optimize agent 1's bid. Agent 1 with value v_1 should maximize $2v_1b_1 - 2b_1^2$ as a function of b_1, and to do so, can differentiate the function and set its derivative equal to zero. The result is $\frac{\mathrm{d}}{\mathrm{d}b_1}(2v_1b_1 - 2b_1^2) = 2v_1 - 4b_1 = 0$ and we can conclude that the optimal bid is $b_1 = v_1/2$.

This proves that agent 1 should bid as prescribed if agent 2 does; and vice versa. Thus, we conclude that the guessed strategy profile is in BNE.

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

2.4 Single-dimensional Games

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2.5 Characterization of Bayes-Nash Equilibrium

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

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

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

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

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

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

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

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

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

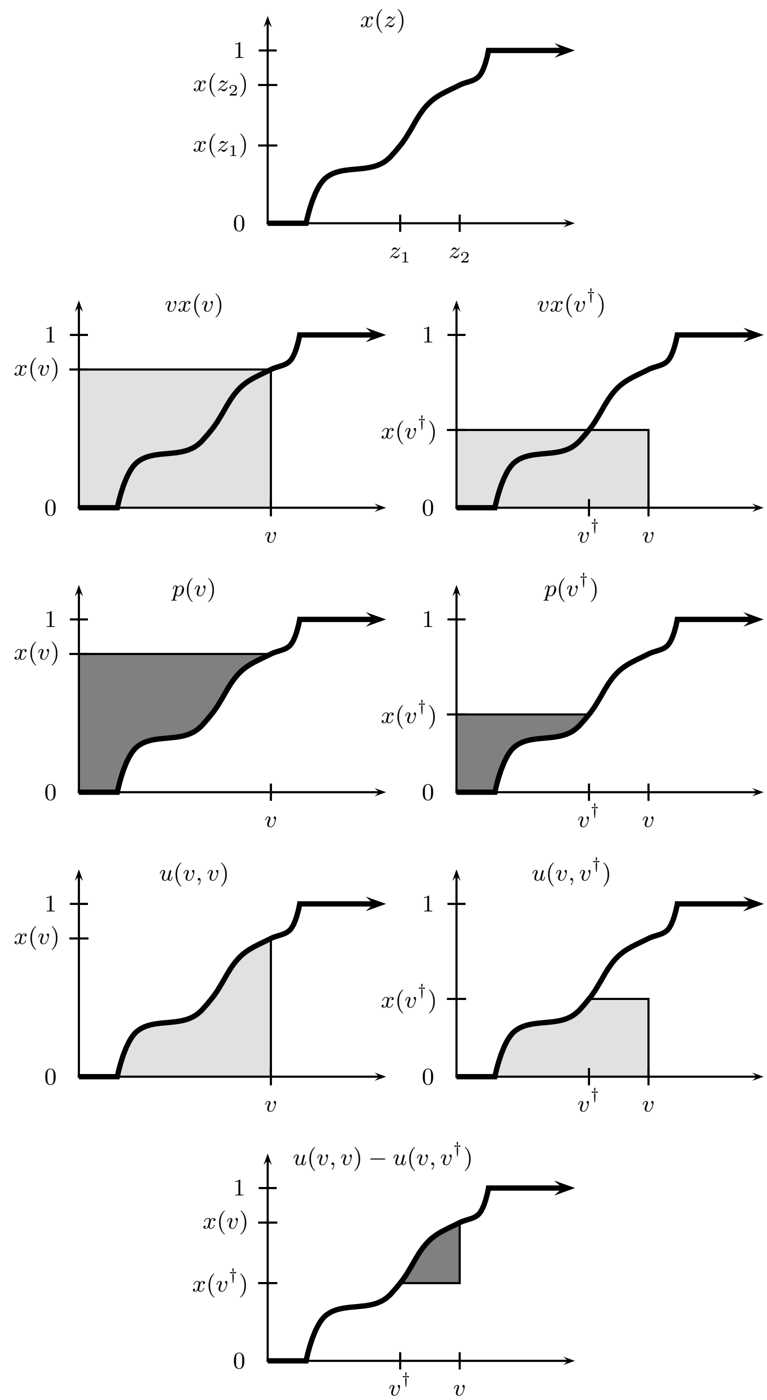

Figure 2.1 The left column shows (shaded) the surplus, payment, and utility of Alice playing action $s(v = z_2)$. The right column shows (shaded) the same for Alice playing action $s(v^\dagger = z_1)$. The final diagram shows (shaded) the difference between Alice's utility for these strategies. Monotonicity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

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

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

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

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

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

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

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

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

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

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

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

Adding these inequalities and canceling the payment terms we have,

$$z_2 x(z_2) + z_1 x(z_1) \geq z_2 x(z_1) + z_1 x(z_2).$$

Rearranging,

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

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

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

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

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

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

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

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

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

$$p'(z) = zx'(z), \text{ so}$$
$$\int_0^v p'(z)\mathrm{d}z = \int_0^v zx'(z)\,\mathrm{d}z.$$

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

$$p(v) = \int_0^v zx'(z)\,\mathrm{d}z + p(0).$$

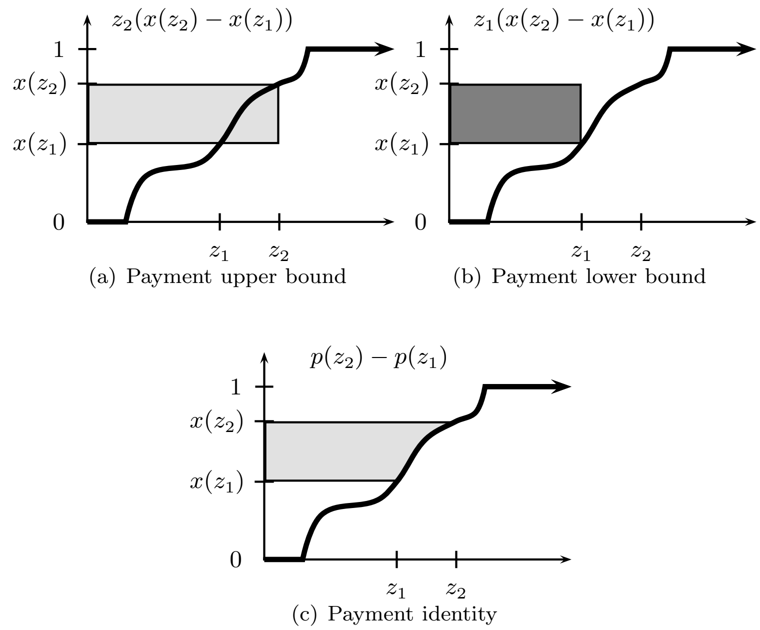

(a) Payment upper bound (b) Payment lower bound

(c) Payment identity

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

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

$$\begin{aligned} p(v) &= \Big[zx(z)\Big]_0^v - \int_0^v x(z)\,\mathrm{d}z + p(0) \\ &= vx(v) - \int_0^v x(z)\,\mathrm{d}z + p(0). \end{aligned}$$

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

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

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

As we conclude the proof of the BNE characterization theorem, it is important to note how little we have assumed of the underlying game. We did not assume it was a single-round, sealed-bid auction. We did not assume that only a winner will make payments. Therefore, we conclude for any potentially wacky, multi-round game the outcomes of all Bayes-Nash equilibria have a nice form.

2.6 Characterization of Dominant Strategy Equilibrium

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

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

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

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

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

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

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

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

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

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

2.7 Revenue Equivalence

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

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

We can now quantitatively compare the second-price and first-price auctions from a revenue standpoint. Consider the case where the agent's values are distributed independently and identically. What is the equilibrium outcome of the second-price auction? The agent with the highest valuation wins. What is the equilibrium outcome of the first-price auction? This question requires a little more thought. Since the distributions are identical, it is reasonable to expect that there is a symmetric equilibrium, i.e., one where $s_i = s_{i'}$ for all i and i'. Furthermore, it is reasonable to expect that the strategies are monotone, i.e., an agent with a higher value will out bid an agent with a lower value. Under these assumptions, the agent with the highest value wins. Of course, in both auctions a 0-valued agent will pay nothing. Therefore, we can conclude that the two auctions obtain the same expected revenue.

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

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

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

2.8 Solving for Bayes-Nash Equilibrium

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

Suppose we are to solve for the BNE strategies of mechanism M. The approach is to express an agent's payment in M as a function of the agent's action, then to calculate the agent's expected payment in a strategically-simple mechanism M' that is revenue equivalent to M (usually a "second-price implementation" of M). Setting these terms equal and solving for the agents action gives the equilibrium strategy.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

We conclude by observing that if we fail to exhibit a Bayes-Nash equilibrium via this approach then our original guess is contracted and there is no equilibrium of the given mechanism that corresponds to the guess. Conversely, if the approach succeeds then the equilibrium found is the only equilibrium consistent with the guess. As an example, we can conclude the following for first-price auctions.[2]

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

2.9 Uniqueness of Equilibria

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

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

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

In contrast, the first-price auction for independent and identical prior distributions does not suffer from multiplicity of Bayes-Nash equilibria. Specifically, the method described in the previous section for solving for the symmetric equilibrium in symmetric auction-like games gives the unique BNE. We describe this result as two parts. First, we exclude the possibility of multiple symmetric equilibria. Second, we exclude the existence of asymmetric equilibria.

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

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

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

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

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

They are,

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

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

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

Lemma 2.9 *For $n = 2$ agents with values drawn independently and identically from a continuous distribution F, the first-price auction with an unknown random reserve from known distribution G admits no asymmetric Bayes-Nash equilibrium.*

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

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

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

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

A first claim for strictly-increasing continuous strategies is that $s_1(v) > s_2(v)$ if any only if $x_1(v) > x_2(v)$. See Figure 2.3 for a graphical representation of the following argument. Since the strategies are continuous and strictly increasing, the inverses of the strategies are well defined.

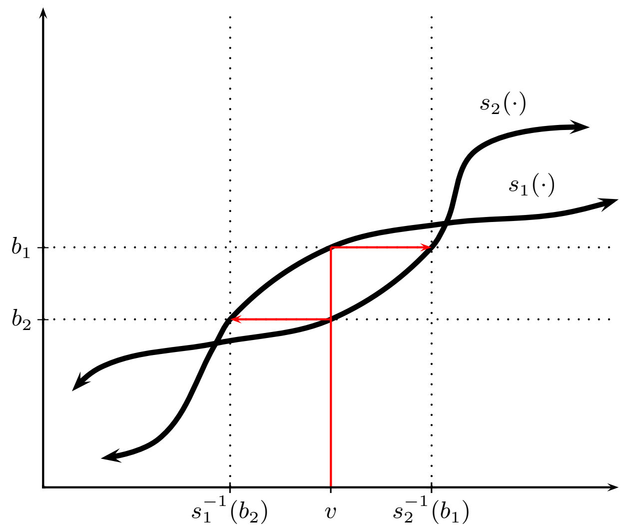

Figure 2.3 Graphical depiction of the first claim in the proof of Lemma 2.9 with $b_i = s_i(v)$. Clearly, $s_2^{-1}(b_1) > s_1^{-1}(b_2)$. Strict monotonicity of the distribution function $F(\cdot)$ then implies that $F(s_2^{-1}(b_1)) > F(s_1^{-1}(b_2))$.

Calculate Alice's interim allocation probability x_1 at value v, for Bob's value $v_2 \sim F$ and reserve bid $\hat{b} \sim G$, as:

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

Likewise, Bob's interim allocation probability is

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

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

A second claim is that the low-bidding Bob on the interval $I = (a, b)$

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

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

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

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

$$\int_a^b x_1(z)\,\mathrm{d}z > \int_a^b x_2(z)\,\mathrm{d}z$$

These observations give a contradiction. □

2.10 The Revelation Principle

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

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

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

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

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

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

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

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

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

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

We defined the allocation and payment rules $\boldsymbol{x}(\cdot)$ and $\boldsymbol{p}(\cdot)$ as functions of the valuation profile for an implicit game G and strategy profile $\boldsymbol{s}$. When the strategy profile is truthtelling, the allocation and payment rules are identical the original mappings of the game from actions to allocations and prices, denoted $\boldsymbol{x}^G(\cdot)$ and $\boldsymbol{p}^G(\cdot)$. Additionally, let

$x_i^G(v_i) = \mathbf{E}\Big[x_i^G(\boldsymbol{v}) \mid v_i\Big]$ and $p_i^G(v_i) = \mathbf{E}\Big[p_i^G(\boldsymbol{v}) \mid v_i\Big]$ for $\boldsymbol{v} \sim \boldsymbol{F}$. Furthermore, the truthtelling strategy profile in a direct-revelation game is onto.

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

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

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

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

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

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

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

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

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

Exercises

2.1 Find a symmetric mixed strategy equilibrium in the chicken game described in Section 2.1. I.e., find a probability ρ such that if James Dean stays with probability ρ and swerves with probability $1-\rho$ then Buzz is happy to do the same.

2.2 Give a characterization of Bayes-Nash equilibrium for discrete single-dimensional type spaces for agents with linear utility. Assume that $T = \{v^0, \ldots, v^N\}$ with the probability that an agent's value is $v \in T$ given by probability mass function $f(v)$. Assume $v^0 = 0$. You will not get a payment identity; instead characterize for any BNE allocation rule, the maximum payments.

(a) Give a characterization for the special case where the values are uniform, i.e., $v^j = j$ for all j.

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

(c) Give a characterization for the general case.

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

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

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

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

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

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

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

(a) Find the envy-free (see Exercise 2.3) outcome and payments with the maximum social surplus. Give a description and formula for the envy-free outcome and payments for each advertiser. (Feel free to specify your payment formula with a comprehensive picture.)

(b) In the real AdWords problem, advertisers only pay if they receive a click, whereas the payments calculated, i.e., $\boldsymbol{p}$, are in expected over all outcomes, click or no click. If we are going to charge advertisers only if they are clicked on, give a formula for calculating these payments $\boldsymbol{p}'$ from $\boldsymbol{p}$.

(c) The real AdWords problem is solved by auction. Design an auction that maximizes the social surplus in dominant strategy equilibrium. Give a formula for the payment rule of your auction (again, a comprehensive picture is fine). Compare your DSE payment rule to the envy-free payment rule. Draw some informal conclusions.

2.5 Consider the first-price auction for selling a single item to two agents whose values are independent but not identical. In each of the settings below prove or disprove the claim that there is a Bayes-Nash equilibrium wherein the item is always allocated to the agent with the highest value.

(a) Agent 1 has value $U[0, 1]$ and agent 2 has value $U[0, 1/2]$.

(b) Agent 1 has value $U[0, 1]$ and agent 2 has value $U[1/2, 1]$.

2.6 Consider the first-price auction for selling k units of an item to n unit-demand agents. This auction solicits bids and allocates one unit to each of the k highest-bidding agents. These winners are charged their bids. This auction is revenue equivalent to the k-unit "second-price" auction where the winners are charged the $(k+1)$st highest bid, $b_{(k+1)}$. Solve for the symmetric Bayes-Nash equilibrium strategies in the first-price auction when the agent values are i.i.d. $U[0,1]$.

2.7 Consider the position auction environment with $n = m = 2$ (see Exercise 2.4). Consider running the following first-price auction: The advertisers submit bids $\boldsymbol{b} = (b_1, b_2)$. The advertisers are assigned to slots in order of their bids. Advertisers pay their bid when clicked. Use revenue equivalence to solve for BNE strategies $\boldsymbol{s}$ when the values of the advertisers are drawn independent and identically from $U[0,1]$.

2.8 Prove that in a two-agent second-price auction for a single-item, that the best Bayes-Nash equilibrium can have a social surplus (i.e., the expected value of the winner) that is arbitrarily larger than the worst Bayes-Nash equilibrium. (Hint: Show that for any fixed β that there is a value distribution $\boldsymbol{F}$ and two BNE where the social surplus in one BNE is strictly larger than a β fraction of the social surplus of the other BNE.)

2.9 Show that with independent, identical, and continuously distributed values, the two-agent all-pay auction (where agents bid, the highest-bidder wins, and all agents pay their bids) admits exactly one strictly continuous Bayes-Nash equilibrium.

2.10 Show that with independent, identical, and continuously distributed values, the two-agent first-price position auction (cf. Exercise 2.4; where agents bid, the highest bidder is served with given probability w_1, the second-highest bidder is served with given probability $w_2 \leq w_1$, and all agents pay their bids when they are served) admits exactly one strictly continuous Bayes-Nash equilibrium.

2.11 Consider the following auction with first-price payment semantics. Agents bid, any agent whose bid is (weakly) higher than all other bids wins, all winners are charged their bids. Notice that in the case of a tie in the highest bid, all of the tied agents win. Prove that there are multiple Bayes-Nash equilibria when agents have values that are independently, identically, and continuously distributed.

2.12 Prove Lemma 2.9: For two agents with values drawn independently and identically from a continuous distribution F with support $[0,1]$, the first-price auction with an unknown random reserve from known distribution G admits no asymmetric Bayes-Nash equilibrium. I.e., remove the assumption of strictly-increasing and continuous strategies from the proof given in the text.

Chapter Notes

The formulation of Bayesian games is due to Harsanyi (1967). The characterization of Bayes-Nash equilibrium, revenue equivalence, and the revelation principle come from Myerson (1981). Parts of the BNE characterization proof presented here come from Archer and Tardos (2001). Amann and Leininger (1996), Bajari (2001), Maskin and Riley (2003), and Lebrun (2006) studied the uniqueness of equilibrium in the first-price and all-pay auctions. The revenue-equivalence-based uniqueness proof presented here is from Chawla and Hartline (2013).

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