Finite Extensive Form Games

4

Most games derived from economic or political situations have in common with most parlor games (like card games and board games) that they are not 'one-shot': players move sequentially, and one and the same player may move more often than once. Such games are best described by drawing a decision tree which tells us whose move it is and what a player's information is when that player has to make a move.

In this chapter these 'games in extensive form' are studied. Attention is restricted to games with finitely many players (usually two), finitely many decision moments and finitely many moves. See Sect. 1.3.3 for a few examples. We also assume that each player has 'complete information': this means that either there is no chance move in the game or, if there is one, there is no player who still has to move after the chance move and is not informed about the outcome of the chance move. This excludes, for instance, the game of entry deterrence with incomplete information in Sect. 1.3.3, but it also excludes most card games—most card games start with cards being dealt, which is a chance move. For the analysis of games with incomplete information see Chap. 5. Chapter 14 extends the analysis of the present and the next chapter.

The first section of this chapter introduces games in extensive form. In order to avoid a load of cumbersome notation the treatment will be somewhat informal but—hopefully—not imprecise. In Sect. 4.2 we define strategies and the 'strategic form' of a game: the definition of Nash equilibrium for extensive form games is then practically implied. The focus in this chapter is on pure Nash equilibrium.

In the third section the concept of Nash equilibrium is refined by considering backward induction and subgame perfection. A further important refinement, called 'perfect Bayesian equilibrium', is treated in the fourth section.

H. Peters, *Game Theory*, Springer Texts in Business and Economics,
DOI 10.1007/978-3-662-46950-7_4

4.1 The Extensive Form

A *game in extensive form* is described by a *game tree*. Such a game tree is characterized by *nodes* and *edges*. Each node is either a *decision node* of a player, or a *chance node*, or an *end node*. Each edge corresponds to either an *action* of a player or a choice made by chance, sometimes called a 'move of Nature'.

Figure 4.1 illustrates these and other concepts.

The upper node in the tree, the *root* of the tree, is a decision node of player 1 and the starting point of the game. Player 1 chooses between three actions, namely A, B, and C. Player 2 learns that player 1 has chosen either one of the actions A and B, or action C. The first event is indicated by the dashed line connecting the two left decision nodes of player 2. In that case, player 2 has two actions, namely l and r. We call the two connected nodes an *information set* of player 2: player 2 knows that the play of the game has arrived at one of these nodes but he does not know at which one. The fact that player 2 has the same set of actions at each of the two nodes in this information set is a necessary consequence: if this were not the case, player 2 would know at which node he was (i.e., would know whether player 1 would have played A or B) by simply examining the set of available actions, which would go against the interpretation of an information set. This last argument is one consequence of the more general assumption that the whole game tree is common knowledge between the players: each player knows it, knows that the other player(s) know(s) it, knows that the other player(s) know(s) that he knows it, and so on.

If player 1 plays C, then there is a chance move, resulting with probability $1/4$ in a decision node of player 2 (following U) and with probability $3/4$ in a decision node of player 1 (following D). At player 2's decision node this player has two actions, namely L and R. At player 1's decision node this player also has two actions, namely a and b. All the remaining nodes are end nodes, indicated by payoff pairs, where the upper number is the payoff to player 1 and the lower number the payoff to player 2. In this diagram, the payoffs are written as column vectors, but we also write them as row vectors, whatever is convenient in a given situation.

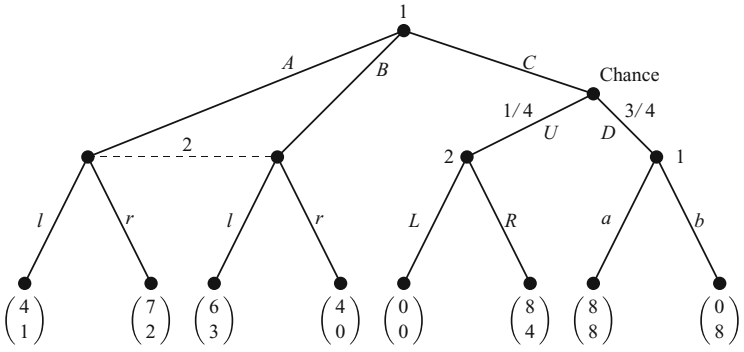

Fig. 4.1 A game in extensive form

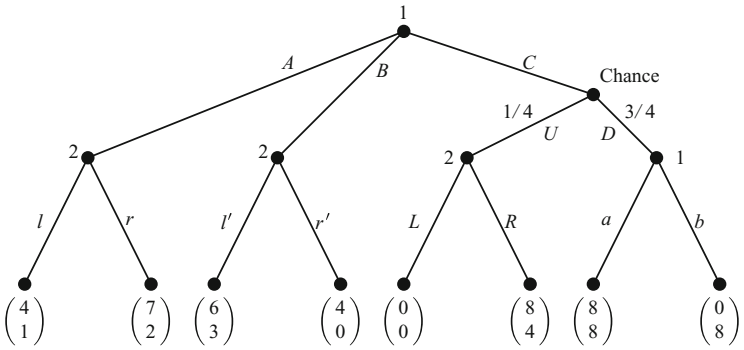

Fig. 4.2 The game of Fig. 4.1, now with perfect information

Also the singleton decision nodes are called information sets. Thus, in this game, each player has two different information sets. An information set is *nontrivial* if it consists of at least two nodes. Games with nontrivial information sets are called games with *imperfect information*. If a game has only trivial information sets, then we say that it has *perfect information*. If we change the present example by assuming that player 2 observes whether player 1 chooses A or B, then the game has perfect information. See Fig. 4.2. We have renamed the actions of player 2 after action B of player 1 for later convenience: these actions can now be regarded as different from player 2's actions following A, since player 2 knows that player 1 has chosen B.

The chance move in our example is not a particularly interesting one, since the players learn what the outcome of the chance move is. (The situation is different if at least one player is not completely informed about the outcome of a chance move and if this lack of information has strategic consequences. In that case, we talk about games with 'incomplete' information, see Chap. 5.)

As mentioned before, we do not give a formal definition of a *game in extensive form*: the examples in Figs. 4.1 and 4.2 illustrate the main ingredients of such a game.[1] An important condition is that the game tree should be a *tree* indeed: it should have a single root and no 'cycles'. This means that a situation as for instance in Fig. 4.3a is not allowed.

We also restrict attention to games in extensive form that have *perfect recall*: each player remembers what he did in the past. For instance, the situation in Fig. 4.3b, where player 1 at his lower information set does not recall which action he took earlier, is not considered.[2]

[1] See Chap. 14 for a formal definition.

[2] The assumption of perfect recall plays a particular role for the relation between mixed and behavioral strategies, see Chap. 14.

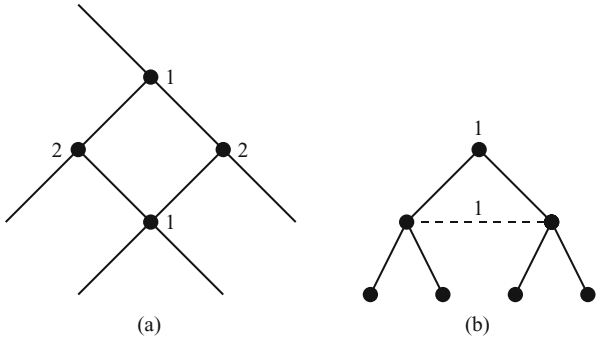

Fig. 4.3 An example of a cycle (**a**), and of a game without perfect recall (**b**)

4.2 The Strategic Form

In a game in extensive form, it is extremely important to distinguish between actions and strategies. An *action* is a possible move of a player at an information set. In the games in Figs. 4.1 and 4.2 player 1 has the actions A, B, and C, and a and b; and player 2 has the actions l and r, and L and R. In contrast,

a *strategy* is a *complete plan to play the game*.

This is one of the most important concepts in game theory. In the games in Figs. 4.1 and 4.2, a possible strategy for player 1 is:

Start by playing C; if the chance move of the game results in D, then play b.

Another strategy of player 1 is:

Start by playing A; if the chance move of the game results in D, then play b.

The last strategy might look strange since player 1's first action A precludes him having to take any further action. Nevertheless, also this plan is regarded as a possible strategy.[3]

A possible strategy for player 2 in the game of Fig. 4.1 is:

Play l if player 1 plays A or B, and play L if player 1 plays C and the chance move results in U.

Note that player 2 cannot make his action contingent on whether player 1 plays A or B, since player 2 does not have that information. In the perfect information game

[3] Although there is not much lost if we would exclude such strategies—as some authors do.

of Fig. 4.2, however, player 2's strategy should tell what player 2 plays after A and what he plays after B. A possible strategy would then be:

> Play l if player 1 plays A, play r' if player 1 plays B, and play L if player 1 plays C and the chance move results in U.

A formal definition of a strategy of a player is:

> a *strategy* is *a list of actions, exactly one at each information set of that player*.

In both our examples, a strategy of player 1 is therefore a list of two actions since player 1 has two information sets. The number of possible strategies of player 1 is the number of different lists of actions. Since player 1 has three possible actions at his first information set and two possible actions at his second information set, this number is equal to $3 \times 2 = 6$. The strategy set of player 1 can be denoted as

$$\{Aa, Ab, Ba, Bb, Ca, Cb\} .$$

Similarly, in the imperfect information game in Fig. 4.1 player 2 has $2 \times 2 = 4$ different strategies, and his strategy set can be denoted as

$$\{lL, lR, rL, rR\} .$$

In the perfect information game in Fig. 4.2 player 2 has three information sets and two actions at each information set, so $2 \times 2 \times 2 = 8$ different strategies, and his strategy set can be denoted as

$$\{ll'L, ll'R, lr'L, lr'R, rl'L, rl'R, rr'L, rr'R\} .$$

In general, it is important to distinguish between a strategy combination and the associated *outcome*. The outcome is the induced play of the game, that is, the path followed in the game tree. For instance, the strategy combination $(Aa, ll'L)$ in the game in Fig. 4.2 induces the outcome (A, l) and the payoffs $(4, 1)$.

There are several reasons why we are interested in strategies. The main reason is that by considering strategies the extensive form game is effectively reduced to a one-shot game. Once we fix a profile (in the present example, pair) of strategies we can compute the payoffs by following the path followed in the game tree. Consider for instance the strategy pair (Cb, rL) in the game in Fig. 4.1. Then player 1 starts by playing C, and this is followed by a chance move; if the result of this move is U, then player 2 plays L; if the result is D, then player 1 plays b. Hence, with probability $1/4$ the resulting payoff pair is $(0, 0)$ and with probability $3/4$ the resulting payoff pair is $(0, 8)$. So the expected payoffs are 0 for player 1 and 6 for player 2. In this way, we can compute the payoffs in the game of Fig. 4.1 resulting from each of the 6×4 possible strategy combinations. Similarly, for the game in Fig. 4.2 we compute

$$\begin{array}{c} \\ Aa \\ Ab \\ Ba \\ Bb \\ Ca \\ Cb \end{array} \begin{array}{c} lL \\ \left(\begin{array}{c} 4,1 \\ 4,1 \\ \underline{6},\underline{3} \\ \underline{6},\underline{3} \\ \underline{6},6 \\ 0,6 \end{array}\right. \end{array} \begin{array}{c} lR \\ \begin{array}{c} 4,1 \\ 4,1 \\ 6,\underline{3} \\ 6,\underline{3} \\ \underline{8},\underline{7} \\ 2,\underline{7} \end{array} \end{array} \begin{array}{c} rL \\ \begin{array}{c} \underline{7},\underline{2} \\ \underline{7},\underline{2} \\ 4,0 \\ 4,0 \\ 6,6 \\ 0,6 \end{array} \end{array} \begin{array}{c} rR \\ \left.\begin{array}{c} 7,\underline{2} \\ 7,\underline{2} \\ 4,0 \\ 4,0 \\ \underline{8},\underline{7} \\ 2,\underline{7} \end{array}\right) \end{array}$$

$$\begin{array}{c} \\ Aa \\ Ab \\ Ba \\ Bb \\ Ca \\ Cb \end{array} \begin{array}{c} ll'L \\ \left(\begin{array}{c} 4,1 \\ 4,1 \\ \underline{6},\underline{3} \\ \underline{6},\underline{3} \\ \underline{6},6 \\ 0,6 \end{array}\right. \end{array} \begin{array}{c} ll'R \\ \begin{array}{c} 4,1 \\ 4,1 \\ 6,\underline{3} \\ 6,\underline{3} \\ \underline{8},\underline{7} \\ 2,\underline{7} \end{array} \end{array} \begin{array}{c} lr'L \\ \begin{array}{c} 4,1 \\ 4,1 \\ 4,0 \\ 4,0 \\ \underline{6},6 \\ 0,6 \end{array} \end{array} \begin{array}{c} lr'R \\ \begin{array}{c} 4,1 \\ 4,1 \\ 4,0 \\ 4,0 \\ \underline{8},\underline{7} \\ 2,\underline{7} \end{array} \end{array} \begin{array}{c} rl'L \\ \begin{array}{c} \underline{7},\underline{2} \\ \underline{7},\underline{2} \\ 6,\underline{3} \\ 6,\underline{3} \\ 6,6 \\ 0,6 \end{array} \end{array} \begin{array}{c} rl'R \\ \begin{array}{c} 7,\underline{2} \\ 7,\underline{2} \\ 6,\underline{3} \\ 6,\underline{3} \\ \underline{8},\underline{7} \\ 2,\underline{7} \end{array} \end{array} \begin{array}{c} rr'L \\ \begin{array}{c} \underline{7},\underline{2} \\ \underline{7},\underline{2} \\ 4,0 \\ 4,0 \\ 6,6 \\ 0,6 \end{array} \end{array} \begin{array}{c} rr'R \\ \left.\begin{array}{c} 7,\underline{2} \\ 7,\underline{2} \\ 4,0 \\ 4,0 \\ \underline{8},\underline{7} \\ 2,\underline{7} \end{array}\right) \end{array}$$

Fig. 4.4 The 6×4 strategic form of the game in Fig. 4.1 and the 6×8 strategic form of the game in Fig. 4.2

6×8 payoff pairs. We next write these payoff pairs in a bimatrix, as in Chap. 3. The resulting bimatrix games are presented in Fig. 4.4.

Such a bimatrix game is called the *strategic form* of the extensive form game. The definition of Nash equilibrium of an extensive form game is then almost implied:

> A *Nash equilibrium* of a game in extensive form is a Nash equilibrium of the strategic form.

This definition holds for pure Nash equilibria and, more generally, Nash equilibria in mixed strategies, but in this chapter we restrict attention to pure strategies and pure strategy Nash equilibria.

The pure strategy Nash equilibria of the bimatrix games in Fig. 4.4 can be found by using the method of Sect. 3.2.1. The equilibria correspond to the double underlined entries. Thus, the imperfect information game has six different Nash equilibria in pure strategies, and the perfect information game has ten different Nash equilibria in pure strategies.

In the next two sections we examine these Nash equilibria more closely and discuss ways to distinguish between them.

4.3 Backward Induction and Subgame Perfection

We first consider the perfect information game of Fig. 4.2. This game can be analyzed using the principle of *backward induction*. This means that we start with the nodes preceding the end nodes, and turn them into end nodes with payoffs resulting from choosing the optimal action(s). Specifically, player 2 chooses r after

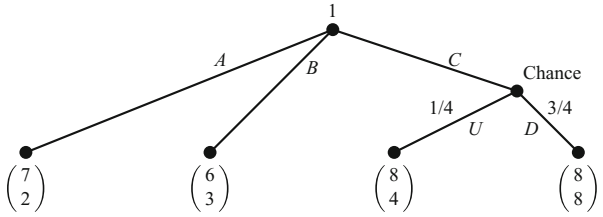

Fig. 4.5 The reduced game of Fig. 4.2

A of player 1 and l' after B of player 1; and R after C and U. Player 1 chooses a after D. Thus, we obtain the reduced game of Fig. 4.5. Note that player 2's strategy has already been completely determined by this first step: it is the strategy $rl'R$. Player 1 has chosen a at his lower information set. Next, in this reduced game, player 1 chooses the action(s) that yield(s) the highest payoff. Since A yields a payoff of 7, B a payoff of 6, and C a(n expected) payoff of $\frac{1}{4} \cdot 8 + \frac{3}{4} \cdot 8 = 8$, it is optimal for player 1 to choose C. Hence, we obtain the strategy combination $(Ca, rl'R)$ with payoffs $\frac{1}{4}(8, 4) + \frac{3}{4}(8, 8) = (8, 7)$. This is one of the ten Nash equilibria of the game (see Fig. 4.4). It is called *backward induction equilibrium.*

It can be shown that, in a game of perfect information, *applying the backward induction principle always results in a (pure) Nash equilibrium.* As a by-product, we obtain that *a game of perfect information has at least one Nash equilibrium in pure strategies, which can be obtained by backward induction.*

It is illustrative to consider backward induction equilibrium—in this game: $(Ca, rl'R)$—as opposed to the induced backward induction outcome. The latter refers to the actual play of the game or, equivalently, the equilibrium path, in this case (Ca, R). Observe that there are other Nash equilibria in this game that generate the same outcome or path, namely $(Ca, ll'R)$, $(Ca, lr'R)$, and $(Ca, rr'R)$: they all generate the path (Ca, R), but differ in the left part of the tree, where player 2 makes at least one decision that is not optimal. Hence, the principle of backward induction ensures that every player always takes an optimal action, even in parts of the game tree that are not actually reached when the game is played.

A generalization of backward induction is *subgame perfection*. The definition of a subgame is as follows:

> A *subgame* is any part of the game tree, starting at an information set consisting of a single decision node (trivial information set) of a player or at a chance node, and which is not connected to the rest of the tree by any later information set.

The last part of this definition is illustrated in Fig. 4.6. Although player 2's information set consists of a single decision node, no subgame starts there: this would not make sense since player 3 does not observe any earlier move and in particular does not know at which of the three nodes in his information set he is. Formally, the subtree starting at player 2's decision node is still connected to the whole tree by player 3's information set.

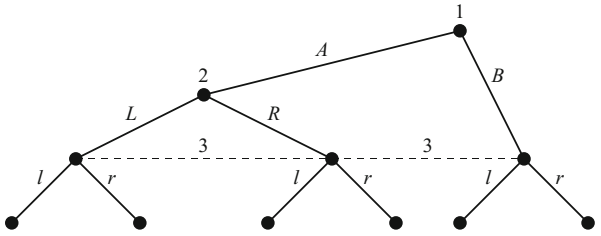

Fig. 4.6 No subgame starts at player 2's decision node

The game in Fig. 4.2 has six different subgames, namely: the entire game; the game starting from the chance move; and the four games starting from the four nodes preceding the end nodes. The definition of a subgame perfect equilibrium is as follows:

> A *subgame perfect equilibrium* is a strategy combination that induces a Nash equilibrium in every subgame.

To see what this means, consider again the game in Fig. 4.2. In order for a strategy combination to be a subgame perfect equilibrium, it has to induce a Nash equilibrium in every subgame. Since the entire game is a subgame, a subgame perfect equilibrium has to be a Nash equilibrium in the entire game, and, thus, the ten Nash equilibria in this game are the candidates for a subgame perfect equilibrium. This is the case for any arbitrary game, and therefore *a subgame perfect equilibrium is always a Nash equilibrium.*

Returning to the game in Fig. 4.2, a subgame perfect equilibrium also has to induce an equilibrium in each of the four one-player subgames preceding the end nodes: although we have not defined Nash equilibria for one-person games, the only reasonable definition is that a player should choose the action that is optimal. In the example, this means that (from left to right) the actions r, l', R, and a, should be chosen. This implies that the players choose optimally also in the subgame starting from the chance node. Summarizing, we look for the Nash equilibrium or equilibria that generate the mentioned actions, and the only Nash equilibrium that does this is again $(Ca, rl'R)$. Hence, the unique subgame perfect equilibrium in this game is $(Ca, rl'R)$. It is not surprising that this is also the backward induction equilibrium: *in games of perfect information, backward induction equilibria and subgame perfect equilibria coincide.*

Let us now consider the imperfect information version of the game in Fig. 4.1. In this case, backward induction cannot be applied to the left part of the game tree: since player 2 does not know whether player 1 has played A or B when he has to choose an action in his left information set, he cannot optimally choose between l and r: l is better if player 1 has played B, but r is better if player 1 has played A. Considering subgame perfection, the only subgames are now: the entire game; the two subgames following U and D; and the subgame starting from the chance move.

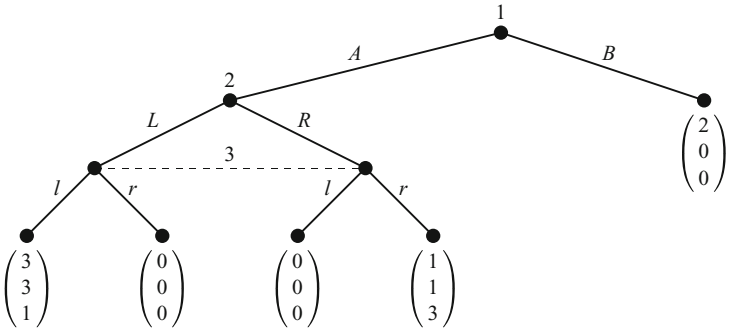

Fig. 4.7 A three-player game

Hence, the restrictions imposed by subgame perfection are that player 1 should play a, player 2 should play R, and the strategy combination should be a Nash equilibrium in the entire game. Of the six Nash equilibria of the game (see Fig. 4.4), this leaves the two equilibria (Ca, lR) and (Ca, rR). So these are the subgame perfect equilibria of the game in Fig. 4.1.

We conclude this section with an example which shows more explicitly than the preceding example that subgame perfection can be more generally applied than the backward induction principle. Consider the game in Fig. 4.7, which is a three player game (for a change). Clearly, backward induction cannot be applied here: player 3 cannot unambiguously determine his optimal action since he does not know whether player 2 has played L or has played R. For subgame perfection, notice that this game has two subgames: the entire game; and the game starting with player 2's decision node. The latter game is a game between players 2 and 3 with strategic form

$$\begin{array}{c} \\ L \\ R \end{array} \begin{array}{c} \begin{array}{cc} l & r \end{array} \\ \begin{pmatrix} 3,1 & 0,0 \\ 0,0 & 1,3 \end{pmatrix} \end{array}$$

where player 2 is the row player and player 3 the column player. This game has two pure Nash equilibria, namely (L, l) and (R, r). Hence, a subgame perfect equilibrium has to induce one of these two equilibria in this subgame. Note that if the first equilibrium is played, then player 1 should play A, yielding him a payoff of 3 rather than the payoff of 2 obtained by playing B. If the other equilibrium is played in the subgame, then player 1 should obviously play B since A now yields only 1. So the two subgame perfect equilibria are (A, L, l) and (B, R, r).

Alternatively, one can first compute the (pure) Nash equilibria of the entire game. The strategic form of the game can be represented as follows, where the left matrix results from player 1 playing A and the right matrix from player 1 playing B.

$$1:A \quad \begin{matrix} & l & r \\ L & \underline{3},\underline{3},\underline{1} & 0,0,0 \\ R & 0,0,0 & 1,\underline{1},\underline{3} \end{matrix} \qquad 1:B \quad \begin{matrix} & l & r \\ L & 2,\underline{0},\underline{0} & \underline{2},\underline{0},\underline{0} \\ R & \underline{2},\underline{0},\underline{0} & \underline{2},\underline{0},\underline{0} \end{matrix}.$$

Best replies are underlined (for player 1 one has to compare the corresponding payoffs over the two matrices), and the pure Nash equilibria are (A, L, l), (B, L, r), (B, R, l), and (B, R, r). The subgame perfect equilibria are those where the combination (L, l) or (R, r) is played, resulting in the two equilibria found above.

4.4 Perfect Bayesian Equilibrium

A further refinement of Nash equilibrium and of subgame perfect equilibrium is provided by the concept of 'perfect Bayesian equilibrium'. Consider an information set of a player in an extensive form game. A *belief* of that player on that information set is a probability distribution over the nodes of that information set or, equivalently, over the actions leading to that information set. Of course, if the information set is trivial (consists of a single node) then also the belief is trivial, namely attaching probability 1 to the unique node. Our somewhat informal definition of a perfect Bayesian equilibrium is as follows.

> A *perfect Bayesian equilibrium* in an extensive form game is a combination of strategies and a specification of beliefs such that the following two conditions are satisfied:
>
> (i) *Bayesian consistency*: the beliefs are consistent with the strategies under consideration;
> (ii) *sequential rationality*: the players choose optimally given their beliefs.

The first condition says that the beliefs should satisfy Bayesian updating with respect to the strategies whenever possible. The second condition says that a player should maximize his expected payoff given his beliefs. In order to see what these conditions mean exactly, we consider some examples. (Formal definitions are provided in Chap. 14.)

Consider the game in Fig. 4.8, which is identical to the game in Fig. 4.1. This game has one nontrivial information set. Suppose player 2's belief at this information set is given by the probabilities α at the left node and $1-\alpha$ at the right node, where $0 \leq \alpha \leq 1$. That is, if this information set is reached then player 2 attaches probability α to player 1 having played A and probability $1-\alpha$ to player 1 having played B. All the other information sets are trivial and therefore the beliefs attach probability 1 to each of the corresponding nodes. Condition (ii), sequential rationality, means that player 2 should choose R and player 1 should choose a at the corresponding (trivial) information sets. At the nontrivial information set, player 2 should choose the action that maximizes his expected payoff, *given his belief at this information set*. The expected payoff from l is equal to $\alpha \cdot 1 + (1-\alpha) \cdot 3 = 3 - 2\alpha$ and the expected payoff from r is $\alpha \cdot 2 + (1-\alpha) \cdot 0 = 2\alpha$. Hence, l is optimal if $3 - 2\alpha \geq 2\alpha$, i.e., if $\alpha \leq 3/4$ and r is optimal if $\alpha \geq 3/4$.

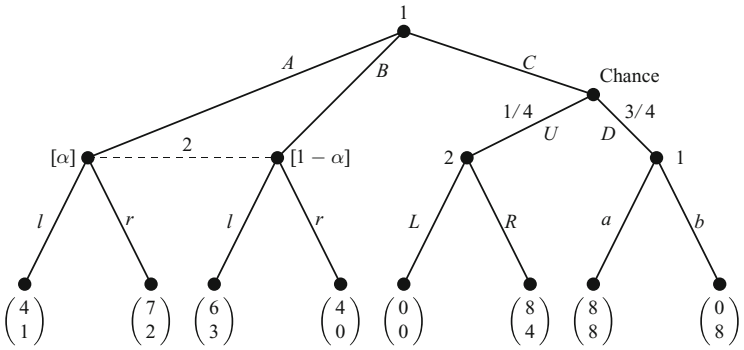

Fig. 4.8 The extensive form game of Fig. 4.1

What does condition (i), Bayesian consistency, imply for this game? It says that, whenever possible, the belief of player 2 at the left, nontrivial information set should be computed by Bayesian updating using the strategy of player 1, specifically, using the initial action of player 1. Thus, if the strategy of player 1 prescribes action A, then player 2 should indeed believe this, so $\alpha = 1$. Similarly, if the strategy of player 1 prescribes action B, then $1 - \alpha = 1$ so $\alpha = 0$. The formula behind this is the formula for conditional probability:

$$\begin{aligned}\alpha &= \text{Prob}[A \text{ is played} \mid A \text{ or } B \text{ is played}] \\ &= \frac{\text{Prob}[A \text{ and } (A \text{ or } B) \text{ is played}]}{\text{Prob}[A \text{ or } B \text{ is played}]} \\ &= \frac{\text{Prob}[A \text{ is played}]}{\text{Prob}[A \text{ or } B \text{ is played}]}\end{aligned}$$

where these probabilities should be computed given the strategy of player 1. Hence, indeed, if the strategy of player 1 prescribes A then $\alpha = 1/1 = 1$ and if the strategy of player 1 prescribes B then $\alpha = 0/1 = 0$. If, however, player 1's strategy prescribes C, then Prob[A or B is played] $= 0$ and we cannot use the formula for computing α: the belief of player 2 on his nontrivial information set is undetermined or *free*.

Remark 4.1 We do not consider mixed strategies in this chapter, but suppose we did and consider, for instance, the mixed strategy where player 1 plays Aa with probability $\frac{1}{2}$, Ab with probability $\frac{1}{6}$, and Ba, Bb, Ca and Cb each with probability $\frac{1}{12}$. Then we would have $\alpha = (\frac{1}{2} + \frac{1}{6})/(\frac{1}{2} + \frac{1}{6} + 2 \cdot \frac{1}{12}) = \frac{4}{5}$. In words, given player 1's strategy, if it turns out that player 1 has played A or B, then player 2 should believe that player 1 has played A with probability $\frac{4}{5}$ and B with probability $\frac{1}{5}$. □

To complete the analysis of this game, note that it is always optimal for player 1 to play C, given the actions R and a following the chance move: C yields 8 to player 1 whereas A or B yield at most 7. But if player 1 does not play A or B, then, as we have just seen, Bayesian consistency [condition (i)] does not put any restriction on the belief α of player 2, since the nontrivial information set of player 2 is reached with zero probability. This means that α is free, but given α player 2 should choose optimally, as computed before. Hence, we have essentially two perfect Bayesian equilibria, namely (Ca, lR) with beliefs $\alpha \leq 3/4$ and (Ca, rR) with beliefs $\alpha \geq 3/4$. Note that these are also the subgame perfect equilibria, now 'backed up' by a belief of player 2 on his nontrivial information set.

Remark 4.2 As already indicated, a perfect Bayesian equilibrium is also a subgame perfect equilibrium, but in order to show this a more detailed formal definition is required. See Chap. 14. □

In order to show that the perfect Bayesian equilibrium requirement can have an additional impact compared to subgame perfection, consider the variation on the game of Fig. 4.1, obtained by replacing the payoffs $(4, 1)$ after A and l by the payoffs $(4, 3)$, resulting in the game in Fig. 4.9.

One can check that the subgame perfect equilibria are still (Ca, lR) and (Ca, rR). Obviously, a rational player 2 would never play r at his nontrivial information set since l is always better, but subgame perfection does not rule this out. But clearly, there is no belief that player 2 could have at this information set that would make r optimal: if we denote player 2's belief by $(\alpha, 1-\alpha)$ as before, then r yields 2α whereas l yields $3\alpha + 3(1-\alpha) = 3$, which is always larger than 2α. Hence, the only perfect Bayesian equilibrium is (Ca, lR), with arbitrary free belief of player 2 at his nontrivial information set.

We conclude with another example. Consider again the game of Fig. 4.7, reproduced in Fig. 4.10 with belief $(\alpha, 1-\alpha)$ attached to the nodes in the information set of player 3. There are two ways to find the perfect Bayesian equilibria of this

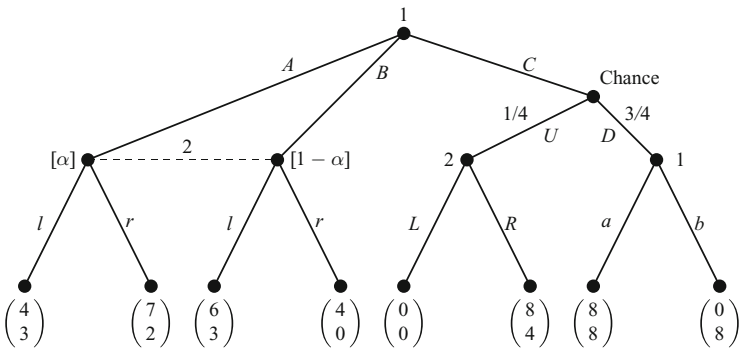

Fig. 4.9 The extensive form game of Fig. 4.1 with payoffs (4,3) after A and l

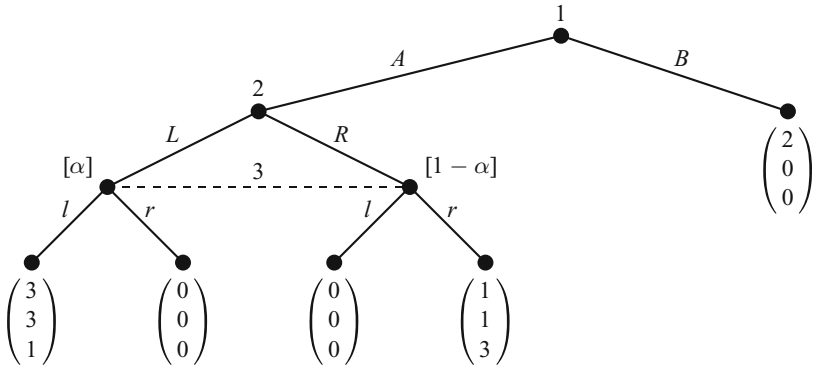

Fig. 4.10 The three-player game of Fig. 4.7 with belief of player 3

game. One can consider the subgame perfect equilibria and find appropriate beliefs. Alternatively, one can start from scratch and apply a form of backward induction. To illustrate the last method, start with player 3. If player 3 plays l then his (expected) payoff is $1\alpha + 0(1-\alpha) = \alpha$. If player 3 plays r then his (expected) payoff is $0\alpha + 3(1-\alpha) = 3 - 3\alpha$. Therefore, l is optimal if $\alpha \geq 3/4$ and r is optimal if $\alpha \leq 3/4$. In fact, we have just applied condition (ii), sequential rationality, in the definition of a perfect Bayesian equilibrium.

Now suppose player 3 plays l. Then it is optimal for player 2 to play L. If player 2 plays L, then condition (i) in the definition of a perfect Bayesian equilibrium, Bayesian consistency, implies $\alpha = 1$: that is, player 3 should indeed believe that player 2 has played L. Since $1 \geq 3/4$, l is the optimal action for player 3. Player 1, finally, should play A, yielding payoff 3 instead of the payoff 2 resulting from B. So we have a perfect Bayesian equilibrium (A, L, l) with belief $\alpha = 1$.

If player 3 plays r, then it is optimal for player 2 to play R, resulting in $\alpha = 0$ by Bayesian consistency; since $0 \leq 3/4$, r is the optimal action for player 3. In this case, player 1 should play B: this yields payoff 2, whereas A yields only 1. Hence, we have a perfect Bayesian equilibrium (B, R, r) with belief $\alpha = 0$.

4.5 Problems

4.1. *Counting Strategies*
Consider the following simplified chess game. White moves first (in accordance with the usual rules). Black observes White's move and then makes its move. Then the game is over and ends in a draw. Determine the strategy sets of White and Black. How many strategies does Black have? [In case you do not know the rules of chess: the only information you need is that both players have 20 possible actions.]

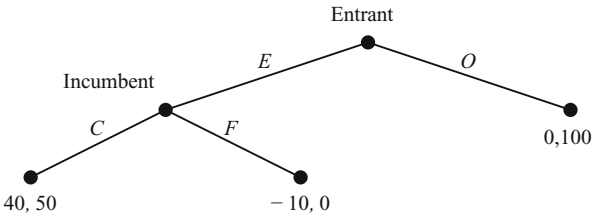

Fig. 4.11 Entry deterrence, Problem 4.3

4.2. *Extensive vs. Strategic Form*
Each game in extensive form leads to a unique game in strategic form. The converse, however, is not true. Consider the following bimatrix game:

$$\begin{pmatrix} 1,0 & 1,0 & 0,1 & 0,1 \\ 2,0 & 0,2 & 2,0 & 0,2 \end{pmatrix}.$$

Find an extensive form game with perfect information and an extensive form game with imperfect information, both having this bimatrix game as their strategic form game.

4.3. *Entry Deterrence*
Consider the entry deterrence game of Chap. 1, of which the extensive form is reproduced in Fig. 4.11.

(a) Write down the strategic form of this game.
(b) Determine the Nash equilibria (in pure strategies). Which one is the backward induction equilibrium? Which one is subgame perfect? In which sense is the other equilibrium based on an 'incredible threat'?

4.4. *Choosing Objects*
Four objects O_1, O_2, O_3, and O_4 have different worths for two players 1 and 2, given by the following table:

	O_1	O_2	O_3	O_4
Worth for player 1	1	2	3	4
Worth for player 2	2	3	4	1

Player 1 starts with choosing an object. After him player 2 chooses an object, then player 1 takes his second object, and finally player 2 gets the object that is left. The payoff for a player is the sum of the worths of the objects he obtains.

(a) Draw the decision tree for this extensive form game.
(b) How many strategies does each player have?

(c) Determine the backward induction or subgame perfect equilibria (in pure strategies). How many different subgame perfect equilibria are there? What are the associated outcomes, and the resulting divisions of the objects?
(d) Is there a Nash equilibrium in this game resulting in a division of the objects different from the division in a subgame perfect equilibrium?

4.5. *A Bidding Game*
Players 1 and 2 bid for an object that has value 2 for each of them. They both have wealth 3 and are not allowed to bid higher than this amount. Each bid must be a nonnegative integer amount. Besides bidding, each player, when it is his turn, has the options to pass (P) or to match (M) the last bid, where the last bid is set at zero at the beginning of the game. If a player passes (P), then the game is over and the other player gets the object and pays the last bid. If a player matches (M), then the game is over and each player gets the object and pays the last bid with probability $\frac{1}{2}$. Player 1 starts, and the players alternate until the game is over. Each new bid must be higher than the last bid.

(a) Draw the game tree of this extensive form game.
(b) How many strategies does player 1 have? Player 2?
(c) How many subgame perfect equilibria does this game have? What is (are) the possible subgame perfect equilibrium outcome(s)?
(d) Describe all (pure strategy) Nash equilibria of the game (do not make the strategic form). Is there any Nash equilibrium that does not result in a subgame perfect equilibrium outcome?

4.6. *An Extensive Form Game*
For the game in Fig. 4.12, write down the strategic form and compute all Nash equilibria, subgame perfect equilibria, and perfect Bayesian equilibria in pure strategies.

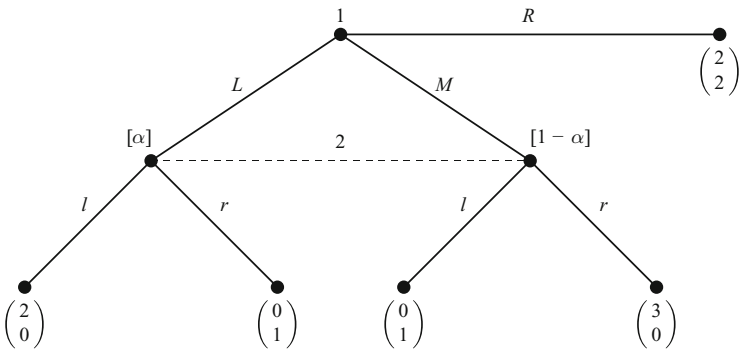

Fig. 4.12 Extensive form game of Problem 4.6

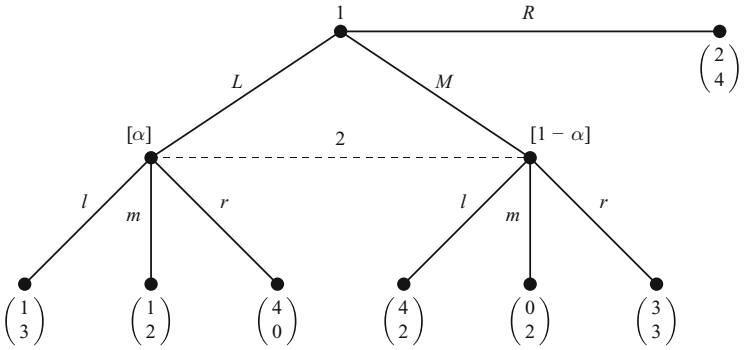

Fig. 4.13 Extensive form game of Problem 4.7

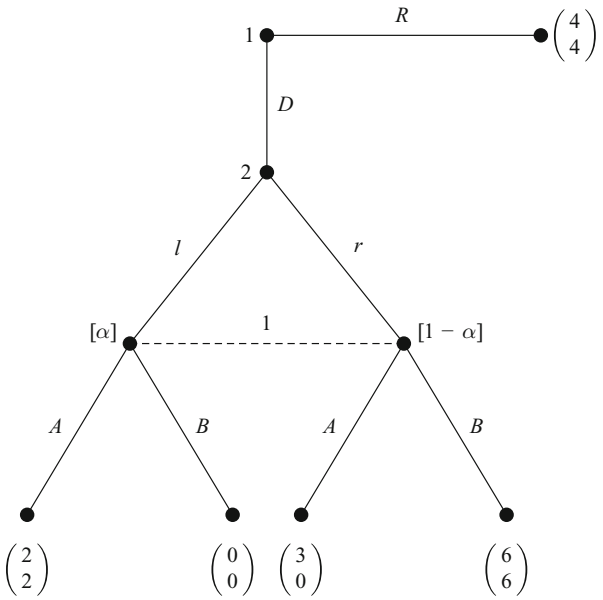

Fig. 4.14 Extensive form game of Problem 4.8

4.7. *Another Extensive Form Game*
For the game in Fig. 4.13, write down the strategic form and compute all Nash equilibria, subgame perfect equilibria, and perfect Bayesian equilibria in pure strategies.

4.8. *Still Another Extensive Form Game*
Consider the extensive form game in Fig. 4.14.

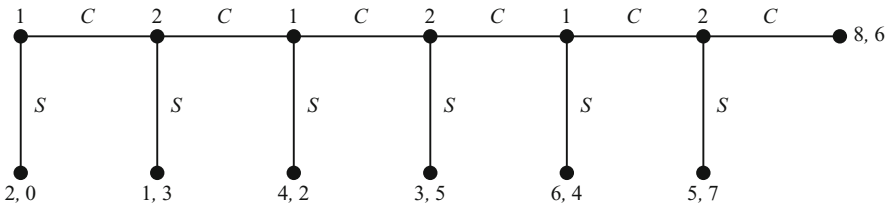

Fig. 4.15 The centipede game of Problem 4.9

(a) Determine the strategic form of this game.
(b) Determine all Nash equilibria in pure strategies.
(c) Determine all subgame perfect equilibria in pure strategies.
(d) Determine all perfect Bayesian equilibria in pure strategies.

4.9. *A Centipede Game*
In the centipede game, the two players move alternatingly. On each move, a player can stop (S) or continue (C). On any move, a player is better off stopping the game than continuing if the other player stops immediately afterward, but is worse off stopping than continuing if the other player continues, regardless of the subsequent actions. The game ends after a finite number of periods. Consider an example of this game in Fig. 4.15.

(a) Determine the backward induction or subgame perfect equilibrium of this game. What is the associated outcome?
(b) Show that there are other Nash equilibria, but that these always result in the same outcome as the subgame perfect equilibrium.

4.10. *Finitely Repeated Prisoners' Dilemma*
Consider the prisoners' dilemma game of Chap. 1:

$$\begin{array}{c} \\ C \\ D \end{array} \begin{array}{c} \begin{array}{cc} C & D \end{array} \\ \begin{pmatrix} -1,-1 & -10,0 \\ 0,-10 & -9,-9 \end{pmatrix} \end{array}.$$

Suppose that this game is played twice. After the first play of the game the players learn the outcome of that play. The final payoff for each player is the sum of the payoffs of the two stages.

(a) Write down the extensive form of this game. How many strategies does each player have? How many subgames does this game have?
(b) Determine the subgame perfect equilibrium or equilibria of this game. What if the game is repeated more than twice but still finitely many times?

4.11. *A Twice Repeated* 2×2 *Bimatrix Game*
Consider the following bimatrix game:

$$\begin{array}{c} \\ T \\ B \end{array} \begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} 2,1 & 1,0 \\ 5,2 & 4,4 \end{pmatrix} \end{array}.$$

Suppose that the game is played twice, and that after the first play of the game the players learn the outcome of that play. The final payoff for each player is the sum of the payoffs of the two stages.

(a) Determine the subgame perfect equilibrium or equilibria of this game. What if the game is repeated more than twice but still finitely many times?
(b) Exhibit a Nash equilibrium (of the twice repeated game) where (B, L) is played in the first round.

4.12. *Twice Repeated* 3×3 *Bimatrix Games*
Consider the following bimatrix game:

$$\begin{array}{c} \\ T \\ C \\ B \end{array} \begin{array}{c} \begin{array}{ccc} L & M & R \end{array} \\ \begin{pmatrix} 8,8 & 0,9 & 0,0 \\ 9,0 & 0,0 & 3,1 \\ 0,0 & 1,3 & 3,3 \end{pmatrix} \end{array}.$$

Suppose that the game is played twice, and that after the first play of the game the players learn the outcome of that play. The final payoff for each player is the sum of the payoffs of the two stages.

(a) How many strategies does each player have in this game? How many subgames does the game have?
(b) Describe a subgame perfect equilibrium in which (T, L) is played in the first round.

For question (c), consider the bimatrix game

$$\begin{array}{c} \\ T \\ C \\ B \end{array} \begin{array}{c} \begin{array}{ccc} L & M & R \end{array} \\ \begin{pmatrix} 5,3 & 0,0 & 2,0 \\ 0,0 & 2,2 & 0,0 \\ 0,0 & 0,0 & 0,0 \end{pmatrix} \end{array}.$$

(c) For the twice repeated version of this game, describe a subgame perfect equilibrium in which (B, R) is played in the first round.

4.6 Notes

The concept of subgame perfection was first formally introduced by Selten (1965, 1975). The result that each finite extensive form game of perfect information has a backward induction equilibrium in pure strategies is intuitive but nevertheless somewhat cumbersome to prove, see for instance Perea (2001, Chap. 3).

There is also a stronger version of Bayesian consistency, resulting in the concept of 'sequential equilibrium', see Kreps and Wilson (1982) and Chap. 14.

The game in Problem 4.12(c) is taken from Benoit and Krishna (1985).

References

Benoit, J.-P., & Krishna, V. (1985). Finitely repeatd games. *Econometrica, 53*, 905–922.

Kreps, D. M., & Wilson, R. B. (1982). Sequential equilibria. *Econometrica, 50*, 863–894.

Perea, A. (2001). *Rationality in extensive form games*. Boston: Kluwer Academic.

Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodels mit Nachfragezeit. *Zeitschrift für Gesammte Staatswissenschaft, 121*, 301–324.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory, 4*, 25–55.