5 Finite Games with Incomplete Information

In a game of *imperfect* information players may be uninformed about the moves made by other players. Every one-shot, simultaneous move game is a game of imperfect information. In a game of *incomplete* information players may be uninformed about certain characteristics of the game or of the players. For instance, a player may have incomplete information about actions available to some other player, or about payoffs of other players. Incomplete information is modelled by assuming that every player can be of a number of different types. A type of a player summarizes all relevant information (in particular, actions and payoffs) about that player. Furthermore, it is assumed that each player knows his own type and, given his own type, has a probability distribution over the types of the other players. Often, these probability distributions are assumed to be consistent in the sense that they are the marginal probability distributions derived from a basic commonly known distribution over all combinations of player types.

In this chapter we consider games with finitely many players, finitely many types, and finitely many strategies. These games can be either static (simultaneous, one-shot) or dynamic (extensive form games). A Nash equilibrium in this context is also called 'Bayesian equilibrium', and in games in extensive form an appropriate refinement is perfect Bayesian equilibrium. As will become clear, in essence the concepts studied in Chaps. 3 and 4 are applied again. Throughout this chapter we restrict attention to pure strategies and pure strategy Nash equilibria.

In Sect. 5.1 we present a brief introduction to the concept of player types in a game, but the remainder of the chapter can also be understood without this general introduction. Section 5.2 considers static games of incomplete information, and Sect. 5.3 discusses so-called signaling games, which is the most widely applied class of extensive form games with incomplete information. Both Sects. 5.2 and 5.3 are based on examples, rather than general definitions. For the latter, see Chap. 14.

H. Peters, *Game Theory*, Springer Texts in Business and Economics,
DOI 10.1007/978-3-662-46950-7_5

5.1 Player Types

Consider the set of players $N = \{1, \ldots, n\}$. For each player $i \in N$, there is a finite set T_i of *types* which that player can have. If we denote by $T = T_1 \times \ldots \times T_n$ the set

$$T = \{(t_1, \ldots, t_n) \mid t_1 \in T_1, t_2 \in T_2, \ldots, t_n \in T_n\} ,$$

i.e., the set of all possible combinations of types, then a *game with incomplete information* specifies a separate game for every possible combination $t = (t_1, \ldots, t_n) \in T$, in a way to be explained in the next sections. We assume that each player i knows his own type t_i and, given t_i, attaches probabilities

$$p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i)$$

to all type combinations $t_1 \in T_1$, …, $t_{i-1} \in T_{i-1}$, $t_{i+1} \in T_{i+1}$, …, $t_n \in T_n$ of the other players.

Often, these probabilities are derived from a common probability distribution p over T, where $p(t)$ is the probability that the type combination is t. This is also what we assume in this chapter. Moreover, we assume that every player i, apart from his own type t_i, also knows the probability distribution p. Hence, if player i has type t_i, then he can compute the probability that the type combination of the other players is the vector $(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n)$. Formally, this probability is equal to the conditional probability

$$p(t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n | t_i) = \frac{p(t_1, \ldots, t_{i-1}, t_i, t_{i+1}, \ldots, t_n)}{\sum p(t'_1, \ldots, t'_{i-1}, t_i, t'_{i+1}, \ldots, t'_n)}$$

where the sum in the denominator is taken over all possible types of the other players, i.e., over all possible $t'_1 \in T_1, \ldots, t'_{i-1} \in T_{i-1}, t'_{i+1} \in T_{i+1}, \ldots, t'_n \in T_n$. Hence, the sum in the denominator is the probability that player i has type t_i.

Thus, a player in a game of incomplete information can make his actions dependent on his own type but not on the types of the other players. However, since he knows the probabilities of the other players' types, he can compute the expected payoffs from taking specific actions. In the next two sections we will see how this works in static and in extensive form games.

5.2 Static Games of Incomplete Information

We discuss a few examples.

5.2.1 Battle-of-the-Sexes with One-Sided Incomplete Information

The first example is a variant of the Battle-of-the-Sexes (see Sect. 1.3.2) in which player 1 (the man) does not know whether player 2 (the woman) wants to go out with him or avoid him. More precisely, player 1 does not know whether he plays the game y (from 'yes') or the game n (from 'no'), where these games are as follows:

$$y: \begin{array}{c} \\ F \\ B \end{array} \begin{array}{c} \begin{array}{cc} F & B \end{array} \\ \begin{pmatrix} 2,1 & 0,0 \\ 0,0 & 1,2 \end{pmatrix} \end{array} \qquad n: \begin{array}{c} \\ F \\ B \end{array} \begin{array}{c} \begin{array}{cc} F & B \end{array} \\ \begin{pmatrix} 2,0 & 0,2 \\ 0,1 & 1,0 \end{pmatrix} \end{array}.$$

Player 1 attaches probability $1/2$ to each of these games, and player 2 knows this. In the terminology of types, this means that player 1 has only one type, simply indicated by '1', and that player 2 has two types, namely y and n. So there are two type combinations, namely $(1, y)$ and $(1, n)$, each occurring with probability $1/2$. Player 2 knows player 1's type with certainty, and also knows her own type, that is, she knows which game is actually being played. Player 1 attaches probability $1/2$ to each type of player 2.

What would be a Nash equilibrium in this game? To see this, it is helpful to model the game as a game in extensive form, using the tree representation of Chap. 4. Such a tree is drawn in Fig. 5.1.

The game starts with a chance move which selects which of the two bimatrix games is going to be played. In the terminology of types, it selects the type of player 2. Player 2 is informed but player 1 is not. Player 2 has four different strategies but player 1 only two.

From this extensive form it is apparent that every Nash equilibrium is subgame perfect, since there are no nontrivial subgames.

Also, every Nash equilibrium is perfect Bayesian, since the only nontrivial information set (namely, that of player 1) is reached with positive probability (namely, equal to 1) for any strategy of player 2, and thus the beliefs are completely

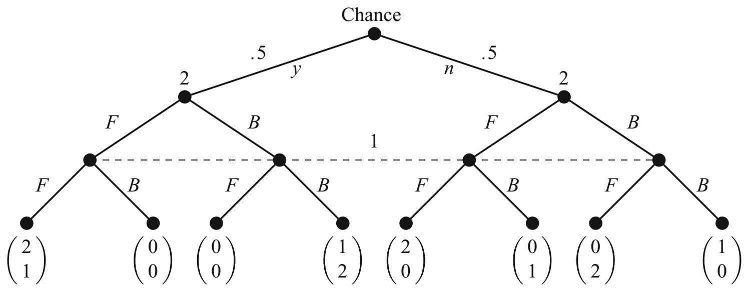

Fig. 5.1 An extensive form representation of the Battle-of-the-Sexes game with incomplete information. The upper numbers are the payoffs for player 1

determined by player 2's strategy through Bayesian updating (see Sect. 4.4). More precisely, suppose the belief of player 1 is denoted by the nonnegative vector $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$, where $\alpha_1 + \ldots + \alpha_4 = 1$, from left to right in Fig. 5.1. For instance, α_3 is the probability attached by player 1 to player 2 having type n and playing strategy F. Suppose, for instance, that player 2 plays F if she has type y and B if she has type n, and let E denote the event that player 1's information set is reached. Then

$$\begin{aligned}
\alpha_1 &= \text{Prob}\,[2 \text{ plays } F \text{ and has type } y \mid E] \\
&= \frac{\text{Prob}\,[2 \text{ plays } F, \text{ has type } y, \text{ and } E]}{\text{Prob}\,[E]} \\
&= \frac{\text{Prob}\,[2 \text{ plays } F \mid 2 \text{ has type } y]\,\text{Prob}\,[2 \text{ has type } y]\,\text{Prob}\,[E]}{\text{Prob}\,[E]} \\
&= \frac{1 \cdot 0.5 \cdot 1}{1} = 0.5\,.
\end{aligned}$$

By a similar computation we find $\alpha_2 = 0$, $\alpha_3 = 0$, and $\alpha_4 = 0.5$. Such computations can also be made if player 2 would play mixed but we restrict attention here to pure strategies. Important is that the beliefs of player 1 are always determined by computing the conditional probabilities since his (only) information set is always reached with positive probability, namely 1. Hence, there are no *free* beliefs, in the terminology of Sect. 4.4.

The strategic form of the game is given in Fig. 5.2. There, the first letter in a strategy of player 2 says what player 2 plays if y is chosen by the Chance move, and the second letter says what player 2 plays if n is chosen. For instance, if player 1 plays F and player 2 plays FF, then the expected payoffs are equal to $0.5 \cdot (2, 1) + 0.5 \cdot (2, 0) = (2, 0.5)$. Or if player 1 plays B and player 2 plays FB, the expected payoffs are equal to $0.5 \cdot (0, 0) + 0.5 \cdot (1, 0) = (0.5, 0)$, etc. Also the best replies are indicated.

From the strategic form it is apparent that the game has a unique Nash equilibrium in pure strategies, namely (F, FB). In this equilibrium player 1 plays F, type y of player 2 plays F and type n of player 2 plays B.

Another equilibrium concept, appropriate for static games with incomplete information, is that of a *Bayesian equilibrium*. In a Bayesian equilibrium, each type of each player plays a best reply against the other players. That is, each type of a player plays an action that maximizes his expected payoff, where the expectation is taken over the type combinations of the other players and their actions. The Nash equilibrium (F, FB) in the game above is such a Bayesian

Fig. 5.2 The strategic form of the game in Fig. 5.1. Player 1 is the row player

$$\begin{array}{c} \\ F \\ B \end{array}
\begin{array}{c}
\begin{array}{cccc} FF & FB & BF & BB \end{array} \\
\begin{pmatrix}
\underline{2}, 0.5 & \underline{1}, \underline{1.5} & \underline{1}, 0 & 0, 1 \\
0, 0.5 & 0.5, 0 & 0.5, \underline{1.5} & \underline{1}, 1
\end{pmatrix}
\end{array}$$

equilibrium: the action F of type y of player 2 is a best reply against F of player 1 (player 1 has only one type), the action B of type n of player 2 is a best reply against F of player 1, and action F of player 1 maximizes player 1's expected payoff against the strategy FB of player 2. In fact, if every type of every player has positive probability (as will be the case throughout this chapter), then the Nash equilibria of the strategic form coincide with the Bayesian equilibria.

The (pure) Nash equilibrium or equilibria of a game like this can also be found without drawing the extensive form and computing the strategic form, as follows. Suppose first that player 1 plays F in an equilibrium. Then the best reply of player 2 is to play F if her type is y and B if here type is n. The expected payoff to player 1 is then 1; playing B against this strategy FB of player 2 yields only 0.5. So (F, FB) is a Nash equilibrium. If, on the other hand, player 1 plays B, then the best reply of player 2 if her type is y, is B and if her type is n it is F. This yields a payoff of 0.5 to player 1, whereas playing F against this strategy BF of player 2 yields payoff 1. Hence, there is no equilibrium where player 1 plays B. Of course, this is also apparent from the strategic form, but the argument can be made without complete computation of the strategic form.

5.2.2 Battle-of-the-Sexes with Two-Sided Incomplete Information

The next example is a further variation on the Battle-of-the-Sexes game in which neither player knows whether the other player wants to be together with him/her or not. It is based on the four bimatrix games in Fig. 5.3. These four bimatrix games correspond to the four possible type combinations of players 1 and 2. The probabilities of these four different combinations are given in Table 5.1. One way to find the Nash equilibria of this game is to draw the extensive form and compute the associated strategic form: see Problem 5.1. Alternatively, we can systematically examine the possible strategy pairs, as follows.

First observe that each player now has four strategies, namely FF, FB, BF, and BB, where the first letter is the action taken by the yes-type (y_1 or y_2), and the second letter is the action taken by the no-type (n_1 or n_2).

Fig. 5.3 Payoffs for Battle-of-the-Sexes with two types per player

$$y_1y_2: \quad \begin{matrix} & F & B \\ F & 2,1 & 0,0 \\ B & 0,0 & 1,2 \end{matrix} \qquad y_1n_2: \quad \begin{matrix} & F & B \\ F & 2,0 & 0,2 \\ B & 0,1 & 1,0 \end{matrix}$$

$$n_1y_2: \quad \begin{matrix} & F & B \\ F & 0,1 & 2,0 \\ B & 1,0 & 0,2 \end{matrix} \qquad n_1n_2: \quad \begin{matrix} & F & B \\ F & 0,0 & 2,2 \\ B & 1,1 & 0,0 \end{matrix}$$

Table 5.1 Type probabilities for Battle-of-the-Sexes with two types per player

t	y_1y_2	y_1n_2	n_1y_2	n_1n_2
$p(t)$	$2/6$	$2/6$	$1/6$	$1/6$

Next, the conditional type probabilities can easily be computed from Table 5.1. For instance,

$$p(y_2|y_1) = \frac{p(y_1y_2)}{p(y_1)} = \frac{p(y_1y_2)}{p(y_1y_2) + p(y_1n_2)} = \frac{2/6}{(2/6) + (2/6)} = 1/2 \,.$$

The other conditional probabilities are computed in the same way, yielding:

$$p(n_2|y_1) = 1/2,\ p(y_2|n_1) = 1/2,\ p(n_2|n_1) = 1/2 \,,$$

$$p(y_1|y_2) = 2/3,\ p(n_1|y_2) = 1/3,\ p(y_1|n_2) = 2/3,\ p(n_1|n_2) = 1/3 \,.$$

We now consider the four pure strategies of player 1 one by one.

(i) Suppose that player 1 plays the strategy FF, meaning that he plays F (the first letter) if his type is y_1 and also F (the second letter) if his type is n_1. Then the expected payoff for type y_2 of player 2 if she plays F is

$$p(y_1|y_2) \cdot 1 + p(n_1|y_2) \cdot 1 = (2/3) \cdot 1 + (1/3) \cdot 1 = 1 \,.$$

If type y_2 of player 2 plays B her expected payoff is

$$p(y_1|y_2) \cdot 0 + p(n_1|y_2) \cdot 0 = (2/3) \cdot 0 + (1/3) \cdot 0 = 0 \,.$$

Hence the best reply of type y_2 of player 2 is F. Similarly, for type n_2 of player 2, playing F yields

$$p(y_1|n_2) \cdot 0 + p(n_1|n_2) \cdot 0 = (2/3) \cdot 0 + (1/3) \cdot 0 = 0$$

and playing B yields

$$p(y_1|n_2) \cdot 2 + p(n_1|n_2) \cdot 2 = (2/3) \cdot 2 + (1/3) \cdot 2 = 2 \,,$$

so that B is the best reply. Hence, player 2's best reply against FF is FB. Suppose, now, that player 2 plays FB, so type y_2 plays F and type n_2 plays B. Then playing F yields type y_1 of player 1 an expected payoff of

$$p(y_2|y_1) \cdot 2 + p(n_2|y_1) \cdot 0 = (1/2) \cdot 2 + (1/2) \cdot 0 = 1$$

and playing B yields

$$p(y_2|y_1) \cdot 0 + p(n_2|y_1) \cdot 1 = (1/2) \cdot 0 + (1/2) \cdot 1 = 1/2 \,,$$

so that F is the best reply for type y_1 of player 1. Similarly, for type n_1 playing F yields

$$p(y_1|n_1) \cdot 0 + p(n_2|n_1) \cdot 2 = (1/2) \cdot 0 + (1/2) \cdot 2 = 1$$

whereas playing B yields

$$p(y_1|n_1) \cdot 1 + p(n_2|n_1) \cdot 0 = (1/2) \cdot 1 + (1/2) \cdot 0 = 1/2 \; .$$

Hence, F is the best reply for type n_1. Hence, player 1's best reply against FB is FF. We conclude that (FF, FB) is a Nash equilibrium.

(ii) Suppose player 1 plays FB. Playing F yields type y_2 of player 2 a payoff of

$$p(y_1|y_2) \cdot 1 + p(n_1|y_2) \cdot 0 = 2/3 \cdot 1 + 1/3 \cdot 0 = 2/3$$

and playing B yields

$$p(y_1|y_2) \cdot 0 + p(n_1|y_2) \cdot 2 = 2/3 \cdot 0 + 1/3 \cdot 2 = 2/3$$

so that both F and B are best replies. Playing F yields type n_2 a payoff of

$$p(y_1|n_2) \cdot 0 + p(n_1|n_2) \cdot 1 = 2/3 \cdot 0 + 1/3 \cdot 1 = 1/3$$

and playing B yields

$$p(y_1|n_2) \cdot 2 + p(n_1|n_2) \cdot 0 = 2/3 \cdot 2 + 1/3 \cdot 0 = 4/3$$

so that B is the best reply. Hence, player 2 has two best replies, namely FB and BB. Against FB player 1's best reply is FF [as established in case (i)] and not FB, so this does not result in a Nash equilibrium. Against BB one can compute in the same way as hitherto that player 1's best reply is BF and not FB, so also this combination is not a Nash equilibrium.

(iii) Suppose that player 1 plays BF. Then player 2 has two best replies, namely BF and BB. Against BF the best reply of player 1 is FF and not BF, so this combination is not a Nash equilibrium. Against BB, player 1's best reply is BF, so the combination (BF, BB) is a Nash equilibrium.

(iv) Finally, suppose player 1 plays BB. Then player 2's best reply is BF. Against this, player 1's best reply is FF and not BB. So BB of player 1 is not part of a Nash equilibrium.

We conclude that the game has two Nash equilibria in pure strategies, namely: (i) both types of player 1 play F, type y_2 of player 2 also plays F but type n_2 of player 2 plays B; (ii) type y_1 of player 1 plays B, type n_1 plays F, and both types of player 2 play B. Again, these equilibria are also called Bayesian Nash equilibria.

5.3 Signaling Games

The extensive form can be used to examine a static game of incomplete information, usually by letting the game start with a chance move that picks the types of the players (see Sect. 5.2). More generally, the extensive form can be used to describe incomplete information games where players move sequentially. An important class of such games is the class of signaling games.

A (finite) signaling game starts with a chance move that picks the type of player 1. Player 1 is informed about his type but player 2 is not. Player 1 moves first, player 2 observes player 1's action and moves next, and then the game ends. Such a game is called a *signaling game* because the action of player 1 may be a signal about his type: that is, from the action of player 1 player 2 may be able to infer something about the type of player 1.

5.3.1 An Example

Consider the example in Fig. 5.4. (The numbers between square brackets at player 2's decision nodes are the beliefs of player 2, which are used in a perfect Bayesian equilibrium below.) In this game, player 1 learns the result of the chance move but player 2 does not. In the terminology of Sect. 5.1, there are two type combinations, namely $(t, 2)$ and $(\tilde{t}, 2)$, each one occurring with probability $1/2$: these notations express the fact that player 2 has only one type (called '2'). Both types of player 1 can choose between L and R. Player 2 only observes the action, L or R, and not the type of player 1. For this reason we use the same letter L for the 'left' action of each type: player 2 cannot distinguish between them. Similar for the 'right' action.

In order to analyze this game and find the (pure strategy) Nash equilibria, one possibility is to first compute the strategic form. Both players have four strategies. Player 1 has strategy set

$$\{LL, LR, RL, RR\} ,$$

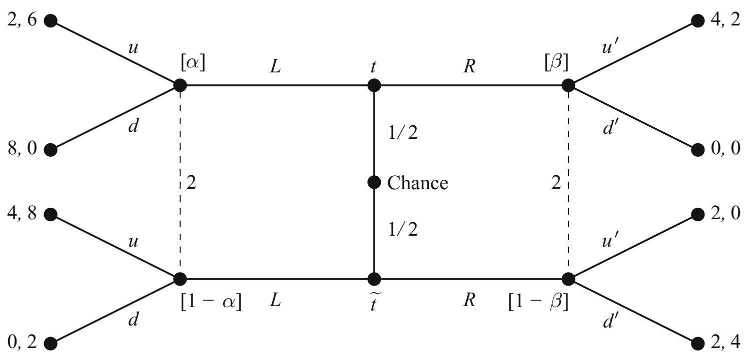

Fig. 5.4 A signaling game

Fig. 5.5 The strategic form of the game in Fig. 5.4

$$
\begin{array}{c} \\ LL \\ LR \\ RL \\ RR \end{array}
\begin{array}{c} \begin{array}{cccc} uu' & ud' & du' & dd' \end{array} \\
\begin{pmatrix} 3,\underline{7} & \underline{3},\underline{7} & 4,1 & 4,1 \\ 2,3 & 2,\underline{5} & \underline{5},0 & \underline{5},2 \\ \underline{4},\underline{5} & 2,4 & 2,2 & 0,1 \\ 3,1 & 1,\underline{2} & 3,1 & 1,\underline{2} \end{pmatrix} \end{array}
$$

where the first letter refers to the action of type t and the second letter to the action of type $\tilde{t}$. Player 2 has strategy set

$$\{uu', ud', du', dd'\} .$$

The (expected) strategic form of the game can be computed in the usual way and is presented in Fig. 5.5. For instance, consider the strategy combination (LR, ud'). Then the expected payoffs are $1/2 \cdot (2, 6) + 1/2 \cdot (2, 4) = (2, 5)$, etc. The (pure) best replies are underlined. This shows that the game has two Nash equilibria, namely (RL, uu') and (LL, ud'). What else can be said about these equilibria? Observe that the only subgame of the game is the entire game, so that both equilibria are trivially subgame perfect. Are they also perfect Bayesian? That is, do they satisfy the conditions of Bayesian consistency and sequential rationality (Sect. 4.4)?

First consider the equilibrium (RL, uu'). Bayesian consistency requires

$$
\begin{aligned}
\alpha &= \text{Prob}\,[\text{player 1 has type } t \mid \text{player 1 plays } L] \\
&= \frac{\text{Prob}\,[\text{player 1 has type } t \text{ and plays } L]}{\text{Prob}\,[\text{player 1 plays } L]} \\
&= \frac{1/2 \cdot 0}{1/2} \\
&= 0
\end{aligned}
$$

and, similarly,

$$
\begin{aligned}
\beta &= \text{Prob}\,[\text{player 1 has type } t \mid \text{player 1 plays } R] \\
&= \frac{\text{Prob}\,[\text{player 1 has type } t \text{ and plays } R]}{\text{Prob}\,[\text{player 1 plays } R]} \\
&= \frac{1/2 \cdot 1}{1/2} \\
&= 1\ .
\end{aligned}
$$

Given these beliefs, playing u yields a payoff of 8 to player 2 (at the left information set), whereas playing d yields only 2, so u is optimal. Playing u' yields a payoff of 2 to player 2 (at the right information set), whereas playing d' yields 0, so u' is optimal. Hence, uu' is indeed the best reply of player 2—as we already knew from the strategic form. Thus, the pair (RL, uu') is a perfect Bayesian equilibrium with beliefs $\alpha = 0$ and $\beta = 1$.

Note that, in this case and, more generally, in case of a Nash equilibrium where each information set of player 2 is reached with positive probability (i.e., every possible action of player 1 is played by some type of player 1), the perfect Bayesian equilibrium requirement does not add anything that is essential: the beliefs of player 2 are completely determined by Bayesian consistency, and sequential rationality is automatically satisfied for these beliefs, given that we already have a Nash equilibrium.

The perfect Bayesian equilibrium (RL, uu') is called *separating*: it separates the two types of player 1, since these types play different actions. In this equilibrium, the action of player 1 is a signal for his type, and the equilibrium is 'information revealing'.

Next, consider the Nash equilibrium (LL, ud'). In this case, by Bayesian consistency we obtain

$$\begin{aligned}\alpha &= \text{Prob}\,[\text{player 1 has type } t \mid \text{player 1 plays } L] \\ &= \frac{\text{Prob}\,[\text{player 1 has type } t \text{ and plays } L]}{\text{Prob}\,[\text{player 1 plays } L]} \\ &= \frac{1/2 \cdot 1}{1} \\ &= 1/2\,.\end{aligned}$$

In words: since each type of player 1 plays L, the conditional probabilities of the two decision nodes in the left information set of player 2 are both equal to $1/2$. Given $\alpha = 1/2$ it follows that u is optimal at player 2's left information set (in fact, in this particular game u is optimal for any α): we know this already from the strategic form, where both uu' and ud' are best replies of player 2 against LL.

How about the belief $(\beta, 1 - \beta)$? Since player 1 always plays L, the right information set of player 2 is reached with zero probability, and therefore we cannot compute β by the formula for conditional probability: Bayesian consistency has no bite, the belief $(\beta, 1 - \beta)$ is *free* (in the terminology of Sect. 4.4). Formally,

$$\begin{aligned}\beta &= \text{Prob}\,[\text{player 1 has type } t \mid \text{player 1 plays } R] \\ &= \frac{\text{Prob}\,[\text{player 1 has type } t \text{ and plays } R]}{\text{Prob}\,[\text{player 1 plays } R]}\,,\end{aligned}$$

but the probability in the denominator of the last expression is zero if player 1 plays LL. However, we still have the sequential rationality requirement: in order for (LL, ud') to be a perfect Bayesian equilibrium the belief $(\beta, 1 - \beta)$ should be such that player 2's action d' is optimal. Hence, the expected payoff to player 2 from playing d' should be at least as large as the expected payoff from playing u', so $4(1 - \beta) \geq 2\beta$, which is equivalent to $\beta \leq 2/3$. Thus, (LL, ud') is a perfect Bayesian equilibrium with beliefs $\alpha = 1/2$ and $\beta \leq 2/3$.

In this case, and more generally, in cases where in a Nash equilibrium *not* every information set of player 2 is reached with positive probability, i.e., there is some action of player 1 which is played by *no* type of player 1, the perfect Bayesian equilibrium requirement does have an impact.

The equilibrium (LL, ud') is called *pooling*, since it 'pools' the two types of player 1: both types play the same action, L in this case. In this equilibrium, the action of player 1 does not reveal any information about his type.

5.3.2 Computing Perfect Bayesian Equilibria in the Extensive Form

Perfect Bayesian equilibria can also be found without first computing the strategic form. We consider again the signaling game in Fig. 5.4.

First, assume that there is an equilibrium where player 1 plays LL. Then $\alpha = 1/2$ by Bayesian consistency, and player 2's optimal action at the left information set (following L) is u. At the right information set, player 2's optimal action is u' if $\beta \geq 2/3$ and d' if $\beta \leq 2/3$. If player 2 would play u' after R, then type t of player 1 would improve by playing R instead of L, so this cannot be an equilibrium. If player 2 plays d' after R, then no type of player 1 would want to play R instead of L. We have established that (LL, ud') with beliefs $\alpha = 1/2$ and $\beta \leq 2/3$ is a (pooling) perfect Bayesian equilibrium.

Second, assume player 1 plays LR in equilibrium. Then player 2's beliefs are given by $\alpha = 1$ and $\beta = 0$, and player 2's best reply is ud'. But then type $\tilde{t}$ of player 1 would gain by playing L instead of R, so this cannot be an equilibrium.

Third, assume player 1 plays RL in equilibrium. Then $\alpha = 0$, $\beta = 1$, and player 2's best reply is uu'. Against uu', RL is player 1's best reply, so that (RL, uu) is a (separating) perfect Bayesian equilibrium with beliefs $\alpha = 0$ and $\beta = 1$.

Fourth, suppose player 1 plays RR in equilibrium. Then $\beta = 1/2$ and player 2's best reply after R is d'. After L, player 2's best reply is u for any value of α. Against ud', however, type t of payer 1 would gain by playing L instead of R. So RR is not part of an equilibrium.

Of course, these considerations can also be based on the strategic form, but we do not need the entire strategic form to find the perfect Bayesian equilibria.

5.3.3 The Intuitive Criterion

In a perfect Bayesian equilibrium, if an information set of player 2 is reached with zero probability, then the belief of player 2 on that information set is free—the only requirement is the sequential rationality requirement demanding that, given this belief, player 2 should choose the optimal action. The question is whether such a free belief is always plausible or reasonable. The so-called *intuitive criterion* (IC) puts a restriction on the plausibility of free beliefs.

It works as follows. Consider a perfect Bayesian equilibrium in a signaling game and suppose that there is an information set which is reached in the equilibrium with zero probability. In other words, there is an action, say A, of player 1 which is played by no type of player 1. Consider a type t of player 1. Suppose this type t obtains payoff x in the equilibrium under consideration. Then consider the maximal possible payoff for this type t attainable by playing A, say m. If $m < x$, then the belief of player 2 on the information set following A should assign zero probability to type t of player 1. The reason is indeed intuitive: type t could never possibly gain by playing A instead of his equilibrium action, since *at best* he would obtain m from doing so, but m is less than what t can get in equilibrium, namely x. Therefore, player 2 should not believe that type t would ever deviate to A. This comparison should be made for every *type* of player 1. This way, we may obtain some restrictions on the belief of player 2 at the information set following action A. If the original belief of player 2 on this information set, corresponding to the perfect Bayesian equilibrium under consideration, satisfies these restrictions, then we say that this perfect Bayesian equilibrium *survives* the IC. However, it could happen that this way all types of player 1 get assigned zero probability: in that case, the IC simply does *not* apply, since the probabilities in a belief have to sum up to 1 [cf. Problem 5.9(a)].

Let us apply the IC to the perfect Bayesian equilibrium (LL, ud') with $\beta \leq 2/3$ in the game in Fig. 5.4. The equilibrium payoff to type t of player 1 is equal to 2. If type t of player 1 deviates to R, then he could get maximally 4, namely if player 2 would play u' following R. It is important to notice that we consider the *maximal possible* payoff after such a deviation—of course, the payoff in the equilibrium after a deviation (in this case 0) can never be higher by the mere definition of Nash equilibrium. Since $4 \not< 2$, type t could have a reason to deviate, and so the IC puts no restriction on the belief of player 2 that player 1, if he would deviate to R, is of type t. The equilibrium payoff to type $\tilde{t}$ of player 1 is 4. Type $\tilde{t}$, however, could get at most 2 (in fact, would always get 2 in this game) by deviating to R. Since $2 < 4$, the IC now says that it is not reasonable for player 2 to assume that type $\tilde{t}$ would ever deviate to R. Thus, $1 - \beta = 0$, so that $\beta = 1$. With this belief, however, (LL, ud') can no longer be sustained as a perfect Bayesian equilibrium, since for this we need $\beta \leq 2/3$. Hence, the perfect Bayesian equilibrium with $\beta \leq 2/3$ does not survive the IC.

5.3.4 Another Example

Consider the signaling game in Fig. 5.6. We compute the pure strategy perfect Bayesian equilibria of this game by considering the strategies of player 1 one by one.

Suppose player 1 plays LL (the first letter refers to type t and the second to type $\tilde{t}$). Then $\alpha = 1/2$ by Bayesian consistency, so player 2 is indifferent between u and d. If player 2 plays u, however, then type $\tilde{t}$ obtains 0 and will therefore deviate to R, so that he obtains at least 1. Hence, player 2 should play d following L in order to get an equilibrium. Then, again to keep type $\tilde{t}$ from deviating, player 2 should

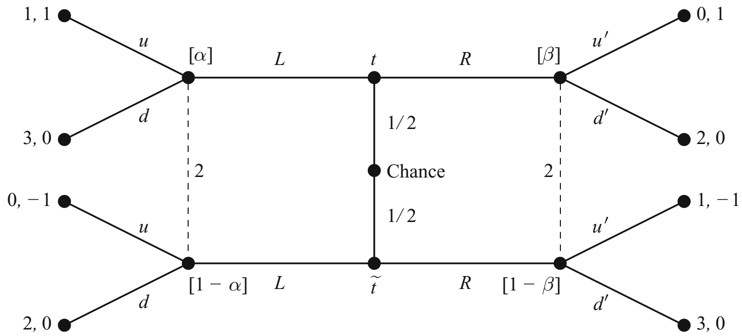

Fig. 5.6 Another signaling game

play u' following R. This is indeed optimal for player 2 for $\beta \geq 1/2$. Obviously, type t will not deviate to R, so (LL, du') with $\alpha = 1/2$ and $\beta \geq 1/2$ is a (pooling) perfect Bayesian equilibrium. In this equilibrium, type t obtains 3 and by deviating to R obtains maximally 2. Type $\tilde{t}$ obtains 2 in equilibrium and by deviating to R maximally 3. Hence, the IC prescribes $\beta = 0$, but for this belief (LL, du') does not result in a perfect Bayesian equilibrium. Hence, (LL, du') with $\alpha = 1/2$ and $\beta \geq 1/2$ does not survive the IC.

Suppose player 1 plays LR. Then the best reply of player 2 is ud', but then type t will deviate to R. So there is no perfect Bayesian equilibrium (or even Nash equilibrium) in which player 1 plays LR.

Suppose player 1 plays RL. Then the best reply of player 2 is du', but then type t will deviate to L. So there is no perfect Bayesian equilibrium (or even Nash equilibrium) in which player 1 plays RL.

Finally, suppose player 1 plays RR. Then $\beta = 1/2$ and player 2 is indifferent between u' and d' following R. If player 2 plays u' then type t of player 1 will deviate to R, where he obtains always more than 0. Hence in an equilibrium player 2 should play d'. Again to keep type t from deviating, player 2 should play u following L. This is indeed optimal for player 2 if $\alpha \geq 1/2$. Hence, (RR, ud') with beliefs $\beta = 1/2$ and $\alpha \geq 1/2$ is a (pooling) perfect Bayesian equilibrium. Type t obtains 2 in equilibrium and maximally 3 by deviating to L. Type $\tilde{t}$ obtains 3 in equilibrium and maximally 2 by deviating to L. Thus, the IC prescribes $1 - \alpha = 0$ or $\alpha = 1$, and therefore this perfect Bayesian equilibrium survives the IC for $\alpha = 1$.

5.4 Problems

5.1. *Battle-of-the-Sexes*
Draw the extensive form of the Battle-of-the-Sexes game in Sect. 5.2 with payoffs in Fig. 5.3 and type probabilities in Table 5.1. Compute the strategic form and find the pure strategy Nash equilibria of the game.

5.2. *A Static Game of Incomplete Information*
Compute all pure strategy Nash equilibria in the following static game of incomplete information:

1. Chance determines whether the payoffs are as in Game 1 or as in Game 2, each game being equally likely.
2. Player 1 learns which game has been chosen but player 2 does not.

The two bimatrix games are:

$$\text{Game 1:}\quad \begin{array}{c} \\ T \\ B \end{array}\!\!\begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} 1,1 & 0,0 \\ 0,0 & 0,0 \end{pmatrix} \end{array} \qquad \text{Game 2:}\quad \begin{array}{c} \\ T \\ B \end{array}\!\!\begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} 0,0 & 0,0 \\ 0,0 & 2,2 \end{pmatrix} \end{array}$$

5.3. *Another Static Game of Incomplete Information*
Player 1 has two types, t_1 and t_1', and player 2 has two types, t_2 and t_2'. The conditional probabilities of these types are:

$$p(t_2|t_1) = 1,\; p(t_2|t_1') = 3/4,\; p(t_1|t_2) = 3/4,\; p(t_1|t_2') = 0\,.$$

(a) Show that these conditional probabilities can be derived from a common distribution p over the four type combinations, and determine p.

As usual suppose that each player learns his own type and knows the conditional probabilities above. Then player 1 chooses between T and B and player 2 between L and R, where these actions may be contingent on the information a player has. The payoffs for the different type combinations are given by the bimatrix games

$$t_1t_2:\; \begin{array}{c} \\ T \\ B \end{array}\!\!\begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} 2,2 & 0,0 \\ 3,0 & 1,1 \end{pmatrix} \end{array} \quad t_1't_2:\; \begin{array}{c} \\ T \\ B \end{array}\!\!\begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} 2,2 & 0,0 \\ 0,0 & 1,1 \end{pmatrix} \end{array} \quad t_1't_2':\; \begin{array}{c} \\ T \\ B \end{array}\!\!\begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} 2,2 & 0,0 \\ 0,0 & 1,1 \end{pmatrix} \end{array},$$

where the type combination (t_1, t_2') is left out since it has zero probability.

(b) Compute all pure strategy Nash equilibria for this game.

5.4. *Job-Market Signaling*
A worker can have either high or low ability, where the probability of high ability is equal to $2/5$. A worker knows his ability, but a firm which wants to hire the worker does not. The worker, whether a high or a low ability type, can choose between additional education or not. Choosing additional education does not enlarge the worker's productivity but may serve as a signal to the firm: a high ability worker can choose education without additional costs, whereas for a low ability worker the cost of education equals $e > 0$. The firm chooses either a high or a low wage, having observed whether the worker took additional education

or not. The payoff to the firm equals the productivity of the worker minus the wage. The payoff to the worker equals the wage minus the cost of education; if, however, this payoff is lower than the worker's reservation utility, he chooses not to work at all and to receive his reservation utility, leaving the firm with 0 payoff. Denote the productivities of the high and low ability worker by p^H and p^L, respectively, and denote the high and low wages by w^h and w^l. Finally, let r^H and r^L denote the reservation utilities of both worker types. (All these numbers are fixed.)

(a) Determine the extensive form of this game.
(b) Choose $p^H = 10$, $p^L = 8$, $w^h = 6$, $w^l = 3$, $r^H = 4$, $r^L = 0$, $e = 4$. Compute the strategic form of this game, and determine the pure strategy Nash equilibria. Also compute the perfect Bayesian equilibrium or equilibria in pure strategies, determine whether they are separating or pooling and whether they survive the IC.

5.5. *A Joint Venture*
Software Inc. and Hardware Inc. are in a joint venture together. The parts used in the joint product can be defective or not; the probability of defective parts is 0.7, and this is commonly known before the start of the game. Each can exert either high or low effort, which is equivalent to costs of 20 and 0. Hardware moves first, but software cannot observe his effort. Revenues are split equally at the end. If both firms exert low effort, total profits are 100. If the parts are defective, the total profit is 100; otherwise (i.e., if the parts are not defective), if both exert high effort, profit is 200, but if only one player does, profit is 100 with probability 0.9 and 200 with probability 0.1. Hardware discovers the truth about the parts by observation before he chooses effort, but software does not.

(a) Determine the extensive form of this game. Is this a signaling game?
(b) Determine the strategic form of this game.
(c) Compute the (pure) Nash equilibria? Which one(s) is (are) subgame perfect? Perfect Bayesian?

5.6. *Entry Deterrence*
The entry deterrence game of Chap. 1 is reproduced in Fig. 5.7. For this game, compute the pure strategy perfect Bayesian equilibria for every value of $x \in \mathbb{R}$. Which one(s) is (are) pooling or separating? Satisfy the intuitive criterion?

5.7. *The Beer-Quiche Game*
Consider the following two-player signaling game. Player 1 is either 'weak' or 'strong'. This is determined by a chance move, resulting in player 1 being 'weak' with probability 1/10. Player 1 is informed about the outcome of this chance move but player 2 is not; but the probabilities of either type of player 1 are

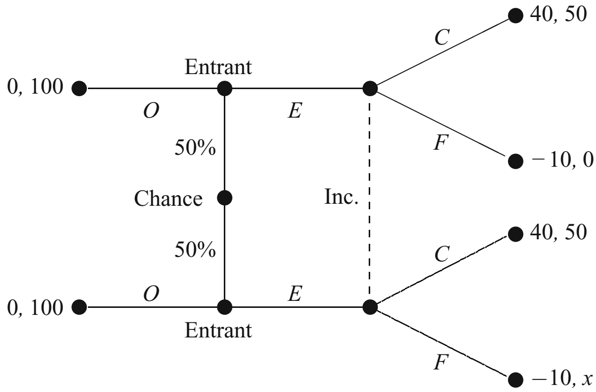

Fig. 5.7 The entry deterrence game of Problem 5.6

common knowledge among the two players. Player 1 has two actions: either have quiche (Q) or have beer (B) for breakfast. Player 2 observes the breakfast of player 1 and then decides to duel (D) or not to duel (N) with player 1. The payoffs are as follows. If player 1 is weak and eats quiche then D and N give him payoffs of 1 and 3, respectively; if he is weak and drinks beer, then these payoffs are 0 and 2, respectively. If player 1 is strong, then the payoffs are 0 and 2 from D and N, respectively, if he eats quiche; and 1 and 3 from D and N, respectively, if he drinks beer. Player 2 has payoff 0 from not duelling, payoff 1 from duelling with the weak player 1, and payoff -1 from duelling with the strong player 1.

(a) Draw a diagram modelling this situation.
(b) Compute all the pure strategy Nash equilibria of the game. Find out which of these Nash equilibria are perfect Bayesian equilibria. Give the corresponding beliefs and determine whether these equilibria are pooling or separating, and which ones satisfy the intuitive criterion.

5.8. *Issuing Stock*
In this story the players are a manager (M) and an existing shareholder (O). The manager is informed about the current value of the firm, a, and the NPV (net present value) of a potential investment opportunity, b, but the shareholder only knows that high values and low values each have probability $1/2$. More precisely, either $(a, b) = (\overline{a}, \overline{b})$ or $(a, b) = (\underline{a}, \underline{b})$, each with probability $1/2$, where $\underline{a} < \overline{a}$ and $\underline{b} < \overline{b}$. The manager moves first and either proposes to issue new stock E (where E is fixed) to undertake the investment opportunity, or decides not to issue new stock. The existing shareholder decides whether to approve of the new stock issue or not. The manager always acts in the interest of the existing shareholder: their payoffs in the game are always equal.

If the manager decides not to issue new stock, then the investment opportunity is foregone, and the payoff is either $\overline{a}$ or $\underline{a}$. If the manager proposes to issue new stock but this is not approved by the existing shareholder, then again the investment opportunity is foregone and the payoff is either $\overline{a}$ or $\underline{a}$. If the manager proposes to issue new stock E and the existing shareholder approves of this, then the payoff to the existing shareholder is equal to $[M/(M+E)](\overline{a}+\overline{b}+E)$ in the good state $(\overline{a},\overline{b})$ and $[M/(M+E)](\underline{a}+\underline{b}+E)$ in the bad state $(\underline{a},\underline{b})$; here, $M = (1/2)[\overline{a}+\overline{b}] + (1/2)[\underline{a}+\underline{b}]$ is the price of the existing shares if the investment is undertaken.

(a) Set up the extensive form of this signaling game.
(b) Take $\overline{a} = 150$, $\underline{a} = 50$, $\overline{b} = 20$, $\underline{b} = 10$, and $E = 100$. Compute the pure strategy perfect Bayesian equilibria of this game. Are they pooling, separating? How about the intuitive criterion? Try to interpret the results from an economic point of view.
(c) Repeat the analysis of (b) for $\overline{b} = 100$.

5.9. *More Signaling Games*

(a) Consider the signaling game in Fig. 5.4, but with payoffs $(1,2)$ instead of $(4,2)$ if type t of player 1 plays R and player 2 plays u'. Show that this game has only pooling equilibria. Which ones survive the IC?
(b) Compute the pure strategy perfect Bayesian equilibria and test for the intuitive criterion in the signaling game in Fig. 5.8.
(c) Consider the signaling game in Fig. 5.9, where the chance move is not explicitly drawn in order to keep the diagram simple. Compute the pure strategy perfect Bayesian equilibria and test for the intuitive criterion.

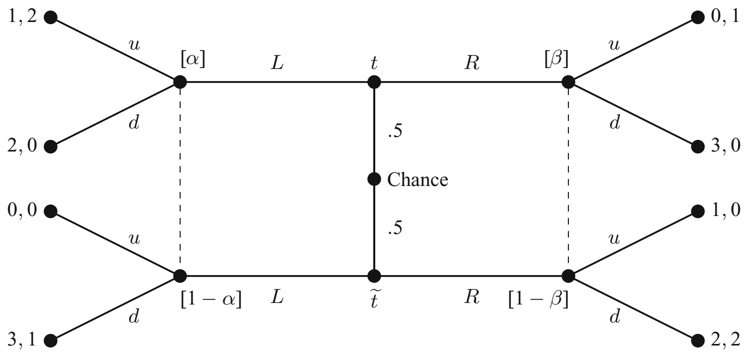

Fig. 5.8 The signaling game of Problem 5.9(b)

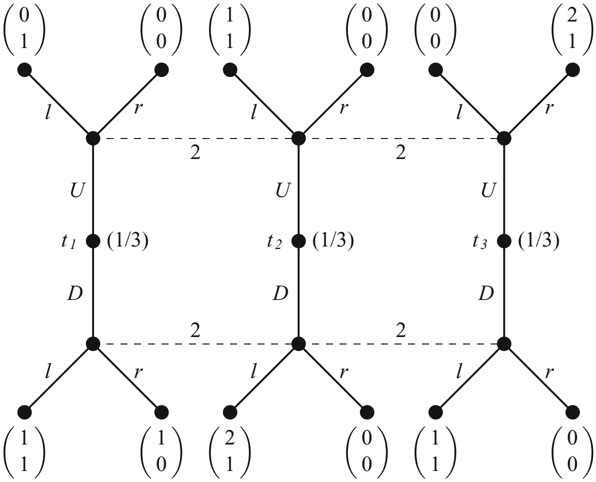

Fig. 5.9 The signaling game of Problem 5.9(c). Each type of player 1 has probability $1/3$

5.5 Notes

The type terminology and corresponding theory is due to Harsanyi (1967/1968). The Battle of the Sexes examples in Sect. 5.2 are taken from Osborne (2004). One of the first examples of a signaling game is the Spence (1973) job market signaling model (see Problem 5.4). The intuitive criterion is due to Cho and Kreps (1987).

Problem 5.5 is taken from Rasmusen (1989). The beer-quiche game of Problem 5.7 is from Cho and Kreps (1987). Problem 5.8 is based on Myers and Majluf (1984).

References

Cho, I. K., & Kreps, D. M. (1987). Signalling games and stable equilibria. *Quarterly Journal of Economics, 102*, 179–221.

Harsanyi, J. C. (1967/1968). Games with incomplete information played by “Bayesian” players I, II, and III. *Management Science, 14*, 159–182, 320–334, 486–502.

Myers, S. C., & Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics, 13*, 187–221.

Osborne, M. J. (2004). *An introduction to game theory*. New York: Oxford University Press.

Rasmusen, E. (1989). *Games and information: An introduction to game theory* (2nd ed., 1994/1995). Oxford: Basil Blackwell.

Spence, A. M. (1973). Job market signalling. *Quarterly Journal of Economics, 87*, 355–374.