Chapter 3
Extensive Games

While in a game in strategic form, the players' strategies are given explicitly, many interactions allow for a succinct representation of the strategies. In this chapter, we study extensive games. For these games, we are given a tree, called the game tree. In each non-leaf vertex of the tree, one of the players chooses an outgoing edge. A strategy vector thus determines a unique path for the source to a leaf of the tree, and the utilities of all players depend only on the leaf chosen this way.

3.1 Extensive Games with Perfect Information

Let $T = (V, E)$ be a rooted directed tree. For a vertex $v \in V$, we denote by $\delta^+(v) = \{e \in E : e = (v, w) \text{ with } w \in V\}$ the set of directed edge starting in v. A vertex v is a leaf if $\delta^+(v) = \emptyset$. We denote the set of leaves by L.
Let $(V_i)_{i \in N}$ be a partition of $V \setminus L$. We will refer to V_i as the decision vertices of player i. For a leave $v \in L$, let $\hat{u}_i(v)$ be the utility of player i when leave v is reached. We refer to the function $\hat{\mathbf{u}} : L \to \mathbb{R}^n$ as the leaf utilities.

Definition 3.1 (Extensive Game with Perfect Information). For a rooted tree T with vertex partition $(V_i)_{i \in N}$ and leaf utilities $\hat{\mathbf{u}}$, the corresponding extensive game with perfect information is the strategic game $G(T) = (N, \mathbf{S}, \mathbf{u})$, where
- N is a finite and non-empty set of n players,
- $S_i = \{E_i \subseteq \cup_{v \in V_i} \delta^+(v) : \delta^+(v) \cap E_i = 1 \text{ for all } v \in V_i\}$ for all players i,
- $u_i(\mathbf{s}) = \hat{u}_i(v(\mathbf{s}))$ for all players i, where $v(\mathbf{s})$ is the unique leaf that is connected to the root in $(V, \cup_{i \in N} S_i)$.

An extensive form is finite if and only if T is finite. To identify the strategic decisions of the players, we will often associate names with the edges of the tree T.

Example 3.2 (Bach-or-Stravinsky, Extensive Game Version). Reconsider the situation of the Bach-or-Stravinsky game from Example 1.5. Imagine that now player 2 first announces irrevocably whether to go to the Bach or the Stravinsky concert. After hearing this decision, player 2 can decide to go to Bach or Stravinsky. As player 2 moves first, they face the same decision as before, i.e., their strategies

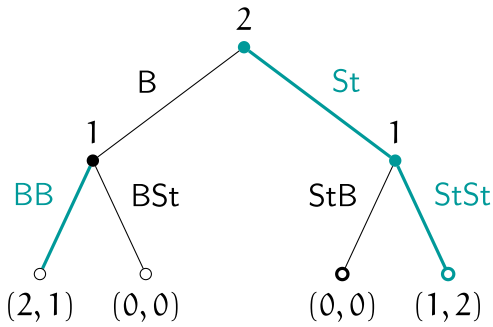

(a) Tree representation

	B	St
{BB, StB}	(2, 1)	(0, 0)
{BB, StSt}	(2, 1)	(1, 2)
{BSt, StB}	(0, 0)	(0, 0)
{BSt, StSt}	(0, 0)	(1, 2)

(b) Strategic representation

Figure 3.1: (a) Tree representation and **(a)** strategic representation of the extensive game version of the Bach-or-Stravinsky game of Exercise 1.5 where player 2 plays first. The strategy vector marked on the right corresponds to the play on the left and is a Nash equilibrium.

remain Bach and Stravinsky. In contrast, player 1 may now condition their decision on the decision of player 2, i.e., their strategies is the set of functions from {Bach, Stravinsky} to {Bach, Stravinsky}, see Figure 3.1. One of the Nash equilibria of this game is coloured.

It is worth noting that the strategy of a player specifies an outgoing edge for each vertex in that player's action set, regardless whether this node is connected to the root, or not.

Example 3.3. Consider the game in Figure 3.2. Player 1 plays before and after player 2. Player 1 has four strategies, i.e., {A, E}, {A, F}, {B, E} and {B, E}, even though for the outcome of the game it is immaterial whether edge E or F is choses when edge B is already contained in the strategy.
Consider the strategy vector {{A, E}, D} marked in Figure 3.3a and Figure 3.3b. This strategy is a Nash equilibrium. It is stable because player 1 threatens to play E when the game has reached to vertex where edges E and F originate. This threat is not credible because player 1 hurts themselves by playing the threat. In fact, player 1 would be better of playing F instead.

To rule out the possibility of players playing non-credible threats, we proceed to introduce subgame perfect equilibria. To this end, we first introduce the notion of a subgame.

Definition 3.4 (Subgame). Let G be an extensive game with game tree $T = (V, E)$, decision vertices $(V_i)_{i \in N}$ and leaf probabilities $\hat{\mathbf{u}}$, and let $v \in V$. The subgame $G^v = (N, \mathbf{S}^v, \mathbf{u}^v)$ is the extensive game to the subtree $T^v(V^v, E^v)$ rooted in v with action set $V_i^v = V_i \cap V^v$ and the same leaf utilities restricted to $L \cap V^v$.

As an example, a subgame of the game of Figure 3.2 is given in Figure 3.3.

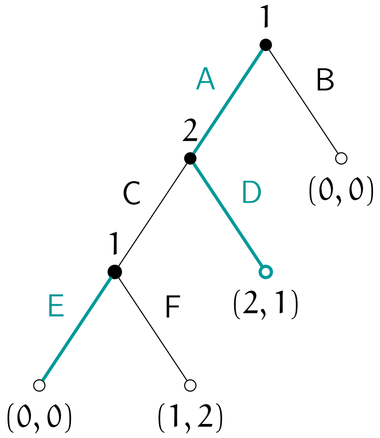

(a) Tree representation

	C	D
{A, E}	(0, 0)	(2, 1)
{A, F}	(1, 2)	(2, 1)
{B, E}	(0, 0)	(0, 0)
{B, F}	(0, 0)	(0, 0)

(b) Strategic representation

Figure 3.2: (a) Tree representation **(b)** strategic representation of an extensive game in which player 1 plays twice.

There is a natural restriction of strategies s_i of G to strategies s_i^v of G^v by setting $s_i^v = s_i \cap E^v$. For a strategy profile $\mathbf{s}$ of G and a vertex v, we also set $\mathbf{s}^v = \bigtimes_{i \in N} s_i^v$. We are now in position to define subgame perfect equilibria.

Definition 3.5 (Subgame-Perfect Equilibrium). A subgame-perfect equilibrium is a strategy profile $\mathbf{s}$ such that for every player i and every vertex $v \in V_i$ we have

$$u_i^v(\mathbf{s}^v) \geq u_i^v(\mathbf{s}_{-i}^v, s_i)$$

for every strategy s_i of player i in the subgame G^v.

The notation of a subgame perfect equilibrium eliminates situations in which players use non-credible threats. For the game in Figure 3.3a, e.g., the unique subgame-perfect equilibrium is the strategy profile $(\{A, F\}, C)$.

We proceed to show that every finite extensive game with perfect information has a subgame-perfect equilibrium.

Theorem 3.6 (Kuhn, 1950). Every finite extensive game with perfect information has a subgame-perfect equilibrium.

Proof. By induction over the height $h(T)$ of T. The theorem is trivial for trees of height 0. For a game with a tree of height k, let $v_1, \dots, v_l$ be the children of the root v_0. By induction, each of the subgames $G^{v_1}, \dots, G^{v_d}$ rooted in $v_1, \dots, v_d$ has a subgame-perfect equilibrium $\mathbf{s}^{v_l}$, $l = 1, \dots, d$. Let i be the player with $v_0 \in V_i$ and let $l^* \in \arg\max_{l \in \{1,\dots,d\}} u_i^{v_l}(\mathbf{s}^{v_l})$. We claim that the strategy vector $\mathbf{s}$ defined

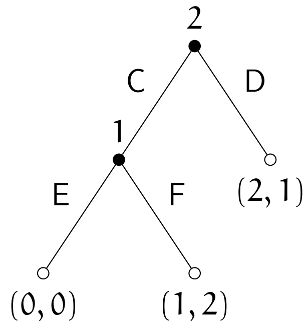

(a) Tree representation

	C	D
E	(0, 0)	(2, 1)
F	(1, 2)	(2, 1)

(b) Strategic representation

Figure 3.3: Subgame of the extensive game in Figure 3.2.

as

$$s_j = \begin{cases} \bigcup_{l=1,\dots,d} s_j^{v_l}, & \text{if } j \neq i \\ \bigcup_{l=1,\dots,d} s_j^{v_l} \cup \{(v_0, v_{l^*})\}, & \text{if } j = i \end{cases}$$

is a subgame-perfect equilibrium. As $\mathbf{s}^{v_{l^*}}$ is a subgame perfect equilibrium for $G^{v_{l^*}}$ there is no profitable deviation in subgames G^v with $v \neq v_0$. We claim that there is no deviation in G as well. For the sake of a contradiction, suppose player i can deviate profitably, i.e., there is $t_i \in S_i$ such that

$$u_i(t_i, \mathbf{s}_{-i}) > u_i(\mathbf{s}).$$

Let $l' \in \{1, \dots, d\}$ be such that $(v_0, v_{l'}) \in t_i$. Then, the strategy $t_i^{v_{l'}} = t_i \setminus \{(v_0, v_{l'})\}$ is a valid strategy of player i in the subgame $G^{v_{l'}}$. This implies

$$u_i(\mathbf{s}) < u_i(t_i, \mathbf{s}_{-i}) = u_i^{v_{l'}}(t_i^{v_{l'}}, \mathbf{s}_{-i}^{v_{l'}}) \leq u_i^{v_{l'}}(\mathbf{s}^{v_{l'}}) \leq \max_{l \in \{1,\dots,d\}} u_i^{v_l}(\mathbf{s}^{v_l}) = u_i(\mathbf{s}),$$

where the second inequality uses that $\mathbf{s}^{v_{l'}}$ is a subgame-perfect equilibrium. □

The proof of Theorem 3.6 implicitly gives rise to an algorithm that computes a subgame-perfect equilibrium, which is known as backward induction, see Algroithm 3.1. The algorithm starts with the non-leaf vertices of maximal height. For each such vertex v, an outgoing edge (v, w^*) is chosen that maximizes the utility $\hat{u}_i(w)$ of the player i with $v \in V_i$. Then, the leaf utility $\hat{\mathbf{u}}(w^*)$ is set to $\hat{\mathbf{u}}(w^*)$ and the algorithm proceeds with vertices of lower hight.
For win-lose-games with two players, where in all outcomes either one of the players wins, or there is a draw, we obtain the following direct corollary.

Input: Extensive game of full information G given by the game tree T, decision vertices $(V_i)_{i\in N}$ and leaf utilities $\hat{\mathbf{u}}$
Output: Subgame-perfect equilibrium $\mathbf{s}$ of G
$s_i = \emptyset$ **for all** $i \in N$
for $k = h(T) - 1, \dots, 1, 0$ **do**
  foreach $v \in V$ with $h(v) = k$ **do**
    $i \leftarrow$ player j such that $v \in V_j$
    choose $w^* \in \arg\max_{w\in V:(v,w)\in E} \hat{u}_i(w)$ arbitrarily
    $s_i \leftarrow s_i \cup \{(v, w^*)\}$
    $\hat{\mathbf{u}}(v) \leftarrow \hat{\mathbf{u}}(w^*)$
  end
end

Algorithm 3.1: Backward induction

Corollary 3.7 (Zermelo, 1913). In finite two-player extensive games of perfect information where all outcomes are such that exactly one of the players wins or there is a draw, either one of the players has a strategy that guarantees a win, or both players have a strategy that guarantees a draw.

Proof. By Theorem 3.6, each such game has a Nash equilibrium. By Theorem 1.25, in an equilibrium, the strategies of player 1 and 2 are maximin and minimax strategies, respectively. □

Chess is essentially is a game as in Corollary 3.7 since each player can enforce a draw one a certain position of the bord is reached three times. Thus, one of the players can enforce to win, or both players can enforce a draw. Even though it seems plausible that playing first is not a disadvantage for white, there is no formal proof for this intuition. So, we cannot rule out the possibility that black can guarantee to win.
For many other games, however, we have a formal proof, that the starting player does not have a disadvantage. As an example, consider the Poisoned-Chocolate-Game.

Example 3.8 (Poisoned-Chocolate-Game). The Poisoned-Chocolate-Game played on a rectangular $m \times n$ grid. The rectangles of the grid correspond to pieces of chocolate and the piece in the lower left is poisoned, see Figure 3.4. Two players take turns choosing a remaining piece of chocolate of the grid. The chosen piece and all pieces that are not lower and not left of that piece are removed. The player who is forced to take the poisoned piece loses.

The following argument has been formalized by John F. Nash for the Hex game. We here give it for the Poisoned-Chocolate-Game.

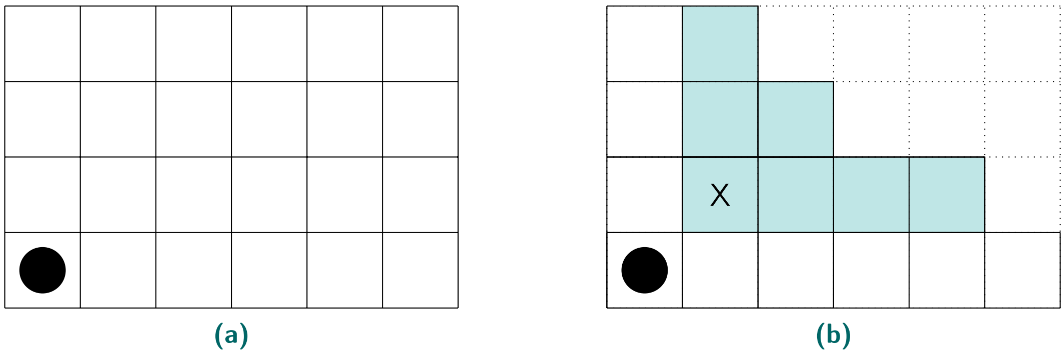

Figure 3.4: **(a)** Starting configuration and **(b)** intermediate configuration of the Poisoned-Chocolate-Game (Example 3.8). Choosing the piece of chocolate marked with X results in removal of all colored pieces of chocolate.

Theorem 3.9 (Strategy Stealing Argument). In the Poisoned-Chocolate-Game with $m + n > 2$, the starting player can guarantee to win.

Proof. The game obviously has no draw, so by Corollary 3.7 either the starting player, say player 1, or the other player, say player 2, can guarantee to win. For the sake of a contradiction, let us assume that player 2 can guarantee to win, i.e., for every first move by player 1 there is a second move by player 2, so that for every third move by player 1, there is a fourth move by player 2, and so on, so that player 1 takes the lower left piece. Imagine player 1 taking the piece that player 2 would take in reaction to player 1 taking the upper right piece. Then, the situation to player 2 prior to the second move looks exactly like the situation to player 1 in the winning strategy of player 2. Hence, player 1 can copy the winning strategy of player 2. This contradicts the assumption that player 2 has a winning strategy. □

For the game Go, the strategy stealing argument can not be applied because the second player typically receives compensation for playing second.

The strategy stealing argument can also be applied to games where a turn of a player is to mark an element from a ground set and the player wins who first marked symbols as in a given set of winning configurations. A famous example of such games are Tic-Tac-Toe games and their generalizations.

Example 3.10 (2×2 Tic-Tac-Toe). Consider the Tic-Tac-Toe game played on a 2×2 grid, see Figure 3.5. Two players X and O alternatingly mark a cell of the 2×2 grid with their sign, starting with player X. The game ends when one of the players has marked a row or a column of the grid.

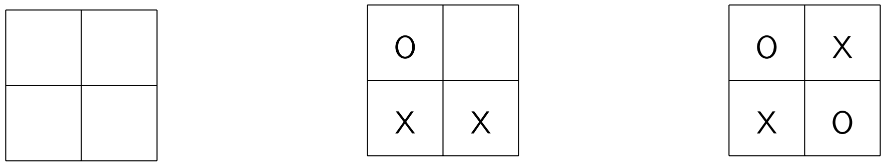

(a) (b) (c)

Figure 3.5: A very simple 2×2 Tic-Tac-Toe game. (a) Starting configuration; (b) A winning configuration for player X; (c) A draw configuration.

3.2 Extensive Games with Imperfect Information

Extensive games with imperfect information differ from games with perfect information in the sense that the players need to make decisions without being informed over all actions taken previously by other players. In an extensive game with perfect information, the vertices of the game tree are partitioned into the players' decision vertices, and each player is fully informed about the position of the vertex in the game tree when making their decision. This is in contrast to a game with incomplete information where players only know that the state of the game is in some vertex $v \in I$ for a so-called information set $I \subseteq V_i$. Formally, we are given an information partition $\mathfrak{I} = (\mathcal{I}_i)_{i \in N}$ where, for each player i, the set $\mathcal{I}_i$ is a partition of V_i, i.e., $\bigcup_{I \in \mathcal{I}_i} I = V_i$ and $I \cap J = \emptyset$ for all $I, J \in \mathcal{I}_i$ with $I \neq J$. In order to make well-defined moves even though the player does not know the current vertex, we need a way to identify the outgoing edges of different vertices v, v' in the same information set I. To this end, we assume that each edge e has a label $L(e)$, that the labels of $\delta^+(v)$ are distinct for all $v \in V$, and that $\bigcup_{e \in \delta^+(v)} L(e) = \bigcup_{e \in \delta^+(v')}$ for all $i \in N$, $I \in \mathcal{I}_i$ and $v, v' \in I$. For a subset $F \subseteq E$ of edges, we write $L(F) = \cup_{e \in F} L(e)$.

Definition 3.11 (Extensive Game with Imperfect Information). For a rooted tree T with vertex partition $(V_i)_{i \in N}$, information partition $\mathfrak{I}$, and leaf utilities $\hat{\mathbf{u}}$, the corresponding extensive game with imperfect information is the strategic game $G(T) = (N, \mathbf{S}, \mathbf{u})$, where

- N is a finite and non-empty set of n players,
- for all players i, we have

$$S_i = \Big\{E_i \subseteq \cup_{v \in V_i} \delta^+(v) : \delta^+(v) \cap E_i = 1 \text{ for all } v \in V_i \\ \text{and } L(\delta^+(v) \cap E_i) = L(\delta^+(v') \cap E_i) \text{ for all } I \in \mathcal{I}_i, v, v' \in I\Big\}.$$

- $u_i(\mathbf{s}) = \hat{u}_i(v(\mathbf{s}))$ for all players i, where $v(\mathbf{s})$ is the unique leaf that is connected to the root in $(V, \cup_{i \in N} S_i)$.

Example 3.12 (Bach-or-Stravinsky with Incomplete Information). In this variant of the Bach-or-Stravinsky game, player 1 first chooses whether the players meet in

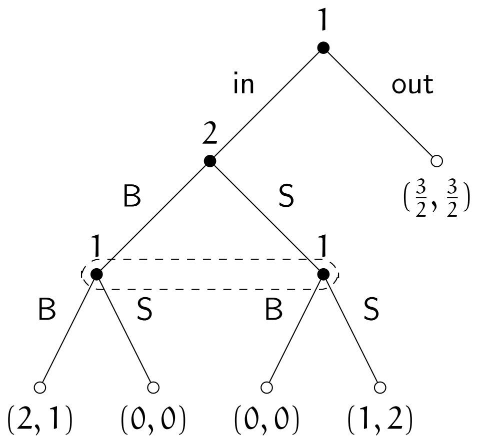

(a) Tree representation

	B	S
{in, B}	(2, 1)	(0, 0)
{in, S}	(0, 0)	(1, 2)
{out, B}	$(\frac{3}{2}, \frac{3}{2})$	$(\frac{3}{2}, \frac{3}{2})$
{out, S}	$(\frac{3}{2}, \frac{3}{2})$	$(\frac{3}{2}, \frac{3}{2})$

(b) Strategic representation

Figure 3.6: An extensive game with incomplete information.

a concert or do not play the game at all giving both players a utility of $3/2$. If player 1 decides that the game is played, they play the simultaneous move Bach-or-Stravinsky game (Example 1.5), see Figure 3.6

Using the information partitions, the class of extensive games with incomplete information is rich enough to encode every strategic game including games without a pure Nash equilibrium like the Matching Pennies Game (Example 1.8). We will thus resort to mixed strategies. There are two natural ways o define mixed strategies.

Definition 3.13 (Mixed Strategy). A mixed strategy of player i in an extensive game is a probability distribution over the set of their pure strategies.

Definition 3.14 (Behavioral Strategy). A behavioral strategy of player i in an extensive game is a collection of independent probability distributions over $L(\delta^+(I))$ for all $I \in \mathcal{I}_i$.

For the game in Figure 3.6, a mixed strategy of player 1 is a probability distribution over the finite set $\big\{\{\text{in}, \text{B}\}, \{\text{in}, \text{S}\}, \{\text{out}, \text{B}\}, \{\text{out}, \text{S}\}\big\}$. A behavioral strategy of player 1 is a pair of probability distributions, one for each information set, i.e., one probability distribution over the finite set $\big\{\text{in}, \text{out}\big\}$, and one over the finite set $\big\{\text{B}, \text{S}\big\}$. Thus, the space of mixed strategies has dimension 3 while the space of behavioral strategies has dimension 2, and the strategy in which $\{\text{out}, \text{B}\}$ and $\{\text{in}, \text{S}\}$ each are played with probability $1/2$ is a mixed strategy that cannot be realized by a behavioral strategy.
We proceed to show that for a large class of extensive games, called games with prefect recall, behavioral strategies are as expressive as mixed strategies.

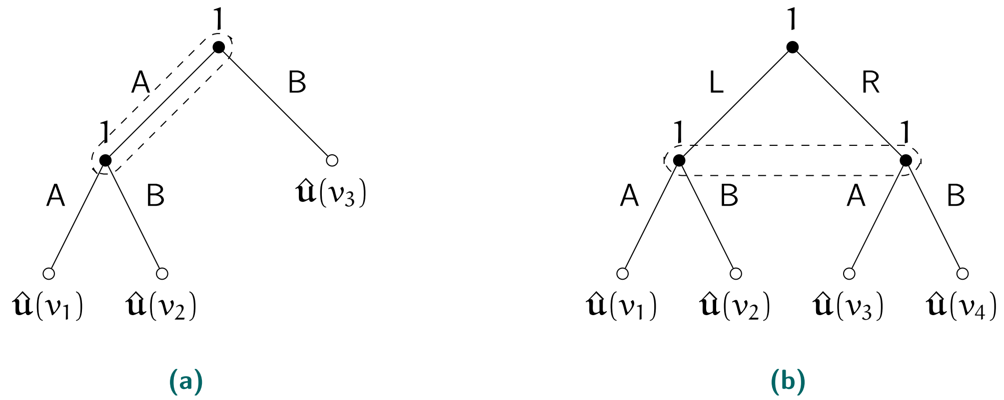

Figure 3.7: Two extensive games in which mixed and behavioral strategies are not outcome equivalent.

> **Definition 3.15** (Perfect Recall). An extensive game with incomplete information has perfect recall, if for each two vertices $v, v' \in I$, $I \in \mathcal{I}$, $i \in N$, the sequences of information sets and actions taken at information sets by player i on $P[v_0, v]$ and $P[v_0, v']$ are the same.

Two games with imperfect recall are given in Figure 3.7. The game in Figure 3.7a, any behavioral strategy can be identified with the value q with with player 1 plays A. For such a strategy, leaves v_1, v_2, and v_3 are reached with probability q^2, $q(1-q)$, and $1-q$, respectively. In contrast, the pure strategies of player 1 are A and B. For the pure strategy A, the leaf v_1 is reached while for the pure strategy B, the leaf v_3 is reached. Thus any mixed strategy corresponds to a respective mixing of outcomes v_1 and v_3.
In the game in Figure 3.7b, there is a mixed strategy in which leafs v_1 and v_4 are reached with probability 1/2 each while any behavioral strategy must also assign positive probability for the leaves v_2 and v_3 when v_1 and v_4 are reached with positive probability.

> **Theorem 3.16** (Kuhn, 1950). For any mixed strategy of a player in a finite extensive game with perfect recall there is an outcome-equivalent behavioral strategy, i.e., each leaf is reached with the same probability.

Proof. Let x_i be a mixed strategy of player i. For a vertex v, let $p_i(v)$ be the sum of the probabilities of pure strategies of player i according to x_i such that vertex v may be reached (for appropriate pure strategies of the other players), i.e.,

$$p_i(v) = \sum\nolimits_{s_i \in S_i : P[v_0,v] \cap \bigcup_{v \in V_i} \delta^+(v) \subseteq s_i} x_{s_i},$$

where we denote by $P[v_0, v]$ the unique path from v_0 to v in T. Let v, v' be two vertices in the same information set I of player i and let $l \in L(I)$ be a label of a corresponding outgoing edge. Since the game has perfect recall, $p_i(v) = p_i(v')$ and since in any pure strategy the edge (v, w) with label l is played if and only if the edge (v', w') with label l is played, we have $p_i(w) = p_i(w')$ for all $w, w' \in V$ such that $L(v, w) = L(v', w') = l$. We define a probability distribution $(q_l)_{l \in L(I)}$ over $L(I)$ by setting

$$q_l = \frac{p_i(w)}{p_i(v)}$$

where $v \in I$ and $w \in V$ with $(v, w) \in E$ and $L(v, w) = l$ are chosen arbitrary. Doing so for all information sets I, we obtain a behavioral strategy q_i. (For information sets I with $p_i(v) = 0$ for all $v \in V$, we may define the corresponding behavioral strategy at I arbitrarily.)

We proceed to show that the behavioral strategy q_i is outcome-equivalent to x_i. To this end, let s_{-i} be a vector of pure strategies of the players other than i, and let $v \in L$ be a leaf of the game tree. consider the subtree $T' = (V, \bigcup_{j \in N \setminus \{i\}} s_i \cup \bigcup_{v \in V_i} \delta^+(v))$. If there is no path from v_0 to v in T', then the probability of reaching v is zero both in x_i and q_i. Next, assume there is such a path in T'. If $p_i(v) = 0$, let v' be the first vertex on the path from v to v_0 with $p_i(v') \neq 0$. Then, the behavioral strategy at the information set I with $v' \in I$ assigns zero probability to the path leading to v as well. Otherwise, if $p_i(v) > 0$, then, $p_i(w) > 0$ for all vertices on $P[v, v_0]$ and the probability of reaching v in q is

$$\prod_{k=0}^{h(v)-1} \frac{p_i(v_{k+1})}{p_i(v_k)} = \frac{p_i(v)}{p_i(v_0)} = p_i(v)$$

where $P[v_0, v] = (v_0, v_1, \ldots, v_{h(v)-1})$ and $v_{v_{h(v)-1}} = v$. □

3.3 Exogenous Uncertainty

Definition 3.17 (Extensive Game with Imperfect Information and Exogenous Uncertainty). An extensive game with imperfect information and exogenous uncertainty is an extensive game with imperfect information in which a special player 0 receives 0 utility from all outcomes and plays according to a fixed behavioral strategy.

Example 3.18 (Beer-Quiche-Game). The Beer-Quiche-Game is a signaling game in which first nature decides whether is of the weak or strong type. After being informed about the type, player 1 decides to have either beer or quiche for breakfast. Player 2 observes the breakfast taken by player 1, but not the type, and chooses to either fight or negotiate with player 1.

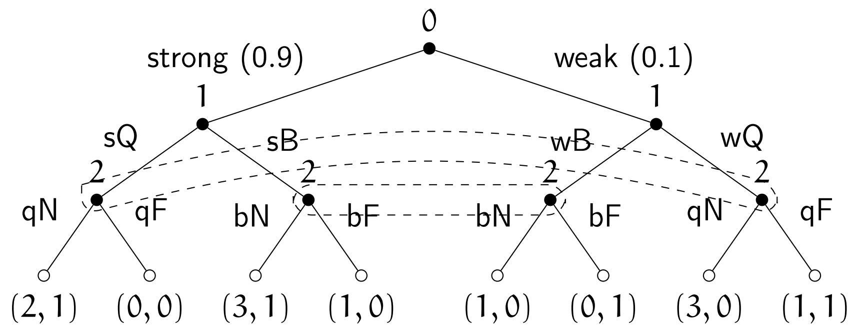

(a) Tree representation

	{bN, qN}	{bN, qF}	{bF, qN}	{bFqF}
{sB, wB}	(2.9, 0.9)	(2.9, 0.9)	(0.9, 0.1)	(0.9, 0.1)
{sB, wQ}	(3.0, 0.9)	(2.8, 1.0)	(1.2, 0)	(1.0, 0.1)
{sQ, wB}	(2.0, 0.9)	(1.2, 0.0)	(1.8, 1)	(0.0, 0.1)
{sQ, wQ}	(2.1, 0.9)	(0.1, 0.1)	(2.1, 0.9)	(0.1, 0.1)

(b) Strategic representation

Figure 3.8: Beer-Quiche-Game

The utility of player 1 is the sum of two components. They receive 1 units of utility for choosing the prefered breakfast (beer for the strong type and quiche for the weak type), and 2 units if player 2 does not fight. Player 2 receives utility 1 if they guessed the type of player 1 correctly (fighting the weak type and negotiating with the strong type).
The game has two Nash equilibria, marked in Figure 3.8b.

Example 3.19 (Kuhn Poker). Two players play a poker game with a deck consisting of the cards King, Queen and Jack. A game consists of the following steps. First, each player puts a blind of 1 coin into the pot. Then, each player is dealt a card where wach player sees only their own card, but neither the card dealt to the other player, nor the remaining card. Then, player 1 has the choice to either bet 1 coin or pass. In both cases, player 2 has to the choice to either bet 1 coin or pass. If both bet or both pass, the game ends in a showdown where the player with the higher card wins the pot. In the case that player 1 bets and player 2 passes, the pot goes to player 1. If player 1 passes and player 2 bets, the third turn is by player 1 again in which they can be for a showdown or pass and forfeit to player 2.
By removing dominated strategies of the players, one can show that player 1 always passes with a queen, player 2 always bets with a king, player 1 never calls with a jack, and always calls with a king.

A behavioral strategy profile is then governed by the following parameters:
– x_J: the probability of player 1 to place a bet with a jack,
– x_Q: the probability of player 1 to call a bet of player 2 when holding a queen,
– x_K: the probability of player 1 to place a bet with a king,
– y_B: the probability of player 2 to place a bet with a jack when player 1 passed,
– y_Q: the probability of player 2 to call with a queen when player 2 placed a bet.
To determine the optimal strategies, we consider all card deals and compute

$$
\begin{aligned}
&\mathbb{P}[-2|JQ] = x_J y_Q && \mathbb{P}[\ \ 1|JQ] = x_J(1-y_Q) && \mathbb{P}[-1|JQ] = 1 - x_J \\
&\mathbb{P}[-2|JK] = x_J && \mathbb{P}[-1|JK] = 1 - x_J && \\
&\mathbb{P}[\ \ 1|QJ] = 1 - y_J && \mathbb{P}[\ \ 2|QJ] = y_J x_Q && \mathbb{P}[-1|QJ] = y_J(1-x_Q) \\
&\mathbb{P}[-2|QK] = x_Q && \mathbb{P}[-1|QK] = 1 - x_Q && \\
&\mathbb{P}[\ \ 2|KJ] = (1-x_K)y_J && \mathbb{P}[\ \ 1|KJ] = 1-(1-x_K)y_J && \\
&\mathbb{P}[\ \ 2|KQ] = x_K y_Q && \mathbb{P}[\ \ 1|KJ] = 1 - x_K y_Q &&
\end{aligned}
$$

Thus, player 1 loses 2 with probability $P_{-2} = \frac{x_J y_Q}{6} + \frac{x_J}{6} + \frac{x_Q}{6}$, wins one with probability $P_1 = \frac{x_J(1-y_Q)}{6} + \frac{1-y_J}{6} + \frac{1-(1-x_K)y_J}{6} + \frac{1-x_K y_Q}{6}$, and wins 2 with probability $P_2 = \frac{y_J x_Q}{6} + \frac{(1-x_K)y_J}{6} + \frac{x_K y_Q}{6}$, and loses 1 in all other cases. We obtain

$$
\begin{aligned}
\mathbb{E}[u_1] &= -2P_{-2} + P_1 + 2P_2 - (1 - P_1 - P_2 - P_{-2}) \\
&= 2P_1 + 3P_2 - P_{-2} - 1 \\
&= \frac{1}{6}\Big(2x_J(1-y_Q) + 2 - 2y_J + 2 - 2(1-x_K)y_J + 2 - 2x_K y_Q \\
&\quad + 3y_J x_Q + 3(1-x_K)y_J + 3x_K y_Q - x_J y_Q - x_J - x_Q\Big) - 1 \\
&= \frac{1}{6}\Big(2x_J(1-y_Q) - 2y_J + 3y_J x_Q + (1-x_K)y_J + x_K y_Q - x_J y_Q - x_J - x_Q\Big) \\
&= \frac{1}{6}\Big(x_J - 3x_J y_Q - y_Q + 3y_J x_Q - x_K y_J + x_K y_Q - x_Q\Big) \\
&= \frac{1}{6}\Big(x_J(1-3y_Q) + x_Q(3y_J - 1) + x_K(y_Q - y_J) - y_Q\Big).
\end{aligned}
$$

This implies that player 1 is indifferent for the choice $y_J = y_Q = 1/3$. The expected utility of player 1 is then $-1/18$. Reformulating this equation, we obtain

$$
\mathbb{E}[u_1] = \frac{1}{6}\Big(y_J(3x_Q - x_K - 1) + y_Q(x_K - 3x_J) + x_J - x_Q\Big).
$$

For arbitrary x_K, we can make player 2 indifferent for $x_J = \frac{x_K}{3}$ and $x_Q = x_J + \frac{1}{3}$. Thus, we obtain a continuum of Nash equilibria where player 2 chooses $y_J = y_Q = 1/3$, and a parameter $\gamma \in [0, 1]$, player 1 chooses $x_J = \frac{\gamma}{3}$, $x_Q = \frac{\gamma+1}{3}$, and $x_K = \gamma$.

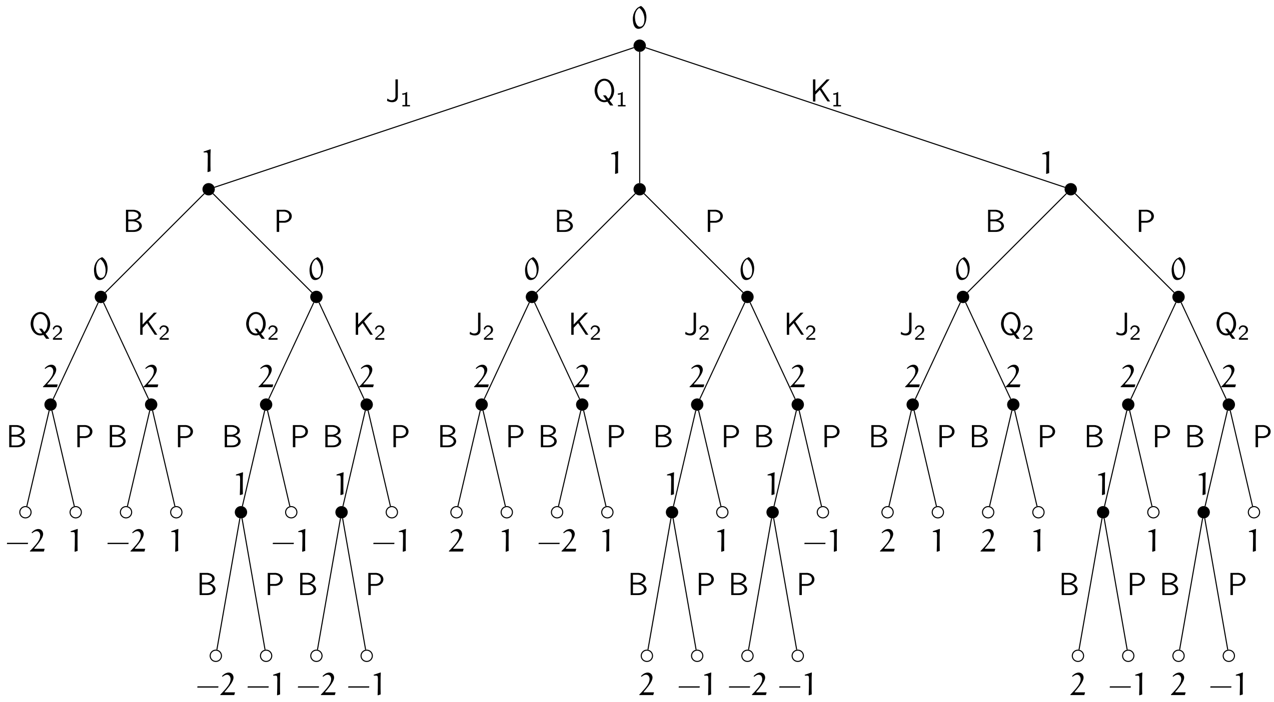

Figure 3.9: Game Tree of a two-player three-card poker game.