CHAPTER 2

Extensive Form Games

Again, we begin our discussion of extensive form games without defining what one is, but giving some examples.

1. Examples of extensive form games

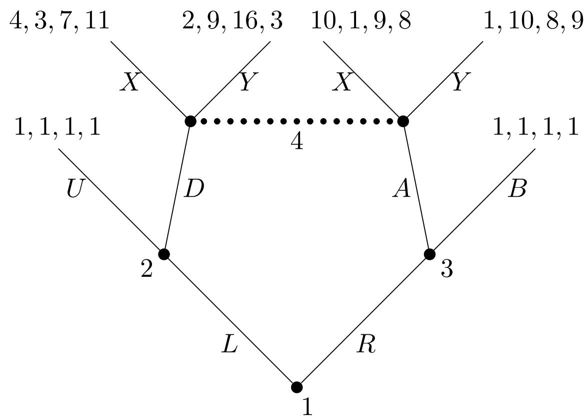

FIGURE 1. An extensive form game.

Look at Figure 1. We interpret this as follows. Each point where a player gets to move in the game or at which the game ends is called a *node*. Nodes at which players move are shown by small black dots in Figure 1 and are called *decision nodes*. The game starts at a particular node, called the *initial node* or *root*. In this case we assume that the lowest node where Player 1 moves is the initial node. Player 1 chooses between L or R. If Player 1 chooses L then Player 2 moves and chooses between U or D. If Player 1 chooses R then Player 3 moves and chooses between A and B. If Player 2 chooses U then the game ends. If Player 2 chooses D then player 4 moves. If player 3 chooses B then the game ends. If player 3 chooses A then Player 4 moves. When it's Player 4's turn to move, he doesn't see whether he is called upon to move because Player 1 chose L and Player 2 chose D or because Player 1 chose R and Player 3 chose A. We say that the two nodes at which player 4 moves are in the same *information set* and we represent this by joining them with a dotted line as in Figure 1. If it's player 4's turn to move, he chooses between X and Y, after which the game ends. The nodes at which the game ends are called *terminal nodes*. To each terminal node we associate a *payoff* to each player. These payoffs tell us how the player evaluates the game ending at that particular node, that is they tell us the players' preferences over the terminal nodes, as well as their preferences over randomisations over those nodes.

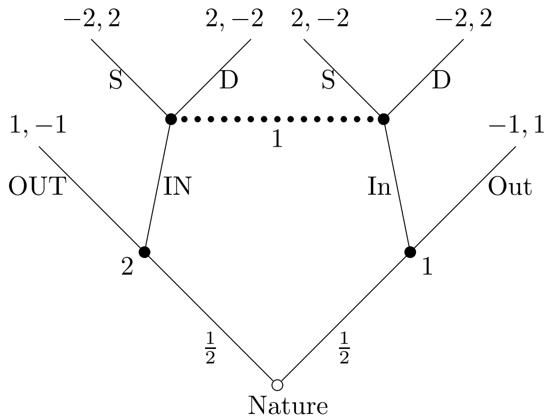

Figure 2. Another extensive form game. This game does not have perfect recall.

As another example, consider the extensive form game shown in Figure 2. In this game, the first mover is not a player but "Nature". That is, at the beginning of the game, there is a random selection of whether Player 1 or Player 2 gets to move, each being chosen with probability $\frac{1}{2}$. (I shall indicate such moves of Nature by an open circle, while nodes at which real players move are indicated by closed dots. The reader is warned that such a convention is not universally followed elsewhere.) If Player 1 is chosen, Player 1 gets to choose between *In* and *Out*, and if Player 2 is chosen, Player 2 gets to choose between *IN* and *OUT*. If either Player 1 chooses *Out* or Player 2 chooses OUT if chosen to move then the game ends. If Player 1 made the choice then Player 1 gets -1 and Player 2 gets 1 while if Player 2 made the choice, Player 1 gets 1 and Player 2 gets -1. On the other hand, if either Player 1 chose In or Player 2 chose *IN* then Player 1 gets to move again, but Player 1 does not know if it was Player 1 or Player 2 that moved previously. This may seem a bit bizarre because it means that Player 1 forgets whether he moved previously in the game or not. Thus we say that this is a game without *perfect recall* (see Section 3 below). One way to rationalise this is to think of Player 1 as a 'team'. So the first time Player 1 gets to move (if Player 1 was chosen by Nature), it is one person or *agent* in the "Player 1 team" moving, and the second time it is another agent in the "Player 1 team", and neither agent in the Player 1 team can communicate with the other. Finally, if Player 1 gets to move for the second time, he gets to choose between S or D, and the appropriate payoffs are given as in Figure 2.

Let's summarise the terminology we have introduced. In a game such as those in Figures 1 or 1, we call nodes at which a player moves *decision nodes.* A node at which some exogenous randomisation occurs is also called a decision node and we say that "Nature" moves at this node. The first node of the game is called the *initial node* or simply the *root.* Nodes at which the game ends and payoffs are specified are called *terminal nodes.* When a player cannot distinguish between two nodes at which he or she gets to move, as is the case for the second move of Player 1 in the game in Figure 2, we say that these two nodes are in the same *information set* of that player.

2. Definition of an extensive form game

One of the essential building blocks of an extensive form game is the *game tree*, g. Consider Figure 3. Formally, a game tree is a finite connected graph with no loops and a distinguished initial node. What does this mean? A finite graph is a finite set of nodes, $X = \{x_1, x_2, \ldots, x_K\}$ and a set of branches connecting them. A branch is a set of two different nodes, $\{x_i, x_j\}$ where $x_i \neq x_j$. In Figure 3, the set of nodes is $X = \{x_1, \ldots, x_9\}$ and the branches are $\{x_1, x_2\}, \{x_2, x_5\}, \{x_6, x_2\}, \{x_1, x_9\}, \{x_9, x_3\}, \{x_9, x_4\}, \{x_3, x_7\},$ and $\{x_3, x_8\}$. The initial node, or root, is x_1.

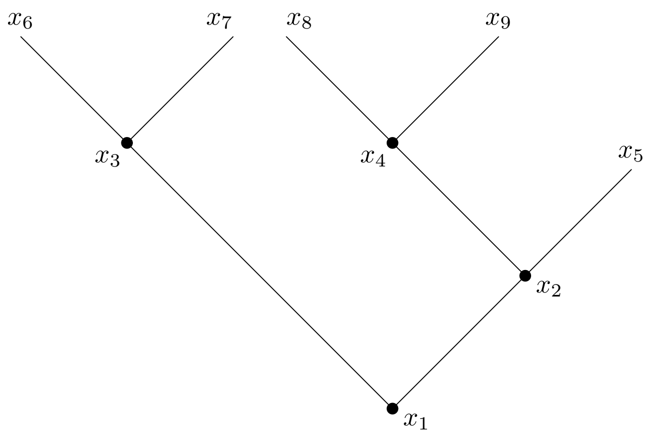

FIGURE 3. A game tree.

In a game tree, we want to rule out trees that look like those in Figures 4 and 5. That is, we don't want to have trees with loops (as in Figure 4) or trees with more than one initial node (as in Figure 4). A *path* is a sequence of branches $(\{x_1, x_2\}, \{x_2, x_3\}, \ldots, \{x_{T-1}, x_T\})$. For example, in Figure 3, one path is $(\{x_7, x_3\}, \{x_3, x_8\})$.[1] Then, to avoid problem trees like those in Figures 4 and 5, we make the following assumption: *Every node except the root is connected to the root by one and only one path.*

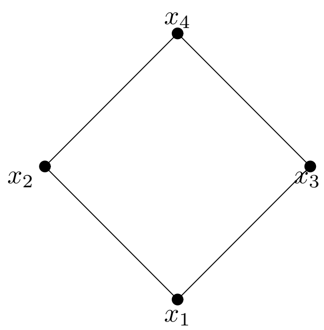

FIGURE 4. Not a game tree.

[1] Notice that in our definition of a path, we didn't say that we are only moving in one direction in the game tree.

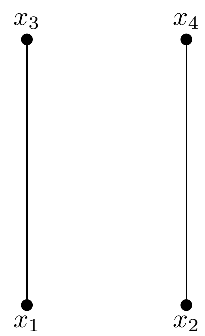

FIGURE 5. Not a game tree either.

As an alternative way of avoiding game trees like those in Figures 4 and 5, given a finite set of nodes X, we define the *immediate predecessor function* $p : X \to X \cup \{\emptyset\}$, to be the function that gives the node that comes immediately before any given node in the game tree. For example, in Figure 3, $p(x_2) = x_1$, $p(x_3) = x_9$, $p(x_7) = x_3$, and so on. To make sure that a game tree has only one initial node, we require that there is a unique $x_0 \in X$ such that $p(x_0) = \emptyset$, that is, that there is only one node with no immediate predecessor. To prevent loops in the game tree, we require that for every $x \neq x_0$, either $p(x) = x_0$ or $p(p(x)) = x_0$ or $\ldots p(p(p \ldots p(x))) = x_0$. That is, by applying the immediate predecessor function p to any node except for the initial node, eventually we get to the initial node x_0.

The terminal nodes in Figure 3 are x_5, x_6, x_7, x_8, and x_9. We denote the set of terminal nodes by T. A terminal node is a node such that there is a path $(\{x_0, x_1\}, \{x_1, x_2\}, \ldots, \{x_{t-1}, x_t\})$ and there are no paths that extend this. Alternatively, we can say that x_t is a terminal node if there is no node $y \in X$ with $p(y) = x_t$.

Now we can give a formal definition of an extensive form game.

DEFINITION 2.1. An *extensive form game* consists of:

(1) $N = \{1, 2, \ldots, N\}$ a finite set of players.
(2) X a set of nodes.
(3) X is a game tree.
(4) A set of actions, A, and a labelling function $\alpha : X \backslash \{x_0\} \to A$ where $\alpha(x)$ is the action at the predecessor of x that leads to x. If $p(x) = p(x')$ and $x \neq x'$ then $\alpha(x) \neq \alpha(x')$.
(5) $\mathcal{H}$ a collection of information sets and $H : X \backslash T \to \mathcal{H}$ a function that assigns for every node, except the terminal ones, which information set the node is in.
(6) A function $n : \mathcal{H} \to N \cup \{0\}$, where player 0 denotes Nature. That is, $n(H)$ is the player who moves at information set H. Let $\mathcal{H}_n = \{H \in \mathcal{H} | n(H) = n\}$ be the information sets controlled by player n.
(7) $\rho : \mathcal{H}_0 \times A \to [0, 1]$ giving the probability that action a is taken at the information set H of Nature. That is, $\rho(H, a)$.
(8) $(u_1, \ldots, u_N)$ with $u_n : T \to \mathbb{R}$ being the payoff to player n.

This completes the definition of an extensive form game. Just a few more comments are in order at this stage.

First, let $C(x) = \{a \in A | a = \alpha x' \text{ for some } x' \text{ with } p(x') = x\}$. That is, $C(x)$ is the set of choices that are available at node x. Note that if x is a terminal node then $C(x) = \emptyset$. If two nodes x and x' are in the same information set, that is, if $H(x) = H(x')$, then the same choices must be available at x and x', that is, $C(x) = C(x')$.

To illustrate this, consider the "game" shown in Figure 6. This figure illustrates the "forgetful driver": A student is driving home after spending the evening at a pub. He reaches a set of traffic lights and can choose to go left or right. If he goes left, he falls off a cliff. If he goes right, he reaches another set of traffic lights. However when he gets to this second set of traffic lights, since he's had a few drinks this evening, he cannot remember if he passed a set of traffic lights already or not. At the second traffic lights he can again either go left or right. If he goes left he gets home and if he goes right he reaches a rooming house. The fact that he forgets at the second set of traffic lights whether he's already passed the first set is indicated in Figure 6 by the nodes x_0 and x_1 being in the same information set. That is, $H(x_0) = H(x_1)$. Under our definition, this is *not* a proper extensive form game. Under our definition, if two nodes are in the same information set, neither should be a predecessor of the other.

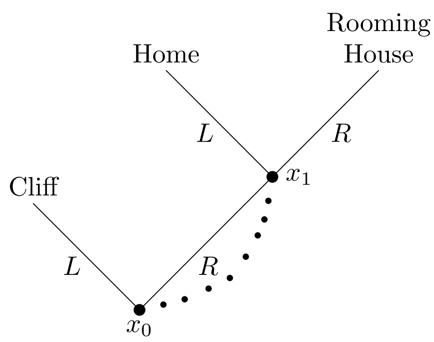

FIGURE 6. The forgetful driver.

Finally, we'll assume that all $H \in \mathcal{H}_0$ are singletons, that is, sets consisting of only one element. In other words, if $H \in \mathcal{H}_0$ and $H(x) = H(x')$ then $x = x'$. This says that Nature's information sets always have only one node in them.

3. Perfect recall

An informal definition is that a player has perfect recall if he remembers everything that he knew and everything that he did in the past. This has a few implications:

- If the player has only one information set then the player has perfect recall (because he has no past!).
- If it is never the case that two different information sets of the player both occur in a single play of the game then the player has perfect recall.
- If every information set is a single node we say the game has *perfect information*. In this case a player when called upon to move sees exactly what has happened in the past. If the game has perfect information then each player has perfect recall.

Now let us give a formal definition of perfect recall. We define it in terms of a player, and if every player in the game has perfect recall, we say that the game itself has perfect recall.

DEFINITION 2.2. Given an extensive form game, we say that player n in that game has *perfect recall* if whenever $H(x) = H(x') \in \mathcal{H}_n$, that is, whenever x and x' are in the same information set and player n moves at that information set, with x'' a predecessor of x with $H(x'') \in \mathcal{H}_n$, that is, x'' comes before x and player n moves at x'', and a'' is the action at x'' on the path to x, then there is some $x''' \in H(x'')$ a predecessor of x' with a'' the action at x''' on the path to x'.

4. Equilibria of extensive form games

To find the Nash equilibria of an extensive form game, we have two choices. First, we could find the equivalent normal form game and find all the equilibria from that game, using the methods that we learned in the previous section. The only trouble with this is that two different extensive form games can give the same normal form. Alternatively, we can find the equilibria directly from the extensive form using a concept (to be explained in subsection 9 below) called *subgame perfect equilibrium* (SPE). Note that we *cannot* find the SPE of an extensive form game from its associated normal form; we must find it directly from the extensive form. The two approaches that we can take to finding the equilibria of an extensive form game are shown in Figure 7.

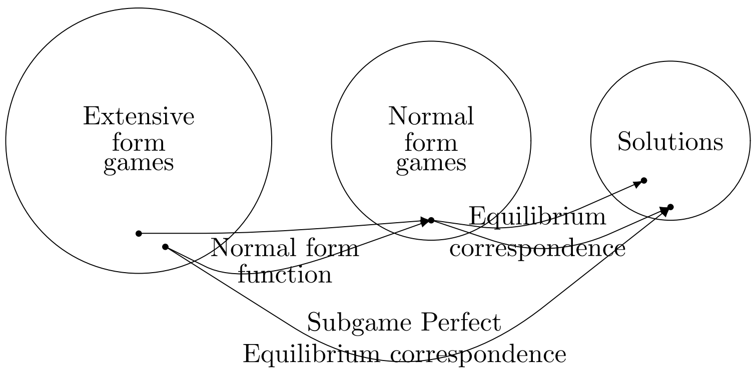

FIGURE 7. Methods for defining equilibria of extensive form games. Note that two different extensive form games may have the same normal form and that a single game may have multiple equilibria, some of which may be subgame perfect and some of which may not.

5. The associated normal form

Let us first consider the method for finding equilibria of extensive form games whereby we find the Nash equilibria of the associated normal form. Consider the extensive form game shown in Figure 8. To find the associated normal form of this game, we first need to know what a strategy of a player is. As we said before, a strategy for a player is a complete set of instructions as to how to play in the game. More formally, we have the following definition.

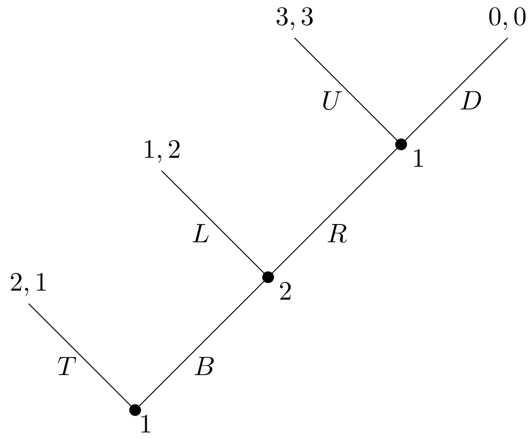

FIGURE 8. Yet another extensive form game.

DEFINITION 2.3. A *strategy for player n* is a function that assigns to each information set of that player a choice at that information set.

Note that this definition doesn't take account of the fact that certain choices by the players may render some of their information sets irrelevant. For example, in Figure 8, if Player 1 plays T at her first information set, her second information set is never reached. However, a strategy for Player 1 must still specify what she would do if her second information set were reached. Thus, strictly speaking, TU and TD are distinct strategies for Player 1. However, two strategies that for any fixed choice of the other player lead to the same outcome are said to be *equivalent.* In the game in Figure 8, the strategies TU and TD are equivalent for Player 1.

Now that we know what a strategy in an extensive form game is, we can set about deriving the associated normal form. Figure 9 shows the associated normal form game of the extensive form game in Figure 8. Player 1's strategies are TU, TD, BU and BD and Player 2's strategies are L and R. We then construct the associated normal form by simply writing all the strategies of Player 1 down the left hand side and all the strategies of Player 2 across the top, and then filling in each cell in the matrix with the appropriate payoff according to the strategies that are played by each player.

EXERCISE 2.1. Find all the Nash equilibria of the game in Figure 9.

As another example, consider the extensive form game in Figure 2. In this game, Player 2's strategy set is $S_2 = \{IN, OUT\}$ and Player 1's strategy set is $S_1 = \{(In,S), (In,D), (Out,S), (Out,D)\}$. In this game, because there is a random move by Nature at the beginning of the game, a profile of pure strategies of the player generates a *probability distribution* over the terminal nodes. To generate the associated normal form of the game, we then use this probability distribution to calculate the expected value of the terminal payoffs. Thus the associated normal form of the extensive form game in Figure 2 is shown in Figure 10. As an example of how the payoffs were calculated, suppose Player 1 plays (*In,S*) and Player 2 plays *IN*. In this case, the payoff is $(-2, 2)$ with probability $\frac{1}{2}$ and $(2, -2)$ with

Player 2

		L	R
	TU	$2,1$	$2,1$
Player 1	TD	$2,1$	$2,1$
	BU	$1,2$	$3,3$
	BD	$1,2$	$0,0$

FIGURE 9. The normal form game corresponding to the extensive form game in Figure 8.

probability $\frac{1}{2}$. Thus the expected payoffs are $\frac{1}{2} \cdot (-2, 2) + \frac{1}{2} \cdot (2, -2) = (0, 0)$. If instead Player 1 plays (*In,S*) and Player 2 plays OUT, the expected payoffs are $\frac{1}{2} \cdot (1, -1) + \frac{1}{2} \cdot (2, -2) = \left(1\frac{1}{2}, -1\frac{1}{2}\right)$. The other payoffs are calculated similarly.

Player 2

		IN	OUT
	In, S	$0,0$	$\frac{3}{2}, -\frac{3}{2}$
Player 1	In, D	$0,0$	$-\frac{1}{2}, \frac{1}{2}$
	Out, S	$-\frac{3}{2}, \frac{3}{2}$	$0,0$
	Out, D	$\frac{1}{2}, -\frac{1}{2}$	$0,0$

FIGURE 10. The normal form game corresponding to the extensive form game in Figure 2.

Let us complete this example by finding all the Nash equilibria of the normal form game in Figure 10. Note that for Player 1, the strategies (In, D) and (Out, S) are strictly dominated by, for example, the strategy $\left(\frac{1}{2}, 0, 0, \frac{1}{2}\right)$. To see this, suppose that Player 1 plays (In, D) and Player 2 plays a mixed strategy $(y, 1-y)$ where $0 \leq y \leq 1$. Then Player 1's expected payoff from (In, D) is $0 \cdot y + \left(-\frac{1}{2}\right) \cdot (1-y) = -\frac{1}{2} + \frac{1}{2}y$. Suppose instead that Player 1 plays the mixed strategy $\left(\frac{1}{2}, 0, 0, \frac{1}{2}\right)$. Player 1's expected payoff from this strategy is $\frac{1}{2}\left(0 \cdot y + 1\frac{1}{2} \cdot (1-y)\right) + \frac{1}{2}\left(\frac{1}{2} \cdot y + 0 \cdot (1-y)\right) = \frac{3}{4} - \frac{1}{2}y$. Since $\frac{3}{4} - \frac{1}{2}y - \left(-\frac{1}{2} + \frac{1}{2}y\right) = \frac{5}{4} - y \geq 0$ for any $0 \leq y \leq 1$, $\left(\frac{1}{2}, 0, 0, \frac{1}{2}\right)$ gives Player 1 a higher expected payoff than (*In,D*) whatever Player 2 does. Similar calculations show that (*Out,S*) is also strictly dominated for Player 1. Thus we know that Player 1 will never play (*In,D*) or (*Out,S*) in a Nash equilibrium. This means that we can reduce the game to the one found in Figure 11.

From Figure 11, we see that the game looks rather like that of matching pennies from Figure 6. We know that in such a game there are no pure strategy equilibria. In fact, there is no equilibrium in which either player plays a pure strategy. So we only need to check for mixed strategy equilibria. Suppose that Player 1 plays (*In,S*) with probability x and (*Out,D*) with probability $1-x$, and Player 2 plays *IN* with probability y and *OUT* with probability $1-y$. Then x must be such that Player 2 is indifferent between *IN* and *OUT*, which implies

$$0 \cdot x + \left(-\tfrac{1}{2}\right) \cdot (1-x) = \left(-1\tfrac{1}{2}\right) \cdot x + 0.\,(1-x)$$

		Player 2	
		IN	OUT
Player 1	In, S	$0, 0$	$\frac{3}{2}, -\frac{3}{2}$
	Out, D	$\frac{1}{2}, -\frac{1}{2}$	$0, 0$

FIGURE 11. The normal form game corresponding to the extensive form game in Figure 2, after eliminating Player 1's strictly dominated strategies.

which implies $x = \frac{1}{4}$. Similarly, y must be such that Player 1 is indifferent between (*In,S*) and (*Out,D*). This implies that

$$0 \cdot y + 1\tfrac{1}{2} \cdot (1 - y) = \tfrac{1}{2} \cdot y + 0 \cdot (1 - y)$$

which implies $y = \frac{3}{4}$. So, the only Nash equilibrium of the game in Figure 11, and hence of the game in Figure 10 is

$$\left\{ \left(\tfrac{1}{4}, 0, 0, \tfrac{3}{4}\right), \left(\tfrac{3}{4}, \tfrac{1}{4}\right) \right\}.$$

EXERCISE 2.2. Consider the normal form game in Figure 10 and the Nash equilibrium strategy that we have just found.

(1) If the players play their equilibrium strategies, what is the expected payoff to each player?
(2) If Player 1 plays his equilibrium strategy, what is the *worst* payoff that he can get (whatever Player 2 does)?

6. Behaviour strategies

Consider the extensive form game shown in Figure 12 and its associated normal form shown in Figure 13. In this game, we need three independent numbers to describe a mixed strategy of Player 2, i.e., $(x, y, z, 1 - x - y - z)$. Suppose that instead Player 2 puts off her decision about which strategy to use until she is called upon to move. In this case we only need two independent numbers to describe the uncertainty about what Player 2 is going to do. That is, we could say that at Player 2's left-hand information set she would choose L with probability x and R with probability $1 - x$, and at her right-hand information set she would choose W with probability y and E with probability $1 - y$. We can see that by describing Player 2's strategies in this way, we can save ourselves some work. This efficiency increases very quickly depending on the number of information sets and the number of choices at each information set. For example, suppose that the game is similar to that of Figure 12, except that Player 1's choice is between four strategies, each of which leads to a choice of Player 2 between three strategies. In this case Player 2 has $3 \cdot 3 \cdot 3 \cdot 3 = 81$ different pure strategies and we'd need 80 independent numbers to describe the uncertainty about what Player 2 is going to do! If instead we suppose that Player 2 puts off her decision about which strategy to use until she is called upon to move, we only need $4 \cdot 2 = 8$ independent numbers. This is really great!

When we say that a player "puts off" his or her decision about which strategy to use until he or she is called upon to move, what we really mean is that we are using what is called a *behaviour strategy* to describe what the player is doing. Formally, we have the following definition.

DEFINITION 2.4. In a given extensive form game with player set N, a *behaviour strategy* for player $n \in N$ is a rule, or function, that assigns to each information set of

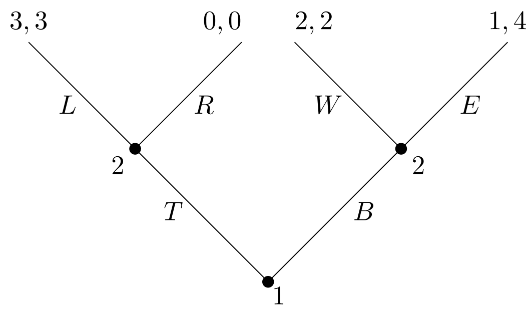

Figure 12. An extensive form game.

		Player 2			
		LW	LE	RW	RE
Player 1	T	3, 3	3, 3	0, 0	0, 0
	B	2, 2	1, 4	2, 2	1, 4

Figure 13. The normal form game corresponding to the extensive form game in Figure 12.

that player a probability distribution over the choices available at that information set.

7. Kuhn's theorem

Remember we motivated behaviour strategies in subsection 6 as a way of reducing the amount of numbers we need compared to using mixed strategies. You might be wondering whether we can always do this, that is, if we can represent any arbitrary mixed strategy by a behaviour strategy. The answer is provided by Kuhn's theorem.

Theorem 2.1. *(Kuhn) Given an extensive form game, a player n who has perfect recall in that game, and a mixed strategy σ_n of player n, there exists a behaviour strategy b_n of player n such that for any profile of strategies of the other players $(x_1, \ldots, x_{n-1}, x_{n+1}, \ldots, x_N)$ where x_m, $m \neq n$, is either a mixed strategy of a behaviour strategy of player m, the strategy profiles $(x_1, \ldots, x_{n-1}, \sigma_n, x_{n+1}, \ldots, x_N)$ and $(x_1, \ldots, x_{n-1}, b_n, x_{n+1}, \ldots, x_N)$ give the same distribution over terminal nodes.*

In other words, Kuhn's theorem says that given what the other players are doing, we can get the same distribution over terminal nodes from σ_n and b_n, as long as player n has perfect recall. To see Kuhn's theorem in action, consider the game in Figure 14. Suppose Player 1 plays T with probability p and B with probability $1 - p$. Player 2's pure strategies are XA, XB, YA, and YB. Suppose Player 2 plays a mixed strategy of playing XA and YB with probability $\frac{1}{2}$ and XB and YA with probability 0. Thus, with probability $\frac{1}{2}$, Player 2 plays XA and we get to terminal node x with probability p and to terminal node a with probability $1-p$, and with probability $\frac{1}{2}$, Player 2 plays YB and we get to terminal node y with probability p and terminal node b with probability $1 - p$. This gives a distribution over terminal nodes as shown in the table in Figure 15.

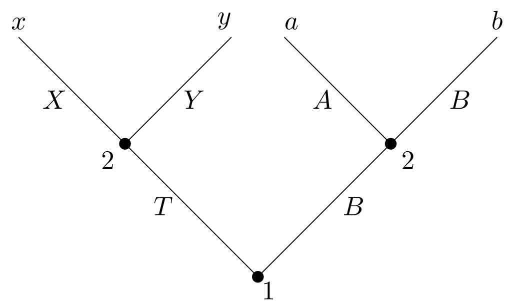

FIGURE 14. An extensive form game.

Terminal node	Probability
x	$p/2$
y	$p/2$
a	$(1-p)/2$
b	$(1-p)/2$

FIGURE 15. Distributions over terminal nodes in the game of Figure 14 for the strategy profile given in the text.

Now suppose that Player 2 plays a behaviour strategy of playing X with probability $\frac{1}{2}$ at his left-hand information set and A with probability $\frac{1}{2}$ at his right-hand information set. Thus we get to terminal node x with probability $p \cdot \frac{1}{2}$, to terminal node y with probability $p \cdot \frac{1}{2}$, to terminal node a with probability $(1-p) \cdot \frac{1}{2}$ and to terminal node b with probability $(1-p) \cdot \frac{1}{2}$. Just as Kuhn's theorem predicts, there is a behaviour strategy that gives the same distribution over terminal nodes as the mixed strategy.

8. More than two players

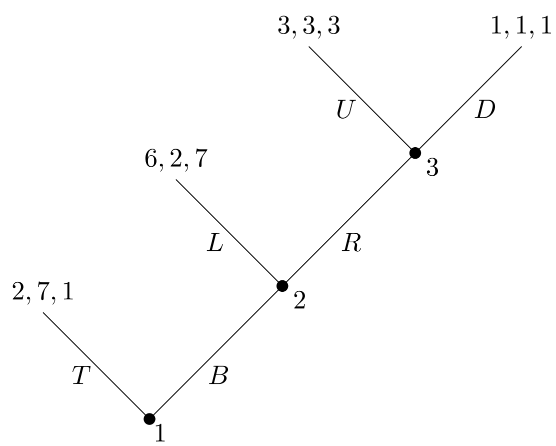

FIGURE 16. An extensive form game with three players.

It is easy to draw extensive form games that have more than two players, such as the one shown in Figure 16, which has three players. How would we find the associated normal form of such a game? Recall that a normal form game is given by (N, S, u). For the game in Figure 16, we have $N = \{1, 2, 3\}$, $S_1 = \{T, B\}$, $S_2 = \{L, R\}$ and $S_3 = \{U, D\}$, hence

$$S = \{(T, L, U), (T, L, D), (T, R, U), (T, R, D), (B, L, U), (B, L, D),$$
$$(B, R, U), (B, R, D)\}.$$

Finally, the payoff functions of each player are shown in Figure 17.

S	T, L, U	T, L, D	T, R, U	T, R, D	B, L, U	B, L, D	B, R, U	B, R, D
u_1	2	2	2	2	6	6	3	1
u_2	7	7	7	7	2	2	3	1
u_3	1	1	1	1	7	7	3	1

FIGURE 17. One way of presenting the normal form game associated with the extensive form game of Figure 16.

The form shown in Figure 17 is a perfectly acceptable exposition of the associated normal form. However, it's more convenient to represent this three player game as shown in Figure 18, where Player 3 chooses the matrix.

Player 3

U

Player 2

Player 1

	L	R
T	$2, 7, 1$	$2, 7, 1$
B	$6, 2, 7$	$3, 3, 3$

D

Player 2

Player 1

	L	R
T	$2, 7, 1$	$2, 7, 1$
B	$6, 2, 7$	$1, 1, 1$

FIGURE 18. A more convenient representation of the normal form of the game in Figure 16.

9. Subgame perfect equilibrium

Subgame perfect equilibrium is an equilibrium concept that relates directly to the extensive form of a game. The basic idea is that equilibrium strategies should continue to be an equilibrium in each subgame (subgames will be defined formally below).

DEFINITION 2.5. A *subgame perfect equilibrium* is a profile of strategies such that for every subgame the parts of the profile relevant to the subgame constitute an equilibrium of the subgame.

To understand subgame perfect equilibrium, we first need to know what a subgame is. Consider the game in Figure 16. This game has two proper subgames. Strictly speaking, the whole game is a subgame of itself, so we call subgames that are not the whole game *proper* subgames. The first is the game that begins from Player 2's decision node and the second is the game that begins from player 3's decision node. However, it would be a mistake to think that every game has proper subgames. Consider the extensive form game in Figure 19, which is the extensive form version of matching pennies. In this game it doesn't make sense to say there is a subgame that begins at either of Player 2's decision nodes. This is because both of these nodes are in the same information set. If we said that a subgame started at either one of them, then we would be saying that Player 2 knew which node she is at, which violates the original structure of the game.

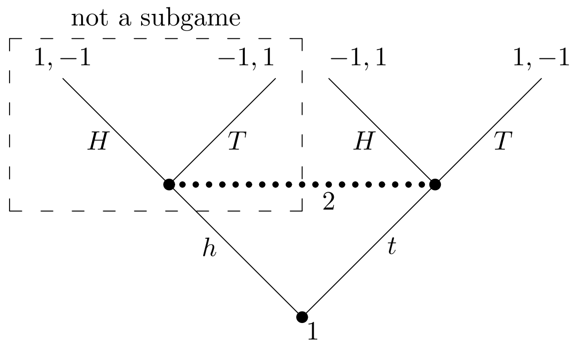

FIGURE 19. The extensive form version of matching pennies. This game has no proper subgames.

As another example, consider the extensive form game shown in Figure 20. In this game, the game starting from Player 2's decision node is not a subgame, because it splits player 3's information set.

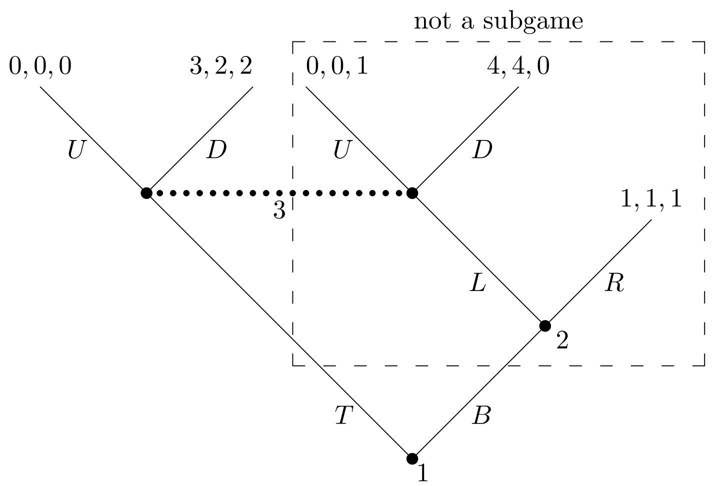

FIGURE 20. The game beginning at Player 2's decision node is not a subgame because it splits Player 3's information set.

Let us now give a formal definition of a subgame.

Definition 2.6. A *subgame* of an extensive form game Γ is some node in the tree of Γ and all the nodes that follow it, with the original tree structure but restricted to this subset of nodes, with the property that any information set of Γ is either completely in the subgame or completely outside the subgame. The rest of the structure is the same as in Γ but restricted to the new (smaller) tree.

Notice that under this definition, Γ is a subgame of Γ. We call subgames that are not Γ itself *proper* subgames of Γ.

And we are now in a position to give a more formal definition of subgame perfect equilibrium.

Definition 2.7. A *subgame perfect equilibrium* of an extensive form game Γ with perfect recall is a profile of behaviour strategies $(b_1, b_2, \ldots, b_N)$ such that, for every subgame Γ' of Γ, $(b'_1, b'_2, \ldots, b'_N)$ is a Nash equilibrium of Γ' where b'_n is b_n restricted to Γ'.

9.1. Finding subgame perfect equilibria. An example of the simplest type of extensive form game in which subgame perfect equilibrium is interesting is shown in Figure 21. Notice that in this game, (T, R) is a Nash equilibrium (of the whole game). Given that Player 1 plays T, Player 2 is indifferent between L and R. And given that Player 2 plays R, Player 1 prefers to play T. There is something strange about this equilibrium, however. Surely, if Player 2 was actually called upon to move, she would play L rather than R. The idea of subgame perfect equilibrium is to get rid of this silly type of Nash equilibrium. In a subgame perfect equilibrium, Player 2 should behave rationally if her information set is reached. Thus Player 2 playing R cannot form part of a subgame perfect equilibrium. Rather, if Player 2 is called upon to move, the only rational thing for her to do is to play L. This means that Player 1 will prefer B over T, and the only subgame perfect equilibrium of this game is (B, L). Note that (B, L) is also a Nash equilibrium of the whole game. This is true in general: Subgame perfect equilibria are always Nash equilibria, but Nash equilibria are not necessarily subgame perfect.

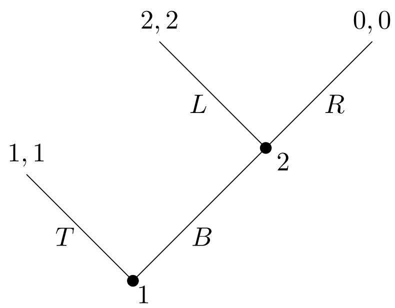

Figure 21. An example of the simplest type of game in which subgame perfect equilibrium is interesting.

As another example, consider the extensive form game in Figure 22. In this game, players 1 and 2 are playing a prisoners' dilemma, while at the beginning of the game player 3 gets to choose whether players 1 and 2 will actually play the prisoners' dilemma or whether the game will end. The only proper subgame of this game begins at Player 1's node. We claim that (C, C, L) is a Nash equilibrium of the whole game. Given that player 3 is playing L, players 1 and 2 can do whatever they like without affecting their payoffs and do not care what the other is playing.

And given that players 1 and 2 are both playing C, player 3 does best by playing L. However, although (C, C, L) is a Nash equilibrium of the game, it is not a subgame perfect equilibrium. This is because (C, C) is not a Nash equilibrium of the subgame beginning at Player 1's decision node. Given that Player 1 is playing C, Player 2 would be better off playing D, and similarly for Player 1. The only subgame perfect equilibrium is (D, D, A).

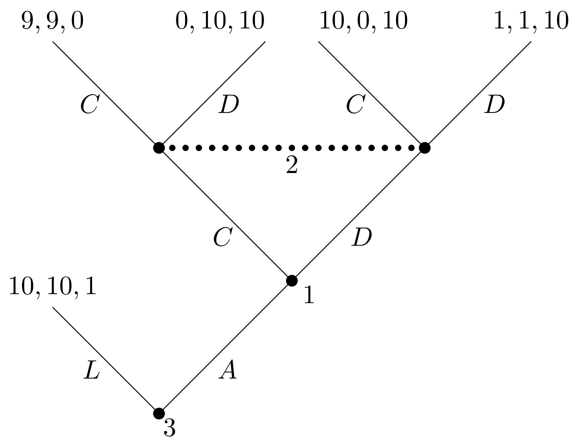

FIGURE 22. Player 3 chooses whether or not Players 1 and 2 will play a prisoners' dilemma game.

In the previous example, we have seen that a extensive form games can have equilibria that are Nash but not subgame perfect. You might then be wondering whether it's possible to have an extensive form game that has no subgame perfect equilibria. The answer is no. Selten in 1965 proved that every finite extensive form game with perfect recall has at least one subgame perfect equilibrium.

9.2. Backwards induction. Backwards induction is a convenient way of finding subgame perfect equilibria of extensive form games. We simply proceed backwards through the game tree, starting from the subgames that have no proper subgames of themselves, and pick an equilibrium. We then replace the subgame by a terminal node with payoff equal to the expected payoff in the equilibrium of the subgame. Then repeat, as necessary.

As an illustration, consider the game in Figure 23. First Player 1 moves and chooses whether she and Player 2 will play a battle of the sexes game or a coordination game. There are two proper subgames of this game, the battle of the sexes subgame and the coordination subgame. Following the backwards induction procedure, one possible equilibrium of the battle of the sexes subgame is (F, F), and one possible equilibrium of the coordination subgame is (B, b). We then replace each subgame by the appropriate expected payoffs, as shown in Figure 24. We can then see that at the initial node, Player 1 will choose S. Thus one subgame perfect equilibrium of this game is $((S, F, B), (F, b))$.

10. Sequential equilibrium

Consider the game shown in Figure 25. This game is known as "Selten's Horse". In this game, there are no proper subgames, hence all Nash equilibria are subgame perfect. Note that (T, R, D) is a Nash equilibrium (and hence subgame perfect). But in this equilibrium, Player 2 isn't really being rational, because if player 3 is

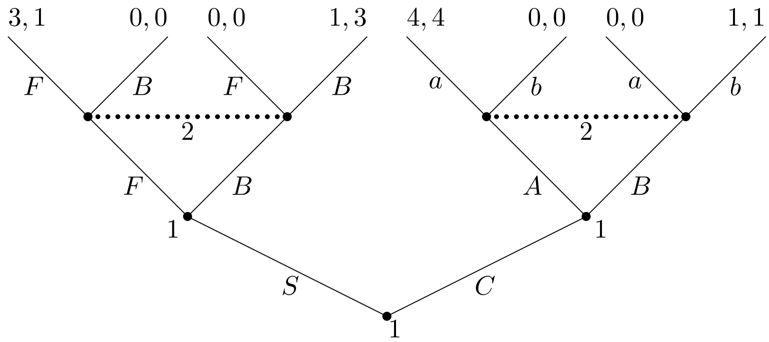

FIGURE 23. Player 1 chooses battle of the sexes or coordination game.

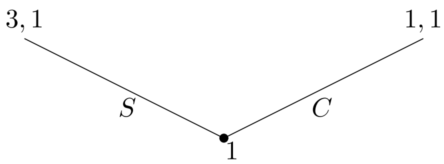

FIGURE 24. Replacing the subgames with the equilibrium expected payoff in each subgame.

really playing D then if Player 2 actually got to move, she would be better off playing L rather R. Thus we have another "silly" Nash equilibrium, and subgame perfection is no help to us here to eliminate it. This kind of game lead to the development of another equilibrium concept called sequential equilibrium.

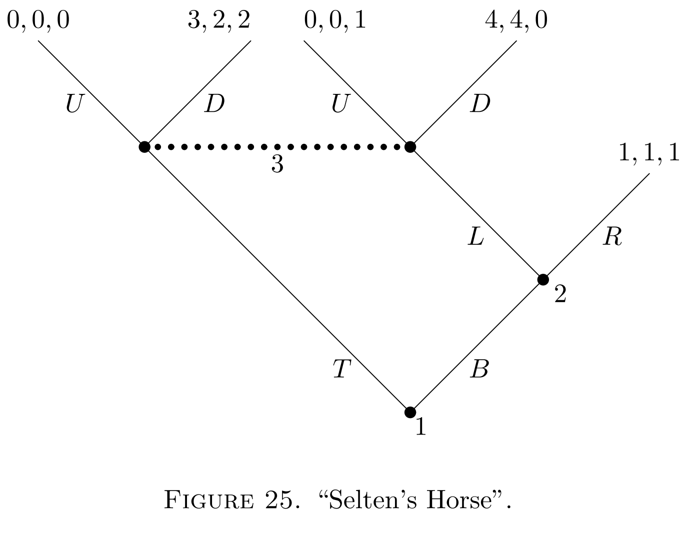

FIGURE 25. "Selten's Horse".

EXERCISE 2.3. Find all the Nash equilibria of the game in Figure 25.

A sequential equilibrium is a pair, (b, μ) where b is a profile of behaviour strategies and μ is a *system of beliefs*. That is, μ is a profile of probability distributions, one for each information set, over the nodes in the information set. The system of beliefs summarises the probabilities with which each player believes he is at each of the nodes within each of his information sets.

To be a sequential equilibrium, the pair (b, μ) must satisfy two properties:

(1) At each information set, b puts positive weight only on those actions that are optimal given b and μ. This is called *sequential rationality*.
(2) μ and b should be *consistent*. This means that if anything can be inferred about μ from b then μ should be that.

If we drop the consistency requirement, we get another equilibrium concept called *Perfect Bayesian Equilibrium*, which we won't talk about in this course but you might come across in other books.

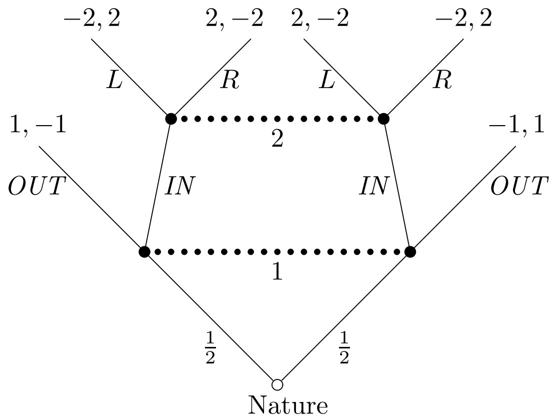

FIGURE 26. Illustrating consistent beliefs.

More formally, consistency requires that there is a sequence of completely mixed behaviour strategy profiles $b^t \to b$ and $\mu^t \to \mu$ with μ^t the conditional distribution over the nodes induced by b^t. To see how this works, consider the game in Figure 26. Each player has two nodes in their (only) information set. Under sequential equilibrium we must specify the probabilities with which each player believes they are at each node in each of their information sets. Furthermore, these beliefs must be consistent with the equilibrium behaviour strategy. Suppose that in this game, $b = \{(0,1), (1,0)\}$, that is, that Player 1 plays OUT and Player 2 plays L. The profile of beliefs is $\mu = ((\mu(x), \mu(y)), (\mu(a), \mu(b)))$. Clearly, since the only move before Player 1's information set is that of Nature, Player 1's beliefs must be that $\mu(x) = \frac{1}{2}$ and $\mu(y) = \frac{1}{2}$. These beliefs will be consistent with any behaviour strategy, since whatever behaviour strategy is being played will not affect the probability that Player 1 is at either node in his information set. To find Player 2's beliefs $\mu(a)$ and $\mu(b)$ we need to find a completely mixed behaviour strategy close to b. Actually, we only need to worry about Player 1's behaviour strategy, since there are no moves of Player 2 before Player 2's information set. Suppose we use $b^t = \left(\frac{1}{t}, \frac{t-1}{t}\right)$ for Player 1. Then,

$$\Pr(a) = \frac{1}{2} \cdot \frac{1}{t} = \frac{1}{2t}$$

and

$$\Pr(b) = \frac{1}{2} \cdot \frac{1}{t} = \frac{1}{2t}.$$

Thus,

$$\mu^t(a) = \frac{\frac{1}{2t}}{\frac{1}{2t} + \frac{1}{2t}} = \frac{1}{2}.$$

Similarly, we can show that $\mu^t(b) = \frac{1}{2}$ also.

Why couldn't we just use Player 1's behaviour strategy $(0, 1)$? Because in this case Player 1 plays OUT and thus Player 2's information set is not reached. Mathematically, we would have

$$\mu(a) = \frac{\frac{1}{2} \cdot 0}{\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 0}$$

which is not defined. So we must use completely mixed behaviour strategies close to b to find the (consistent) beliefs, since if we use a completely mixed behaviour strategy, every information set in the game will be reached with strictly positive probability (even if the probability is very small).

11. Signaling games

Signaling games are a special type of extensive form game that arise in many economic models. The basic structure of a signaling game is as follows:

(1) Nature moves first, and picks one "state of nature" or "type of Player 1".
(2) Player 1 sees the result of Nature's move, and makes his choice. It is assumed that Player 1's available actions do not depend on Nature's choice.
(3) Player 2 sees Player 1's choice but not Nature's, and makes her choice.
(4) Payoffs are determined based on the choices of the two players and Nature, and the game ends.

The extensive form of a simple signaling game is shown in Figure 27. In this game, Player 1 is either type a or b with probability $\frac{1}{2}$. Player 1 can choose T or B. Player 2 then observes Player 1's choice but not Nature's choice. If Player 1 chose T then Player 2 chooses between L and R and if Player 1 chose B then Player 2 chooses between X, Y, and Z (this illustrates the fact that Player 2's available choices can depend on Player 1's choice, but neither player's available choices can depend on Nature's choice).

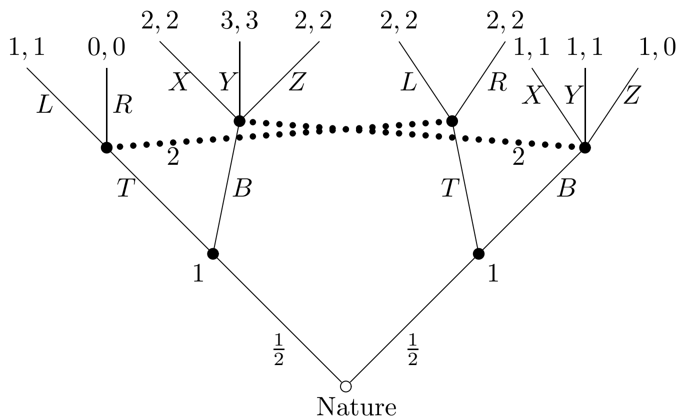

FIGURE 27. The extensive form of a simple signaling game.

Typically, signaling games are not presented in the form of Figure 27, however. Instead, a presentation due originally to Kreps and Wilson (1982) is used, in which we suppress Nature's move and instead have a number of initial nodes with associated probabilities. This alternative presentation of the game in Figure 27 is shown in Figure 28. As you can see, it is much neater to present signaling games in this way.

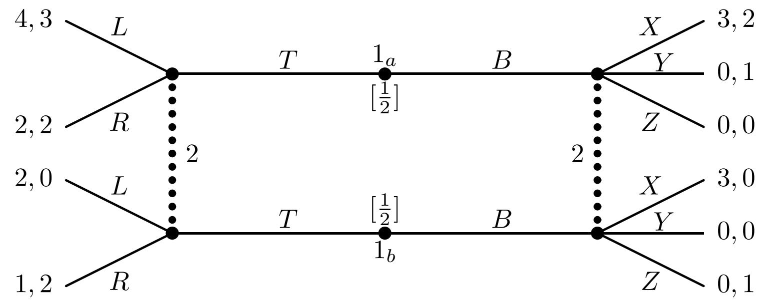

Figure 28. An alternative, and tidier, presentation of the extensive form signaling game in Figure 27.

One economic example of a signaling game is as follows. In some market there is an incumbent firm that sets its price in the time period before a potential entrant enters the market. The incumbent knows if the market demand conditions are good or bad. The potential entrant does not know the demand conditions but does see the price chosen by the incumbent. The entrant then decides whether to enter or not. If the entrant enters, it finds out the market conditions and then there is some competition in the second period. The fixed costs are such that the entrant will make positive profits only if market conditions are good. A representation of this game is shown in Figure 29, where it is assumed that the incumbent can only choose a 'high' or 'low' price.

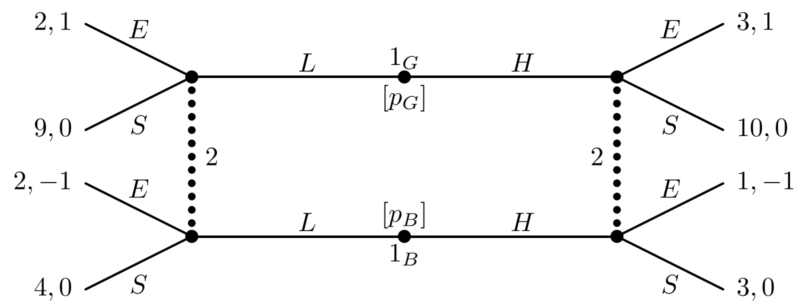

Figure 29. A simple price signaling game.

A famous signaling game is shown in Figure 30. The story behind this game is as follows. There are two players, Player A and Player B. At the beginning of the game, Nature selects Player A to be either a *wimp* (with probability 0.1) or *surly* (with probability 0.9). At the start of the game, Player A knows whether he is a wimp or surly. Player A then has to choose what to have for breakfast. He has two choices: beer or quiche.[2] If Player A is surly, he prefers beer for breakfast and if Player A is a wimp he prefers quiche for breakfast, everything else equal. That is, if Player A has his preferred breakfast, his incremental payoff is 1, otherwise 0. After breakfast, Player A meets Player B. At the meeting, Player B observes what Player A had for breakfast (perhaps by seeing the bits of quiche stuck in Player A's beard if Player A had quiche for breakfast or smelling the alcohol on Player A's breath if A had beer for breakfast), but Player B does not know whether Player A is a wimp or surly. Having observed Player A's breakfast, Player B must then choose whether or not to duel (fight) with Player A. Then the game ends and payoffs are decided. Regardless of whether Player A is a wimp or surly, Player A dislikes fighting. Thus Player A's incremental payoff is 2 if Player B chooses not to duel and 0 if Player B chooses to duel. So, for example, if Player A is surly and has his preferred breakfast of beer and then Player B chooses not to duel, Player A's payoff is $1 + 2 = 3$. Or, if Player A is a wimp and has his less preferred breakfast of beer and Player B chooses to duel, Player A's payoff is $0 + 0 = 0$. Player B, on the other hand, only prefers to duel if Player A is a wimp. If Player A is surly and Player B chooses to duel, Player B's payoff is 0, while if Player B chooses not to duel Player B's payoff is 1. If Player A is a wimp and Player B chooses to duel, Player B's payoff is 2, while if Player B chooses not to duel his payoff is 1.[3]

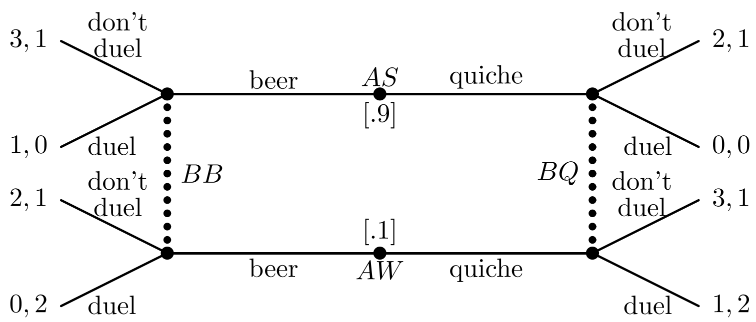

FIGURE 30. The beer quiche game of Cho and Kreps.

EXERCISE 2.4. Consider the signaling game in Figure 30.
(1) Find the normal form of this game.
(2) Find all the pure strategy Nash equilibria (from the normal form).
(3) Find all the mixed strategy Nash equilibria (from the normal form).

Since the beer-quiche game in Figure 30 is a game of asymmetric information (that is, Player B doesn't know whether Player A is surly or a wimp), it makes sense to look for a sequential equilibrium. In this game a sequential equilibrium will be an assessment $((b_A, b_B), \mu)$ that is consistent and sequentially rational, where

$$b_A = ((b_{AS}(\text{beer}), b_{AS}(\text{quiche})), (b_{AW}(\text{beer}), b_{AW}(\text{quiche})))$$
$$b_B = ((b_{BB}(\text{don't}), b_{BB}(\text{duel})), b_{BQ}(\text{don't}), b_{BQ}(\text{duel}))$$
$$\mu = ((\mu_{\text{beer}}(S), \mu_{\text{beer}}(W)), (\mu_{\text{quiche}}(S), \mu_{\text{quiche}}(W))).$$

[2]For those of you who don't know, quiche is a sort-of egg pie type thing that only wimps and Australian rugby players eat.

[3]Note that Player B's payoff only depends on his own decision about whether to duel or not, and whether Player A is surly or a wimp. That is, Player B's payoff does *not* depend on what Player A had for breakfast.

In such a game as this, there are two basic types of equilibrium that can arise. In the first type of equilibrium, called a *pooling* equilibrium, both types of Player A choose the same probabilities for having beer and quiche for breakfast. In this case, Player A's breakfast gives Player B no information about Player A's type and thus Player B's best estimate of Player A's type is that Player A is surly with probability 0.9 and a wimp with probability 0.1. In the second type of equilibrium, called a *separating* equilibrium, the two types of Player A choose different probabilities for the breakfasts. In this case, Player B gets some information about Player A's type from Player A's breakfast, and will be able to modify his estimate of Player A's type.

You will analyse this game more fully in tutorials. So as not to spoil the fun, we shall content ourselves with just answering one question about this game here: Are there any equilibria in which both beer and quiche are played with positive probability? Before answering this question, note the following facts about the game:

(1) If b_B is such that the surly Player A is willing to have quiche for breakfast then the wimp Player A strictly wants to have quiche for breakfast (since quiche is the wimp's preferred breakfast).
(2) Similarly, if b_B is such that the wimp Player A wants to have beer for breakfast then the surly Player A strictly wants to have beer for breakfast.

Now suppose that in equilibrium the surly pPlayer A plays quiche with positive probability. Then by fact (1), the wimp Player A plays quiche with probability 1. Thus either surly Player A plays quiche with probability 1, in which case beer is not played with positive probability, or surly Player A plays beer with positive probability. But then $\mu_{\text{beer}}(S) = 1$ and thus Player B chooses not to duel after observing Player A having beer for breakfast, that is, $b_{BB}(\text{duel}) = 0$. But then surly Player A will not want to play quiche with positive probability.

Suppose, on the other hand, that surly Player A plays beer with probability 1, i.e., $b_{AS}(\text{beer}) = 1$. Then $\mu_{\text{beer}}(S) \geq 0.9$. Then sequential rationality implies $b_{BB}(\text{ duel}) = 0$. So if wimp Player A plays quiche with positive probability then $\mu_{\text{quiche}}(S) = 0$ and $\mu_{\text{quiche}}(W) = 1$ and so $b_{BQ}(\text{duel}) = 1$. Thus wimp Player A chooses beer for sure and quiche is not played with positive probability.

So, we have established that there are no sequential equilibria in which both beer and quiche are played with strictly positive probability. As you will discover in tutorials, there are sequential equilibria in which both types of Player A play beer and there are sequential equilibria in which both types of Player A play quiche.