9 Bayesian Games

Prerequisite: Chapter 2 and Section 4.1.3; Section 9.7 requires Chapter 4

AN ASSUMPTION underlying the notion of Nash equilibrium is that each player holds the correct belief about the other players' actions. To do so, a player must know the game she is playing; in particular, she must know the other players' preferences. In many situations the participants are not perfectly informed about their opponents' characteristics: bargainers may not know each others' valuations of the object of negotiation, firms may not know each others' cost functions, combatants may not know each others' strengths, and jurors may not know their colleagues' interpretations of the evidence in a trial. In some situations, a participant may be well informed about her opponents' characteristics, but may not know how well these opponents are informed about her own characteristics. In this chapter I describe the model of a "Bayesian game", which generalizes the notion of a strategic game to allow us to analyze any situation in which each player is imperfectly informed about an aspect of her environment that is relevant to her choice of an action.

9.1 Motivational examples

I start with two examples that illustrate the main ideas in the model of a Bayesian game. I define the notion of Nash equilibrium separately for each game. In the next section I define the general model of a Bayesian game and the notion of Nash equilibrium for such a game.

EXAMPLE 273.1 (Variant of *BoS* with imperfect information) Consider a variant of the situation modeled by *BoS* (Figure 19.1) in which player 1 is unsure whether

Prob. $\frac{1}{2}$ — 1 — Prob. $\frac{1}{2}$

2	B	S
B	2,1	0,0
S	0,0	1,2

2 wishes to meet 1

2	B	S
B	2,0	0,2
S	0,1	1,0

2 wishes to avoid 1

Figure 274.1 A variant of *BoS* in which player 1 is unsure whether player 2 wants to meet her or to avoid her. The frame labeled 2 enclosing each table indicates that player 2 knows the relevant table. The frame labeled 1 enclosing both tables indicates that player 1 does not know the relevant table; the probabilities she assigns to the two tables appear on the frame.

player 2 prefers to go out with her or prefers to avoid her, whereas player 2, as before, knows player 1's preferences. Specifically, suppose player 1 thinks that with probability $\frac{1}{2}$ player 2 wants to go out with her, and with probability $\frac{1}{2}$ player 2 wants to avoid her. (Presumably this assessment comes from player 1's experience: half of the time she is involved in this situation she faces a player who wants to go out with her, and half of the time she faces a player who wants to avoid her.) That is, player 1 thinks that with probability $\frac{1}{2}$ she is playing the game on the left of Figure 274.1 and with probability $\frac{1}{2}$ she is playing the game on the right. Because probabilities are involved, an analysis of the situation requires us to know the players' preferences over lotteries, even if we are interested only in pure strategy equilibria; thus the numbers in the tables are Bernoulli payoffs.

We can think of there being two *states*, one in which the players' Bernoulli payoffs are given in the left table and one in which these payoffs are given in the right table. Player 2 knows the state—she knows whether she wishes to meet or avoid player 1—whereas player 1 does not; player 1 assigns probability $\frac{1}{2}$ to each state.

The notion of Nash equilibrium for a strategic game models a steady state in which each player's beliefs about the other players' actions are correct, and each player acts optimally, given her beliefs. We wish to generalize this notion to the current situation.

From player 1's point of view, player 2 has two possible *types*, one whose preferences are given in the left table of Figure 274.1 and one whose preferences are given in the right table. Player 1 does not know player 2's type, so to choose an action rationally she needs to form a belief about the action of each type. Given these beliefs and her belief about the likelihood of each type, she can calculate her expected payoff to each of her actions. For example, if she thinks that the type who wishes to meet her will choose B and the type who wishes to avoid her will choose S, then she thinks that B will yield her a payoff of 2 with probability $\frac{1}{2}$ and a payoff of 0 with probability $\frac{1}{2}$, so that her expected payoff is $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$, and S will yield her an expected payoff of $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 = \frac{1}{2}$. Similar calculations for the other combinations of actions for the two types of player 2 yield the expected payoffs in Figure 275.1. Each column of the table is a pair of actions for the two types of

	(B,B)	(B,S)	(S,B)	(S,S)
B	2	1	1	0
S	0	$\frac{1}{2}$	$\frac{1}{2}$	1

Figure 275.1 The expected payoffs of player 1 for the four possible pairs of actions of the two types of player 2 in Example 273.1. Each row corresponds to an action of player 1, and each column corresponds to a pair of actions of the two types of player 2; the action of the type who wishes to meet player 1 is listed first.

player 2, the first member of each pair being the action of the type who wishes to meet player 1 and the second member being the action of the type who wishes to avoid player 1.

For this situation we define a pure strategy *Nash equilibrium* to be a triple of actions, one for player 1 and one for each type of player 2, with the property that

- the action of player 1 is optimal, given the actions of the two types of player 2 (and player 1's belief about the state)
- the action of each type of player 2 is optimal, given the action of player 1.

That is, we treat the two types of player 2 as separate players and analyze the situation as a three-player strategic game in which player 1's payoffs as a function of the actions of the two other players (i.e. the two types of player 2) are given in Figure 275.1, and the payoff of each type of player 2 is independent of the actions of the other type and depends on the action of player 1 as given in the tables in Figure 274.1 (the left table for the type who wishes to meet player 1, and the right table for the type who wishes to avoid player 1). In a Nash equilibrium, player 1's action is a best response in Figure 275.1 to the pair of actions of the two types of player 2, the action of the type of player 2 who wishes to meet player 1 is a best response in the left table of Figure 274.1 to the action of player 1, and the action of the type of player 2 who wishes to avoid player 1 is a best response in the right table of Figure 274.1 to the action of player 1.

Why should player 2, who knows whether she wants to meet or avoid player 1, have to plan what to do in both cases? She does not have to do so! But we, as analysts, need to consider what she would do in both cases. The reason is that to determine her best action, player 1, who does not know player 2's type, needs to form a belief about the action each type would take, and we wish to impose the equilibrium condition that these beliefs be correct, in the sense that for each type of player 2 they specify a best response to player 1's equilibrium action. Thus the equilibrium action of player 2 for each of her possible types may be interpreted as player 1's (correct) belief about the action that each type of player 2 would take, not as a plan of action for player 2.

I claim that $(B,(B,S))$, where the first component is the action of player 1 and the other component is the pair of actions of the two types of player 2, is a Nash equilibrium. Given that the actions of the two types of player 2 are (B,S), player 1's action B is optimal, from Figure 275.1; given that player 1 chooses B, B is optimal for the type who wishes to meet player 1 and S is optimal for the type who wishes to avoid player 1, from Figure 274.1. Suppose that in fact player 2 wishes to meet player 1. Then we interpret the equilibrium as follows. Both player 1 and player 2 choose B; player 1, who does not know if player 2 wants to meet her or avoid her, believes that if player 2 wishes to meet her she will choose B, and if she wishes to avoid her she will choose S.

? EXERCISE 276.1 (Equilibria of a variant of *BoS* with imperfect information) Show that there is no pure strategy Nash equilibrium of this game in which player 1 chooses S. If you have studied mixed strategy Nash equilibrium (Chapter 4), find the mixed strategy Nash equilibria of the game. (First check whether there is an equilibrium in which both types of player 2 use pure strategies; then look for equilibria in which one or both of these types randomize.)

We can interpret the actions of the two types of player 2 to reflect player 2's intentions in the hypothetical situation *before* she knows the state. We can tell the following story. Initially player 2 does not know the state; she is informed of the state by a *signal* that depends on the state. Before receiving this signal, she plans an action for each possible signal. After receiving the signal she carries out her planned action for that signal. We can tell a similar story for player 1. To be consistent with her not knowing the state when she takes an action, her signal must be uninformative: it must be the same in each state. Given her signal, she is unsure of the state; when choosing an action she takes into account her belief about the likelihood of each state, given her signal. The framework of states, beliefs, and signals, which is unnecessarily baroque in this simple example, comes into its own in the analysis of more complex situations.

EXAMPLE 276.2 (Variant of *BoS* with imperfect information) Consider another variant of the situation modeled by *BoS*, in which neither player knows whether the other wants to go out with her. Specifically, suppose that player 1 thinks that with probability $\frac{1}{2}$ player 2 wants to go out with her, and with probability $\frac{1}{2}$ player 2 wants to avoid her, and player 2 thinks that with probability $\frac{2}{3}$ player 1 wants to go out with her and with probability $\frac{1}{3}$ player 1 wants to avoid her. As before, assume that each player knows her own preferences.

We can model this situation by introducing four states, one for each of the possible configurations of preferences. I refer to these states as yy (each player wants to go out with the other), yn (player 1 wants to go out with player 2, but player 2 wants to avoid player 1), ny, and nn.

The fact that player 1 does not know player 2's preferences means that she cannot distinguish between states yy and yn, or between states ny and nn. Similarly, player 2 cannot distinguish between states yy and ny, or between states yn and nn. We can model the players' information by assuming that each player receives a *signal* before choosing an action. Player 1 receives the same signal, say y_1, in states yy and yn, and a different signal, say n_1, in states ny and nn; player 2 receives the same signal, say y_2, in states yy and ny, and a different signal, say n_2, in states yn and nn. After player 1 receives the signal y_1, she is referred to as *type* y_1

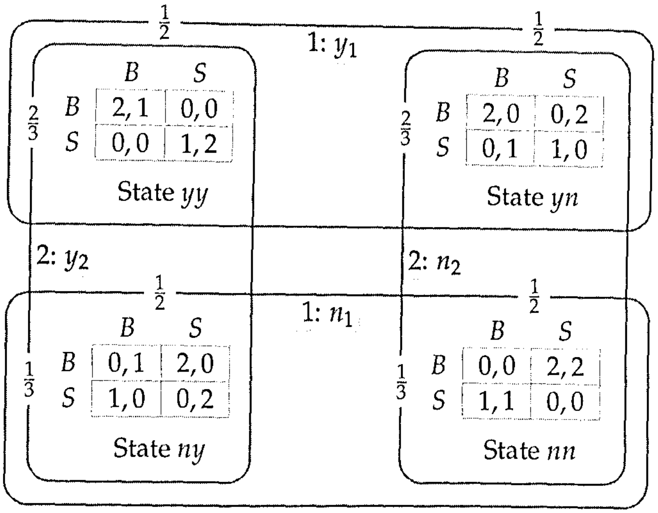

Figure 277.1 A variant of *BoS* in which each player is unsure of the other player's preferences. The frame labeled i: x encloses the states that generate the signal x for player i; the numbers appearing over this frame next to each table are the probabilities assigned by type x of player i to each state she regards to be possible.

of player 1 (who wishes to go out with player 2); after she receives the signal n_1 she is referred to as *type* n_1 of player 1 (who wishes to avoid player 2). Similarly, player 2 has two *types*, y_2 and n_2.

Type y_1 of player 1 believes that the probability of each of the states yy and yn is $\frac{1}{2}$; type n_1 of player 1 believes that the probability of each of the states ny and nn is $\frac{1}{2}$. Similarly, type y_2 of player 2 believes that the probability of state yy is $\frac{2}{3}$ and that of state ny is $\frac{1}{3}$; type n_2 of player 2 believes that the probability of state yn is $\frac{2}{3}$ and that of state nn is $\frac{1}{3}$. This model of the situation is illustrated in Figure 277.1.

As in Example 273.1, to study the equilibria of this model we consider the players' plans of action before they receive their signals. That is, each player plans an action for each of the two possible signals she may receive. We may think of there being four players: the two types of player 1 and the two types of player 2. A *Nash equilibrium* consists of four actions, one for each of these players, such that the action of each type of each original player is optimal, given her belief about the state after observing her signal, and given the actions of each type of the other original player.

Consider type y_1 of player 1. Her beliefs about the probabilities of states yy and yn and her payoffs in these states are the same as the beliefs and payoffs of player 1 for the two states in Example 273.1. Thus her expected payoffs for the four pairs of actions of the two types of player 2 are given in Figure 275.1, interpreting each column to represent a pair of actions of types y_2 and n_2 of player 2.

? EXERCISE 277.1 (Expected payoffs in a variant of *BoS* with imperfect information) Construct tables like the one in Figure 275.1 for type n_1 of player 1, and for types y_2 and n_2 of player 2.

I claim that $((B, B), (B, S))$ and $((S, B), (S, S))$ are Nash equilibria of the game, where in each case the first component gives the actions of the two types of player 1 and the second component gives the actions of the two types of player 2. You may use Figure 275.1 to verify that B is a best response of type y_1 of player 1 to the pair (B, S) of actions of player 2, and S is a best response to the pair of actions (S, S). You may use your answer to Exercise 277.1 to verify that in each of the claimed Nash equilibria the action of type n_1 of player 1 and the action of each type of player 2 is a best response to the other players' actions.

In each of these examples a Nash equilibrium is a list of actions, one for each type of each player, such that the action of each type of each player is a best response to the actions of all the types of the other player, given the player's beliefs about the state after she observes her signal. The actions planned by the various types of player i are not relevant to the decision problem of any type of player i, but there is no harm in taking them, as well as the actions of the types of the *other* player, as given when player i is choosing an action. Thus we may define a Nash equilibrium in each example to be a Nash equilibrium of the strategic game in which the set of players is the set of all types of all players in the original situation.

In the next section I define the general notion of a Bayesian game and the notion of Nash equilibrium in such a game. These definitions require significant theoretical development. If you find the theory in the next section heavy going, you may be able to skim the section and then study the subsequent illustrations, relying on the intuition developed in the examples in this section, and returning to the theory only as necessary for clarification.

9.2 General definitions

9.2.1 Bayesian games

A strategic game with imperfect information is called a "Bayesian game". (The reason for this nomenclature will become apparent.) As in a strategic game, the decision-makers are called *players*, and each player is endowed with a set of *actions*.

A key component in the specification of the imperfect information is the set of *states*. Each state is a complete description of one collection of the players' relevant characteristics, including both their preferences and their information. For every collection of characteristics that some player believes to be possible, there must be a state. For instance, suppose in Example 273.1 that player 2 wishes to meet player 1. In this case, the reason for including in the model the state in which player 2 wishes to avoid player 1 is that player 1 believes such a preference to be possible.

At the start of the game a state is realized. The players do not observe this state. Rather, each player receives a *signal* that may give her some information about the state. Denote the signal player i receives in state ω by $\tau_i(\omega)$. The function τ_i is called player i's *signal function*. (Note that the signal is a *deterministic* function of the state: for each state, a definite signal is received.) The states that generate any given signal t_i are said to be *consistent* with t_i. The sizes of the sets of states consistent with each of player i's signals reflect the quality of player i's information. If, for example, $\tau_i(\omega)$ is different for each value of ω, then player i knows, given her signal, the state that has occurred; after receiving her signal, she is perfectly informed about all the players' relevant characteristics. At the other extreme, if $\tau_i(\omega)$ is the same for all states, then player i's signal conveys no information about the state. If $\tau_i(\omega)$ is constant over some subsets of the set of states, but is not the same for all states, then player i's signal conveys partial information. For example, if there are three states, ω_1, ω_2, and ω_3, and $\tau_i(\omega_1) \neq \tau_i(\omega_2) = \tau_i(\omega_3)$, then when the state is ω_1 player i knows that it is ω_1, whereas when it is either ω_2 or ω_3 she knows only that it is one of these two states. (In Figures 274.1 and 277.1, each frame encloses a set of states that yield the same signal for one of the players.)

We refer to player i in the event that she receives the signal t_i as *type* t_i of player i. Each type of each player holds a *belief* about the likelihood of the states consistent with her signal. If, for example, $t_i = \tau_i(\omega_1) = \tau_i(\omega_2)$, then type t_i of player i assigns probabilities to ω_1 and ω_2. (A player who receives a signal consistent with only one state naturally assigns probability 1 to that state.)

Each player may care about the actions chosen by the other players, as in a strategic game with perfect information, and also about the state. The players may be uncertain about the state, so we need to specify their preferences regarding probability distributions over pairs (a, ω) consisting of an action profile a and a state ω. I assume that each player's preferences over such probability distributions are represented by the expected value of a *Bernoulli payoff function.* Thus I specify each player i's preferences by giving a Bernoulli payoff function u_i over pairs (a, ω). (Note that in Examples 273.1 and 276.2, both players care only about the other player's action, not independently about the state.)

In summary, a Bayesian game is defined as follows.

DEFINITION 279.1 (*Bayesian game*) A **Bayesian game** consists of

- a set of **players**
- a set of **states**

and for each player

- a set of **actions**
- a set of **signals** that she may receive and a **signal function** that associates a signal with each state
- for each signal that she may receive, a **belief** about the states consistent with the signal (a probability distribution over the set of states with which the signal is associated)
- a **Bernoulli payoff function** over pairs (a, ω), where a is an action profile and ω is a state, the expected value of which represents the player's preferences among lotteries over the set of such pairs.

Note that the set of actions of each player is independent of the state. Each player may care about the state, but the set of actions available to her is the same in every state.

The eponymous Thomas Bayes (1702–61) first showed how probabilities should be changed in the light of new information. His formula (discussed in Section 17.6.5) is needed when one works with a variant of Definition 279.1 in which each player is endowed with a "prior" belief about the states, from which the belief of each of her types is derived. For the purposes of this chapter, the belief of each type of each player is more conveniently taken as a primitive, rather than being derived from a prior belief.

The game in Example 273.1 fits into this general definition as follows.

Players The pair of people.

States The set of states is $\{meet, avoid\}$.

Actions The set of actions of each player is $\{B, S\}$.

Signals Player 1 may receive a single signal, say z; her signal function τ_1 satisfies $\tau_1(meet) = \tau_1(avoid) = z$. Player 2 receives one of two signals, say m and v; her signal function τ_2 satisfies $\tau_2(meet) = m$ and $\tau_2(avoid) = v$.

Beliefs Player 1 assigns probability $\frac{1}{2}$ to each state after receiving the signal z. Player 2 assigns probability 1 to the state *meet* after receiving the signal m, and probability 1 to the state *avoid* after receiving the signal v.

Payoffs The payoffs $u_i(a, meet)$ of each player i for all possible action pairs are given in the left panel of Figure 274.1, and the payoffs $u_i(a, avoid)$ are given in the right panel.

Similarly, the game in Example 276.2 fits into the definition as follows.

Players The pair of people.

States The set of states is $\{yy, yn, ny, nn\}$.

Actions The set of actions of each player is $\{B, S\}$.

Signals Player 1 receives one of two signals, y_1 and n_1; her signal function τ_1 satisfies $\tau_1(yy) = \tau_1(yn) = y_1$ and $\tau_1(ny) = \tau_1(nn) = n_1$. Player 2 receives one of two signals, y_2 and n_2; her signal function τ_2 satisfies $\tau_2(yy) = \tau_2(ny) = y_2$ and $\tau_2(yn) = \tau_2(nn) = n_2$.

Beliefs Player 1 assigns probability $\frac{1}{2}$ to each of the states yy and yn after receiving the signal y_1 and probability $\frac{1}{2}$ to each of the states ny and nn after receiving the signal n_1. Player 2 assigns probability $\frac{2}{3}$ to the state yy and probability $\frac{1}{3}$ to the state ny after receiving the signal y_2, and probability $\frac{2}{3}$ to the state yn and probability $\frac{1}{3}$ to the state nn after receiving the signal n_2.

Payoffs The payoffs $u_i(a, \omega)$ of each player i for all possible action pairs and states are given in Figure 277.1.

9.2.2 Nash equilibrium

In a strategic game, each player chooses an action. In a Bayesian game, each player chooses a collection of actions, one for each signal she may receive. That is, in a Bayesian game *each type of each player* chooses an action. In a Nash equilibrium of such a game, the action chosen by each type of each player is optimal, given the actions chosen by every type of every other player. (In a steady state, each player's experience teaches her these actions.) Any given type of player i is not affected by the actions chosen by the other types of player i, so there is no harm in thinking that player i takes as given these actions, as well as those of the other players. Thus we may define a Nash equilibrium of a Bayesian game to be a Nash equilibrium of a strategic game in which each player is one of the types of one of the players in the Bayesian game. What is each player's payoff function in this strategic game?

Consider type t_i of player i. For each state ω she knows every other player's type (i.e. she knows the signal received by every other player). This information, together with her belief about the states, allows her to calculate her expected payoff for each of her actions and each collection of actions for the various types of the other players. For instance, in Example 273.1, player 1's belief is that the probability of each state is $\frac{1}{2}$, and she knows that player 2 is type m in the state *meet* and type v in the state *avoid*. Thus if type m of player 2 chooses B and type v of player 2 chooses S, player 1 thinks that if she chooses B then her expected payoff is

$$\tfrac{1}{2}u_1((B,B),\mathit{meet}) + \tfrac{1}{2}u_1((B,S),\mathit{avoid}),$$

where u_1 is her payoff function in the Bayesian game. (In general her payoff may depend on the state, though in this example it does not.) The top box of the second column in Figure 275.1 gives this payoff; the other boxes give player 1's payoffs for her other action and the other combinations of actions for the two types of player 2.

In a general game, denote the probability assigned by the belief of type t_i of player i to state ω by $\Pr(\omega \mid t_i)$. Denote the action taken by each type t_j of each player j by $a(j,t_j)$. Player j's signal in state ω is $\tau_j(\omega)$, so her action in this state is $a(j,\tau_j(\omega))$. For each state ω and each player j, let $\hat{a}_j(\omega) = a(j,\tau_j(\omega))$. Then the expected payoff of type t_i of player i when she chooses the action a_i is

$$\sum_{\omega\in\Omega} \Pr(\omega \mid t_i)u_i((a_i,\hat{a}_{-i}(\omega)),\omega), \tag{281.1}$$

where Ω is the set of states and $(a_i,\hat{a}_{-i}(\omega))$ is the action profile in which player i chooses the action a_i and every other player j chooses $\hat{a}_j(\omega)$. (Note that this expected payoff does not depend on the actions of any other types of player i, but only on the actions of the various types of the *other* players.)

We may now define precisely a Nash equilibrium of a Bayesian game.

DEFINITION 281.2 (*Nash equilibrium of Bayesian game*) A **Nash equilibrium of a Bayesian game** is a Nash equilibrium of the strategic game (with vNM preferences) defined as follows.

Players The set of all pairs (i, t_i) in which i is a player in the Bayesian game and t_i is one of the signals that i may receive.

Actions The set of actions of each player (i, t_i) is the set of actions of player i in the Bayesian game.

Preferences The Bernoulli payoff function of each player (i, t_i) is given by (281.1).

? EXERCISE 282.1 (Fighting an opponent of unknown strength) Two people are involved in a dispute. Person 1 does not know whether person 2 is strong or weak; she assigns probability α to person 2's being strong. Person 2 is fully informed. Each person can either fight or yield. Each person's preferences are represented by the expected value of a Bernoulli payoff function that assigns the payoff of 0 if she yields (regardless of the other person's action) and a payoff of 1 if she fights and her opponent yields; if both people fight, then their payoffs are $(-1, 1)$ if person 2 is strong and $(1, -1)$ if person 2 is weak. Formulate this situation as a Bayesian game and find its Nash equilibria if $\alpha < \frac{1}{2}$ and if $\alpha > \frac{1}{2}$.

? EXERCISE 282.2 (An exchange game) Each of two individuals receives a ticket on which there is an integer from 1 to m indicating the size of a prize she may receive. The individuals' tickets are assigned randomly and independently; the probability of an individual's receiving each possible number is positive. Each individual is given the option of exchanging her prize for the other individual's prize; the individuals are given this option simultaneously. If both individuals wish to exchange, then the prizes are exchanged; otherwise each individual receives her own prize. Each individual's objective is to maximize her expected monetary payoff. Model this situation as a Bayesian game and show that in any Nash equilibrium the highest prize that either individual is willing to exchange is the smallest possible prize.

? EXERCISE 282.3 (Adverse selection) Firm A (the "acquirer") is considering taking over firm T (the "target"). It does not know firm T's value; it believes that this value, when firm T is controlled by its own management, is at least \$0 and at most \$100, and assigns equal probability to each of the 101 dollar values in this range. Firm T will be worth 50% more under firm A's management than it is under its own management. Suppose that firm A bids y to take over firm T, and firm T is worth x (under its own management). Then if T accepts A's offer, A's payoff is $\frac{3}{2}x - y$ and T's payoff is y; if T rejects A's offer, A's payoff is 0 and T's payoff is x. Model this situation as a Bayesian game in which firm A chooses how much to offer and firm T decides the lowest offer to accept. Find the Nash equilibrium (equilibria?) of this game. Explain why the logic behind the equilibrium is called *adverse selection*.

9.3 Two examples concerning information

The notion of a Bayesian game may be used to study how information patterns affect the outcome of strategic interaction. Here are two examples.

$\frac{1}{2}$ 1 $\frac{1}{2}$

$\frac{1}{2}$ 2 $\frac{1}{2}$

	L	M	R
T	$1,2\epsilon$	$1,0$	$1,3\epsilon$
B	$2,2$	$0,0$	$0,3$

State ω_1

	L	M	R
T	$1,2\epsilon$	$1,3\epsilon$	$1,0$
B	$2,2$	$0,3$	$0,0$

State ω_2

Figure 283.1 The first Bayesian game considered in Section 9.3.1.

9.3.1 *More information may hurt*

A decision-maker in a single-person decision problem cannot be worse off if she has more information: if she wishes, she can ignore the information. In a game the same is not true: if a player has more information and the other players know that she has more information, then she may be worse off.

Consider, for example, the two-player Bayesian game in Figure 283.1, where $0 < \epsilon < \frac{1}{2}$. In this game there are two states, and neither player knows the state. Player 2's unique best response to each action of player 1 is L. (If player 1 chooses T, L yields 2ϵ, whereas M and R each yield $\frac{3}{2}\epsilon$; if player 1 chooses B, L yields 2, whereas M and R each yield $\frac{3}{2}$.) Further, player 1's unique best response to L is B. Thus (B, L) is the unique Nash equilibrium of the game; it yields each player a payoff of 2. (If you have studied Chapter 4, you will be able to verify that, moreover, the game has no other mixed strategy equilibrium.)

Now consider the variant of this game in which player 2 is informed of the state: player 2's signal function τ_2 satisfies $\tau_2(\omega_1) \neq \tau_2(\omega_2)$. In this game $(T, (R, M))$ is the unique Nash equilibrium. (Each type of player 2 has a strictly dominant action, to which T is player 1's unique best response.)

Player 2's payoff in the unique Nash equilibrium of the original game is 2, whereas her payoff in the unique Nash equilibrium of the game in which she knows the state is 3ϵ in each state. Thus she is worse off when she knows the state than when she does not. To understand this result, notice that among player 2's actions, R is good only in state ω_1, M is good only in state ω_2, and L is a compromise. When she does not know the state she chooses L, inducing player 1 to choose B. When she is fully informed she tailors her action to the state, choosing R in state ω_1 and M in state ω_2, inducing player 1 to choose T. The game has no steady state in which she ignores her information and chooses L because this action leads player 1 to choose B, making R better for player 2 in state ω_1 and M better in state ω_2.

9.3.2 *Infection*

The notion of a Bayesian game may be used to model not only situations in which players are uncertain about each others' preferences, but also situations in which they are uncertain about each others' *knowledge*. Consider, for example, the Bayesian game in Figure 284.1.

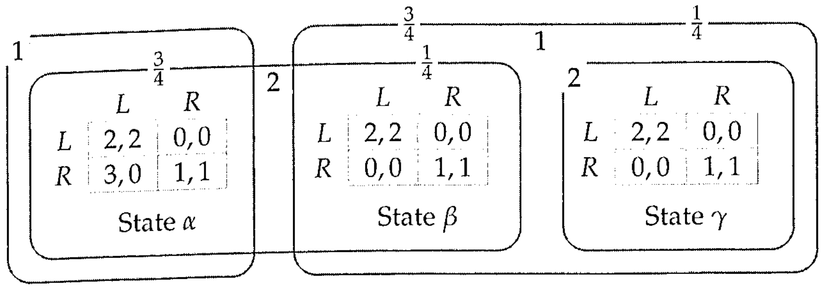

Figure 284.1 The first Bayesian game in Section 9.3.2. In the unique Nash equilibrium of this game, each type of each player chooses R.

Notice that player 2's preferences are the same in all three states, and player 1's preferences are the same in states β and γ. In particular, in state γ, each player knows the other player's preferences, and player 2 knows that player 1 knows her preferences. The shortcoming in the players' information in state γ is that player 1 does not know that player 2 knows her preferences: player 1 knows only that the state is either β or γ, and in state β player 2 does not know whether the state is α or β, and hence does not know player 1's preferences (because player 1's preferences in these two states differ).

This imperfection in player 1's knowledge of player 2's information significantly affects the equilibria of the game. If information were perfect in state γ, then both (L, L) and (R, R) would be Nash equilibria. However, the whole game has a *unique* Nash equilibrium, in which the outcome in state γ is (R, R), as you are asked to show in the next exercise. The argument shows that the incentives faced by player 1 in state α "infect" the remainder of the game.

? EXERCISE 284.1 (Infection) Show that the Bayesian game in Figure 284.1 has a unique Nash equilibrium, in which each player chooses R regardless of her signal. (Start by considering player 1's action in state α. Next consider player 2's action when she gets the signal that the state is α or β. Then consider player 1's action when she gets the signal that the state is β or γ. Finally consider player 2's action in state γ.)

Now extend the game as in Figure 285.1. Consider state δ. In this state, player 2 knows player 1's preferences (because she knows that the state is either γ or δ, and in both states player 1's preferences are the same). What player 2 does not know is whether player 1 knows that player 2 knows player 1's preferences. The reason is that player 2 does not know whether the state is γ or δ; and in state γ player 1 does not know that player 2 knows her preferences, because she does not know whether the state is β or γ, and in state β player 2 (who does not know whether the state is α or β) does not know her preferences. Thus the level of the shortcoming in the players' information is higher than it is in the game in Figure 284.1. Nevertheless, the incentives faced by player 1 in state α again "infect" the remainder of the game, and in the only Nash equilibrium every type of each player chooses R.

The game may be further extended. As it is extended, the level of the imperfection in the players' information in the last state increases. When the number of

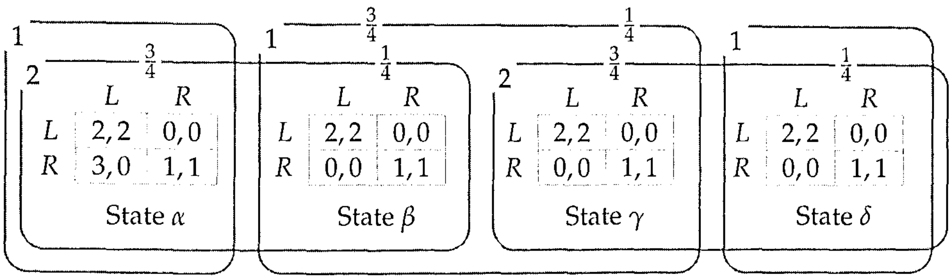

Figure 285.1 The second Bayesian game in Section 9.3.2.

states is large, the players' information in the last state is only very slightly imperfect. Nevertheless, the incentives of player 1 in state α still cause the game to have a unique Nash equilibrium, in which every type of each player chooses R.

In each of these examples, the equilibrium induces an outcome in every state that is worse for both players than another outcome (namely (L, L)); in all states but the first, the alternative outcome is a Nash equilibrium in the game with perfect information. For some other specifications of the payoffs in state α and the players' beliefs, the game has a unique equilibrium in which the "good" outcome (L, L) occurs in every state; the point is only that one of the two Nash equilibria is selected, not that the "bad" equilibrium is necessarily selected. (Modify the payoffs of player 1 in state α so that L strictly dominates R, and change the beliefs to assign probability $\frac{1}{2}$ to each state compatible with each signal.)

9.4 Illustration: Cournot's duopoly game with imperfect information

9.4.1 *Imperfect information about cost*

Two firms compete in selling a good; one firm does not know the other firm's cost function. How does this lack of information affect the firms' behavior?

Assume that both firms can produce the good at constant unit cost. Assume also that they both know that firm 1's unit cost is c, but only firm 2 knows its own unit cost; firm 1 believes that firm 2's cost is c_L with probability θ and c_H with probability $1 - \theta$, where $0 < \theta < 1$ and $c_L < c_H$.

We may model this situation as a Bayesian game that is a variant of Cournot's game (Section 3.1).

Players Firm 1 and firm 2.

States $\{L, H\}$.

Actions Each firm's set of actions is the set of its possible outputs (nonnegative numbers).

Signals Firm 1's signal function τ_1 satisfies $\tau_1(H) = \tau_1(L)$ (its signal is the same in both states); firm 2's signal function τ_2 satisfies $\tau_2(H) \neq \tau_2(L)$ (its signal is perfectly informative of the state).

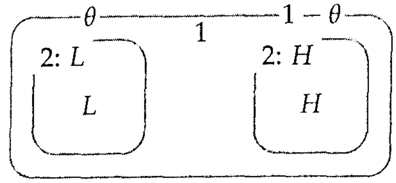

Figure 286.1 The information structure for the model in the variant of Cournot's model in Section 9.4.1, in which firm 1 does not know firm 2's cost. The frame labeled 2: x, for $x = L$ and $x = H$, encloses the state that generates the signal x for firm 2.

Beliefs The single type of firm 1 assigns probability θ to state L and probability $1-\theta$ to state H. Each type of firm 2 assigns probability 1 to the single state consistent with its signal.

Payoff functions The firms' Bernoulli payoffs are their profits; if the actions chosen are (q_1, q_2) and the state is I (either L or H), then firm 1's profit is $q_1(P(q_1+q_2)-c)$ and firm 2's profit is $q_2(P(q_1+q_2)-c_I)$, where $P(q_1+q_2)$ is the market price when the firms' outputs are q_1 and q_2.

The information structure in this game is similar to that in Example 273.1; it is illustrated in Figure 286.1.

A Nash equilibrium of this game is a triple (q_1^*, q_L^*, q_H^*), where q_1^* is the output of firm 1, q_L^* is the output of type L of firm 2 (i.e. firm 2 when it receives the signal $\tau_2(L)$), and q_H^* is the output of type H of firm 2 (i.e. firm 2 when it receives the signal $\tau_2(H)$), such that

- q_1^* maximizes firm 1's profit given the output q_L^* of type L of firm 2 and the output q_H^* of type H of firm 2
- q_L^* maximizes the profit of type L of firm 2 given the output q_1^* of firm 1
- q_H^* maximizes the profit of type H of firm 2 given the output q_1^* of firm 1.

To find an equilibrium, we first find the firms' best response functions. Given firm 1's beliefs, its best response $b_1(q_L, q_H)$ to (q_L, q_H) solves

$$\max_{q_1} \left[\theta(P(q_1+q_L)-c)q_1 + (1-\theta)(P(q_1+q_H)-c)q_1\right].$$

Firm 2's best response $b_L(q_1)$ to q_1 when its cost is c_L solves

$$\max_{q_L} \left[(P(q_1+q_L)-c_L)q_L\right],$$

and its best response $b_H(q_1)$ to q_1 when its cost is c_H solves

$$\max_{q_H} \left[(P(q_1+q_H)-c_H)q_H\right].$$

A Nash equilibrium is a triple (q_1^*, q_L^*, q_H^*) such that

$$q_1^* = b_1(q_L^*, q_H^*),\ q_L^* = b_L(q_1^*), \text{ and } q_H^* = b_H(q_1^*).$$

? EXERCISE 287.1 (Cournot's duopoly game with imperfect information) Consider the game when the inverse demand function is given by $P(Q) = \alpha - Q$ for $Q \leq \alpha$ and $P(Q) = 0$ for $Q > \alpha$ (see (56.2)). For values of c_H and c_L close enough that there is a Nash equilibrium in which all outputs are positive, find this equilibrium. Compare this equilibrium with the Nash equilibrium of the game in which firm 1 knows that firm 2's unit cost is c_L, and with the Nash equilibrium of the game in which firm 1 knows that firm 2's unit cost is c_H.

9.4.2 *Imperfect information about both cost and information*

Now suppose that firm 2 does not know whether firm 1 knows firm 2's cost. That is, suppose that one circumstance that firm 2 believes to be possible is that firm 1 knows its cost (although in fact it does not). Because firm 2 thinks this circumstance to be possible, we need *four* states to model the situation, which I call $L0$, $H0$, $L1$, and $H1$, with the following interpretations.

$L0$: firm 2's cost is low and firm 1 does not know whether it is low or high

$H0$: firm 2's cost is high and firm 1 does not know whether it is low or high

$L1$: firm 2's cost is low and firm 1 knows it is low

$H1$: firm 2's cost is high and firm 1 knows it is high.

Firm 1 receives one of three possible signals, 0, L, and H. The states $L0$ and $H0$ generate the signal 0 (firm 1 does not know firm 2's cost), the state $L1$ generates the signal L (firm 1 knows firm 2's cost is low), and the state $H1$ generates the signal H (firm 1 knows firm 2's cost is high). Firm 2 receives one of two possible signals, L, in states $L0$ and $L1$, and H, in states $H0$ and $H1$. Denote by θ (as before) the probability assigned by type 0 of firm 1 to firm 2's cost being c_L, and by π the probability assigned by each type of firm 2 to firm 1's knowing firm 2's cost. (The case $\pi = 0$ is equivalent to the one considered in Section 9.4.1.) A Bayesian game that models the situation is defined as follows.

Players Firm 1 and firm 2.

States $\{L0, L1, H0, H1\}$, where the first letter in the name of the state indicates firm 2's cost and the second letter indicates whether firm 1 does (1) or does not (0) know firm 2's cost.

Actions Each firm's set of actions is the set of its possible outputs (nonnegative numbers).

Signals Firm 1 gets one of the signals 0, L, and H, and her signal function τ_1 satisfies $\tau_1(L0) = \tau_1(H0) = 0$, $\tau_1(L1) = L$, and $\tau_1(H1) = H$. Firm 2 gets the signal L or H and her signal function τ_2 satisfies $\tau_2(L0) = \tau_2(L1) = L$ and $\tau_2(H0) = \tau_2(H1) = H$.

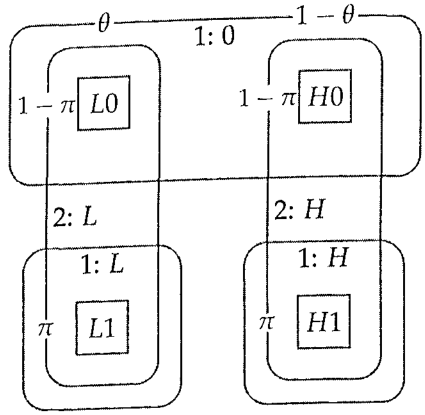

Figure 288.1 The information structure for the model in Section 9.4.2, in which firm 2 does not know whether firm 1 knows its cost. The frame labeled i: x encloses the states that generate the signal x for firm i.

Beliefs Firm 1: type 0 assigns probability θ to state $L0$ and probability $1-\theta$ to state $H0$; type L assigns probability 1 to state $L1$; type H assigns probability 1 to state H. Firm 2: type L assigns probability π to state $L1$ and probability $1-\pi$ to state $L0$; type H assigns probability π to state $H1$ and probability $1-\pi$ to state $H0$.

Payoff functions The firms' Bernoulli payoffs are their profits; if the actions chosen are (q_1, q_2), then firm 1's profit is $q_1(P(q_1+q_2)-c)$ and firm 2's profit is $q_2(P(q_1+q_2)-c_L)$ in states $L0$ and $L1$, and $q_2(P(q_1+q_2)-c_L)$ in states $H0$ and $H1$.

The information structure in this game is illustrated in Figure 288.1. You are asked to investigate its Nash equilibria in the following exercise.

? EXERCISE 288.1 (Cournot's duopoly game with imperfect information) Write down the maximization problems that determine the best response function of each type of each player. (Denote by q_0, q_ℓ, and q_h the outputs of types 0, ℓ, and h of firm 1, and by q_L and q_H the outputs of types L and H of firm 2.) Now suppose that the inverse demand function is given by $P(Q) = \alpha - Q$ for $Q \leq \alpha$ and $P(Q) = 0$ for $Q > \alpha$. For values of c_H and c_L close enough that there is a Nash equilibrium in which all outputs are positive, find this equilibrium. Check that when $\pi = 0$ the equilibrium output of type 0 of firm 1 is equal to the equilibrium output of firm 1 you found in Exercise 287.1, and that the equilibrium outputs of the two types of firm 2 are the same as the ones you found in that exercise. Check also that when $\pi = 1$ the equilibrium outputs of type ℓ of firm 1 and type L of firm 2 are the same as the equilibrium outputs when there is perfect information and the costs are c and c_L, and that the equilibrium outputs of type h of firm 1 and type H of firm 2 are the same as the equilibrium outputs when there is perfect information and the costs are c and c_H. Show that for $0 < \pi < 1$, the equilibrium outputs of types L and H of firm 2 lie between their values when $\pi = 0$ and when $\pi = 1$.

9.5 Illustration: providing a public good

Suppose that a public good is provided to a group of people if at least one person is willing to pay the cost of the good (as in the model of crime reporting in Section 4.8). Assume that the people differ in their valuations of the good, and each person knows only her own valuation. Who, if anyone, will pay the cost?

Denote the number of individuals by n, the cost of the good by $c > 0$, and individual i's payoff if the good is provided by v_i. If the good is not provided, then each individual's payoff is 0. Each individual i knows her own valuation v_i. She does not know anyone else's valuation, but she knows that all valuations are at least $\underline{v}$ and at most $\overline{v}$, where $0 \le \underline{v} < c < \overline{v}$. She believes that the probability that any one individual's valuation is at most v is $F(v)$, independent of all other individuals' valuations, where F is a continuous increasing function. The fact that F is increasing means that the individual does not assign zero probability to any range of values between $\underline{v}$ and $\overline{v}$; the fact that it is continuous means that she does not assign positive probability to any single valuation. (An example of the function F is shown in Figure 289.1.)

The following mechanism determines whether the good is provided. All n individuals simultaneously submit envelopes; the envelope of any individual i may contain either a contribution of c or nothing (no intermediate contributions are allowed). If all individuals submit 0, then the good is not provided and each individual's payoff is 0. If at least one individual submits c, then the good is provided, each individual i who submits c obtains the payoff $v_i - c$, and each individual i who submits 0 obtains the payoff v_i. (The pure strategy Nash equilibria of a variant of this model, in which more than one contribution is needed to provide the good, are considered in Exercise 33.1.)

We can formulate this situation as a Bayesian game as follows.

Players The set of n individuals.

States The set of all profiles $(v_1, \ldots, v_n)$ of valuations, where $\underline{v} \le v_i \le \overline{v}$ for all i.

Actions Each player's set of actions is $\{0, c\}$.

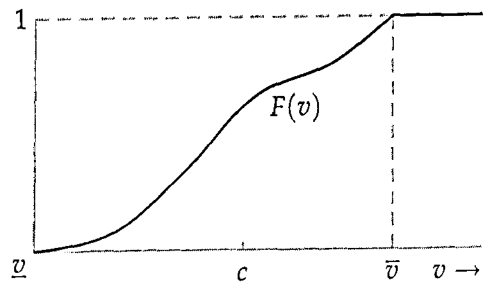

Figure 289.1 An example of the function F for the model in Section 9.5. For each value of v, $F(v)$ is the probability that any given individual's valuation is at most v.

Signals The set of signals that each player may observe is the set of possible valuations. The signal function τ_i of each player i is given by $\tau_i(v_1, \ldots, v_n) = v_i$ (each player knows her own valuation).

Beliefs Every type of player i assigns probability $F(v_1)F(v_2)\cdots F(v_{i-1}) \times F(v_{i+1})\cdots F(v_n)$ to the event that the valuation of every other player j is at most v_j.

Payoff functions Player i's Bernoulli payoff in state $(v_1, \ldots, v_n)$ is

$$\begin{cases} 0 & \text{if no one contributes} \\ v_i & \text{if } i \text{ does not contribute but some other player does} \\ v_i - c & \text{if } i \text{ contributes.} \end{cases}$$

? EXERCISE 290.1 (Nash equilibria of game of contributing to a public good) Find conditions under which for each value of i this game has a pure strategy Nash equilibrium in which each type v_i of player i with $v_i \geq c$ contributes, whereas every other type of player i, and all types of every other player, do not contribute.

In addition to the Nash equilibria identified in this exercise, the game has a symmetric Nash equilibrium in which every player contributes if and only if her valuation exceeds some critical amount v^*. For such a strategy profile to be an equilibrium, a player whose valuation is less than v^* must optimally not contribute, and a player whose valuation is at least v^* must optimally contribute. Consider player i. Suppose that every other player contributes if and only if her valuation is at least v^*. The probability that at least one of the other players contributes is the probability that at least one of the other players' valuations is at least v^*, which is $1 - (F(v^*))^{n-1}$. (Note that $(F(v^*))^{n-1}$ is the probability that all the other valuations are at most v^*.) Thus if player i's valuation is v_i, her expected payoff is $(1 - (F(v^*))^{n-1})v_i$ if she does not contribute and $v_i - c$ if she does contribute. Hence the conditions for player i to optimally not contribute when $v_i < v^*$ and optimally contribute when $v_i \geq v^*$ are $(1 - (F(v^*))^{n-1})v_i \geq v_i - c$ if $v_i < v^*$, and $(1 - (F(v^*))^{n-1})v_i \leq v_i - c$ if $v_i \geq v^*$, or equivalently

$$\begin{aligned} v_i(F(v^*))^{n-1} &\leq c \quad \text{if } v_i < v^* \\ v_i(F(v^*))^{n-1} &\geq c \quad \text{if } v_i \geq v^*. \end{aligned} \tag{290.2}$$

If these inequalities are satisfied, then

$$v^*(F(v^*))^{n-1} = c. \tag{290.3}$$

Conversely, if v^* satisfies (290.3), then it satisfies the two inequalities in (290.2). Thus the game has a Nash equilibrium in which every player contributes whenever her valuation is at least v^* if and only if v^* satisfies (290.3).

Note that because $F(v) = 1$ only if $v \geq \overline{v}$, and $\overline{v} > c$, we have $v^* > c$. That is, every player's cutoff for contributing exceeds the cost of the public good. When at least one player's valuation exceeds c, all players are better off if the public good is provided and the high-valuation player contributes than if the good is not provided. But in the equilibrium, the good is provided only if at least one player's valuation exceeds v^*, which exceeds c.

As the number of individuals increases, is the good more or less likely to be provided in this equilibrium? The probability that the good is provided is the probability that at least one player's valuation is at least v^*, which is equal to $1 - (F(v^*))^n$. (Note that $(F(v^*))^n$ is the probability that every player's valuation is less than v^*.) From (290.3) this probability is equal to $1 - cF(v^*)/v^*$. How does v^* vary with n? As n increases, for any given value of v^* the value of $(F(v^*))^{n-1}$ decreases, and thus the value of $v^*(F(v^*))^{n-1}$ decreases. Thus to maintain the equality (290.3), the value of v^* must increase as n increases. We conclude that as n increases, the change in the probability that the good is provided depends on the change in $F(v^*)/v^*$ as v^* increases: the probability increases if $F(v^*)/v^*$ is a decreasing function of v^*, whereas it decreases if $F(v^*)/v^*$ is an increasing function of v^*. If F is uniform and $\underline{v} > 0$, for example, $F(v^*)/v^*$ is an increasing function of v^*, so that the probability that the good is provided decreases as the population size increases.

The notion of a Bayesian game may be used to model a situation in which each player is uncertain of the number of other players. In the next exercise you are asked to study another variant of the crime-reporting model of Section 4.8 in which each of the two players does not know whether she is the only witness or whether there is another witness (in which case she knows that witness's valuation). (The exercise requires a knowledge of mixed strategy Nash equilibrium (Chapter 4).)

? EXERCISE 291.1 (Reporting a crime with an unknown number of witnesses) Consider the variant of the model of Section 4.8 in which each of two players does not know whether she is the only witness or whether there is another witness. Denote by π the probability each player assigns to being the sole witness. Model this situation as a Bayesian game with three states: one in which player 1 is the only witness, one in which player 2 is the only witness, and one in which both players are witnesses. Find a condition on π under which the game has a pure Nash equilibrium in which each player chooses *Call* (given the signal that she is a witness). When the condition is violated, find the symmetric mixed strategy Nash equilibrium of the game, and check that when $\pi = 0$ this equilibrium coincides with the one found in Section 4.8 for $n = 2$.

9.6 Illustration: auctions

9.6.1 Introduction

In the analysis of auctions in Section 3.5, every bidder knows every other bidder's valuation of the object for sale. Here I use the notion of a Bayesian game to analyze auctions in which bidders are not perfectly informed about each others' valuations.

Assume that a single object is for sale, and that each bidder independently receives some information—a "signal"—about the value of the object to her. If each bidder's signal is simply her valuation of the object, as assumed in Section 3.5, we say that the bidders' valuations are *private*. If each bidder's valuation depends on other bidders' signals as well as her own, we say that the valuations are *common*.

The assumption of private values is appropriate, for example, for a work of art whose beauty rather than resale value interests the buyers. Each bidder knows her valuation of the object, but not that of any other bidder; the other bidders' valuations have no bearing on her valuation. The assumption of common values is appropriate, for example, for an oil tract containing unknown reserves on which each bidder has conducted a test. Each bidder i's test result gives her some information about the size of the reserves, and hence her valuation of these reserves, but the other bidders' test results, if known to bidder i, would typically improve this information.

As in the analysis of auctions in which the bidders are perfectly informed about each others' valuations, I study models in which bids for a single object are submitted simultaneously (bids are *sealed*), and the participant who submits the highest bid obtains the object. As before I consider both *first-price* auctions, in which the winner pays the price she bid, and *second-price* auctions, in which the winner pays the highest of the remaining bids.

(In Section 3.5 I argue that the first-price rule models an open descending ("Dutch") auction, and the second-price rule models an open ascending ("English") auction. Note that the argument that the second-price rule corresponds to an open ascending auction depends upon the bidders' valuations being private. If a bidder is uncertain of her valuation, which is related to that of other bidders, then in an open ascending auction she may obtain information about her valuation from other participants' bids, information not available in a sealed-bid auction.)

I first consider the case in which the bidders' valuations are private, then the case in which they are common.

9.6.2 *Independent private values*

In the case in which the bidders' valuations are private, the assumptions about these valuations are similar to those in the previous section (on the provision of a public good). Each bidder knows that all other bidders' valuations are at least $\underline{v}$, where $\underline{v} \geq 0$, and at most $\overline{v}$. She believes that the probability that any given bidder's valuation is at most v is $F(v)$, independent of all other bidders' valuations, where F is a continuous increasing function (as in Figure 289.1).

The preferences of a bidder whose valuation is v are represented by the expected value of the Bernoulli payoff function that assigns 0 to the outcome in which she does not win the object and $v - p$ to the outcome in which she wins the object and pays the price p. (That is, each bidder is risk neutral.) I assume that the expected payoff of a bidder whose bid is tied for first place is $(v - p)/m$, where m is the number of tied winning bids. (The assumption about the outcome when bids are tied for first place has mainly "technical" significance; in Section 3.5, it was convenient to make an assumption different from the one here.)

Denote by $P(b)$ the price paid by the winner of the auction when the profile of bids is b. For a first-price auction $P(b)$ is the winning bid (the largest b_i), whereas for a second-price auction it is the highest bid made by a bidder different from the winner. Given the appropriate specification of P, the following Bayesian game models **first- and second-price auctions with independent private valuations** (and imperfect information about valuations).

Players The set of bidders, say $1, \ldots, n$.

States The set of all profiles $(v_1, \ldots, v_n)$ of valuations, where $\underline{v} \leq v_i \leq \overline{v}$ for all i.

Actions Each player's set of actions is the set of possible bids (nonnegative numbers).

Signals The set of signals that each player may observe is the set of possible valuations. The signal function τ_i of each player i is given by $\tau_i(v_1, \ldots, v_n) = v_i$ (each player knows her own valuation).

Beliefs Every type of player i assigns probability $F(v_1)F(v_2) \cdots F(v_{i-1}) \times F(v_{i+1}) \cdots F(v_n)$ to the event that the valuation of every other player j is at most v_j.

Payoff functions Player i's Bernoulli payoff in state $(v_1, \ldots, v_n)$ is 0 if her bid b_i is not the highest bid, and $(v_i - P(b))/m$ if no bid is higher than b_i and m bids (including b_i) are equal to b_i:

$$u_i(b, (v_1, \ldots, v_n)) = \begin{cases} (v_i - P(b))/m & \text{if } b_j \leq b_i \text{ for all } j \neq i \text{ and} \\ & \quad b_j = b_i \text{ for } m \text{ players} \\ 0 & \text{if } b_j > b_i \text{ for some } j \neq i. \end{cases}$$

Nash equilibrium in a second-price sealed-bid auction As in a second-price sealed-bid auction in which every bidder knows every other bidder's valuation,

> *in a second-price sealed-bid auction with imperfect information about valuations, a player's bid equal to her valuation weakly dominates all her other bids.*

Precisely, consider some type v_i of some player i, and let b_i be a bid not equal to v_i. Then for all bids by all types of all the other players, the expected payoff of type v_i of player i is at least as high when she bids v_i as it is when she bids b_i, and for some bids by the various types of the other players, her expected payoff is greater when she bids v_i than it is when she bids b_i.

The argument for this result is similar to the argument in Section 3.5.2 in the case in which the players know each other's valuations. The main difference between the arguments arises because in the case in which the players do not know each others' valuations, any given bids for every type of every player but i leave player i uncertain about the highest of the remaining bids, because she is uncertain of the other players' types. (The difference in the tie-breaking rules between the two cases also necessitates a small change in the argument.) In the next exercise you are asked to fill in the details.

? EXERCISE 294.1 (Weak domination in a second-price sealed-bid auction) Show that for each type v_i of each player i in a second-price sealed-bid auction with imperfect information about valuations the bid v_i weakly dominates all other bids.

We conclude, in particular, that a second-price sealed-bid auction with imperfect information about valuations has a Nash equilibrium in which every type of every player bids her valuation. The game has also other equilibria, some of which you are asked to find in the next exercise.

? EXERCISE 294.2 (Nash equilibria of a second-price sealed-bid auction) For every player i, find a Nash equilibrium of a second-price sealed-bid auction in which player i wins. (Think about the Nash equilibria when the players know each others' valuations, studied in Section 3.5.)

Nash equilibrium in a first-price sealed-bid auction As when the players are perfectly informed about each other's valuations, the bid of v_i by type v_i of player i weakly dominates any bid greater than v_i, does not weakly dominate bids less than v_i, and is itself weakly dominated by any such lower bid. (If type v_i of player i bids v_i, her payoff is certainly 0 (either she wins and pays her valuation, or she loses), whereas if she bids less than v_i, she may win and obtain a positive payoff.)

These facts suggest that the game may have a Nash equilibrium in which each player bids less than her valuation. An analysis of the game for an arbitrary distribution F of valuations requires calculus and is relegated to an appendix (Section 9.8). Here I consider the case in which there are two bidders and each player's valuation is distributed "uniformly" between 0 and 1. This assumption on the distribution of valuations means that the fraction of valuations less than v is exactly v, so that $F(v) = v$ for all v with $0 \leq v \leq 1$.

Denote by $\beta_i(v)$ the bid of type v of player i. I claim that if there are two bidders and the distribution of valuations is uniform between 0 and 1, the game has a (symmetric) Nash equilibrium in which the function β_i is the same for both players, with $\beta_i(v) = \frac{1}{2}v$ for all v. That is, each type of each player bids exactly half her valuation.

To verify this claim, suppose that each type of player 2 bids in this way. Then as far as player 1 is concerned, player 2's bids are distributed uniformly between 0 and $\frac{1}{2}$. Thus if player 1 bids more than $\frac{1}{2}$ she surely wins, whereas if she bids

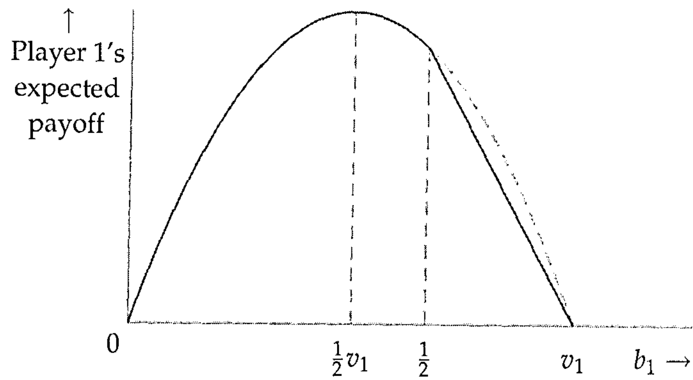

Figure 295.1 Player 1's expected payoff as a function of its bid in a first-price sealed-bid auction in which there are two bidders and the valuations are uniformly distributed from 0 to 1, given that player 2 bids $\frac{1}{2}v_2$.

$b_1 \leq \frac{1}{2}$ the probability that she wins is the probability that player 2's valuation is less than $2b_1$ (in which case player 2 bids less than b_1), which is $2b_1$. Consequently her payoff as a function of her bid b_1 is

$$\begin{cases} 2b_1(v_1 - b_1) & \text{if } 0 \leq b_1 \leq \frac{1}{2} \\ v_1 - b_1 & \text{if } b_1 > \frac{1}{2}. \end{cases}$$

This function is shown in Figure 295.1. Its maximizer is $\frac{1}{2}v_1$ (see Exercise 497.1), so that player 1's optimal bid is half her valuation. Both players are identical, so this argument shows also that given $\beta_1(v) = \frac{1}{2}v$, player 2's optimal bid is half her valuation. Thus, as claimed, the game has a Nash equilibrium in which each type of each player bids half her valuation.

When the number n of bidders exceeds 2, a similar analysis shows that the game has a (symmetric) Nash equilibrium in which every player bids the fraction $1 - 1/n$ of her valuation: $\beta_i(v) = (1 - 1/n)v$ for every player i and every valuation v. (You are asked to verify a claim more general than this one in Exercise 296.1.)

In this example—and, it turns out, for any distribution F satisfying the conditions given at the start of this section—the players' common bidding function in a symmetric Nash equilibrium may be given an illuminating interpretation. Choose $n - 1$ valuations randomly and independently, each according to the cumulative distribution function F. The highest of these $n - 1$ valuations is a "random variable": its value depends on the $n - 1$ valuations that were chosen. Denote it by X. Fix a valuation v. Some values of X are less than v; others are greater than v. Consider the distribution of X in those cases in which it is less than v. The expected value of this distribution is denoted $\mathrm{E}[\mathrm{X} \mid \mathrm{X} < v]$: the expected value of X conditional on X being less than v. We may prove the following result. (See the appendix, Section 9.8, for suggestive arguments.)

> *If each bidder's valuation is drawn independently from the same continuous and increasing cumulative distribution, a first-price sealed-bid auction (with imperfect information about valuations) has a (symmetric) Nash equilibrium in which each type v of each player bids* $\mathrm{E}[\mathbf{X} \mid \mathbf{X} < v]$, *the expected value of the highest of the other players' bids conditional on v being higher than all the other valuations.*

Put differently, each bidder asks the following question: Over all the cases in which my valuation is the highest, what is the expectation of the highest of the other players' valuations? This expectation is the amount she bids.

If F is uniform from 0 to 1 and $n = 2$, we may verify that indeed the equilibrium we found may be expressed in this way. For any valuation v of player 1, the cases in which player 2's valuation is less than v are distributed uniformly from 0 to v, so that the expected value of player 2's valuation conditional on its being less than v is $\frac{1}{2}v$, which is equal to the equilibrium bidding function that we found.

Comparing equilibria of first- and second-price auctions Near the end of Section 3.5.3 we saw that first- and second-price auctions are "revenue equivalent" when the players know each others' valuations: their distinguished equilibria yield the same outcome. I now argue that the same is true, under the assumptions of this section, when the players are uncertain of each others' valuations.

Consider the equilibrium of a second-price auction in which every player bids her valuation. In this equilibrium, the expected price paid by a bidder with valuation v who wins is the expectation of the highest of the other $n - 1$ valuations, conditional on this maximum being less than v, or, in the notation above, $\mathrm{E}[\mathbf{X} \mid \mathbf{X} < v]$. We have just seen that a first-price auction has a symmetric Nash equilibrium in which this amount is precisely the bid of a player with valuation v, and hence the amount paid by such a player. Thus a winning bidder pays the same expected price in the equilibrium of each auction. In both cases, the player with the highest valuation submits the winning bid, so both auctions yield the auctioneer the same revenue:

> *If each bidder is risk-neutral and each bidder's valuation is drawn independently from the same continuous and increasing cumulative distribution, then the Nash equilibrium of a second-price sealed-bid auction (with imperfect information about valuations) in which each player bids her valuation yields the same revenue as the symmetric Nash equilibrium of the corresponding first-price sealed-bid auction.*

This result depends on the assumption that each player's preferences are represented by the expected value of a risk-neutral Bernoulli payoff function. The next exercise asks you to study an example in which each player is risk averse. (See page 103 for a discussion of risk neutrality and risk aversion.)

?. EXERCISE 296.1 (Auctions with risk-averse bidders) Consider a variant of the Bayesian game defined earlier in this section in which the players are risk averse.

Specifically, suppose each of the n players' preferences are represented by the expected value of the Bernoulli payoff function $x^{1/m}$, where x is the player's monetary payoff and $m > 1$. Suppose also that each player's valuation is distributed uniformly between 0 and 1, as in the example on page 294. Show that the Bayesian game that models a first-price sealed-bid auction under these assumptions has a (symmetric) Nash equilibrium in which each type v_i of each player i bids $(1 - 1/[m(n-1)+1])v_i$. (You need to use the mathematical fact that the solution of the problem $\max_b[b^k(v-b)^\ell]$ is $kv/(k+\ell)$.) Compare the auctioneer's revenue in this equilibrium with her revenue in the symmetric Nash equilibrium of a second-price sealed-bid auction in which each player bids her valuation. (Note that the equilibrium of the second-price auction does not depend on the players' payoff functions.)

9.6.3 Common valuations

In an auction with common valuations, each player's valuation depends on the other players' signals as well as her own. (As before, I assume that the players' signals are independent.) I denote the function that gives player i's valuation by g_i, and assume that it is increasing in all the signals. Given the appropriate specification of the function P that determines the price $P(b)$ paid by the winner as a function of the profile b of bids, the following Bayesian game models **first- and second-price auctions with common valuations** (and imperfect information about valuations).

Players The set of bidders, say $\{1, \ldots, n\}$.

States The set of all profiles $(t_1, \ldots, t_n)$ of signals that the players may receive.

Actions Each player's set of actions is the set of possible bids (nonnegative numbers).

Signals The signal function τ_i of each player i is given by $\tau_i(t_1, \ldots, t_n) = t_i$ (each player observes her own signal).

Beliefs Each type of each player believes that the signal of every type of every other player is independent of all the other players' signals.

Payoff functions Player i's Bernoulli payoff in state $(t_1, \ldots, t_n)$ is 0 if her bid b_i is not the highest bid, and $(g_i(t_1, \ldots, t_n) - P(b))/m$ if no bid is higher than b_i and m bids (including b_i) are equal to b_i:

$$u_i(b, (t_1, \ldots, t_n)) = \begin{cases} (g_i(t_1, \ldots, t_n) - P(b))/m & \text{if } b_j \le b_i \text{ for all } j \ne i \text{ and} \\ & \quad b_j = b_i \text{ for } m \text{ players} \\ 0 & \text{if } b_j > b_i \text{ for some } j \ne i. \end{cases}$$

Nash equilibrium in a second-price sealed-bid auction The main ideas in the analysis of sealed-bid common value auctions are illustrated by an example in which there are two bidders, each bidder's signal is uniformly distributed from 0 to 1, and the valuation of each bidder i is given by $v_i = \alpha t_i + \gamma t_j$, where j is the other player and $\alpha \geq \gamma \geq 0$. The case in which $\alpha = 1$ and $\gamma = 0$ is exactly the one studied in Section 9.6.2: in this case, the bidders' valuations are private. If $\alpha = \gamma$, then for any given signals, each bidder's valuation is the same—a case of "pure common valuations". If, for example, the signal t_i is the number of barrels of oil in a tract, then the expected valuation of a bidder i who knows the signals t_i and t_j is $p \cdot \frac{1}{2}(t_i + t_j)$, where p is the monetary worth of a barrel of oil. Our assumption, of course, is that a bidder does *not* know any other player's signal. However, a key point in the analysis of common value auctions is that the other players' bids contain *some* information about the other players' signals—information that may profitably be used.

I claim that under these assumptions a second-price sealed-bid auction has a Nash equilibrium in which each type t_i of each player i bids $(\alpha + \gamma)t_i$.

To verify this claim, suppose that each type of player 2 bids in this way and type t_1 of player 1 bids b_1. To determine the expected payoff of type t_1 of player 1, we need to find the probability with which she wins, and both the expected price she pays and the expected value of player 2's signal if she wins.

Probability that player 1 wins: Given that player 2's bidding function is $(\alpha + \gamma)t_2$, player 1's bid of b_1 wins only if $b_1 \geq (\alpha + \gamma)t_2$, or if $t_2 \leq b_1/(\alpha + \gamma)$. Now, t_2 is distributed uniformly from 0 to 1, so the probability that it is at most $b_1/(\alpha + \gamma)$ is $b_1/(\alpha + \gamma)$. Thus a bid of b_1 by player 1 wins with probability $b_1/(\alpha + \gamma)$.

Expected price player 1 pays if she wins: The price she pays is equal to player 2's bid, which, *conditional on its being less than* b_1, is distributed uniformly from 0 to b_1. Thus the expected value of player 2's bid, *given that it is less than* b_1, is $\frac{1}{2}b_1$.

Expected value of player 2's signal if player 1 wins: Player 2's bid, given her signal t_2, is $(\alpha + \gamma)t_2$, so that the expected value of signals that yield a bid of less than b_1 is $\frac{1}{2}b_1/(\alpha + \gamma)$ (because of the uniformity of the distribution of t_2).

Now, player 1's expected payoff if she bids b_1 is the difference between her expected valuation, given her signal t_1 and the fact that she wins, and the expected price she pays, multiplied by her probability of winning. Combining the previous calculations, player 1's expected payoff if she bids b_1 is thus $(\alpha t_1 + \frac{1}{2}\gamma b_1/(\alpha + \gamma) - \frac{1}{2}b_1)b_1/(\alpha + \gamma)$, or

$$\frac{\alpha}{2(\alpha + \gamma)^2} \cdot (2(\alpha + \gamma)t_1 - b_1)b_1.$$

This function is maximized at $b_1 = (\alpha + \gamma)t_1$. That is, if each type t_2 of player 2 bids $(\alpha + \gamma)t_2$, any type t_1 of player 1 optimally bids $(\alpha + \gamma)t_1$. Symmetrically, if each type t_1 of player 1 bids $(\alpha+\gamma)t_1$, any type t_2 of player 2 optimally bids $(\alpha+\gamma)t_2$. Hence, as claimed, the game has a Nash equilibrium in which each type t_i of each player i bids $(\alpha+\gamma)t_i$.

? EXERCISE 299.1 (Asymmetric Nash equilibria of second-price sealed-bid common value auctions) Show that when $\alpha=\gamma=1$, for *any* value of $\lambda>0$ the game has an (asymmetric) Nash equilibrium in which each type t_1 of player 1 bids $(1+\lambda)t_1$ and each type t_2 of player 2 bids $(1+1/\lambda)t_2$.

Note that when player 1 calculates her expected value of the object, she finds the expected value of player 2's signal *given that her bid wins.* She should not base her bid simply on an estimate of the valuation derived from her own signal and the (unconditional) expectation of the other player's signal. She wins precisely when her bid exceeds those of the other players, so if she bids in this way, then over all the cases in which she wins, she more likely than not overvalues the object. A bidder who incorrectly behaves in this way is said to suffer from the *winner's curse.* (Bidders in real auctions know this problem: when a contractor gives you a quotation to renovate your house, she does not base her price simply on an unbiased estimate out how much it will cost her to do the job; rather, she takes into account that you will select her only if her competitors' estimates are all higher than hers, in which case her estimate may be suspiciously low.)

Nash equilibrium in a first-price sealed-bid auction I claim that under the assumptions on the players' signals and valuations in Section 9.6.2, a first-price sealed-bid auction has a Nash equilibrium in which each type t_i of each player i bids $\frac{1}{2}(\alpha+\gamma)t_i$. This claim may be verified by arguments like those in that section. In the next exercise, you are asked to supply the details.

? EXERCISE 299.2 (First-price sealed-bid auction with common valuations) Verify that under the assumptions on signals and valuations in Section 9.6.2, a first-price sealed-bid auction has a Nash equilibrium in which the bid of each type t_i of each player i is $\frac{1}{2}(\alpha+\gamma)t_i$.

Comparing equilibria of first- and second-price auctions We see that the revenue equivalence of first- and second-price auctions that holds when valuations are private holds also for the symmetric equilibria of the examples in which the valuations are common. That is, the expected price paid by a player of any given type is the same in the symmetric equilibrium of the first-price auction as it is in the symmetric equilibrium of the second-price auction: in each case type t_i of player i pays $\frac{1}{2}(\alpha+\gamma)t_i$ if she wins, and wins with the same probability.

In fact, the revenue equivalence principle holds much more generally. Suppose that each bidder is risk neutral and independently receives a signal from the same distribution, and that this distribution satisfies the conditions in Section 9.6.2. Then the expected payment of a bidder of any given type is the same in the symmetric Nash equilibrium of a second-price sealed-bid auction as it is in the symmetric Nash equilibrium of a first-price sealed-bid auction. Further, this revenue equivalence is not restricted to first- and second-price auctions; a general result, encompassing a wider range of auction forms, is stated at the end of the appendix (Section 9.8).

AUCTIONS OF THE RADIO SPECTRUM

In the 1990s several countries started auctioning the right to use parts of the radio spectrum used for wireless communication (by mobile telephones, for example). Spectrum licenses in the United States were originally allocated on the basis of hearings by the Federal Communications Commission (FCC). This procedure was time-consuming, and a large backlog developed, prompting a switch to lotteries. Licenses awarded by the lotteries could be resold at high prices, attracting many participants. In one case that drew attention, the winner of a license to run cellular telephones in Cape Cod sold it to Southwestern Bell for $41.5 million (*New York Times*, May 30, 1991, page A1). In the early 1990s, the U.S. government was persuaded that auctioning licenses would allocate them more efficiently and might raise nontrivial revenue.

For each interval of the spectrum, many licenses were available, each covering a geographic area. A buyer's valuation of a license could be expected to depend on the other licenses it owned, so many interdependent goods were for sale. In designing an auction mechanism, the FCC had many choices: for example, the bidding could be open, or it could be sealed, with the price equal to either the highest bid or the second-highest bid; the licenses could be sold sequentially, or simultaneously, in which case participants could submit bids for individual licenses or for combinations of licenses. Experts in auction theory were consulted on the design of the mechanism. John McMillan (who advised the FCC), writes that "When theorists met the policy-makers, concepts like Bayes-Nash equilibrium, incentive-compatibility constraints, and order-statistic theorems came to be discussed in the corridors of power" (1994, 146). No theoretical analysis fitted the environment of the auction well, but the experts appealed to some principles from the existing theory, the results of laboratory experiments, and experience in auctions held in New Zealand and Australia in the early 1990s in making their recommendations. The mechanism adopted in 1994 was an open ascending auction for which bids were accepted simultaneously for all licenses in each round. Experts argued that the open (as opposed to sealed-bid) format and the simultaneity of the auctions promoted an efficient outcome because at each stage the bidders could see their rivals' previous bids for all licenses.

The FCC has conducted several auctions, starting with "narrowband" licenses (each covering a sliver of the spectrum, used by paging services) and continuing with "broadband" licenses (used for voice and data communications). These auctions have provided more employment for game theorists, many of whom have advised the companies bidding for licenses. In response to growing congestion of the airwaves and the expectation that a significant part of the rapidly growing Internet traffic will move to wireless devices, in 2000 President Bill Clinton ordered further auctions of large parts of the spectrum (*New York Times*, October 14, 2000). Whether the auctions that have been held have allocated licenses efficiently is hard to tell, though it appears that the winners were able to obtain the sets of licenses they wanted. Certainly the auctions have been successful in generating revenue: the first four generated over $18 billion.

9.7 Illustration: juries

9.7.1 Model

In a trial, jurors are presented with evidence concerning the guilt or innocence of a defendant. Their interpretations of the evidence may differ. Each juror assesses the evidence, and on the basis of her interpretation votes either to convict or acquit the defendant. Assume that a unanimous verdict is required for conviction: the defendant is convicted if and only if every juror votes to convict her. (This rule is used in the United States and Canada, for example.) What can we say about the chance of an innocent defendant's being convicted and a guilty defendant's being acquitted?

In deciding how to vote, each juror must consider the costs of convicting an innocent person and of acquitting a guilty person. She must consider also the likely effect of her vote on the outcome, which depends on the other jurors' votes. For example, a juror who thinks that at least one of her colleagues is likely to vote for acquittal may act differently from one who is sure that all her colleagues will vote for conviction. Thus an answer to the question requires us to consider the strategic interaction between the jurors, which we may model as a Bayesian game.

Assume that each juror comes to the trial with the belief that the defendant is guilty with probability π (the same for every juror), a belief modified by the evidence presented. We model the possibility that jurors interpret the evidence differently by assuming that for each of the defendant's true statuses (guilty and innocent), each juror interprets the evidence to point to guilt with positive probability, and to innocence with positive probability, and that the jurors' interpretations are independent (no juror's interpretation depends on any other juror's interpretation). I assume that the probabilities are the same for all jurors and denote the probability of any given juror's interpreting the evidence to point to guilt when the defendant is guilty by p, and the probability of her interpreting the evidence to point to innocence when the defendant is innocent by q. I assume also that a juror is more likely than not to interpret the evidence correctly, so that $p > \frac{1}{2}$ and $q > \frac{1}{2}$, and hence in particular $p > 1 - q$.

Each juror wishes to convict a guilty defendant and acquit an innocent one. She is indifferent between these two outcomes and prefers each of them to one in which an innocent defendant is convicted or a guilty defendant is acquitted.

Assume specifically that each juror's Bernoulli payoffs are:

$$\begin{cases} 0 & \text{if guilty defendant convicted or} \\ & \quad \text{innocent defendant acquitted} \\ -z & \text{if innocent defendant convicted} \\ -(1-z) & \text{if guilty defendant acquitted.} \end{cases} \tag{302.1}$$

The parameter z may be given an appealing interpretation. Denote by r the probability a juror assigns to the defendant's being guilty, given all her information. Then her expected payoff if the defendant is acquitted is $-r(1-z) + (1-r) \cdot 0 = -r(1-z)$ and her expected payoff if the defendant is convicted is $r \cdot 0 - (1-r)z = -(1-r)z$. Thus she prefers the defendant to be acquitted if $-r(1-z) > -(1-r)z$, or $r < z$, and convicted if $r > z$. That is, z is equal to the probability of guilt required for the juror to want the defendant to be convicted. Put differently, for any juror

$$\begin{array}{l} \text{acquittal is at least as good as conviction if and only if} \\ \Pr(\text{defendant is guilty, given juror's information}) \leq z. \end{array} \tag{302.2}$$

We may now formulate a Bayesian game that models the situation. The players are the jurors, and each player's action is a vote to convict (C) or to acquit (Q). We need one state for each configuration of the players' preferences and information. Each player's preferences depend on whether the defendant is guilty or innocent, and each player's information consists of her interpretation of the evidence. Thus we define a state to be a list $(X, s_1, \ldots, s_n)$, where X denotes the defendant's true status, guilty (G) or innocent (I), and s_i represents player i's interpretation of the evidence, which may point to guilt (g) or innocence (b). (I do not use i for "innocence" because I use it to index the players; b stands for "blameless".) The signal that each player i receives is her interpretation of the evidence, s_i. In any state in which $X = G$ (i.e. the defendant is guilty), each player assigns the probability p to any other player's receiving the signal g, and the probability $1-p$ to her receiving the signal b, independently of all other players' signals. Similarly, in any state in which $X = I$ (i.e. the defendant is innocent), each player assigns the probability q to any other player's receiving the signal b, and the probability $1-q$ to her receiving the signal g, independently of all other players' signals.

Each player cares about the verdict, which depends on the players' actions and on the defendant's true status. Given the assumption that unanimity is required to convict the defendant, only the action profile $(C, \ldots, C)$ leads to conviction. Thus (302.1) implies that player i's payoff function in the Bayesian game is defined as follows:

$$u_i(a, \omega) = \begin{cases} 0 & \text{if } a \neq (C, \ldots, C) \text{ and } \omega_1 = I \text{ or} \\ & \quad \text{if } a = (C, \ldots, C) \text{ and } \omega_1 = G \\ -z & \text{if } a = (C, \ldots, C) \text{ and } \omega_1 = I \\ -(1-z) & \text{if } a \neq (C, \ldots, C) \text{ and } \omega_1 = G, \end{cases} \tag{302.3}$$

where ω_1 is the first component of the state, giving the defendant's true status.

In summary, the Bayesian game that models the situation has the following components.

Players A set of n jurors.

States The set of states is the set of all lists $(X, s_1, \ldots, s_n)$ where $X \in \{G, I\}$ and $s_j \in \{g, b\}$ for every juror j, where $X = G$ if the defendant is guilty, $X = I$ if she is innocent, $s_j = g$ if player j receives the signal that she is guilty, and $s_j = b$ if player j receives the signal that she is innocent.

Actions The set of actions of each player is $\{C, Q\}$, where C means vote to convict, and Q means vote to acquit.

Signals The set of signals that each player may receive is $\{g, b\}$ and player j's signal function is defined by $\tau_j(X, s_1, \ldots, s_n) = s_j$ (each juror is informed only of her own signal).

Beliefs Type g of any player i believes that the state is $(G, s_1, \ldots, s_n)$ with probability $\Pr(G \mid g)p^{k-1}(1-p)^{n-k}$ and $(I, s_1, \ldots, s_n)$ with probability $\Pr(I \mid g) \times (1-q)^{k-1}q^{n-k}$, where k is the number of players j (including i) for whom $s_j = g$ in each case. Type b of any player i believes that the state is $(G, s_1, \ldots, s_n)$ with probability $\Pr(G \mid b)p^k(1-p)^{n-k-1}$ and $(I, s_1, \ldots, s_n)$ with probability $\Pr(I \mid b)(1-q)^k q^{n-k-1}$, where k is the number of players j for whom $s_j = g$ in each case.

Payoff functions The Bernoulli payoff function of each player i is given in (302.3).

9.7.2 Nash equilibrium

One juror Start by considering the very simplest case, in which there is a single juror. Suppose that her signal is b. To determine whether she prefers conviction or acquittal, we need to find the probability $\Pr(G \mid b)$ she assigns to the defendant's being guilty, given her signal. We can find this probability by using Bayes' rule (see Section 17.6.5, in particular (504.3)), as follows:

$$\begin{aligned}\Pr(G \mid b) &= \frac{\Pr(b \mid G)\Pr(G)}{\Pr(b \mid G)\Pr(G) + \Pr(b \mid I)\Pr(I)} \\ &= \frac{(1-p)\pi}{(1-p)\pi + q(1-\pi)}.\end{aligned}$$

Thus by (302.2), acquittal yields an expected payoff at least as high as does conviction if and only if

$$z \geq \frac{(1-p)\pi}{(1-p)\pi + q(1-\pi)}.$$

That is, after getting the signal that the defendant is innocent, the juror chooses acquittal as long as z is not too small—as long as she is too concerned about acquitting a guilty defendant. If her signal is g, then a similar calculation leads to the conclusion that conviction yields an expected payoff at least as high as does acquittal if

$$z \leq \frac{p\pi}{p\pi + (1-q)(1-\pi)}.$$

Thus if

$$\frac{(1-p)\pi}{(1-p)\pi + q(1-\pi)} \leq z \leq \frac{p\pi}{p\pi + (1-q)(1-\pi)}, \tag{304.1}$$

then the juror optimally acts according to her signal, acquitting the defendant when her signal is b and convicting her when it is g. (A bit of algebra shows that the term on the left of (304.1) is less than the term on the right, given $p > 1 - q$.)

Two jurors Now suppose there are two jurors. Are there values for z such that the game has a Nash equilibrium in which each juror votes according to her signal? Suppose that juror 2 acts in this way: type b votes to acquit, and type g votes to convict. Consider type b of juror 1. If juror 2's signal is b, juror 1's vote has no effect on the outcome, because juror 2 votes to acquit and unanimity is required for conviction. Thus when deciding how to vote, juror 1 should ignore the possibility that juror 2's signal is b and assume it is g. That is, juror 1 should take as evidence her signal and the fact that juror 2's signal is g. Hence, given (302.2), for type b of juror 1 acquittal is at least as good as conviction if the probability that the defendant is guilty, given that juror 1's signal is b and juror 2's signal is g, is at most z. This probability is

$$\begin{aligned}\Pr(G \mid b, g) &= \frac{\Pr(b, g \mid G)\Pr(G)}{\Pr(b, g \mid G)\Pr(G) + \Pr(b, g \mid I)\Pr(I)} \\ &= \frac{(1-p)p\pi}{(1-p)p\pi + q(1-q)(1-\pi)}.\end{aligned}$$

Thus type b of juror 1 optimally votes for acquittal if

$$z \geq \frac{(1-p)p\pi}{(1-p)p\pi + q(1-q)(1-\pi)}.$$

By a similar argument, for type g of juror 1 conviction is at least as good as acquittal if

$$z \leq \frac{p^2\pi}{p^2\pi + (1-q)^2(1-\pi)}.$$

Thus when there are two jurors, the game has a Nash equilibrium in which each juror acts according to her signal, voting to acquit the defendant when her signal is b and to convict her when it is g, if

$$\frac{(1-p)p\pi}{(1-p)p\pi + q(1-q)(1-\pi)} \leq z \leq \frac{p^2\pi}{p^2\pi + (1-q)^2(1-\pi)}. \tag{304.2}$$

Consider the expressions on the left of (304.1) and (304.2). Divide the numerator and denominator of the expression on the left of (304.1) by $1-p$ and the numerator and denominator of the expression on the left of (304.2) by $(1-p)p$. Then, given $p > 1 - q$, we see that the expression on the left of (304.2) is greater than the expression on the left of (304.1). That is, the lowest value of z for which an equilibrium exists in which each juror votes according to her signal is higher when there are two jurors than when there is only one juror. Why? Because a juror who receives the signal b, knowing that her vote makes a difference only if the other juror votes to convict, makes her decision on the assumption that the other juror's signal is g, and so is less worried about convicting an innocent defendant than is a single juror in isolation.

Many jurors Now suppose the number of jurors is arbitrary, equal to n. Suppose that every juror other than juror 1 votes to acquit when her signal is b and to convict when her signal is g. Consider type b of juror 1. As in the case of two jurors, juror 1's vote has no effect on the outcome unless every other juror's signal is g. Thus when deciding how to vote, juror 1 should assume that all the other signals are g. Hence, given (302.2), for type b of juror 1 acquittal is at least as good as conviction if the probability that the defendant is guilty, given that juror 1's signal is b and every other juror's signal is g, is at most z. This probability is

$$\begin{aligned}\Pr(G \mid b, g, \ldots, g) &= \frac{\Pr(b, g, \ldots, g \mid G)\Pr(G)}{\Pr(b, g, \ldots, g \mid G)\Pr(G) + \Pr(b, g, \ldots, g \mid I)\Pr(I)} \\ &= \frac{(1-p)p^{n-1}\pi}{(1-p)p^{n-1}\pi + q(1-q)^{n-1}(1-\pi)}.\end{aligned}$$

Thus type b of juror 1 optimally votes for acquittal if

$$\begin{aligned}z &\geq \frac{(1-p)p^{n-1}\pi}{(1-p)p^{n-1}\pi + q(1-q)^{n-1}(1-\pi)} \\ &= \frac{1}{1 + \dfrac{q}{1-p}\left(\dfrac{1-q}{p}\right)^{n-1}\dfrac{1-\pi}{\pi}}.\end{aligned}$$

Now, given that $p > 1 - q$, the denominator decreases to 1 as n increases. Thus the lower bound on z for which type b of juror 1 votes for acquittal approaches 1 as n increases. (You may check that if $p = q = 0.8$, $\pi = 0.5$, and $n = 12$, the lower bound on z exceeds 0.999999.) In particular, in a large jury, if jurors care even slightly about acquitting a guilty defendant, then a juror who interprets the evidence to point to innocence will nevertheless vote for conviction. The reason is that the vote of a juror who interprets the evidence to point to innocence makes a difference to the outcome only if every other juror interprets the evidence to point to guilt, in which case the probability that the defendant is in fact guilty is very high.

We conclude that the model of a large jury in which the jurors are concerned about acquitting a guilty defendant has no Nash equilibrium in which every juror votes according to her signal. What *are* its equilibria? You are asked to find the conditions for two equilibria in the next exercise.

? EXERCISE 306.1 (Signal-independent equilibria in a model of a jury) Find conditions under which the game, for an arbitrary number of jurors, has a Nash equilibrium in which every juror votes for acquittal regardless of her signal, and conditions under which every juror votes for conviction regardless of her signal.

Under some conditions on z the game has in addition a symmetric mixed strategy Nash equilibrium in which each type g juror votes for conviction, and each type b juror votes for acquittal and conviction each with positive probability. Denote by β the mixed strategy of each juror of type b. As before, a juror's vote affects the outcome only if all other jurors vote for conviction, so when choosing an action, a juror should assume that all other jurors vote for conviction.

Each type b juror must be indifferent between voting for conviction and voting for acquittal, because she takes each action with positive probability. By (302.2) we thus need the mixed strategy β to be such that the probability that the defendant is guilty, given that all other jurors vote for conviction, is equal to z. Now, the probability of any given juror's voting for conviction is $p + (1-p)\beta(C)$ if the defendant is guilty and $1 - q + q\beta(C)$ if she is innocent. Thus

$$\begin{aligned}
&\Pr(G \mid \text{signal } b \text{ and } n-1 \text{ votes for } C)\\
&\quad = \frac{\Pr(b \mid G)(\Pr(\text{vote } C \mid G))^{n-1}\Pr(G)}{\Pr(b \mid G)(\Pr(\text{vote } C \mid G))^{n-1}\Pr(G) + \Pr(b \mid I)(\Pr(\text{vote } C \mid I))^{n-1}\Pr(I)}\\
&\quad = \frac{(1-p)(p+(1-p)\beta(C))^{n-1}\pi}{(1-p)(p+(1-p)\beta(C))^{n-1}\pi + q(1-q+q\beta(C))^{n-1}(1-\pi)}.
\end{aligned}$$

The condition that this probability equals z implies

$$(1-p)(p+(1-p)\beta(C))^{n-1}\pi(1-z) = q(1-q+q\beta(C))^{n-1}(1-\pi)z, \tag{306.2}$$

and hence

$$\beta(C) = \frac{pX - (1-q)}{q - (1-p)X},$$

where $X = [\pi(1-p)(1-z)/((1-\pi)qz)]^{1/(n-1)}$. For a range of parameter values, $0 \le \beta(C) \le 1$, so that $\beta(C)$ is indeed a probability. Notice that when n is large, X is close to 1, and hence $\beta(C)$ is close to 1: a juror who interprets the evidence as pointing to innocence very likely nonetheless votes for conviction.

Each type g juror votes for conviction, and so must get an expected payoff at least as high from conviction as from acquittal. From an analysis like that for each type b juror, this condition is

$$p(p+(1-p)\beta(C))^{n-1}\pi(1-z) \ge (1-q)(1-q+q\beta(C))^{n-1}(1-\pi)z.$$

Given $p > \frac{1}{2}$ and $q > \frac{1}{2}$, this condition follows from (306.2).

An interesting property of this equilibrium is that the probability that an innocent defendant is convicted *increases* as n increases: the larger the jury, the *more* likely an innocent defendant is to be convicted. (The proof of this result is not simple.)

Variants The key point behind the results is that under unanimity rule a juror's vote makes a difference to the outcome only if every other juror votes for conviction. Consequently, a juror, when deciding how to vote, rationally assesses the defendant's probability of guilt under the assumption that every other juror votes for conviction. The fact that this implication of unanimity rule drives the results suggests that the Nash equilibria might be quite different if less than unanimity were required for conviction. The analysis of such rules is difficult, but indeed the Nash equilibria they generate differ significantly from the Nash equilibria under unanimity rule. In particular, the analogue of the mixed strategy Nash equilibria considered earlier generates a probability that an innocent defendant is convicted that approaches zero as the jury size increases, as Feddersen and Pesendorfer (1998) show.

The idea behind the equilibria of the model in the next exercise is related to the ideas in this section, though the model is different.

? EXERCISE 307.1 (Swing voter's curse) Whether candidate 1 or candidate 2 is elected depends on the votes of two citizens. The economy may be in one of two states, A and B. The citizens agree that candidate 1 is best if the state is A and candidate 2 is best if the state is B. Each citizen's preferences are represented by the expected value of a Bernoulli payoff function that assigns a payoff of 1 if the best candidate for the state wins (obtains more votes than the other candidate), a payoff of 0 if the other candidate wins, and payoff of $\frac{1}{2}$ if the candidates tie. Citizen 1 is informed of the state, whereas citizen 2 believes it is A with probability 0.9 and B with probability 0.1. Each citizen may either vote for candidate 1, vote for candidate 2, or not vote.

a. Formulate this situation as a Bayesian game. (Construct the table of payoffs for each state.)

b. Show that the game has exactly two pure Nash equilibria, in one of which citizen 2 does not vote and in the other of which she votes for 1.

c. Show that an action of one of the players in the second equilibrium is weakly dominated.

d. Why is "swing voter's curse" an appropriate name for the determinant of citizen 2's decision in the first equilibrium?

9.8 Appendix: auctions with an arbitrary distribution of valuations

9.8.1 *First-price sealed-bid auctions*

In this section I construct a symmetric equilibrium of a first-price sealed-bid auction for a distribution F of valuations that satisfies the assumptions in Section 9.6.2 and is differentiable on $(\underline{v}, \overline{v})$. (Unlike the remainder of the book, the section uses calculus.)

As before, denote the bid of type v of player i (i.e. player i when her valuation is v) by $\beta_i(v)$. In a symmetric equilibrium, every player uses the same bidding function: for some function β we have $\beta_i = \beta$ for every player i. A reasonable guess is that if such an equilibrium exists, β is increasing: the higher a player's valuation, the more she bids. Under the additional assumption that β is differentiable, I derive a condition that it must satisfy in any symmetric equilibrium. Exactly one function satisfies this condition, and this function is in fact increasing (as you are asked to show in an exercise).

Suppose that all $n-1$ players other than i bid according to the increasing differentiable function β. Then, given the assumptions on F, the probability of a tie is zero, and hence for any bid b, the expected payoff of player i when her valuation is v and she bids b is

$$(v-b)\,\Pr(\text{Highest bid is } b) = (v-b)\,\Pr(\text{All } n-1 \text{ other bids} \leq b). \qquad (308.1)$$

Now, a player bidding according to the function β bids at most b, for $\beta(\underline{v}) \leq b \leq \beta(\overline{v})$, if her valuation is at most $\beta^{-1}(b)$ (the inverse of β evaluated at b). Thus the probability that the bids of the $n-1$ other players are all at most b is the probability that the highest of $n-1$ randomly selected valuations—a random variable denoted $\mathbf{X}$ in Section 9.6.2—is at most $\beta^{-1}(b)$. Denoting the cumulative distribution function of $\mathbf{X}$ by H, the expected payoff in (308.1) is thus

$$(v-b)H(\beta^{-1}(b)) \text{ if } \beta(\underline{v}) \leq b \leq \beta(\overline{v})$$

(and 0 if $b < \beta(\underline{v})$, $v - b$ if $b > \beta(\overline{v})$).

I now claim that in a symmetric equilibrium in which every player bids according to β, we have $\beta(v) \leq v$ if $v > \underline{v}$ and $\beta(\underline{v}) = \underline{v}$. If $v > \underline{v}$ and $\beta(v) > v$, then a player with valuation v wins with positive probability (players with valuations less than v bid less than $\beta(v)$, because β is increasing) and obtains a negative payoff if she does so. She obtains a payoff of zero by bidding v, so for equilibrium we need $\beta(v) \leq v$ whenever $v > \underline{v}$. Given that β satisfies this condition, if $\beta(\underline{v}) > \underline{v}$ then a player with valuation $\underline{v}$ wins with positive probability, and obtains a negative payoff if she does so. Thus $\beta(\underline{v}) \leq \underline{v}$. If $\beta(\underline{v}) < \underline{v}$, then players with valuations slightly greater than $\underline{v}$ also bid less than $\underline{v}$ (because β is continuous), so that a player with valuation $\underline{v}$ who increases her bid slightly wins with positive probability and obtains a positive payoff if she does so, rather than obtaining the payoff of zero. We conclude that $\beta(\underline{v}) = \underline{v}$.

Now, the expected payoff of a player of type v when every other player uses the bidding function β is differentiable on $(\underline{v}, \beta(\overline{v}))$ (given that β is increasing and differentiable, and $\beta(\underline{v}) = \underline{v}$) and, if $v > \underline{v}$, is increasing at $\underline{v}$. Thus the derivative of this expected payoff with respect to b is zero at any best response less than $\beta(\overline{v})$:

$$-H(\beta^{-1}(b)) + \frac{(v-b)H'(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))} = 0. \qquad (308.2)$$

(The derivative of β^{-1} at the point b is $1/\beta'(\beta^{-1}(b))$.)

In a symmetric equilibrium in which every player bids according to β, the best response of type v of any given player to the other players' strategies is $\beta(v)$. Because β is increasing, we have $\beta(v) < \beta(\overline{v})$ for $v < \overline{v}$, so $\beta(v)$ must satisfy (308.2) whenever $\underline{v} < v < \overline{v}$. If $b = \beta(v)$, then $\beta^{-1}(b) = v$, so that substituting $b = \beta(v)$ into (308.2) and multiplying by $\beta'(v)$ yields

$$\beta'(v)H(v) + \beta(v)H'(v) = vH'(v) \text{ for } \underline{v} < v < \overline{v}.$$

The left-hand side of this differential equation is the derivative with respect to v of $\beta(v)H(v)$. Thus for some constant C we have

$$\beta(v)H(v) = \int_{\underline{v}}^{v} xH'(x)\,dx + C \text{ for } \underline{v} < v < \overline{v}.$$

The function β is bounded, so considering the limit as v approaches $\underline{v}$, we deduce that $C = 0$.

We conclude that if the game has a symmetric Nash equilibrium in which each player's bidding function is increasing and differentiable on $(\underline{v}, \overline{v})$, then this function is defined by

$$\beta^*(v) = \frac{\int_{\underline{v}}^{v} xH'(x)\,dx}{H(v)} \text{ for } \underline{v} < v \leq \overline{v}$$

and $\beta^*(\underline{v}) = \underline{v}$. Now, the function H is the cumulative distribution function of $\mathbf{X}$, the highest of $n-1$ independently drawn valuations. Thus $\beta^*(v)$ is the expected value of $\mathbf{X}$ conditional on its being less than v: $\beta^*(v) = \mathrm{E}[\mathbf{X} \mid \mathbf{X} < v]$, as claimed in Section 9.6.2.

We may alternatively express the numerator in the expression for $\beta^*(v)$ as $vH(v) - \int_{\underline{v}}^{v} H(x)dx$ (using integration by parts), so that given $H(v) = (F(v))^{n-1}$ (the probability that $n-1$ valuations are at most v), we have

$$\beta^*(v) = v - \frac{\int_{\underline{v}}^{v} H(x)\,dx}{H(v)} = v - \frac{\int_{\underline{v}}^{v} (F(x))^{n-1}\,dx}{(F(v))^{n-1}} \text{ for } \underline{v} < v \leq \overline{v}. \qquad (309.1)$$

In particular, $\beta^*(v) < v$ for $\underline{v} < v \leq \overline{v}$.

?? EXERCISE 309.2 (Property of the bidding function in a first-price auction) Show that the bidding function β^* defined in (309.1) is increasing.

? EXERCISE 309.3 (Example of Nash equilibrium in a first-price auction) Verify that for the distribution F uniform from 0 to 1 the bidding function defined by (309.1) is $(1 - 1/n)v$.

9.8.2 Revenue equivalence of auctions

I argue at the end of Section 9.6.2 that, under the assumptions in that section, the expected price paid by the winner is the same in a symmetric equilibrium of a first-price auction as it is in the equilibrium of a second-price auction in which each player bids her valuation. A more general result may be established.

Consider an auction that differs from an auction with independent private valuations as defined in Section 9.6.2 only in that the price is an arbitrary continuous function of the bids (not necessarily the highest or second highest bid). Suppose that the auction has a symmetric equilibrium in which the common bidding function β is increasing. Assume that every player other than i adheres to β. Then as i's bid increases, the probability that it exceeds all the other players' bids increases continuously from 0 to 1. Thus i's choosing a bid is tantamount to her choosing a probability of winning. Denote by $e(p)$ her expected payment when her probability of winning is p. Then her probability of winning in the equilibrium solves the problem $\max_p(p \cdot v - e(p))$, where v is her valuation. Assume that this problem has a unique solution, denoted $p^*(v)$. If e is differentiable, the first-order condition implies that $v = e'(p^*(v))$ if $0 < p^*(v) < 1$. Thus if $0 < p^*(v) < 1$ for all v, we may integrate both sides of this equation, obtaining

$$e(p^*(v)) = e(p^*(\underline{v})) + \int_{\underline{v}}^{v} x\, dp^*(x). \qquad (310.1)$$

Now, the common equilibrium bidding function is increasing, so the object is sold to the bidder with the highest valuation. Thus $p^*(v) = \Pr(\mathbf{X} < v)$ and the expected payment $e(p^*(\underline{v}))$ of a bidder with the lowest possible valuation is zero. Hence (310.1) implies that the expected payment $e(p^*(v))$ of a bidder with any given valuation v is equal to $\Pr(\mathbf{X} < v)\mathrm{E}[\mathbf{X} \mid \mathbf{X} < v]$, independent of the price determination rule.

This result is a special case of the more general *revenue equivalence principle*, which applies to a class of common value auctions, as well as private value auctions, and may be stated as follows.

> *Suppose that each bidder (i) is risk neutral, (ii) independently receives a signal from the same continuous and increasing cumulative distribution, and (iii) has a valuation that depends continuously on all the bidders' signals. Consider auction mechanisms in the symmetric Nash equilibria of which the object is sold to the bidder with the highest signal and the expected payoff of a bidder with the lowest possible valuation is zero. In the symmetric Nash equilibrium of any such mechanism, the expected payment of a bidder of any given type is the same, and hence the auctioneer's expected revenue is the same.*

9.8.3 *Reserve prices*

The following exercise (which requires calculus) asks you to show that, in an example, a seller in a second-price sealed-bid auction can increase the expected selling price by imposing a positive "reserve price". (Note that the revenue equivalence principle does not apply to an auction with a positive reserve price because when the highest valuation is less than the reserve price the object is not sold, and in particular is not sold to the bidder with the highest valuation.)

?, EXERCISE 310.2 (Reserve prices in second-price sealed-bid auction) Consider a second-price sealed-bid auction with two bidders in which each bidder's valuation is drawn independently from the distribution uniform from 0 to 1. Suppose that the seller imposes the reserve price r. That is, if both bids are less than r, the object is not sold (and neither bidder makes any payment), if one bid is less than r and the other is at least r, the object is sold at the price r, and if both bids are at least r, the object is sold at a price equal to the second highest bid. Show that for each bidder (and any value of r), a bid equal to her valuation weakly dominates all her other bids. For the Nash equilibrium in which each bidder submits her valuation, find the reserve price r that maximizes the expected price at which the object is sold.

Notes

The notion of a general Bayesian game was defined and studied by Harsanyi (1967/68). The formulation I describe here is taken (with a minor change) from Osborne and Rubinstein (1994, Section 2.6).

The origin of the observation that more information may hurt (Section 9.3.1) is unclear. The idea of "infection" in Section 9.3.2 was first studied by Rubinstein (1989). The game in Figure 284.1 is a variant suggested by Eddie Dekel of the one analyzed by Morris, Rob, and Shin (1995).

Games modeling voluntary contributions to a public good were first considered by Olson (1965, Section I.D) and have been subsequently much studied. The model in Section 9.5 is a variant of one in an unpublished paper of William F. Samuelson dated 1984.

Vickrey (1961) initiated the study of auctions described in Section 9.6. First-price common value auctions (Section 9.6.3) were first studied by Wilson (1967, 1969, 1977). The "winner's curse" appears to have been first articulated by Capen, Clapp, and Campbell (1971). The general revenue equivalence principle at the end of Section 9.8.2 is due to Myerson (1981). The equilibria in Exercise 299.1 are described by Milgrom (1981, Theorem 6.3). The literature is surveyed by Klemperer (1999). The box on spectrum auctions on page 300 is based on McMillan (1994), Cramton (1995, 1997, 1998), and McAfee and McMillan (1996).

Section 9.7 is based on Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998).

Exercise 282.1 was suggested by Ariel Rubinstein. Exercise 282.2 is based on Brams, Kilgour, and Davis (1993). A model of adverse selection was first studied by Akerlof (1970); the model in Exercise 282.3 is taken from Samuelson and Bazerman (1985). Exercise 307.1 is based on Feddersen and Pesendorfer (1996).

10 Extensive Games with Imperfect Information

Prerequisite: Chapters 4 and 5, and Section 7.6

THE NOTION of an extensive game with perfect information defined in Chapter 5 models situations in which each player, whenever taking an action, is informed of the actions chosen previously by all players. In this chapter I discuss a generalization of this notion that allows us to model situations in which each player, when choosing her action, may not be informed of the other players' previous actions.

10.1 Extensive games with imperfect information

To describe an extensive game with *perfect* information, we need to specify the set of players, the set of terminal histories, the player function, and the players' preferences (see Definition 155.1). To describe an extensive game with *imperfect* information, we need to add a single item to this list: a specification of each player's information about the history at every point at which she moves. Denote by H_i the set of histories after which player i moves. We specify player i's information by partitioning (dividing up) H_i into a collection of *information sets*. (See Section 17.2 for the notion of a partition.) This collection is called player i's *information partition*. When making her decision, player i is informed of the information set that has occurred but not of which history within that set has occurred.

Suppose, for example, that player i moves after the histories C, D, and E (i.e. $H_i = \{C, D, E\}$) and is informed only that the history is C, or that it is either D or E. That is, if the history C has occurred, then she is informed that C has occurred, whereas if either D or E has occurred, then she is informed only that either D or E has occurred and not C—she is not informed whether the history that has occurred was D or E. Then player i's information partition consists of two information sets: $\{C\}$ and $\{D, E\}$. If instead she is not informed at all about which history has occurred, then her information partition consists of a single information set, $\{C, D, E\}$, whereas if she is informed precisely about the history, then her information partition consists of three information sets, $\{C\}$, $\{D\}$, and $\{E\}$.

As before, denote the set of actions available to the player who moves after the history h by $A(h)$ (see (156.1)). We allow two histories h and h' to be in the same information set only if $A(h) = A(h')$. Why? The player who moves after any history must know the set of actions available after that history, so if h and h' are in the same information set and $A(h) \neq A(h')$, then the player who moves at this information set can deduce which of these two histories has occurred by looking at the actions available to her. Only if $A(h) = A(h')$ is the player's knowledge of the set of actions available to her consistent with her not knowing whether the history is h or h'. If the information set that contains h and h' is I_i, the common value of $A(h)$ and $A(h')$ is denoted $A(I_i)$. That is, $A(I_i)$ is the set of actions available to player i at her information set I_i.

Many interesting extensive games with imperfect information contain a move of chance, so my definition of an extensive game, unlike my original definition of an extensive game with perfect information (155.1), allows for such moves. Given the presence of such moves, an outcome is a lottery over the set of terminal histories, so each player's preferences must be defined over such lotteries.

DEFINITION 314.1 (*Extensive game*) An **extensive game** (with imperfect information and chance moves) consists of

- a set of **players**
- a set of sequences (**terminal histories**) having the property that no sequence is a proper subhistory of any other sequence
- a function (the **player function**) that assigns either a player or "chance" to every sequence that is a proper subhistory of some terminal history
- a function that assigns to each history that the player function assigns to chance a probability distribution over the actions available after that history, where each such distribution is independent of every other such distribution
- for each player, a partition (the player's **information partition**) of the set of histories assigned to that player by the player function such that for every history h in any given member of the partition, the set $A(h)$ of actions available is the same
- for each player, **preferences** over the set of lotteries over terminal histories.

The simplest extensive games, in which each player moves once and no player, when moving, is informed of any other player's action, model situations that may alternatively be modeled as strategic games, as illustrated by the next example.

EXAMPLE 314.2 (*BoS* as an extensive game) Section 2.3 considers a situation in which each of two people chooses whether to imbibe the pleasures of Bach or those of Stravinsky, and neither person, when choosing a concert, knows the one chosen by the other person. Example 18.2 models this situation as a strategic game. We may alternatively model it as the following extensive game, in which the players choose their actions sequentially, but the second-mover is not informed of the choice made by the first-mover.

Players The two people, say 1 and 2.

Terminal histories (B, B), (B, S), (S, B), and (S, S).

Player function $P(\varnothing) = 1$, $P(B) = P(S) = 2$.

Chance moves None.

Information partitions Player 1's information partition contains a single information set, $\varnothing$ (player 1 has a single move, and when moving she is informed that the game is beginning); player 2's information partition also contains a single information set, $\{B, S\}$ (player 2 has a single move, and when moving she is not informed whether the history is B or S).

Preferences As given in the description of the situation.

This game is illustrated in Figure 315.1. As before, the small circle at the top of the figure represents the start of the game and each line represents an action; the numbers below each terminal history are Bernoulli payoffs whose expected values represent the players' preferences over lotteries. The dotted line connecting the histories B and S indicates that these histories are in the same information set of player 2. Note that the condition that player 2's set of actions at every history within her information set be the same is satisfied: after each such history, player 2's set of actions is $\{B, S\}$.

The next example shows how an extensive game may model a situation having a richer information structure and random events. The situation considered in the example is an extremely simple version of the card game poker, in which there are two players, one of whom is dealt a single card.

EXAMPLE 315.1 (Card game) Each of two players begins by putting a dollar in the pot. Player 1 is then dealt a card that is equally likely to be *High* or *Low*; she observes her card, but player 2 does not. Player 1 may *see* or *raise*. If she *sees*, she

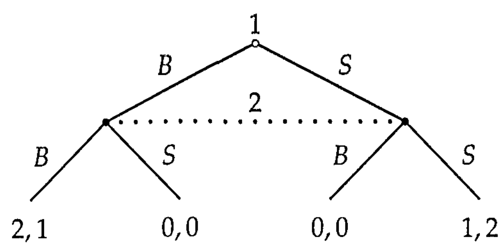

Figure 315.1 The extensive game in Example 314.2. The dotted line (with "2" above it) represents the fact that when moving, player 2 is not informed whether the history is B or S. That is, B and S are in the same information set.

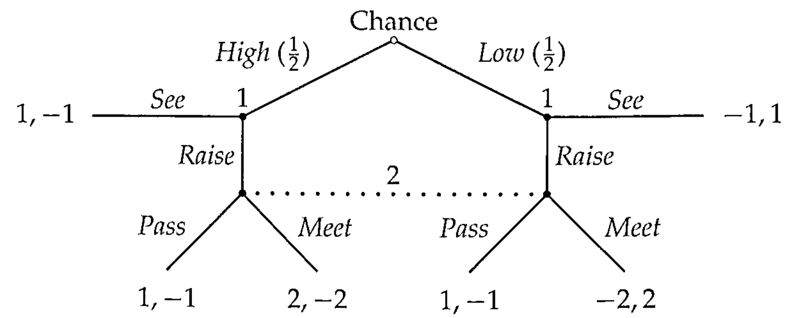

Figure 316.1 An extensive game that models the card game in Example 315.1. The number in parentheses beside each label of the moves of chance is the probability with which that move occurs.

shows her card to player 2. (Player 2 does not have any card for player 1 to see, but you can imagine her holding a fixed card with value between *High* and *Low*. In Exercise 316.1 you are asked to consider a situation in which player 2, like player 1, is dealt a card.) If player 1's card is *High* she takes the money in the pot, and if it is *Low* player 2 takes the money in the pot; in both cases the game ends. If player 1 *raises,* she adds a dollar to the pot and player 2 chooses whether to *pass* or *meet*. If player 2 *passes,* player 1 takes the money in the pot. If player 2 *meets,* she adds a dollar to the pot, and player 1 shows her card. If the card is *High*, player 1 takes the money in the pot, while if it is *Low* player 2 does so.

An extensive game that models this situation is shown in Figure 316.1. Player 1 has two information sets, one containing the single history *High* and one containing the single history *Low*. Player 2 has one information set, consisting of the two histories (*High*, *Raise*) and (*Low*, *Raise*). This information set reflects the fact that player 2 cannot observe player 1's card. Note that again the requirement that the set of actions at every history within an information set be the same is satisfied at player 2's information set.

⑦ EXERCISE 316.1 (Variant of card game) Consider the following variant of the card game in the previous example. Initially each player puts a dollar in the pot. Then *each* player is dealt a card; each player's card is equally likely to be *High* or *Low*, independent of the other player's card. Each player sees only her own card. Player 1 may *see* or *raise*. If she *sees*, then the players compare their cards. The one with the higher card wins the pot; if the cards are the same, then each player takes back the dollar she had put in the pot. If player 1 *raises*, then she adds $\$k$ to the pot (where k is a fixed positive number), and player 2 may *pass* or *meet*. If player 2 *passes*, then player 1 takes the money in the pot. If player 2 *meets*, then she adds $\$k$ to the pot and the players compare cards, the one with the higher card winning the pot; if the cards are the same, then each player takes back the $\$(1+k)$ she had put in the pot. Model this card game as an extensive game. (Drawing a diagram is sufficient; you can avoid the need for information sets to cross histories by putting the initial move in the center of your diagram.)

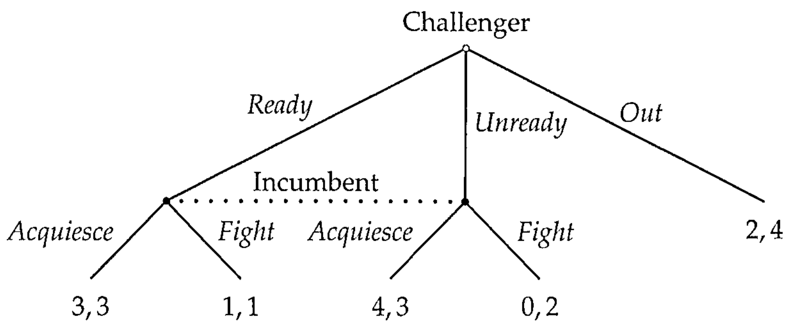

Figure 317.1 An extensive game that models the entry game in Example 317.1. The dotted line indicates that the histories *Ready* and *Unready* are in the same information set. The challenger's payoff is listed first, the incumbent's second.

In these examples, one of the players has an information set that contains more than one history. A game in which every information set of every player contains a single history is equivalent to an extensive game with perfect information. Thus, for instance, the entry game of Example 154.1 can be regarded as a general extensive game in which the challenger has a single information set consisting of the empty history, and the incumbent has a single information set consisting of the history *In*.

The next example is a variant of the entry game in which the challenger, before entering, takes an action that the incumbent does not observe.

EXAMPLE 317.1 (Entry game) An incumbent faces the possibility of entry by a challenger, as in Example 154.1. The challenger has three choices: it may stay out, prepare itself for combat and enter, or enter without making preparations. Preparation is costly but reduces the loss from a fight. The incumbent may either fight or acquiesce to entry. A fight is less costly to the incumbent if the entrant is unprepared; but regardless of the entrant's readiness, the incumbent prefers to acquiesce than to fight. The incumbent observes whether the challenger enters but not whether an entrant is prepared.

An extensive game with imperfect information that models this situation, given some additional assumptions about the players' preferences, is shown in Figure 317.1. Each player has a single information set: the challenger's information set consists of the empty history, and the incumbent's information set consists of the two histories *Ready* (the challenger enters, prepared for combat) and *Unready* (the challenger enters without having made preparations).

10.2 Strategies

As in an extensive game with perfect information, a strategy for a player in a general extensive game specifies the action the player takes whenever it is her turn to move. In a general game, a player takes an action at each of her information sets, so a strategy is defined as follows.

DEFINITION 318.1 (*Strategy in extensive game*) A (pure) **strategy** of player i in an extensive game is a function that assigns to each of i's information sets I_i an action in $A(I_i)$ (the set of actions available to player i at the information set I_i).

Consider, for example, the game in Figure 315.1 (*BoS* as an extensive game). Each player has a single information set, at which two actions, B and S, are available. Thus each player has two strategies, B and S. In the game in Figure 316.1 (the card game), player 1 has two information sets, at each of which she has two actions, *Raise* and *See*. Thus she has four strategies: (*Raise*, *Raise*) (raise whether her card is *High* or *Low*), (*Raise*, *See*) (raise if the card is *High*, see if it is *Low*), (*See*, *Raise*), and (*See*, *See*). Player 2 has a single information set, at which she has two actions, namely *Meet* and *Pass*; thus she has two strategies, *Meet* and *Pass*.

? EXERCISE 318.2 (Strategies in variants of card game and entry game) Find the players' strategies in the games in Exercise 316.1 and Example 317.1.

In some of the games we study, we wish to allow players to choose their actions randomly. One way of doing so is to allow each player to choose a mixed strategy—a probability distribution over her pure strategies—as we did for strategic games (see Definition 107.1).

DEFINITION 318.3 (*Mixed strategy in extensive game*) A **mixed strategy** of a player in an extensive game is a probability distribution over the player's pure strategies.

As before, a pure strategy may be identified with a mixed strategy that assigns probability 1 to the pure strategy. Given this identification, a player's set of mixed strategies contains her pure strategies.

10.3 Nash equilibrium

Having defined a strategy, the definition of a Nash equilibrium is straightforward: a strategy profile is a Nash equilibrium if no player has an alternative strategy that increases her payoff, given the other players' strategies. The following definition closely follows Definition 108.1.

DEFINITION 318.4 (*Nash equilibrium of extensive game*) The mixed strategy profile α^* in an extensive game is a **(mixed strategy) Nash equilibrium** if, for each player i and every mixed strategy α_i of player i, player i's expected payoff to α^* is at least as large as her expected payoff to $(\alpha_i, \alpha^*_{-i})$ according to a payoff function whose expected value represents player i's preferences over lotteries.

As before, I refer to an equilibrium in which no player's strategy entails any randomization (i.e. every player's strategy assigns probability 1 to a single action at each information set) as a *pure Nash equilibrium*.

One way to find a Nash equilibrium of an extensive game is to construct the strategic form of the game and analyze it as a strategic game, as we did for extensive games with perfect information (see Section 5.3).

EXAMPLE 319.1 (*BoS* as an extensive game) Each player in Example 314.2 has two strategies, B and S. The strategic form of the game is exactly the strategic game *BoS*, as given in Figure 19.1. Thus the game has two pure Nash equilibria, (B, B) and (S, S), and a mixed equilibrium in which player 1 uses B with probability $\frac{2}{3}$ and player 2 uses B with probability $\frac{1}{3}$.

In this example player 2, when taking an action, is not informed of the action chosen by player 1. This lacuna in player 2's information is reflected in her information set, which contains both the history B and the history S. However, even though player 2 is not informed of player 1's action, her experience playing the game tells her the history (or probability distribution over histories) to expect. In the steady state in which every person who plays the role of either player chooses B, for example, each player knows that the other player will choose B. She is not *informed* of this fact, but her long experience playing the game leads her to the (correct) conclusion about the other player's action. Similarly, in the steady state in which B is chosen by a third of the people who play the role of player 2, the experience of each person who plays the role of player 1 leads her to expect her adversary to choose B with probability $\frac{1}{3}$. The point is that a player's information partition reflects the information obtained from her observations of the other players' actions during a play of the game; her experience playing the game may yield her more information about the other players' steady state actions.

EXAMPLE 319.2 (Card game) The strategic form of the card game in Example 315.1 is given in Figure 320.1. By inspection, we see that the game has no equilibrium in which either player's strategy is pure. Player 1's strategy (*See*, *See*) is strictly dominated by her mixed strategy that assigns probability $\frac{1}{2}$ to (*Raise*, *Raise*) and probability $\frac{1}{2}$ to (*Raise*, *See*). Thus it is not used with positive probability in any Nash equilibrium. Player 1's strategy (*See*, *Raise*) is not a best response to any mixed strategy of player 2 that assigns positive probability to *Meet*, so given the absence of an equilibrium in pure strategies, in all Nash equilibria player 1 randomizes between (*Raise*, *Raise*) and (*Raise*, *See*). The condition that each player receive the same expected payoff from each strategy to which she assigns positive probability generates the conditions $q = \frac{1}{2}(1 - q)$ and $-p = -\frac{1}{2}(1 - p)$, where p is the probability player 1 assigns to (*Raise*, *Raise*) and q is the probability player 2 assigns to *Pass*. Thus $p = q = \frac{1}{3}$. In conclusion, the game has a unique Nash equilibrium, in which player 1 assigns probability $\frac{1}{3}$ to (*Raise*, *Raise*) and probability $\frac{2}{3}$ to (*Raise*, *See*), and player 2 assigns probability $\frac{1}{3}$ to *Pass* and probability $\frac{2}{3}$ to *Meet*. Player 1's equilibrium strategy implies that she always chooses *Raise* if her card is *High*, and chooses *Raise* with probability $\frac{1}{3}$ if her card is *Low*. That is, player 1 "bluffs" with probability $\frac{1}{3}$.

? EXERCISE 319.3 (Nash equilibria of card game) Consider a generalization of the card game in Example 315.1 in which the amount added to the pot when player 1 raises and player 2 meets is \$$k$. (In Example 315.1, $k = 1$). Find the Nash equilibrium of this game for any $k > 0$. How does player 1's propensity to bluff depend on the value of k?

	Pass	*Meet*
Raise, Raise	$1,-1$	$0,0$
Raise, See	$0,0$	$\frac{1}{2},-\frac{1}{2}$
See, Raise	$1,-1$	$-\frac{1}{2},\frac{1}{2}$
See, See	$0,0$	$0,0$

Figure 320.1 The strategic form of the card game in Example 315.1.

	X	Y
X	$3,2$	$1,1$
Y	$4,3$	$2,4$

Figure 320.2 The payoffs in the situation considered in Example 320.2.

? EXERCISE 320.1 (Nash equilibria of variant of card game) Find the Nash equilibria of the game in Exercise 316.1 for $0 < k < 1$ and for $k > 1$. How does player 1's propensity to bluff depend on the value of k?

EXAMPLE 320.2 (Commitment and observability) Two people each have two actions, X and Y. Their payoffs to the four action pairs are shown in Figure 320.2.

First suppose that they choose actions simultaneously. The strategic game that models this situation has a unique Nash equilibrium, in which both players choose Y. (Note that X is strictly dominated by Y for player 1.)

Now suppose that the players choose their actions sequentially, player 1 first, and that player 2 observes player 1's action before choosing her action. That is, player 1 commits to her action before player 2 chooses her action. The extensive game with perfect information that models this situation has a single subgame perfect equilibrium, in which both players choose X. Player 1 is better off in this equilibrium than she is in the equilibrium of the simultaneous-move game.

Finally suppose that player 1 moves first, but her action is not perfectly observed by player 2. If player 1 chooses X, player 2 may think she chose Y, and vice versa. We may model this situation as an extensive game with imperfect information in which player 1's move is followed by a move of chance that selects a signal observed by player 2. Player 2 observes the signal, but not player 1's action. Assume that the probability that the signal is correct is the same for both actions, and less than 1. Denote this probability by $1 - \epsilon$. (Thus if player 1 chooses X, the signal is X with probability $1 - \epsilon$ and Y with probability ϵ, and if player 1 chooses Y, it is Y with probability $1 - \epsilon$ and X with probability ϵ.) Assume that $0 \leq \epsilon < \frac{1}{4}$.

This extensive game, its strategic form, and its Nash equilibria are shown in Figure 321.1. (The strategy IJ of player 2, where I and J are both either X or Y, means choose I after the signal X and J after the signal Y. The mixed strategy equilibria may be found by using the methods of Section 4.10.) We see, in particular, that the game has a pure strategy Nash equilibrium (Y, YY) for all values of ϵ with $0 \leq \epsilon < \frac{1}{4}$.

Extensive game:

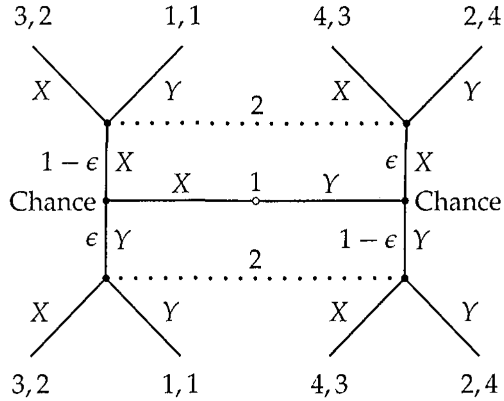

Strategic form:

	XX	XY	YX	YY
X	$3,2$	$3-2\epsilon, 2-\epsilon$	$1+2\epsilon, 1+\epsilon$	$1,1$
Y	$4,3$	$2+2\epsilon, 4-\epsilon$	$4-2\epsilon, 3+\epsilon$	$2,4$

Nash equilibrium strategy pairs:

Player 1		Player 2			
X	Y	XX	XY	YX	YY
0	1	0	0	0	1
ϵ	$1-\epsilon$	0	$\frac{1}{2-4\epsilon}$	0	$\frac{1-4\epsilon}{2-4\epsilon}$
$1-\epsilon$	ϵ	$\frac{1-4\epsilon}{2-4\epsilon}$	$\frac{1}{2-4\epsilon}$	0	0

Figure 321.1 An extensive game (top) that models a situation in which the payoffs are those in Figure 320.2, player 1 moves first, and player 2 observes an imperfect signal of player 1's action. Note that the empty history is in the center of the game. The strategic form of the game is given in the middle of the figure, followed by its Nash equilibria for $0 \le \epsilon < \frac{1}{4}$.

In summary, the game in which player 1 moves first and her action is perfectly observable has a unique subgame perfect equilibrium, the outcome of which is that both players choose X, whereas the game in which player 1's action is observable with error has a pure strategy Nash equilibrium in which the outcome is that both players choose Y, regardless of how small the error.

Thus the advantage gained by the commitment entailed in being the first-mover, as reflected in the subgame perfect equilibrium of the game with perfect information, is completely lost in the pure strategy Nash equilibrium of the game in which the second player's observation of the first-mover's action is even slightly imperfect. Why? Suppose both player 1 and player 2 choose Y, and consider the implications of player 1's switching to X. In the game with perfect information, player 2's observation of X is inconsistent with the equilibrium; she interprets it as a deviation, to which she optimally responds by choosing X, making the deviation worthwhile for player 1. In the game with imperfect information, player 2's observation of X is consistent with the equilibrium; she interprets it as an inaccurate signal (regardless of how unlikely such a signal is) and continues to choose Y, making player 1's deviation undesirable for her.

In Chapter 5 we saw that the notion of Nash equilibrium is not adequate in all extensive games with perfect information, and I developed the notion of subgame perfect equilibrium to deal with the problem. How may we extend the idea behind subgame perfect equilibrium to the larger class of extensive games with (possibly) imperfect information?

EXAMPLE 322.1 (Entry game) The strategic form of the entry game in Example 317.1 is shown in Figure 322.1. We see that the game has two Nash equilibria in pure strategies, (*Unready*, *Acquiesce*) and (*Out*, *Fight*). (You may verify that it has also mixed strategy Nash equilibria in which the challenger uses the pure strategy *Out* and the probability assigned by the incumbent to *Acquiesce* is at most $\frac{1}{2}$.)

	Acquiesce	*Fight*
Ready	3,3	1,1
Unready	4,3	0,2
Out	2,4	2,4

Figure 322.1 The strategic form of the entry game of Example 317.1.

As in the version of the entry game with perfect information studied in Chapter 5, the Nash equilibrium (*Out*, *Fight*) is not plausible. If in fact the challenger enters, the incumbent optimally acquiesces, regardless of the history that has occurred (i.e. regardless of whether the challenger is ready for combat). In games with perfect information we eliminated such equilibria by defining the notion of subgame perfect equilibrium, which requires that each player's strategy be optimal, given the other players' strategies, for every history after which she moves, regardless of whether the history occurs if the players adhere to their strategies.

The natural extension of this idea to games with imperfect information requires that each player's strategy be optimal at each of her information sets. In the entry game we are studying, the implementation of this idea is straightforward. The incumbent's action *Fight* is unambiguously suboptimal at its information set because the incumbent prefers *Acquiesce* if the challenger enters, regardless of whether the challenger is ready. Thus the equilibria in which the incumbent assigns positive probability to *Fight* do not satisfy the added requirement. The remaining equilibrium, (*Unready*, *Acquiesce*), does satisfy the requirement: the incumbent chooses its unambiguously optimal strategy at its information set.

The implementation of the idea in other games is less straightforward because the optimality of an action at an information set may depend on the history that has occurred. Consider a variant of the entry game in which the incumbent prefers to fight than to accommodate an unprepared entrant, as illustrated in Figure 323.1. Like the original game, this game has a Nash equilibrium in which the challenger

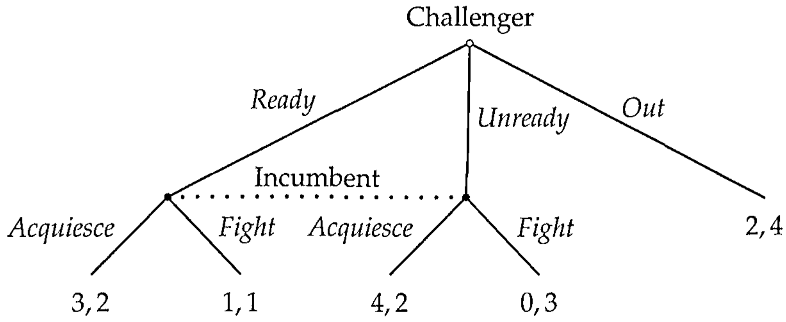

Figure 323.1 A variant of the entry game in Example 317.1 in which the incumbent prefers to fight if the challenger enters unprepared. The challenger's payoff is listed first, the incumbent's second.

stays out and the incumbent fights. But given that fighting is now optimal if the challenger enters unprepared, the reasonableness of the equilibrium in the modified game depends on the history the incumbent believes has occurred if it has to move. The challenger's strategy *Out* gives the incumbent no basis on which to form such a belief, so to pursue the argument we need to consider what this belief might be. In the next section I describe a solution for extensive games that includes the players' beliefs as part of an equilibrium.

10.4 Beliefs and sequential equilibrium

A Nash equilibrium of a strategic game may be characterized by two requirements: that each player choose her best action given her belief about the other players, and that each player's belief be correct (see page 21). The notion of equilibrium I now define for extensive games embodies the same two requirements, and, like the notion of subgame perfect equilibrium for extensive games with perfect information, insists that they hold at each point at which a player has to choose an action. When defining a Nash equilibrium of a strategic game precisely, we do not need to consider the players' beliefs separately from their strategies because the requirement that the beliefs be correct completely determines them: each player's belief about every other player's strategy is simply equal to that strategy. For an extensive game, the players' strategies may not completely determine their beliefs, as we have just seen in the game in Figure 323.1. Thus we are led to define a notion of equilibrium for pairs consisting of a strategy profile and a collection of beliefs.

10.4.1 Beliefs

We assume that at an information set that contains more than one history, the player whose turn it is to move forms a belief about the history that has occurred. We model this belief as a probability distribution over the histories in the information set. (At an information set containing a single history, the only possible belief assigns probability 1 to that history.) We call a collection of beliefs, one for each information set of every player, a *belief system*.

DEFINITION 324.1 A **belief system** in an extensive game is a function that assigns to each information set a probability distribution over the histories in that information set.

Consider, for example, the entry game of Example 317.1. This game has two information sets, one consisting of the empty history, and one consisting of the histories *Ready* and *Unready*. Thus a belief system for the game consists of a pair of probability distributions: one assigns probability 1 to the empty history (the challenger's belief at the start of the game), and the other assigns probabilities to the histories *Ready* and *Unready* (the incumbent's belief after the challenger enters).

10.4.2 Strategies

When studying Nash equilibrium in the previous section we incorporated the possibility of randomization by allowing players to choose mixed strategies. In the current context, a different formulation is convenient. Rather than allowing a player to choose a probability distribution over her pure strategies, we have her assign to each of her information sets a probability distribution over the actions available at that set. We refer to such an assignment as a *behavioral strategy*.

DEFINITION 324.2 (*Behavioral strategy in extensive game*) A **behavioral strategy** of player i in an extensive game is a function that assigns to each of i's information sets I_i a probability distribution over the actions in $A(I_i)$, with the property that each probability distribution is independent of every other distribution.

A behavioral strategy in which every probability distribution assigns probability 1 to a single action is equivalent to a pure strategy. In all the games we study, mixed strategies and behavioral strategies are equivalent; I choose to use behavioral strategies now only because they are easier to work with. The relation between the two types of strategy is illustrated by two examples.

In the game in Example 314.2 (*BoS* as an extensive game), each player has a single information set, so a behavioral strategy for each player is a single probability distribution over her actions. Thus in this game each player's set of behavioral strategies is identical to her set of mixed strategies.

In the game in Example 315.1 (the card game), player 1 has two information sets, so a behavioral strategy for her consists of a pair of probability distributions, each over the set $\{Raise, See\}$. By contrast, a mixed strategy for player 1 in this game is a single probability distribution over the set of her four pure strategies, $\{(Raise, Raise), (Raise, See), (See, Raise), (See, See)\}$. A behavioral strategy is thus specified by two numbers (the probabilities of *Raise* if player 1's card is *High* and if it is *Low*), while a mixed strategy is specified by three numbers. (Remember that a probability distribution over k actions is determined by $k-1$ numbers, because the k probabilities must sum to 1.) In this sense a behavioral strategy is simpler than a mixed strategy. However, restricting a player to behavioral strategies does not limit her options, in the sense that for every mixed strategy and every list of strategies for the other players there is a behavioral strategy that generates the same probability distribution over outcomes. (This result is not limited to this game—it holds in a very wide class of games that includes all the games studied in this chapter.)

10.4.3 Equilibrium

The notion of equilibrium I define applies to pairs consisting of a profile of behavioral strategies and a belief system; the following piece of terminology is convenient.

DEFINITION 325.1 (*Assessment*) An **assessment** in an extensive game is a pair consisting of a profile of behavioral strategies and a belief system.

An assessment is an equilibrium if it satisfies the following two requirements (which I make precise in the subsequent discussion).

Sequential rationality Each player's strategy is optimal whenever she has to move, given her belief and the other players' strategies.

Consistency of beliefs with strategies Each player's belief is consistent with the strategy profile.

The requirement of sequential rationality generalizes the requirement of subgame perfect equilibrium that each player's strategy be optimal in the part of the game that follows each history after which she moves, given the strategy profile, regardless of whether this history occurs if the players follow their strategies. In the more general context of an extensive game, sequential rationality requires that each player's strategy be optimal *in the part of the game that follows each of her information sets*, given the strategy profile *and given the player's belief about the history in the information set that has occurred*, regardless of whether the information set is reached if the players follow their strategies.

Consider, for example, the game in Figure 326.1. Suppose that player 1's strategy (indicated by the black branches) selects E at the start of the game and J after the history (C, F), and player 2's belief at her information set (indicated by the numbers in brackets) is that the history C has occurred with probability $\frac{2}{3}$ and the history D has occurred with probability $\frac{1}{3}$. Sequential rationality requires that player 2's strategy be optimal at her information set, given the subsequent behavior specified by player 1's strategy, even though this set is not reached if player 1 follows her strategy. Player 2's expected payoff in the part of the game starting at her information set, given her belief, is $\frac{2}{3} \cdot 0 + \frac{1}{3} \cdot 1 = \frac{1}{3}$ to her strategy F and $\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3}$ to her strategy G. Thus sequential rationality requires her to select G. Sequential rationality requires also that player 1's strategy be optimal at each of her two (one element) information sets, given player 2's strategy. Player 1's optimal action after the history (C, F) is J; if player 2's strategy is G, her optimal actions at the start of the game are D and E. Thus given player 2's strategy G, player 1 has two optimal strategies, DJ and EJ.

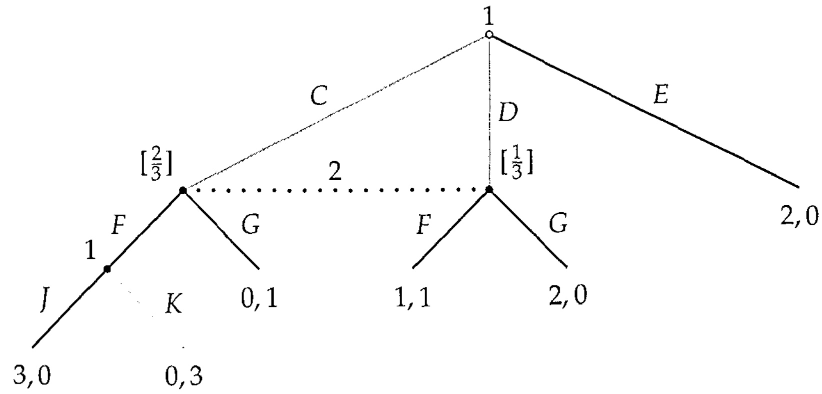

Figure 326.1 A game illustrating the sequential rationality requirement on an assessment. Given player 1's strategy, indicated by the black branches, and player 2's belief, given by the numbers in brackets at the end of the histories C and D, player 2's expected payoff to G exceeds her expected payoff to F, so that her optimal action is G.

Now let (β, μ) be an assessment in an arbitrary game (i.e. let β be a profile of behavioral strategies and μ a belief system), and let I_i be an information set of player i. Denote by $O_{I_i}(\beta, \mu)$ the probability distribution over terminal histories that results if each history in I_i occurs with the probability assigned to it by i's belief μ_i (which is not necessarily the probability with which it occurs if the players adhere to β), and *subsequently* the players adhere to the strategy profile β. (Compare this definition with that of $O_h(s)$ for an extensive game with perfect information (see page 165). For the information set $\{C, D\}$ and the strategy of player 1 and belief of player 2 indicated in the game in Figure 326.1, the probability distribution assigns probability $\frac{2}{3}$ to the terminal history (C, F, J) and probability $\frac{1}{3}$ to (D, F) if player 2 uses the strategy F, and probability $\frac{2}{3}$ to (C, G) and probability $\frac{1}{3}$ to (D, G) if she uses the strategy G.)

We can now state precisely the sequential rationality requirement: for each player i and each of her information sets I_i, her expected payoff to $O_{I_i}(\beta, \mu)$ is at least as large as her expected payoff to $O_{I_i}((\gamma_i, \beta_{-i}), \mu)$ for each of her behavioral strategies γ_i.

The requirement that each player's belief be consistent with the strategy profile is new, though forms of it are implicit in the definitions of Nash equilibrium and subgame perfect equilibrium. The idea is that in a steady state, each player's belief must be correct: the probability it assigns to any history must be the probability with which that history occurs if the players adhere to their strategies. The implementation of this idea is straightforward at an information set reached with positive probability if the players follow their strategies, but is unclear at an information set not reached if the players follow their strategies. *Every* history in such an information set has probability 0 if the players follow their strategies, but if such an information set is reached, the player who moves there must believe that *some* history has occurred. We deal with this difficulty by allowing the player who moves at such an information set to hold *any* belief at that information set.

Thus we formulate the consistency requirement to restrict the belief system only at information sets reached with positive probability if every player adheres to her strategy. Precisely, we require that the probability assigned to every history h^* in such an information set by the belief of the player who moves there be equal to the probability that h^* occurs according to the strategy profile, conditional on the information set's being reached. Denoting the information set by I_i and the strategy profile by β, by Bayes' rule (Section 17.6.5) this probability is

$$\frac{\Pr(h^* \text{ according to } \beta)}{\sum_{h \in I_i} \Pr(h \text{ according to } \beta)}. \tag{327.1}$$

Consider again the game in Figure 326.1. If player 1's behavioral strategy assigns probability 1 to the action E at the start of the game, the consistency requirement places no restriction on player 2's belief because player 2's information set is not reached if player 1 adheres to her strategy. If player 1's action at the start of the game assigns positive probability to C or D, however, the consistency requirement bites. If player 1's strategy assigns probability p to C and probability q to D, so that $\Pr(C \text{ according to } \beta) = p$ and $\Pr(D \text{ according to } \beta) = q$ (where β is the strategy profile), then (327.1) implies that player 2's belief assigns probability $p/(p+q)$ to C and probability $q/(p+q)$ to D. In particular, if player 1 chooses D with probability 1, then player 2's belief must assign probability 0 to C and probability 1 to D.

EXAMPLE 327.2 (Consistent beliefs in *BoS* as an extensive game) In the game in Example 314.2, for every strategy of player 1, player 2's information set is reached with probability 1, so the consistency condition restricts player 2's belief in every assessment. The condition simply requires that player 2's belief assign to B and S the probabilities with which player 1 chooses these actions. That is, in this game consistency requires that player 2's belief always be correct. The same is true for any game in which all the players move simultaneously (no player knows any other player's action when making her move).

EXAMPLE 327.3 (Consistent beliefs in card game) Consider the game in Example 315.1. If player 1's strategy selects *See* whether her card is *High* or *Low*, the consistency condition does not restrict player 2's belief, because player 2's information set is not reached. For every other strategy of player 1, the condition determines player 2's belief. Denote the probability that player 1 chooses *Raise* if her card is *High* by p_H and the probability she chooses *Raise* if her card is *Low* by p_L. Then $\Pr((\textit{High}, \textit{Raise}) \text{ according to } \beta) = \frac{1}{2}p_H$ and $\Pr((\textit{Low}, \textit{Raise}) \text{ according to } \beta) = \frac{1}{2}p_L$, so by (327.1), the consistency requirement is that player 2's belief assign probability $p_H/(p_H + p_L)$ to the history H and probability $p_L/(p_H + p_L)$ to the history L.

EXAMPLE 327.4 (Consistent beliefs in entry game) In the games in Figures 317.1 and 323.1, denote by p_R, p_U, and p_O the probabilities the challenger assigns to *Ready*, *Unready*, and *Out*. If $p_O = 1$, the consistency condition does not restrict the incumbent's belief. Otherwise, the condition requires that the incumbent assign probability $p_R/(p_R + p_U)$ to *Ready* and probability $p_U/(p_R + p_U)$ to *Unready*.

In summary, the notion of equilibrium I use for extensive games is defined precisely as follows.

DEFINITION 328.1 (*Weak sequential equilibrium*) An assessment (β, μ) (consisting of a behavioral strategy profile β and a belief system μ) is a **weak sequential equilibrium** if it satisfies the following two conditions.

Sequential rationality Each player's strategy is optimal in the part of the game that follows each of her information sets, given the strategy profile and her belief about the history in the information set that has occurred. Precisely, for each player i and each information set I_i of player i, player i's expected payoff to the probability distribution $O_{I_i}(\beta, \mu)$ over terminal histories generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile β is at least as large as her expected payoff to the probability distribution $O_{I_i}((\gamma_i, \beta_{-i}), \mu)$ generated by her belief μ_i at I_i and the behavior prescribed subsequently by the strategy profile (γ_i, β_{-i}), for each of her behavioral strategies γ_i.

Weak consistency of beliefs with strategies For every information set I_i reached with positive probability given the strategy profile β, the probability assigned by the belief system to each history h^* in I_i is given by (327.1).

The game in Figure 326.1 illustrates this notion of equilibrium. As we have seen, in this game player 1's strategy EJ is sequentially rational given player 2's strategy G, and player 2's strategy G is sequentially rational given the beliefs indicated in the figure and player 1's strategy EJ. Further, the belief is consistent with the strategy profile (EJ, G), because this profile does not lead to player 2's information set. Thus the game has a weak sequential equilibrium in which the strategy profile is (EJ, G) and player 2's belief is the one indicated in the figure (or any other belief for which G is optimal). Now consider the strategy profile (DJ, G). Player 1's strategy is sequentially rational given player 2's strategy, and player 2's strategy is sequentially rational given her belief and player 1's strategy. But player 2's belief is not consistent with the strategy profile. The only consistent belief assigns probability 1 to D, which makes player 2's action F, rather than G, optimal. Thus the game has no weak sequential equilibrium in which the strategy profile is (DJ, G).

In an extensive game with perfect information, only one belief system is possible—that in which every player believes at each information set that the single compatible history has occurred with probability 1. Thus the condition of sequential rationality in such a game is the same as the optimality condition on a strategy profile in a subgame perfect equilibrium. Thus

> *in an extensive game with perfect information, the strategy profile in any weak sequential equilibrium is a subgame perfect equilibrium.*

In a general extensive game, the requirement of sequential rationality implies, in particular, that each player's strategy is optimal at the beginning of the game, given the other players' strategies and the player's belief about the history at each information set. Further, the consistency requirement implies that each player's belief about the history is correct at any information set reached with positive probability when the players follow their strategies. Thus if an assessment is a weak sequential equilibrium, then each player's strategy in the assessment is optimal at the beginning of the game, given the other players' strategies. That is,

the strategy profile in any weak sequential equilibrium is a Nash equilibrium.

To find the weak sequential equilibria of a game, we may use a combination of the techniques for finding subgame perfect equilibria of extensive games with perfect information and for finding Nash equilibria of strategic games. Consider, for example, the game in Figure 326.1. The sequential rationality requirement implies that in any weak sequential equilibrium player 1 chooses J after the history (C, F). Now we can consider two cases.

- Does the game have a weak sequential equilibrium in which player 1 chooses E? If player 1 chooses E, player 2's belief is not restricted by consistency, so we need to ask (*a*) whether any strategy of player 2 makes E optimal for player 1, and, if so, (*b*) whether there is a belief of player 2 that makes any such strategy optimal. We see that E is optimal if and only if player 2 chooses F with probability at most $\frac{2}{3}$. Any such strategy of player 2 is optimal if player 2 believes the history is C with probability $\frac{1}{2}$, and the strategy of choosing F with probability 0 is optimal if player 2 believes the history is C with any probability of at least $\frac{1}{2}$. We conclude that an assessment is a weak sequential equilibrium if player 1's strategy is EJ and player 2 either chooses F with probability at most $\frac{2}{3}$ and believes that the history is C with probability $\frac{1}{2}$, or chooses G and believes that the history is C with probability at least $\frac{1}{2}$.

- Does the game have a weak sequential equilibrium in which player 1 chooses C and/or D with positive probability? Denote the probability player 1 assigns to C by p and the probability she assigns to D by q. In such an equilibrium, player 2's belief is constrained by consistency to assign probability $p/(p+q)$ to C and probability $q/(p+q)$ to D. Thus G is optimal if $p > q$, F is optimal if $p < q$, and any mixture is optimal if $p = q$. Now, if player 2 chooses G, the only optimal strategy of player 1 that assigns positive probability to C or D assigns probability 1 to D, so that $q = 1$, making G not optimal for player 2. If player 2 chooses F, player 1's only optimal action at the start of the game is C, making F not optimal for player 2. If player 2 chooses a mixture of F and G, then player 1 must assign the same probability to both C and D, and hence must choose D with positive probability, which is incompatible with the fact that her expected payoff to D is less than her expected payoff to E. Thus the game has no weak sequential equilibrium in which player 1 chooses C and/or D with positive probability.

We conclude that in every weak sequential equilibrium of the game, player 1's strategy is EJ and player 2 either chooses F with probability at most $\frac{2}{3}$ and believes that the history is C with probability $\frac{1}{2}$, or chooses G and believes that the history is C with probability at least $\frac{1}{2}$. (Another method of finding these equilibria is to find all the Nash equilibria of the game, and then check which of these equilibria are associated with weak sequential equilibria.)

EXAMPLE 330.1 (Weak sequential equilibria of entry game) As we saw in Example 322.1, the entry game of Example 317.1 has two pure strategy Nash equilibria, (*Unready*, *Acquiesce*) and (*Out*, *Fight*). Consider the first equilibrium. Consistency requires that the incumbent believe that the history is *Unready* at its information set, making *Acquiesce* optimal. Thus the game has a weak sequential equilibrium in which the strategy profile is (*Unready*, *Acquiesce*) and the incumbent's belief is that the history is *Unready*. Now consider the second equilibrium. Regardless of the incumbent's belief at its information set, *Fight* is not an optimal action in the remainder of the game—for every belief, *Acquiesce* yields a higher payoff than does *Fight*. Thus no assessment in which the strategy profile is (*Out*, *Fight*) is both sequentially rational and consistent, so that the game has no weak sequential equilibrium in which the strategy profile is (*Out*, *Fight*).

The games in Examples 314.2 (*BoS* as an extensive game) and 315.1 (the card game) have no Nash equilibria in which some information set is not reached, so that for each game the set of strategy profiles that appear in weak sequential equilibria is equal to the set of Nash equilibria.

Why "weak" sequential equilibrium? The consistency condition's limitation to information sets reached with positive probability generates, in some games, a relatively large set of equilibrium assessments, some of which do not plausibly correspond to steady states. Consider, for example, the game in Figure 331.1, a variant of the entry game in Figure 323.1 in which *Ready* is better than *Unready* for the challenger regardless of the incumbent's action. This game has a weak sequential equilibrium in which the challenger's strategy is *Out*, the incumbent's strategy is F, and the incumbent believes at its information set that the history is *Unready*. In this equilibrium the incumbent believes that the challenger has chosen *Unready*, although this action is dominated (by *Ready*) for the challenger. Thus the incumbent's belief does not seem reasonable.

Notions of equilibrium that impose stronger consistency conditions narrow down the set of equilibrium outcomes in some games. One notion, sequential equilibrium, is widely used. However, the consistency condition it imposes is not straightforward, and the notion does not in any case eliminate all implausible equilibria—for example, it does not eliminate the equilibrium of the game in Figure 331.1 in which the challenger chooses *Out*. Further, in many games the

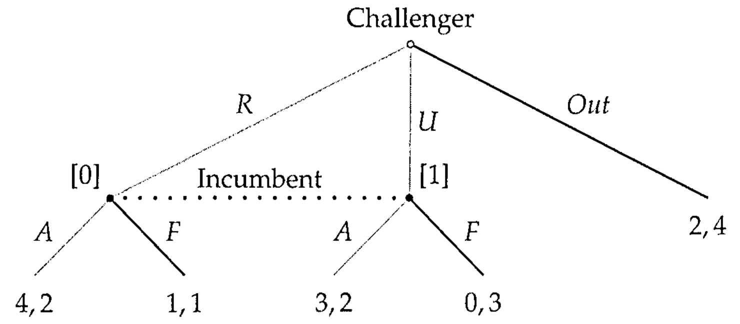

Figure 331.1 A variant of the entry game in Figure 323.1 in which *Ready* is better than *Unready* for the challenger regardless of the incumbent's action and the incumbent prefers *Fight* to *Acquiesce* only if the challenger is unprepared, but there is a weak sequential equilibrium in which the challenger stays out and the incumbent fights. The players' strategies in this equilibrium are indicated by the black lines; the probabilities the incumbent's belief assigns to each history in its information set are shown in brackets.

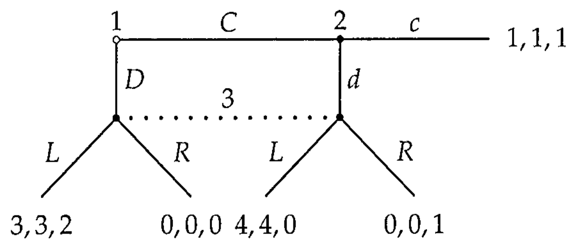

Figure 331.2 The game in Exercise 331.1 (Selten's horse).

set of weak sequential equilibria is not large and does not include implausible assessments. For these reasons I work with weak sequential equilibrium and do not consider "refinements" of this notion.

? EXERCISE 331.1 (Selten's horse) Find the weak sequential equilibria of the game in Figure 331.2 in which each player's strategy is pure. (Find the pure strategy Nash equilibria, then determine which is part of a weak sequential equilibrium. The name of the game comes from the person who first studied it, and its shape.)

? EXERCISE 331.2 (Weak sequential equilibrium and Nash equilibrium in subgames) Consider the variant of the game in Figure 331.1 shown in Figure 332.1, in which the challenger's initial move is broken into two steps. Show that this game has a weak sequential equilibrium in which the players' actions in the subgame following the history *In* do not constitute a Nash equilibrium of the subgame.

10.5 Signaling games

In many interactions, information is "asymmetric": some parties are informed about variables that affect everyone, and some parties are not. In one interesting class of situations, the informed parties have the opportunity to take actions

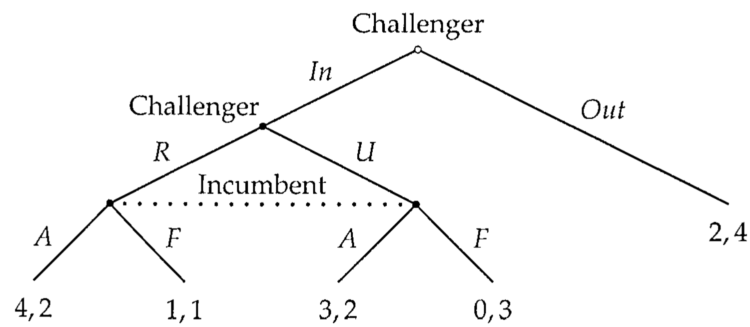

Figure 332.1 A variant of the entry game in Figure 331.1 in which the challenger's decision is broken into two steps.

observed by the uninformed parties before the latter take actions that affect everyone. In some circumstances, the informed parties' actions may "signal" their information.

Suppose, for example, that an employer can observe whether a job applicant has a college degree, but not her ability, and that a person of high ability may obtain a degree at low cost, while one of low ability may do so only at high cost. Then the fact that a person has a degree may signal to an employer that she has high ability—not because college teaches any skills, but because only high-ability individuals find obtaining a degree worthwhile, given the cost. (I return to this example in Section 10.7.)

In a general two-player "signaling game", a *sender* is informed about a variable relevant to both her and a *receiver* (or set of receivers), who is uninformed. The sender takes an action observed by the receiver, who then takes an action that affects them both. Depending on the way in which the message and the receiver's action affect the parties, the sender may want to limit or distort the information her signal conveys.

Such a situation may be modeled as an extensive game in which the sender has several possible "types", each corresponding to a value of the variable about which she is informed. The value she observes, and thus her type, are determined by chance. The receiver does not observe the sender's type, but sees an action she takes, and then herself takes an action. A simple example is given by another variant of the entry game.

EXAMPLE 332.1 (Entry as a signaling game) A challenger contests an incumbent's turf. The challenger is *strong* with probability p and *weak* with probability $1 - p$, where $0 < p < 1$; it knows its type, but the incumbent does not. The challenger may either *ready* itself for battle or remain *unready*. (It does not have the option of staying out.) The incumbent observes the challenger's readiness, but not its type, and chooses whether to *fight* or *acquiesce*. An unready challenger's payoff is 5 if the incumbent acquiesces to its entry. Preparations cost a strong challenger 1 unit of payoff and a weak one 3 units, and fighting entails a loss of 2 units for each type. The incumbent prefers to fight (payoff 1) rather than to acquiesce to (payoff

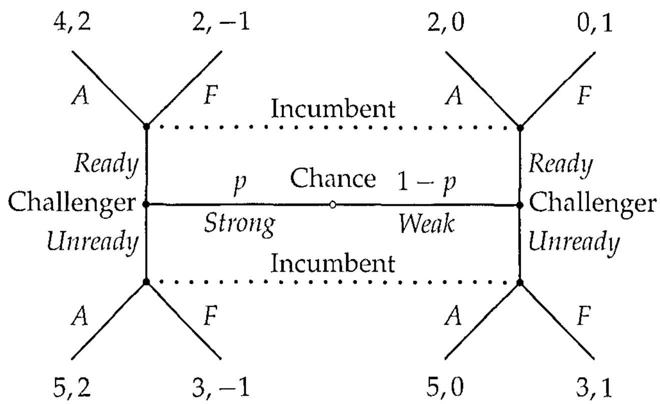

Figure 333.1 The game in Example 332.1. The empty history is in the center of the diagram. The incumbent's actions are A (acquiesce) and F (fight). The challenger's payoff is listed first, the incumbent's second.

0) a weak challenger (who is quickly dispensed with), and prefers to acquiesce to (payoff 2) rather than to fight (payoff -1) a strong one.

The extensive game in Figure 333.1 models this situation. Note that the empty history is in the center of the diagram. The first move is made by chance, which determines the challenger's type. Both types have two actions, *Ready* and *Unready*, so that the challenger has four strategies. The incumbent has two information sets, at each of which it has two actions (A and F), and thus also has four strategies.

I now find the pure weak sequential equilibria of this game. First note that a weak challenger prefers *Unready* to *Ready* regardless of the incumbent's actions—even if the incumbent acquiesces to a ready challenger and fights an unready one, the challenger prefers to be unready. Thus in any weak sequential equilibrium a weak challenger chooses *Unready*. I consider each possible action of a strong challenger in turn.

Strong challenger chooses *Ready* Both the incumbent's information sets are reached, so the consistency condition restricts its belief at each set. At the top information set, the incumbent must believe that the history was (*Strong*, *Ready*), and hence choose A. At the bottom information set it must believe that the history was (*Weak*, *Unready*), and hence choose F. Thus if the challenger deviates and chooses *Unready* if it is strong, then it is worse off—it obtains the payoff of 3 rather than 4. We conclude that the game has a weak sequential equilibrium in which the challenger chooses *Ready* when it is strong and *Unready* when it is weak, and the incumbent acquiesces when it sees *Ready* and fights when it sees *Unready*.

Strong challenger chooses *Unready* At its bottom information set, the incumbent believes, by consistency, that the history was (*Strong*, *Unready*) with probability p and (*Weak*, *Unready*) with probability $1 - p$. Thus its expected payoff to A is $2p$ and to F is $-p + 1 - p = 1 - 2p$. Hence A is optimal if $p \geq \frac{1}{4}$ and F is optimal if $p \leq \frac{1}{4}$.

Suppose that $p \geq \frac{1}{4}$ and the incumbent chooses A in response to *Unready*. A strong challenger that chooses *Unready* obtains the payoff of 5. If it switches to *Ready* its payoff is less than 5 regardless of the incumbent's action. Thus if $p \geq \frac{1}{4}$, then the game has weak sequential equilibria in which both types of challenger choose *Unready* and the incumbent acquiesces to an unready challenger. The incumbent may hold any belief about the type of a ready challenger, and, depending on its belief, may fight or acquiesce.

Now suppose that $p \leq \frac{1}{4}$ and that the incumbent chooses F in response to *Unready*. A strong challenger that chooses *Unready* obtains the payoff of 3. If it switches to *Ready* its payoff is 2 if the incumbent fights and 4 if it acquiesces. Thus for an equilibrium the incumbent must fight a ready challenger. If it believes that a ready challenger is weak with high enough probability (at least $\frac{3}{4}$), fighting is indeed optimal. Is such a belief possible in an equilibrium? Yes: the consistency condition does not restrict the incumbent's belief upon observing *Ready* because this action is not taken when the challenger follows its strategy to choose *Unready* regardless of its type. Thus if $p \leq \frac{1}{4}$ the game has a weak sequential equilibrium in which both types of challenger choose *Unready*, the incumbent fights regardless of the challenger's action, and assigns probability of at least $\frac{3}{4}$ to the challenger's being weak if it observes that the challenger is ready for battle.

In summary, the game has two sorts of weak sequential equilibrium.

- The challenger chooses *Ready* if it is strong and *Unready* if it is weak. The incumbent believes that a ready challenger is strong and an unready one is weak and acquiesces to a ready challenger and fights an unready one.
- The challenger chooses *Unready* regardless of its type. The incumbent believes that an unready challenger is strong with probability p. If $p \geq \frac{1}{4}$ the game has equilibria in which the incumbent acquiesces to an unready challenger, holds any belief about the type of a ready challenger, and takes whatever action is optimal, given its belief, if the challenger is ready. If $p \leq \frac{1}{4}$ the game has equilibria in which the incumbent fights all challengers and believes a ready challenger is strong with probability at most $\frac{1}{4}$.

This example illustrates two kinds of pure strategy equilibrium that may exist in signaling games.

Separating equilibrium Each type of the sender chooses a different action (in the first sort of equilibrium in the example, a strong challenger chooses *Ready* and a weak challenger chooses *Unready*), so that upon observing the sender's action, the receiver knows the sender's type.

Pooling equilibrium All types of the sender choose the same action (in the second sort of equilibrium in the example, both types of challenger choose *Unready*), so that the sender's action gives the receiver no clue to the sender's type.

If the sender has more than two types, mixtures of these types of equilibrium may exist—the set of types may be divided into groups, within each of which all types choose the same action and between which the actions are different.

? EXERCISE 335.1 (Pooling and separating equilibria in a signaling game) Consider a game that differs from the one in Figure 333.1 only in the payoffs to the terminal histories $(Weak, Ready, A)$ and $(Weak, Ready, F)$. For what values of these payoffs, if any, does the game have a weak sequential ("separating") equilibrium in which a strong challenger chooses *Ready* and a weak one chooses *Unready*? For what values of these payoffs, if any, does the game have a weak sequential ("pooling") equilibrium in which both types of challenger choose *Unready*?

? EXERCISE 335.2 (Sir Philip Sydney game) Some young animals expend energy begging for food from their parents—they squawk and bleat and scream, sometimes extravagantly. Can we expect these demands to signal their needs accurately? Consider the following signaling game. (The rationale for its name, due to its originator, the distinguished evolutionary biologist John Maynard Smith, is given in the Notes at the end of the chapter.)

A hungry parent has a piece of food that it may give to its offspring, or keep for itself. It does not detect whether its offspring is hungry. In either case, the offspring may signal that it is hungry to its parent (by squawking, for example). An animal is stronger and thus produces more offspring (i.e. has a higher "biological fitness") if it gets the food than if it does not. Normalize the parent's strength if it keeps the food to be 1, and denote its strength if it gives the food to its offspring by $S < 1$. If the offspring does not squawk, its strength is 1 if it gets the food, $V < 1$ if it is not hungry and does not get the food, and 0 if it is hungry and does not get the food. If the offspring squawks, its strength is multiplied by the factor $1 - t$, where $0 \leq t \leq 1$ (i.e. squawking may be costly). Denote the degree to which the parent and offspring are related by r, and take each player's payoff to be its strength plus r times the other player's strength. Evolutionary pressure will lead to behavior for each player that maximizes that player's payoff, given the other player's behavior. (For more on evolutionary games, see Chapter 13.) The game is shown in Figure 336.1.

Find the conditions on r, in terms of S, V, and t, under which the game has a separating equilibrium in which the offspring squawks if and only if it is hungry and the parent gives it the food if and only if it squawks. Show that if the offspring's payoff from obtaining the food when it is quiet exceeds its payoff from not obtaining it, whether or not it is hungry (which means that $r < (1 - V)/(1 - S)$), then the game has such an equilibrium only if $t > 0$. That is, in this case an equilibrium exists in which the signal is accurate only if the signal is costly. Show that if $r < (1 - S)/(1 - (1 - p)V)$, then the game has a pooling equilibrium in which the offspring is always quiet and the parent always keeps the food. (For some other parameter values, the game has a pooling equilibrium in which the offspring is always quiet and the parent always gives the food.)

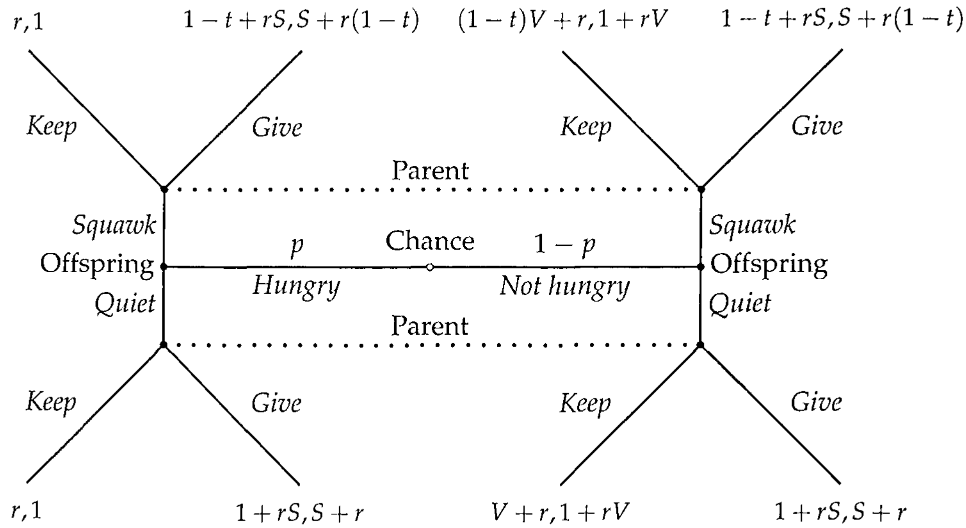

Figure 336.1 The Sir Philip Sydney game (Exercise 335.2). The empty history is in the center of the diagram. The offspring's payoff is listed first, the parent's second.

10.6 Illustration: conspicuous expenditure as a signal of quality

Why do some firms sponsor sporting events? Here is an explanation. The quality of some goods cannot be detected before the goods are consumed, and some firms sell goods of low quality while others sell goods of high quality. A consumer will purchase from a firm repeatedly once she detects goods of high quality, but will not buy low-quality goods more than once. Thus a purveyor of high-quality goods gains more from a consumer's sampling its wares than does a dealer in low-quality merchandise. Consumers may therefore deduce that it is worthwhile only for sellers of high-quality goods to sponsor sporting events—the potential benefit does not exceed the cost for low-quality firms—so that conspicuous expenditure is a sure sign of a high-quality firm, and lack thereof the mark of a low-quality firm. A high-quality firm is thus nudged into sports sponsorship.

Does this argument stand up to careful examination? Suppose that with probability π a firm is one that produces goods of quality $H > 0$ and with probability $1 - \pi$ one that produces goods of quality $L = 0$. (Note that the firm does not choose its quality.) The cost of producing a good of quality H is $c_H > 0$ and that of producing a good of quality 0 is $c_L = 0$.

A consumer interacts with the firm over two periods. In the first period, the firm chooses a price p_1 for its output and an amount $E \geq 0$ to spend conspicuously, for example, sponsoring sporting events. The expenditure E has no effect on the quality of the good. The consumer observes E and p_1, but not the quality of the good, and decides whether to purchase the good. If she makes no purchase, the interaction ends; her payoff is 0 and the firm's is $-E$. If she makes a purchase, she learns the quality of the good, the firm chooses a price p_2 for the second period, and she decides whether to purchase the good again. In any period in which she purchases the good her payoff is $v - p$, where v is the quality of the good and p is the price, and her total payoff is the sum of her payoffs in the two periods. (Thus the consumer is willing to purchase the low-quality good in either period only if its price is zero.) The firm's payoff is $p_1 - E - c_I$, where I is either L or H, if it sells to the consumer only in the first period, and $p_1 + p_2 - E - 2c_I$, if it sells to the consumer in both periods. (Note that the expenditure E occurs only in the first period.)

Assume that $H - c_H > 0$, so that at the highest price the consumer is willing to pay for the high-quality good (namely H), this good yields the firm a positive profit.

The structure of an extensive game that models this situation is shown in Figure 337.1. In this figure only *one* of the possible pairs (p_1, E) of first-period actions for the firm is indicated, and the range of second-period prices p_2 is indicated by the shaded triangles.

Consider the equilibria of this game. In the second period the players are fully informed. Thus the consumer buys the good in this period only if its price is at most equal to its quality. A price equal to H yields the high-quality firm a profit, so in an equilibrium a high-quality firm charges $p_2^H = H$ and the consumer buys the good; the firm's payoff is $H - c_H$ and the consumer's payoff is 0. The profit of a low-quality firm is 0 independent of its price: if the price is positive the consumer does not purchase the good, and if it is zero the firm's profit is zero regardless of whether the consumer purchases the good. Thus any second-period price $p_2^L \geq 0$ for a low-quality firm is compatible with equilibrium.

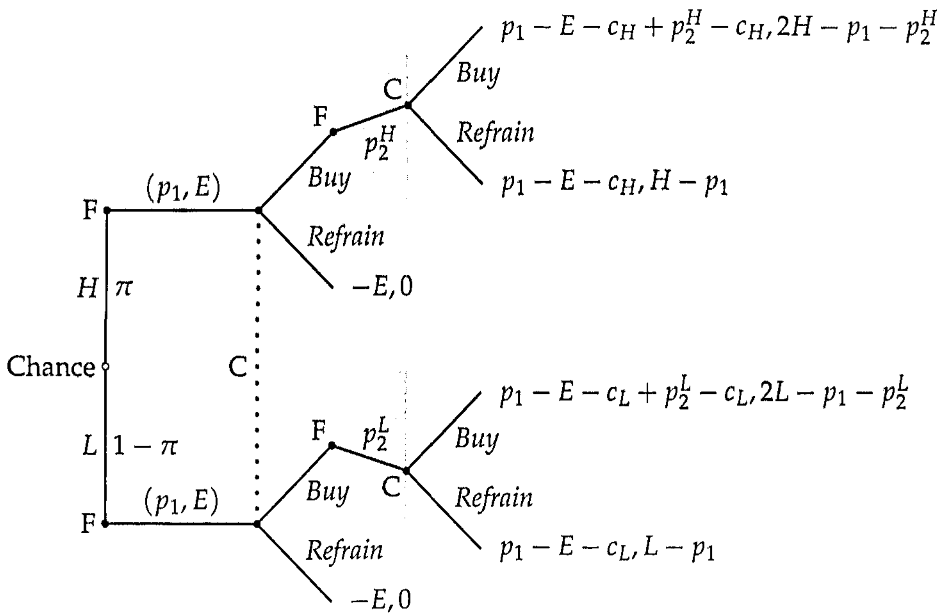

Figure 337.1 An outline of an extensive game that models a firm that signals its quality by sponsoring sporting events. The firm is denoted "F", and the consumer "C". Only *one* of the firm's actions (p_1, E) is shown. The firm's payoffs are listed first, the consumer's second.

I claim that under some conditions the whole game has a (separating) weak sequential equilibrium in which a producer of high-quality goods spends money in the first period, while a producer of low-quality ones does not. Consider the following assessment.

Firm's strategy Choose (p^{H*}, E^*) in the first period and H in the second period after any history in which chance chooses H, and choose $(0,0)$ in the first period and any nonnegative price in the second period after any history in which chance chooses L.

Consumer's belief The consumer believes that the history was H if and only if the (price, expenditure) pair (p, E) observed in the first period satisfies $p \leq p^{H*}$ and $E = E^*$.

Consumer's strategy Buy the good in the first period if and only if the (price, expenditure) pair (p, E) satisfies $p \leq p^{H*}$ and $E = E^*$, and buy it in the second period if and only if either the history starts with H and the second-period price is at most H, or the history starts with L and the second-period price is 0.

Note that the consumer's belief in this assessment is extreme: unless she observes an expenditure of exactly E^* in the first period, she concludes that the firm is low quality. If the firm follows its strategy, the consumer observes a first-period action of either (p^{H*}, E^*) or $(0,0)$. The consistency condition implies that if she observes (p^{H*}, E^*), then she must believe that the firm is high quality, whereas if she observes $(0,0)$, then she must believe that it is low quality but does not restrict her belief after any other observation. The extreme belief I have specified makes a profitable deviation by a high-quality firm as difficult as possible, because any deviation causes the consumer to believe that the firm is low quality.

Under what conditions is this assessment a weak sequential equilibrium? The players' beliefs are consistent by construction, so we need to check only that each player's strategy is sequentially rational.

Firm The firm's strategy is sequentially rational if and only if neither type of firm can increase its expected payoff, given the consumer's strategy and belief.

Type H If the firm chooses (p^{H*}, E^*) its profit is $p^{H*} + H - E^* - 2c_H$, if it chooses (p, E^*) with $p < p^{H*}$ its profit is $p + H - E^* - 2c_H$, and if it chooses any other (price, expenditure) pair its profit is 0 (the consumer believes it is low quality and does not purchase the good). Thus for equilibrium we need

$$p^{H*} + H - E^* - 2c_H \geq 0.$$

Type L If the firm chooses any (p, E) with $E \neq E^*$ or $E = E^*$ and $p > p^{H*}$, then the consumer believes it is low quality, so that its profit is 0. If it chooses any (p, E^*) with $p \leq p^{H*}$, then the consumer believes it is high quality and purchases the good in the first period, yielding it a profit of $p - E^*$. Thus for equilibrium we need

$$0 \geq p^{H*} - E^*.$$

Consumer If she observes (p, E^*) with $p \leq p^{H*}$, she believes the firm is high quality, so that her expected payoff is $H - p^{H*}$ if she buys the good, and 0 if she does not. Thus for an equilibrium in which she buys the good if and only if the (price, expenditure) pair (p, E^*) satisfies $p \leq p^{H*}$, we need

$$H - p^{H*} \geq 0.$$

If she observes any other (price, expenditure) pair (p, E), she believes the firm is low quality, so that her expected payoff is $0 - p$ if she buys the good and 0 if she does not, so that she optimally does not buy the good.

In summary, the assessment is a weak sequential equilibrium if and only if

$$E^* + 2c_H - H \leq p^{H*} \leq \min\{E^*, H\}. \tag{339.1}$$

If $H \geq 2c_H$, there exist pairs (p^{H*}, E^*) with $E^* > 0$ that satisfy this condition (see Figure 339.1), and thus weak sequential equilibria in which a high-quality firm spends conspicuously (on sporting events, for example). The firm does so because if it does not, then the consumer believes it to be low quality, and hence does not buy its good in the first period, does not find out that it is high quality, and thus does not buy its good in the second period. Eliminating the expenditure saves the firm E^*, but it costs it $H - 2c_H$ in lost profit. Imitating the high-quality firm is not worthwhile for the low-quality one because the cost E^* of doing so exceeds the resulting increase in its profit of p^{H*} (the first-period revenue from selling at the price p^{H*}).

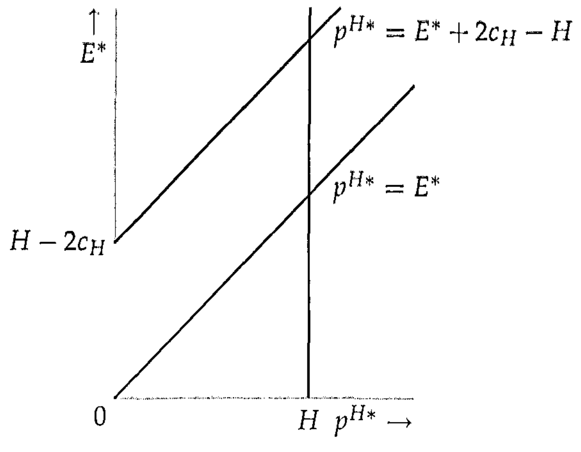

Figure 339.1 The shaded area is the set of pairs (p^{H*}, E^*) satisfying (339.1). For any pair in the set, the game has a weak sequential equilibrium in which a firm of type H chooses (p^{H*}, E^*).

We see that two features of the interaction are important in generating the result. First, a high-quality firm obtains a benefit from attracting a consumer beyond the period in which it makes the expenditure—the consumer purchases the good in the second period. If there were no repeat purchases, then a low-quality firm, by imitating a high-quality one, could induce the consumer to patronize it and thus obtain the same revenue as does a high-quality firm. Given that the low-quality firm's unit cost is lower than that of a high-quality firm, imitation would pay. Second, the consumer does not purchase the good in the first period if it observes the (price, expenditure) pair chosen in equilibrium by a low-quality firm. If it did purchase the good in this case, a high-quality firm could dispense with its conspicuous expenditure and still obtain the benefit of repeat business because the consumer, having purchased its good expecting it to be of low quality, would discover that it is in fact of high quality. If the consumer does not purchase from the low-quality firm, why does this firm exist? The important point is that a high-quality firm loses *some* business when it imitates a low-quality one. If we modify the model by assuming a variety of consumers, some of whom optimally patronize the low-quality firm at the price it charges and some of whom do not, the qualitative feature of the separating equilibria is retained.

Though the game has a separating equilibrium if (339.1) is satisfied, it also has pooling equilibria, which you are invited to study in the following exercise.

? EXERCISE 340.1 (Pooling equilibria of game in which expenditure signals quality) Find weak sequential equilibria of the game in which each type of firm chooses the same (price, expenditure) pair in the first period.

10.7 Illustration: education as a signal of ability

Why are you obtaining a college degree? Because you think that the principles you learn in your courses will prepare you for the day when you run IBM or preside over Italy? Possibly—but there may be another reason. Perhaps nothing you learn in college has any bearing on the job you expect to take, but you need to get a degree to prove to potential employers that your ability is high. How does your obtaining a degree prove this point? Because the cost to persons of low ability of obtaining the degree is much higher than it is for you (they will take longer, and find the process painful), so that such persons cannot profitably imitate you. Thus a college degree signals high ability, even if colleges do nothing to foster that ability: employers know that a recipient of a college degree must have high ability because only for such a person is it worthwhile to obtain a degree. If the cost of achieving high proficiency in freestyle snowboarding were much lower for a person with the skills valued by IBM or the Italian citizenry than for someone without the skills, a certificate attesting to that achievement could be your ticket to a rewarding job. But it is not, so you are in college.

Here is a simple model we can use to study this logic. A worker's ability, which is either H or $L < H$, is known to her but not to either of two potential employers.

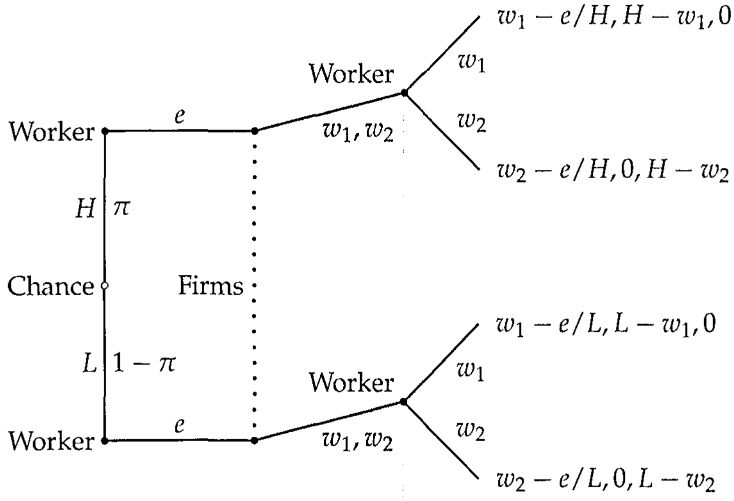

Figure 341.1 An outline of an extensive game that models a worker that signals her ability to a pair of firms by acquiring education. The diagram shows only *one* of the possible education levels e the worker may obtain. The firms' wage offers are made simultaneously. The worker's payoffs are listed first, the firms' second and third.

The worker chooses the amount e of education to obtain, then the firms, observing e, simultaneously offer wages w_1 and w_2, and finally the worker chooses one of the wage offers. Assume that education is less expensive to obtain for a worker of high ability than it is for a worker of low ability. Specifically, assume that the cost to a worker of ability K of obtaining e units of education is e/K, so that the payoff of such a worker who obtains e units of education and takes a job paying w is $w - e/K$. The payoff of a firm that pays a worker of ability K the wage w is $K - w$. The extensive game that corresponds to this model is given in Figure 341.1.

I claim that the game has a weak sequential equilibrium in which a high-ability worker chooses a positive amount of education. Consider the following assessment, in which e^* is a positive number.

Worker's strategy Type H chooses $e = e^*$ and type L chooses $e = 0$; after observing the firms' wage offers, both types choose the highest offer if they differ, and that of firm 1 if they are the same.

Firms' belief Each firm believes that a worker is type H if she chooses e^* and type L otherwise.

Firms' strategies Each firm offers the wage H to a worker who chooses e^* and the wage L to a worker who chooses any other value of e.

The firms' beliefs are consistent with the worker's strategy. (No worker chooses an education level different from e^* and 0, so the consistency condition imposes no restriction on the firms' beliefs after observing such a level.)

I now find conditions on the parameters under which the players' strategies are sequentially rational.

Worker The worker's strategy of accepting the highest wage offer at the end of the game is clearly optimal. Now consider the worker's initial action.

Type H A type H worker obtains the payoff $H - e^*/H$ if she follows her strategy and chooses the education level e^*, and the payoff $L - e/H$ if she chooses any other education level e. The education level 0 achieves the highest payoff, of L, for a deviant, so for equilibrium we need $H - e^*/H \geq L$, or

$$e^* \leq H(H - L).$$

Type L A type L worker obtains the payoff L if she follows her strategy and chooses the education level 0. If she chooses any education level other than e^* she obtains the same wage, and pays a cost, so such a deviation is not profitable. If she chooses the education level e^*, then the firms believe her ability to be H, and she obtains the payoff $H - e^*/L$. Thus for equilibrium we need

$$e^* \geq L(H - L).$$

Firms Each firm's payoff is 0, given its belief and its strategy. If it raises the wage it offers in response to any value of e, its expected profit is negative, given its belief, and if it lowers the wage its expected profit remains zero (its offer is not accepted).

In summary, the assessment is a weak sequential equilibrium if and only if

$$L(H - L) \leq e^* \leq H(H - L).$$

I have assumed that $H > L$, so the left-hand side of this expression is less than the right-hand side, and hence values of e^* satisfying the expression exist. That is, for any values of H and L, the game has separating equilibria in which high-ability workers obtain some education, whereas low-ability ones do not. Education has no effect on the workers' productivity; a high-ability worker obtains it to avoid being labeled low in ability by potential employers.

Like several of the signaling games we have studied, this game has also pooling equilibria, in which both types of worker obtain the same amount of education. That is, the model is consistent both with a steady state in which only high-ability workers obtain education and with one in which all workers do so.

? EXERCISE 342.1 (Pooling equilibria of game in which education signals ability) Find the range of education levels e for which the game has a weak sequential equilibrium in which both types of worker choose the education level e. Compare these levels with those possible in a separating equilibrium.

10.8 Illustration: strategic information transmission

You research the market for new products and submit a report to your boss, who decides which product to develop. Your preferences differ from those of your boss: you are interested in promoting the interests of your division, whereas she is interested in promoting the interests of the whole firm. What do you tell her?

If you report the results of your research without distortion, the product she will choose will not be best for you. Can you do better by distorting or obscuring the fruits of your research? If you *systematically* distort your finding, then your boss will be able to unravel your report and deduce your actual findings, so obfuscation seems a more promising route.

To study the issues precisely, consider the following model. A sender (you) observes the state t, a number from 0 to 1 (the result of your research), that a receiver (your boss) cannot see. Assume that the distribution of the state is uniform: for any number z from 0 to 1, the probability that t is at most z is z. The sender submits a report r, a number, to the receiver, who observes the report and takes an action y, also a number. Both the sender and receiver care about the relation between the state t and the receiver's action y, and neither is affected directly by the value of the sender's report r. Specifically, assume that

$$\begin{aligned} \text{Sender's payoff function:} \quad & -(y-(t+b))^2 \\ \text{Receiver's payoff function:} \quad & -(y-t)^2, \end{aligned}$$

where b (the sender's "bias") is a fixed positive number that reflects the divergence between the sender's and receiver's preferences. These functions are shown in Figure 343.1. Note that a receiver who believes that the state is t optimally chooses $y = t$, whereas the best action for the sender in this case is $y = t + b$. The game is illustrated in Figure 344.1.

10.8.1 Perfect information transmission?

Consider the possibility of an equilibrium in which the sender accurately reports the state she observes—that is, her strategy is $r(t) = t$ for all t. Given this strategy, the consistency condition requires that the receiver believe (correctly) that the state is t when the sender reports t, and hence for any report t optimally chooses the

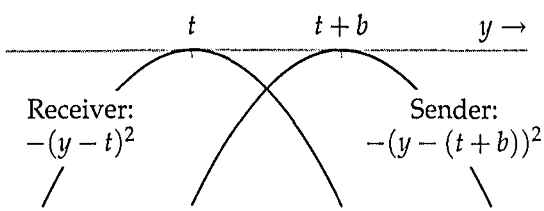

Figure 343.1 The players' payoff functions when the state is t in the game of strategic information transmission. The number b is a positive constant.

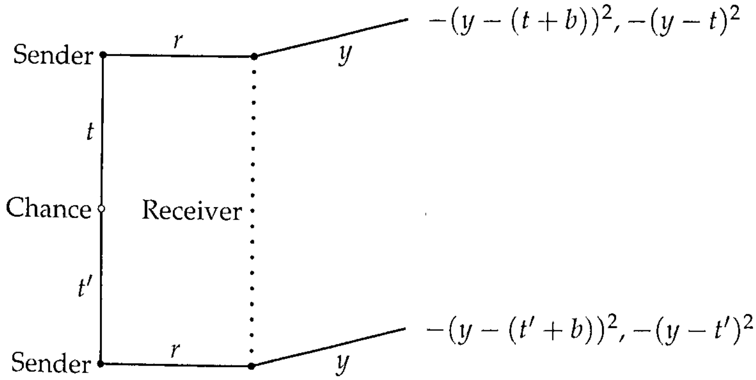

Figure 344.1 An outline of an extensive game that models strategic information transmission. Only *two* of the possible states (t and t') and *one* possible report of the sender are shown. The sender's payoff is listed first, the receiver's second.

action t (the maximizer of $-(y-t)^2$). Is the sender's strategy a best response to this strategy of the receiver? Not if $b > 0$! Suppose the state is t. If the sender reports t, the receiver chooses $y = t$, so that the sender's payoff is $-(t-(t+b))^2 = -b^2$. If instead she reports $t+b$, the receiver believes the state is $t+b$ and chooses $y = t+b$, so that the sender's payoff is $-(t+b-(t+b))^2 = 0$. If $b > 0$ the sender is thus better off reporting $t+b$ when the state is t. So unless the sender's and receiver's preferences are identical, the game has no equilibrium in which the sender accurately reports the state.

A similar argument shows that for $b > 0$ the game has no equilibrium in which the sender's strategy is any increasing function or any decreasing function. If the sender's strategy is such a function r and $r(t) = t'$, the consistency condition requires that upon observing the report t', the receiver believe that the state is t, so that the receiver optimally takes the action t. As before, a sender who observes t is thus better off reporting $t+b$.

10.8.2 *No information transmission?*

Now consider the possibility of an equilibrium in which the sender's report is constant, independent of the value of the variable she observes—say $r(t) = c$ for all t. Such a report conveys no information; the consistency condition requires that if the receiver observes the report c, her belief must remain the same as it was initially, namely that the state is uniformly distributed from 0 to 1. Given this belief, her optimal action is $y = \frac{1}{2}$. (This claim should be plausible, given the symmetry of the belief and the shape of the receiver's payoff function. If you know calculus, you may verify the claim precisely.)

The consistency condition does not constrain the receiver's belief about the state upon her receiving a report different from c, because such a report does not occur if the sender follows her strategy. Assume that the receiver completely ignores the sender's report: her belief remains the same as it was originally for *every* value of the report. Then the receiver's optimal action is $y = \frac{1}{2}$ whatever report she receives. That is, the sender's report has no effect on the action chosen by the receiver, so that the strategy r for which $r(t) = c$ for all t is optimal for the sender.

In summary, for every value of b the game has a weak sequential equilibrium in which the sender's report conveys no information (it is constant, independent of her type) and the receiver ignores the report (she maintains her initial belief about the state, regardless of the report) and takes the action $y = \frac{1}{2}$. If b is small, then this equilibrium is not very attractive. If $b < \frac{1}{2}$, then for some states there is an action different from $\frac{1}{2}$ that is better for both the sender and the receiver, and the smaller is b, the larger the set of states for which such increases in payoffs are possible. (If $b = \frac{1}{4}$, for example, then for any t with $0 \leq t < \frac{1}{4}$, both the sender and the receiver are better off if the receiver's action is $t + b$ than if it is $\frac{1}{2}$.)

10.8.3 Some information transmission?

Does the game have equilibria in which *some* information is transmitted? Suppose that the sender makes one of two reports, depending on the state. Specifically, suppose that if $0 \leq t < t_1$ she reports r_1 and if $t_1 \leq t \leq 1$ she reports $r_2 \neq r_1$.

Consider the receiver's optimal response to this strategy. If she sees the report r_1, she knows that the state t satisfies $0 \leq t < t_1$; given her initial belief that the state is uniformly distributed from 0 to 1, the consistency condition requires that she now believe that the state is uniformly distributed from 0 to t_1. Her optimal action given this belief is $y = \frac{1}{2}t_1$. (As before, this claim should be plausible given the symmetry and shape of the receiver's payoff function; it may be verified using calculus.) Similarly, if the receiver sees the report r_2 she believes that the state is uniformly distributed from t_1 to 1, and chooses the action $y = \frac{1}{2}(t_1 + 1)$. The consistency condition does not restrict the receiver's belief if she sees a report other than r_1 or r_2. Assume that for each such report the receiver's belief is one of the two beliefs she holds if she sees r_1 or r_2, so that her optimal response to every report is either $\frac{1}{2}t_1$ or $\frac{1}{2}(t_1 + 1)$.

Now, for equilibrium we need the sender's report r_1 to be optimal if $0 \leq t < t_1$ and her report r_2 to be optimal if $t_1 \leq t \leq 1$, given the receiver's strategy. By changing her report, the sender can change the receiver's action from $\frac{1}{2}t_1$ to $\frac{1}{2}(t_1 + 1)$, or vice versa. Thus for the report r_1 to be optimal for the sender in every state t with $0 \leq t \leq t_1$, she must like the action $\frac{1}{2}t_1$ at least as much as the action $\frac{1}{2}(t_1 + 1)$, and for the report r_2 to be optimal in every state t for which $t_1 \leq t \leq 1$ she must like the action $\frac{1}{2}(t_1 + 1)$ at least as much as the action $\frac{1}{2}t_1$. In particular, in state t_1 the sender must be exactly indifferent between the two actions, as in Figure 346.1. This indifference implies that $t_1 + b$ is midway $\frac{1}{2}t_1$ and $\frac{1}{2}(t_1 + 1)$, or $t_1 + b = \frac{1}{2}[\frac{1}{2}t_1 + \frac{1}{2}(t_1 + 1)]$, so that

$$t_1 = \tfrac{1}{2} - 2b. \tag{345.1}$$

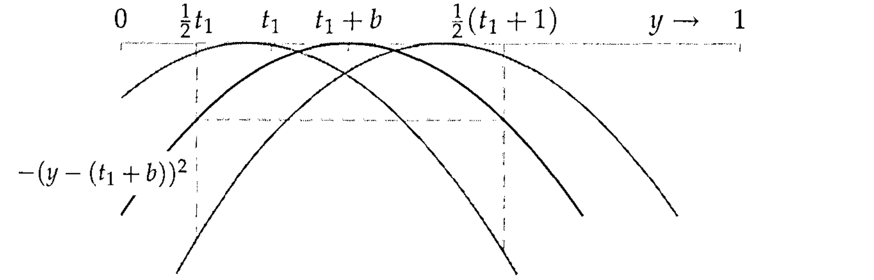

Figure 346.1 A value of t_1 such that in state t_1 the sender is indifferent between the receiver's actions $\frac{1}{2}t_1$ and $\frac{1}{2}(t_1 + 1)$. The black curve is the sender's payoff function in state t_1. The gray curve on the left is the sender's payoff function in a state less than t_1; she prefers $\frac{1}{2}t_1$ to $\frac{1}{2}(t_1 + 1)$. The gray curve on the right is the sender's payoff function in a state greater than t_1; she prefers $\frac{1}{2}(t_1 + 1)$ to $\frac{1}{2}t_1$.

We need $t_1 > 0$, so that this condition can be satisfied only if $b < \frac{1}{4}$. That is, if $b \geq \frac{1}{4}$, then the game has no equilibrium in which the sender makes two different reports, depending on the state. Put differently, there is no point in the receiver's asking the sender to submit a report if her preferences diverge sufficiently from those of the sender—she should simply take the best action for herself given her prior belief.

I claim that (345.1) is not only necessary for an equilibrium, but also sufficient. That is, if t_1 satisfies (345.1), then in every state t with $0 \leq t < t_1$ the sender optimally reports r_1 and in every state t with $t_1 \leq t \leq 1$ she optimally reports $r_2 \neq r_1$. This optimality follows from the shapes of the payoff functions; it is illustrated in Figure 346.1, in which the gray curves are the payoff functions of senders in states less than t_1 (left) and greater than t_1 (right).

In conclusion, if $b < \frac{1}{4}$, then the game has a weak sequential equilibrium in which the sender transmits two different reports, depending on the state: for $0 \leq t < t_1$ she submits one report and the receiver takes the action $\frac{1}{2}t_1$, and for $t_1 \leq t \leq 1$ she submits a different report and the receiver takes the action $\frac{1}{2}(t_1 + 1)$.

This equilibrium is better for both the receiver and the sender (before she observes the state) than the one in which no information is transmitted. Consider the receiver. (An analogous argument may be made for the sender.) In the equilibrium in which no information is transmitted, she takes the action $\frac{1}{2}$ in all states, so that her payoff in each state t is $-(\frac{1}{2} - t)^2$. In the equilibrium we have just found, in which the sender transmits two different reports depending on the state, her payoff is $-(\frac{1}{2}t_1 - t)^2$ for $0 \leq t < t_1$ and $-(\frac{1}{2}(t_1 + 1) - t)^2$ for $t_1 \leq t < 1$.

? EXERCISE 346.1 (Comparing the receiver's expected payoff in two equilibria) Plot the receiver's payoff as a function of t (from 0 to 1) in each equilibrium. Given that the distribution of the state is uniform, the receiver's expected payoff in an equilibrium is the negative of the area between the horizontal axis and the function that gives the payoff in that state. Show that the receiver's payoff is greater in the two-report equilibrium than it is in the equilibrium with no information transmission.

10.8.4 How much information transmission?

For $b < \frac{1}{4}$, does the game have equilibria in which more information is transmitted than in the two-report equilibrium? Consider the possibility of an equilibrium in which the sender makes one of K reports, depending on the state. Specifically, suppose that the sender's report is r_1 if $0 \leq t < t_1$, r_2 if $t_1 \leq t < t_2$, ..., r_K if $t_{K-1} \leq t \leq 1$, where $r_i \neq r_j$ whenever $i \neq j$. For convenience, let $t_0 = 0$ and $t_K = 1$.

The analysis follows the lines of that for the two-report equilibrium. If the receiver observes the report r_k, then the consistency condition requires that she believe the state to be uniformly distributed from t_{k-1} to t_k, so that she optimally takes the action $\frac{1}{2}(t_{k-1} + t_k)$. If she observes a report different from any r_k, the consistency condition does not restrict her belief; assume that her belief in such a case is the belief she holds upon receiving one of the reports r_k.

Now, for equilibrium we need the sender's report r_k to be optimal when the state is t with $t_{k-1} \leq t < t_k$, for $k = 1, \ldots, K$. As before, a sufficient condition for optimality is that in each state t_k, $k = 1, \ldots, K$, the sender be indifferent between the reports r_k and r_{k+1}, and thus between the receiver's actions $\frac{1}{2}(t_{k-1} + t_k)$ and $\frac{1}{2}(t_k + t_{k+1})$. This indifference implies that $t_k + b$ is equal to the average of $\frac{1}{2}(t_{k-1} + t_k)$ and $\frac{1}{2}(t_k + t_{k+1})$:

$$t_k + b = \tfrac{1}{2}\left[\tfrac{1}{2}(t_{k-1} + t_k) + \tfrac{1}{2}(t_k + t_{k+1})\right],$$

which is equivalent to

$$t_{k+1} - t_k = t_k - t_{k-1} + 4b.$$

That is, the interval of states for which the sender's report is r_{k+1} is longer by $4b$ than the interval for which the report is r_k. The length of the first interval, from 0 to t_1, is t_1, and the sum of the lengths of all the intervals must be 1, so we need

$$t_1 + (t_1 + 4b) + \cdots + (t_1 + (K-1)b) = 1$$

or

$$Kt_1 + 4b(1 + 2 + \cdots + (K-1)) = 1.$$

Using the fact that the sum of the first n positive integers is $\frac{1}{2}n(n+1)$, the equation is

$$Kt_1 + 2bK(K-1) = 1. \tag{347.1}$$

If b is small enough that $2bK(K-1) < 1$, there is a positive value of t_1 that satisfies the equation.

Suppose, for example, that $\frac{1}{24} \leq b < \frac{1}{12}$. Then the inequality is satisfied for $K \leq 3$, so that in the equilibrium in which the most information is transmitted the sender chooses one of three reports. From (347.1), we have $t_1 = \frac{1}{3} - 4b$ and hence $t_2 = \frac{2}{3} - 4b$ in this equilibrium. For $b = \frac{1}{24}$, the equilibrium values of t_1 ($\frac{1}{6}$) and t_2 ($\frac{1}{2}$) are illustrated in Figure 348.1, and the equilibrium action y taken by the receiver, as a function of the state t, is shown in Figure 348.2. (I have chosen this

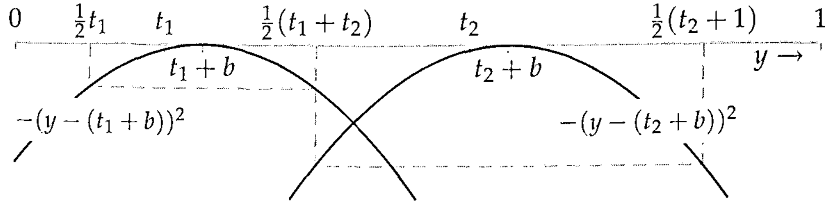

Figure 348.1 The values of t_1 and t_2 in a three-report equilibrium in the model of strategic information transmission for $b = \frac{1}{24}$, the most informative equilibrium for this value of b. In state t_1 the sender is indifferent between the receiver's actions $\frac{1}{2}t_1$ and $\frac{1}{2}(t_1+t_2)$, and in state t_2 she is indifferent between the receiver's actions $\frac{1}{2}(t_1+t_2)$ and $\frac{1}{2}(t_1+1)$.

value of b because it allows the diagram to clearly show the equilibrium. Note that if b were any smaller, a four-report equilibrium would exist.) The values of the reports r_k do not matter, as long as no two are the same. We may think of them as words in a language; the receiver's long experience playing the game teaches her that r_k means "the state is between t_{k-1} and t_k".

In summary, if there is a positive value of t_1 that satisfies (347.1), then the game has a weak sequential equilibrium in which the sender submits one of K different reports, depending on the state. For any given value of b, the largest value of K for which an equilibrium exists is the largest value for which $2bK(K-1) < 1$. If $2bK(K-1) = 1$, then, using the quadratic formula, we have $K = \frac{1}{2}(1+\sqrt{1+2/b})$. Thus, in particular, the larger the value of b, the smaller the largest value of K possible in an equilibrium. That is, the greater the difference between the sender's and receiver's preferences, the coarser the information transmitted in the equilibrium

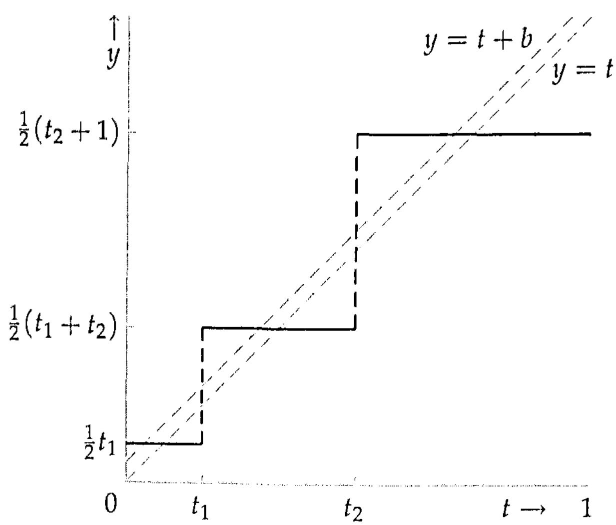

Figure 348.2 The receiver's action y, as a function of the state t, in a three-report equilibrium of the game of strategic information transmission for $b = \frac{1}{24}$, the most informative equilibrium for this value of b. The dashed gray lines show the action optimal in each state for the sender (top) and receiver (bottom).

with the largest number of steps—the "most informative" equilibrium. I claim that the equilibrium that yields the receiver the highest payoff is the most informative one, and this equilibrium is also best for the sender before she observes the state. You should find this claim plausible, though some calculations are needed to verify it.

10.8.5 Summary on strategic information transmission

When a receiver takes an action based on an unverifiable report about the state supplied by a sender, and the sender's and receiver's preferences differ, no equilibrium exists in which the sender accurately reports the state. Regardless of the difference between the sender's and receiver's preferences, an equilibrium exists in which the sender's report is independent of the state, so that no information is transmitted. If the sender's and receiver's preferences are sufficiently similar, there are also equilibria in which the sender makes one of several reports, depending on the state; the maximal number of reports possible in an equilibrium is larger the more similar the parties' preferences. The equilibrium with the maximal number of reports is better for both parties (before the state is known) than all other equilibria.

10.8.6 Delegation

Is there a better way for the receiver to make a decision? Consider the alternative of delegation: the receiver lets the sender choose the action. No reports, no obfuscation; the sender simply chooses the best action for herself—in each state t she chooses the action $t + b$.

Compare this outcome with the outcome of the best equilibrium for the receiver if she acts on the basis of a report by the sender. For $b = \frac{1}{24}$, the outcome, as a function of the state, for the three-report equilibrium (the best equilibrium for this value of b) is given in Figure 348.2, where the dashed line $y = t + b$ indicates the outcome of delegation.

Which outcome does the receiver prefer? Under delegation, the distance between the outcome and the receiver's favorite action is exactly b in every state. Thus the receiver's payoff in every state is $-b^2$. In the three-report equilibrium of the game in which the receiver solicits information from the sender, the outcome for most states is further than b from the receiver's favorite action; only for states within b of $\frac{1}{2}t_1$, $\frac{1}{2}(t_1 + t_2)$, or $\frac{1}{2}(t_2 + 1)$ is it closer to the receiver's favorite action that is the outcome of delegation. If that does not convince you that the receiver prefers the outcome of delegation to the outcome of the three-report equilibrium, take a look at Figure 350.1. This figure plots the receiver's payoffs in the two cases, for each state. The horizontal line labeled $-b^2$, very close to the axis, is the receiver's payoff in each state under delegation. The scalloped curve is her payoff in each state in the three-report equilibrium. The state is distributed uniformly from 0 to 1, so the receiver's expected payoff under delegation is the negative of the area

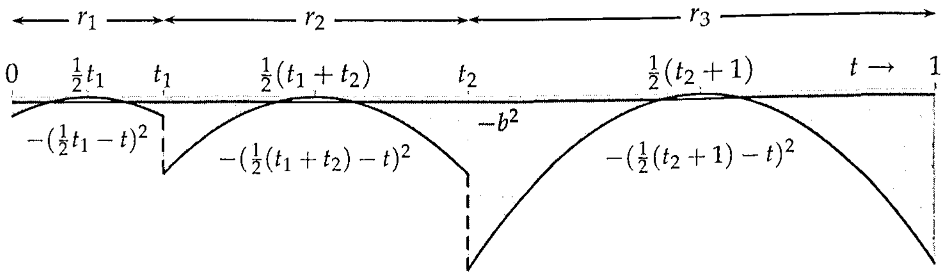

Figure 350.1 The receiver's payoff, as a function of the state t, in a three-report equilibrium of the game of information transmission and under delegation for $b = \frac{1}{24}$. The negative of the shaded area is the receiver's expected payoff in the three-report equilibrium, and the negative of the area between the horizontal line $-b^2$ and the horizontal axis is her expected payoff under delegation.

of the very slim rectangle between the axis and the line $-b^2$, while her expected payoff in the three-report equilibrium is the negative of the shaded area above the scalloped line. Which is larger? You do not need to make an exact calculation to be sure that the receiver prefers the outcome under delegation!

I have analyzed only a single value of b, namely $\frac{1}{24}$. For smaller values of b, the loss from both delegation and the most informative equilibrium of the game of information transmission is smaller, but delegation remains superior. For larger values of b, the loss from both outcomes is larger. Delegation remains better than the most informative equilibrium in the game of information transmission as long as $b < \sqrt{3}/6$ (approximately 0.29). (Calculus is needed to check this cutoff.) As we have seen, for $b \geq \frac{1}{4}$ the most informative equilibrium in the game of information transmission is not informative at all: the sender sends the same report in every state. Thus if the sender's preferences are close enough to the receiver's that the receiver benefits from consulting the sender, she is better off simply delegating the decision rather than soliciting a report. The directive "do what you want" yields a better outcome for the receiver than does "tell me what I should do" because the quality of the information induced by the latter is so low.

⑦ EXERCISE 350.1 (Variant of model with piecewise linear payoff functions) Consider the variant of the game of strategic information transmission in which the sender's payoff function is $-|y - (t + b)|$ and the receiver's payoff function is $-|y - t|$ (where $|x|$ denotes the absolute value of x). That is, in each state t each payoff function has the shape shown in Figure 71.1, with a peak at $t + b$ for the receiver and a peak at t for the sender. How do the k-report weak sequential equilibria of this variant of the model differ from those of the original model?

⑦ EXERCISE 350.2 (Pooling equilibrium in a general model) Consider a generalization of the game of strategic information transmission (Figure 344.1) in which the distribution of states is arbitrary and each player's payoff is an arbitrary function of y and t (but, as before, is independent of r). Assume that there exists an action, say y^*, that maximizes the receiver's expected payoff given the distribution of states. Show that any such game has a weak sequential equilibrium in which the receiver chooses the same action, regardless of the signal.

10.9 Illustration: agenda control with imperfect information

The U.S. House of Representatives assigns to committees the task of formulating modifications of the law. The bills proposed by committees are considered by the legislature under one of several rules, ranging from the "closed rule" (or "gag rule"), under which the legislature may either accept or reject the proposed bill but may not amend it, to the "open rule", under which any amendment may be made. (The Rules Committee determines how each bill is handled.) In general, the preferences of a committee differ from those of the entire legislature. What is the rationale for restricting the actions the legislature may take?

If we treat the legislature and the committee as single players with well-defined preferences, and assume that the legislature is perfectly informed about the environment, an analysis of the open rule is straightforward: the legislature simply chooses the bill that is best according to its preferences, ignoring the committee's proposal. A model of the closed rule is studied in Section 6.1.3. In general, the equilibrium outcome in this case is not the legislature's favorite bill. Thus under perfect information the legislature has no reason to adopt the closed rule.

But of course if the legislature is perfectly informed, it has no reason to assign the drafting of laws to a committee to begin with. If it is not perfectly informed, then a committee has a role: it can discover the "state of the world" and propose legislation to fit. In this environment, a committee whose preferences differ from those of the legislature has an incentive to report distorted information, and the rules under which proposed legislation is considered may affect the degree of distortion. Are there circumstances under which the closed rule produces an outcome better for the legislature than the open rule?

Consider the following model, which follows closely the one in the previous section (10.8). The desirability of a bill depends on the state t, a number from 0 to 1, which the committee, but not the legislature, observes. Assume that the distribution of the state is uniform: for any number z from 0 to 1, the probability that $t \leq z$ is z. After observing the state, the committee recommends a bill r (a number) to the legislature. Under the open rule, the legislature may then choose any bill it wishes. Under the closed rule, it is restricted to either accept r or reject it, in which case the outcome is the status quo y_0. Assume that in state t the legislature's favorite bill is $y = t$, and the committee's favorite bill is $y = t + b$, where b is a fixed positive number. Specifically, assume that

$$\begin{aligned} &\text{Committee's payoff function:} && -(y - (t + b))^2 \\ &\text{Legislature's payoff function:} && -(y - t)^2, \end{aligned}$$

where y is the bill passed by the legislature.

10.9.1 *Open rule*

Under the open rule, the model is identical to the one in the previous section (10.8), with the committee as sender and the legislature as receiver. Thus we know that if $b > 0$, then in an equilibrium the committee's report obscures its information, the more so the larger b. We also know the character of the best (and most informative) equilibrium for the legislature. In this equilibrium for $b = \frac{1}{24}$, for example, the committee sends one of three different reports, as illustrated in Figures 348.1 and 348.2.

10.9.2 Closed rule

Under the closed rule, the legislature retains the power to veto, but not otherwise to amend, the legislation proposed by the committee, in which case the outcome is the fixed status quo y_0. From Section 10.8.6 we know that if the legislature simply delegates the choice of legislation to the committee, the outcome is better for both the legislature and the committee than the outcome of the best equilibrium under the open rule, assuming that the legislature's and committee's preferences do not diverge so much that the legislature would be better off making the decision itself. The closed rule is close to delegation, but the legislature retains the right to veto the proposed legislature in favor of the status quo. Does the legislature benefit from retaining this vestige of control?

The game for the closed rule is illustrated in Figure 353.1. The legislature optimally exercises its option to choose y_0 if it prefers this outcome to the committee's proposal. In particular, if the committee proposes its favorite bill $y = t + b$ for all values of t (as it does under delegation), then the legislature optimally exercises its option whenever $y_0 - b \leq t \leq y_0 + b$, or equivalently whenever $t - b \leq y_0 \leq t + b$. (Refer to Figure 353.2.) Given this response of the legislature, the committee's proposal of $t + b$ is not optimal for t with $y_0 < t < y_0 + b$: the proposal $t + b$ leads to the outcome y_0, and hence the payoff $-(y_0 - (t + b))^2 < -b^2$, but if the committee recommends the bill $t + 2b$, then the legislature accepts it, giving the committee the higher payoff $-(t + 2b - (t + b))^2 = -b^2$. Thus in no equilibrium does the committee propose its favorite bill in all states.

If the status quo y_0 satisfies $b < y_0 < 1 - 3b$, however, I claim that the game has an equilibrium in which the committee recommends its favorite bill in all states except those between $y_0 - b$ and $y_0 + b$, and between $y_0 + b$ and $y_0 + 3b$. The bill proposed by the committee (and accepted by the legislature) in this equilibrium is shown in Figure 354.1. It is defined precisely as follows:

$$\begin{cases} t + b & \text{if } t \leq y_0 - b \\ y_0 & \text{if } y_0 - b < t \leq y_0 \\ y_0 + 2b & \text{if } y_0 < t \leq y_0 + 2b \\ y_0 + 4b & \text{if } y_0 + 2b < t \leq y_0 + 3b \\ t + b & \text{if } y_0 + 3b < t. \end{cases}$$

For any recommendation of less than y_0 or greater than $y_0 + 4b$ the legislature can infer the state precisely. Thus the consistency condition requires that the legislature's belief about the state be correct if the committee makes such a recommendation. For the recommendations y_0, $y_0 + 2b$, and $y_0 + 4b$, the legislature cannot infer the state precisely. The committee makes each of these recommendations for

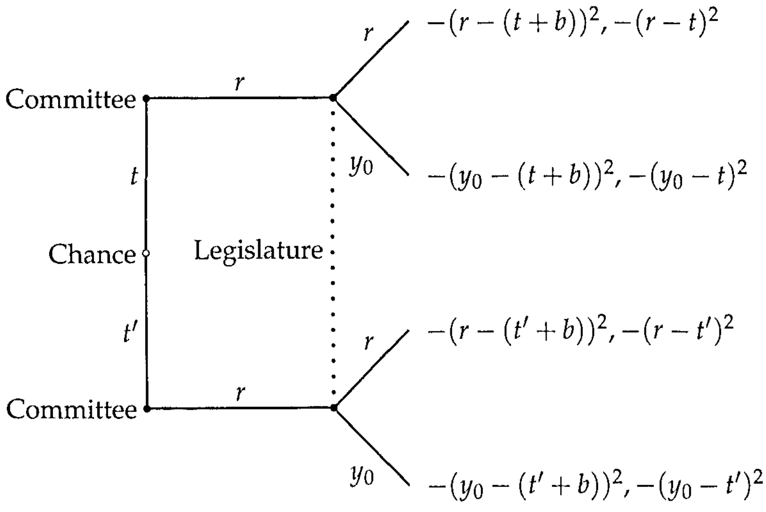

Figure 353.1 An outline of an extensive game that models the closed rule for the consideration of a committee proposal by a legislature. Only *two* of the possible states (t and t') and *one* of the possible reports the committee can make are shown. The committee's payoffs are listed first, the legislature's second.

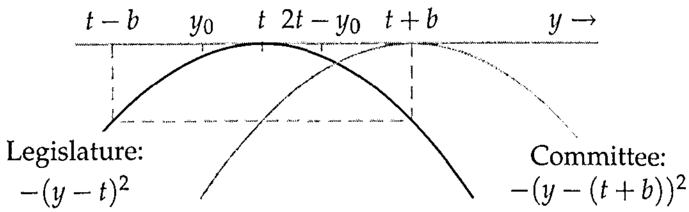

Figure 353.2 For y_0 between $t-b$ and $t+b$, the legislature prefers y_0 to $t+b$, and thus optimally exercises its option to choose y_0 rather than a proposal of $t+b$. For y_0 between $t-b$ and t, it is indifferent between y_0 and $2t-y_0$.

an interval of states, so that given that the initial distribution of states is uniform, the consistency condition requires that the legislature's belief about the state after observing one of these recommendations be uniform on the interval of states that generates the recommendation. In summary, the consistency condition requires that the legislature's belief satisfy the following conditions.

$$\begin{cases} \text{State is } r-b & \text{if } r < y_0 \\ \text{State is uniformly distributed from } y_0-b \text{ to } y_0 & \text{if } r = y_0 \\ \text{State is uniformly distributed from } y_0 \text{ to } y_0+2b & \text{if } r = y_0+2b \\ \text{State is uniformly distributed from } y_0+2b \text{ to } y_0+4b & \text{if } r = y_0+4b \\ \text{State is } r-b & \text{if } y_0+4b < r. \end{cases}$$

The consistency condition does not restrict the legislature's belief about the state when it observes a recommendation between y_0 and y_0+2b, or one between y_0+2b and y_0+4b, because the committee does not make such a recommendation if it follows its strategy. Assume that in these cases it believes the state is y_0. Finally,

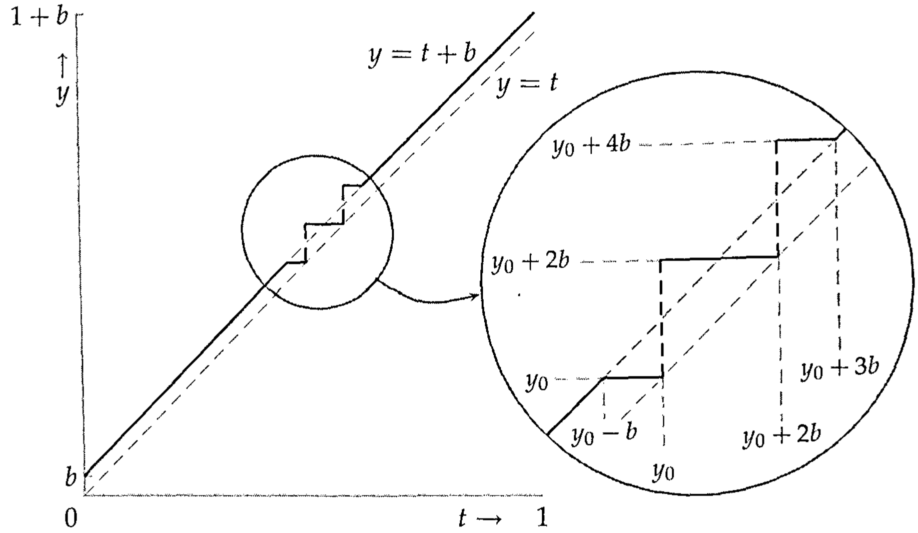

Figure 354.1 The bill y proposed by the committee and accepted by the legislature, as a function of the state t, in an equilibrium of the game of agenda control under the closed rule for $b = \frac{1}{24}$. The 45° dashed gray lines show the action optimal in each state for the committee (top) and legislature (bottom).

the legislature's strategy is to accept all recommendations except those between y_0 and $y_0 + 2b$, and between $y_0 + 2b$ and $y_0 + 4b$ (which are not made if the committee adheres to its equilibrium strategy).

I now argue that this assessment is a weak sequential equilibrium. The legislature's belief satisfies the consistency condition by construction. Consider the optimality of the committee's strategy, given the legislature's strategy. In all states less than $y_0 - b$ or greater than $y_0 + 3b$, the outcome is the committee's favorite bill, so certainly no strategy yields it a higher payoff in these states. Now consider the state y_0, in which its strategy calls for it to recommend the bill y_0, so that it obtains the payoff $-b^2$. If it changes its recommendation, the bill passed by the legislature either remains y_0 (if it recommends between y_0 and $y_0 + 2b$ or between $y_0 + 2b$ and $y_0 + 4b$), changes to less than y_0 (if it recommends less than y_0), or changes to at least $y_0 + 2b$ (if it recommends $y_0 + 2b$ or at least $y_0 + 4b$). None of these changes increases the committee's payoff (refer to Figure 355.1), so its recommendation of y_0 is optimal. In states between $y_0 - b$ and y_0, and between y_0 and $y_0 + 3b$, similar analyses show that the recommendation y_0 is also optimal. (Note that the committee is indifferent between the bills y_0 and $y_0 + 2b$ in state y_0, and between the bills $y_0 + 2b$ and $y_0 + 4b$ in state $y_0 + 2b$.)

Finally consider the legislature's action. A recommendation r from the committee of less than y_0 or greater than $y_0 + 4b$ reveals the state to be $r - b$, so that the legislature prefers r to y_0. Now consider the other possible recommendations the committee may make.

Recommendation of y_0: The outcome is the same, y_0, regardless of the legislature's action.

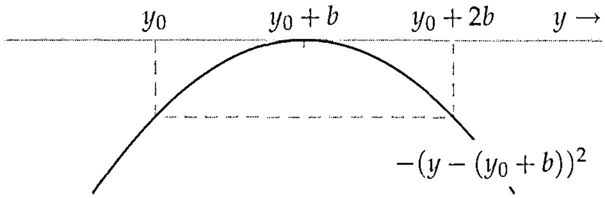

Figure 355.1 In state y_0, the committee prefers the bill y_0 to any bill less than y_0 or greater than $y_0 + 2b$.

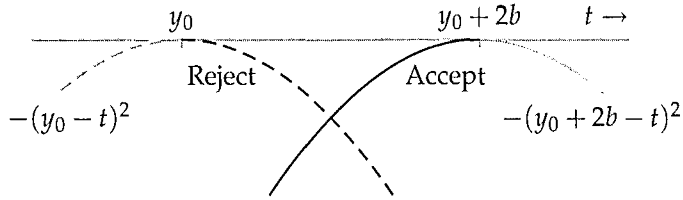

Figure 355.2 The legislature's payoff as a function of the state from y_0 to $y_0 + 2b$ if it accepts the recommended bill $y_0 + 2b$ (solid black curve) and rejects it (dashed black curve).

Recommendation of $y_0 + 2b$: The legislature believes that the state is uniformly distributed from y_0 to $y_0 + 2b$. Its payoff, as a function of the states in this interval, is given by the solid black curve in Figure 355.2 if it accepts the committee's recommendation and by the dashed black curve if it rejects the committee's recommendation. These curves are mirror images of each other, so the legislature is indifferent between accepting and rejecting the committee's recommendation.

Recommendation of $y_0 + 4b$: The legislature believes that the state is uniformly distributed from $y_0 + 2b$ to $y_0 + 3b$. Its payoff, as a function of the states in this interval, is given by the solid black curve in Figure 356.1 if it accepts the committee's recommendation and by the dashed black curve if it rejects the committee's recommendation. (For clarity, the vertical scale in this figure is different from that in Figure 355.2.) The solid curve lies entirely above the dashed one, so the legislature prefers to accept the committee's recommendation than to reject it.

Recommendation between y_0 and $y_0 + 2b$, or between $y_0 + 2b$ and $y_0 + 4b$: The legislature believes that the state is y_0, so it optimally rejects the committee's recommendation.

10.9.3 Comparison of open rule, closed rule, and delegation

The legislature's action, as a function of the state, in the best equilibrium for the legislature under the open rule for $b = \frac{1}{24}$, is given in Figure 348.2. Figure 354.1 is the corresponding figure for the equilibrium found in the previous section (10.9.2) for the closed rule. Which equilibrium does the legislature prefer? Comparing the two

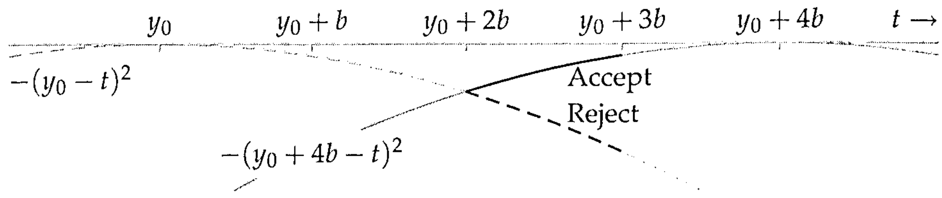

Figure 356.1 The legislature's payoff as a function of the state from $y_0 + 2b$ to $y_0 + 3b$ if it accepts the recommended bill $y_0 + 4b$ (solid black curve) and rejects it (dashed black curve).

figures, it should be plausible that it prefers the equilibrium for the closed rule. In this equilibrium, the difference between the outcome and the legislature's favorite bill is b for most states and is never more than $2b$, whereas for the equilibrium under the open rule the outcome wanders much further afield. The legislature really dislikes large deviations from its favorite bill (its payoff function is quadratic), so the equilibrium under the closed rule is better for it than the equilibrium under the open rule. (Note that I have not argued that the equilibrium found for the closed rule yields the best outcome for the legislature, only that it is better than the best equilibrium under the open rule.)

For other values of b, a comparison of the equilibria under the two rules leads to the same conclusion: if $b < y_0 < 1 - 3b$ and b is small enough that there is an equilibrium under the open rule in which the committee submits at least two different reports, depending on the state, then there is an equilibrium under the closed rule that the legislature prefers to the best equilibrium for it under the open rule.

So in a comparison of the open rule and the closed rule, the analysis provides a rationale for the closed rule: in limiting the legislature's ability to take advantage of the information reported by the committee, the rule gives the committee an incentive to divulge that information, to the ultimate advantage of both the committee and the legislature.

Now compare the outcomes under the closed rule and under delegation. Look carefully at Figure 354.1. In the equilibrium depicted there for the closed rule, the bill proposed by the committee differs from $t + b$ in the region of y_0. The deviations are symmetric about $t + b$, as Figure 357.1 makes clear. Now, the legislature's payoff decreases at an increasing rate away from its favorite bill—its payoff to a bill 2δ from its favorite is more than twice as large a negative number as its payoff to a bill δ from its favorite. Thus it prefers the outcome in which the bill is $t + b$ in every state t to the outcome in the equilibrium we have been studying. That is, for this value of b, delegation is better for it than the equilibrium we found for the closed rule. If it delegates the decision to a committee, the legislature loses control of the outcome; if it retains some control it loses information, because the committee's interest leads to a distorted communication of the state. In the model we have been studying, the loss of information outweighs the loss of control as long as the committee's and legislature's preferences do not differ too much, leading delegation to be the best mechanism for the legislature in this case.

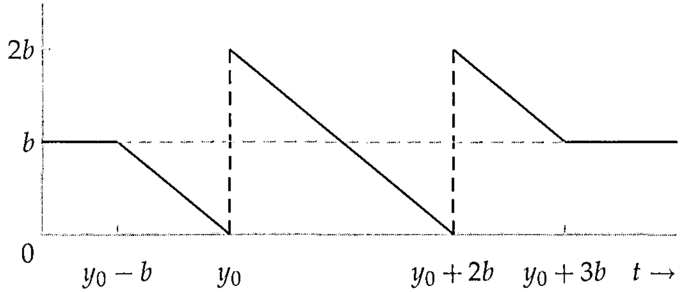

Figure 357.1 The deviation of the bill proposed by the committee from the legislature's favorite bill in an equilibrium in the model of agenda control with the closed rule, as a function of the state around the state y_0.

Thus while the closed rule generates an equilibrium better for the legislature than any equilibrium under the open rule, delegation generates an even better outcome. Assuming the game for the closed rule has no better equilibrium than the one we studied, the conclusion is that if the preferences of the legislature and the committee do not differ so much that the legislature should simply choose the legislation itself, then the legislature's best option is to relinquish all control over the outcome and to cede the choice of legislation to the committee.

Notes

The notion of an extensive game is due to von Neumann and Morgenstern (1944). Kuhn (1950b, 1953) suggested the formulation described in this chapter. The notion of sequential equilibrium, of which weak sequential equilibrium is a variant, is due to Kreps and Wilson (1982); Selten's (1975) notion of "perfect equilibrium" is closely related. (The notion of weak sequential equilibrium is called "assessment equilibrium" by Binmore (1992); it is sometimes called "weak perfect Bayesian equilibrium", though the notion of perfect Bayesian equilibrium (due to Fudenberg and Tirole 1991) is defined only for a restricted set of games.)

The models of poker in Example 315.1 and Exercise 316.1 are variants of models studied by Borel (1938, 91–97) and von Neumann and Morgenstern (1944, 19.14). Von Neumann studied several other models of the game (see von Neumann and Morgenstern 1944, footnote 2 on page 186 and Section 19), as did some of the other pioneers of game theory, including Bellman and Blackwell (1949), Kuhn (1950a), Nash and Shapley (1950), and Karlin (1959b, Chapter 9).

The idea of the model in Section 10.6 is due to Nelson (1970, 1974). The model itself is a simplified version, suggested by Ariel Rubinstein, of the one of Milgrom and Roberts (1986). Section 10.7 is based on Spence (1974), as interpreted by Cho and Kreps (1987). Section 10.8 is based on Crawford and Sobel (1982). The model in Section 10.9 is due to Gilligan and Krehbiel (1987); the equilibrium under the closed rule is taken from Krishna and Morgan (2001). Banks (1991) presents many other applications of signaling games in political science.

Example 320.2 is based on Bagwell (1995). The game in Exercise 331.1 is taken from Selten (1975) and that in Exercise 331.2 is taken from Kohlberg and Mertens (1986). The game in Exercise 335.2 comes from Maynard Smith (1991). In the Battle of Zutphen in 1586, the young British aristocrat and poet Sir Philip Sydney, having been seriously wounded, reportedly gave a water bottle to a dying soldier with the words "Thy necessity is yet greater than mine" (see Greville 1986, 77 and 215). Maynard Smith writes that "This unusual example of altruism by a member of the English upper classes" inspired him to formulate the game. Bergstrom and Lachmann (1997) compare the pooling and separating equilibria of the game. For more discussion of the game, see Godfray and Johnstone (2000).