3

Game Theory: Basic Concepts

> Each discipline of the social sciences rules comfortably within its own chosen domain... so long as it stays largely oblivious of the others.
>
> Edward O. Wilson (1998):191

3.1 Big John and Little John

Big John and Little John eat coconuts, which dangle from a lofty branch of a palm tree. Their favorite coconut palm produces only one fruit per tree. To get the coconut, at least one of them must climb the tree and knock the coconut loose so that it falls to the ground. Careful energy measurements show that a coconut is worth 10 Kc (kilocalories) of energy, the cost of running up the tree, shaking the coconut loose, and running back down to the ground costs 2 Kc for Big John, but is negligible for Little John, who is much smaller. Moreover, if both individuals climb the tree, shake the coconut loose, then climb down the tree and eat the coconut, Big John gets 7 Kc and Little John gets only 3 Kc, because Big John hogs most of it; if only Big John climbs the tree, while Little John waits on the ground for the coconut to fall, Big John gets 6 Kc and Little John gets 4 Kc (Little John eats some before Big John gets back down from the tree); if only Little John climbs the tree, Big John gets 9 Kc and Little John gets 1 Kc (most of the food is gone by the time Little John gets there).

What will Big John and Little John do if each wants to maximize net energy gain? There is one crucial issue that must be resolved: who decides first what to do, Big John or Little John? There are three possibilities: (a) Big John decides first; (b) Little John decides first; (c) both individuals decide simultaneously. We will go through the three cases in turn.

Assuming Big John decides first, we get the situation depicted in Fig. 3.1. We call a figure like this a *game tree*, and we call the game it defines an extensive form game. At the top of the game tree is the *root node* (the little dot labeled "Big John") with two *branches*, labeled w (wait) and c (climb).

This means Big John gets to choose and can go either left (w) or right (c). This brings us to the two nodes labeled "Little John," in each of which Little John can wait (w) or climb (c).

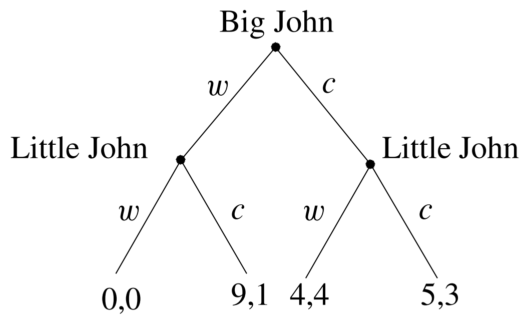

Figure 3.1. Big John and Little John: Big John chooses first.

While Big John has only two strategies, Little John actually has four:

a. Climb no matter what Big John does (cc).
b. Wait no matter what Big John does (ww).
c. Do the same thing Big John does (wc).
d. Do the opposite of what Big John does (cw).

The first letter in parenthesis indicates Little John's move if Big John waits, and the second is Little John's move if Big John climbs.

We call a move taken by a player at a node an *action*, and we call a series of actions that fully define the behavior of a player a *strategy*—actually a *pure strategy*, in contrast to "mixed" and "behavioral" strategies, which we will discuss later, that involve randomizing. Thus, Big John has two strategies, each of which is simply an action, while Little John has four strategies, each of which is two actions—one to be used when Little John goes left, and one when Little John goes right.

At the bottom of the game tree are four nodes, which we variously call *leaf* or *terminal nodes*. At each terminal node is the payoff to the two players, Big John (player 1) first and Little John (player 2) second, if they choose the strategies that take them to that particular leaf. You should check that the payoffs correspond to our description above. For instance, at the leftmost leaf when both wait, with neither John expending or ingesting energy, the payoff is (0,0). At the rightmost leaf both climb the tree, costing Big John 2 Kc, after which Big John gets 7 Kc and Little John gets 3 Kc. Their net payoffs are thus (5,3). And similarly for the other two leaves.

How should Big John decide what to do? Clearly, Big John should figure out how Little John will react to each of Big John's two choices, w and c. If Big John chooses w, then Little John will choose c, because this pays 1 Kc as opposed to 0 Kc. Thus, Big John gets 9 Kc by moving left. If Big John chooses c, Little John will choose w, because this pays 4 Kc as opposed to 3 Kc for choosing c. Thus Big John gets 4 Kc for choosing c, as opposed to 9 Kc for choosing w. We now have answered Big John's problem: choose w.

What about Little John? Clearly, Little John must choose c on the left node, but what should he choose on the right node? Of course it doesn't really matter, because Little John will never *be* at the right node. However, we must specify not only what a player does "along the path of play" (in this case the left branch of the tree), but at *all possible nodes on the game tree*. This is because we can only say for sure that Big John is choosing a best response to Little John if we know what Little John does, and conversely. If Little John makes a wrong choice at the right node, in some games (though not this one) Big John would do better by playing c. In short, Little John must choose one of the four strategies listed above. Clearly, Little John should choose cw (do the opposite of Big John), because this maximizes Little John's payoff no matter what Big John does.

Conclusion: the only reasonable solution to this game is for Big John to wait on the ground, and Little John to do the opposite of what Big John does. Their payoffs are (9,1). We call this a Nash equilibrium (named after John Nash, who invented the concept in about 1950). A Nash equilibrium in a two-player game is a pair of strategies, each of which is a *best response* to the other; i.e., each gives the player using it the highest possible payoff, given the other player's strategy.

There is another way to depict this game, called its *strategic form* or *normal form*. It is common to use both representations and to switch back and forth between them, according to convenience. The normal form corresponding to Fig. 3.1 is in Fig. 3.2. In this example we array strategies of player 1 (Big John) in rows, and the strategies of player 2 (Little John) in columns. Each entry in the resulting matrix represents the payoffs to the two players if they choose the corresponding strategies.

We find a Nash equilibrium from the normal form of the game by trying to pick out a row and a column such that the payoff to their intersection is the highest possible for player 1 down the column, and the highest possible for player 2 across the row (there may be more than one such pair). Note that

		Little John			
		cc	cw	wc	ww
Big John	w	9,1	9,1	0,0	0,0
	c	5,3	4,4	5,3	4,4

Figure 3.2. Normal form of Big John and Little John when Big John moves first.

(w, cw) is indeed a Nash equilibrium of the normal form game, because 9 is better than 4 for Big John down the cw column, and 1 is the best Little John can do across the w row.

Can we find any other Nash equilibria to this game? Clearly (w, cc) is also a Nash equilibrium, because w is a best reply to cc and conversely. But the (w, cc) equilibrium has the drawback that if Big John should happen to make a mistake and play c, Little John gets only 3, whereas with cw, Little John gets 4. We say cc is *weakly dominated* by cw, meaning that cw pays off at least as well for Little John no matter what Big John does, but for at least one move of Big John, cw has a higher payoff than cc for Little John (§4.1).

But what if Little John plays ww? Then Big John should play c, and it is clear that ww is a best response to c. So this gives us another Nash equilibrium, (c, ww), in which Little John does much better, getting 4 instead of 1, and Big John does much worse, getting 4 instead of 9. Why did we not see this Nash equilibrium in our analysis of the extensive form game? The reason is that (c, ww) involves Little John making an *incredible threat* (see §4.2 for a further analysis of Little John's incredible threat).

"I don't care what you do, says Little John—I'm waiting here on the ground—no matter what." The threat is "incredible" because Big John knows that if he plays w, then when it is Little John's turn to carry out the threat to play w, Little John will not in fact do so, simply because 1 is better than 0.[1] We say a Nash equilibrium of an extensive form game is *subgame perfect* if, at any point in the game tree, the play dictated by the Nash equilibrium *remains* a Nash equilibrium of the subgame. The strategy (c, ww) is not subgame perfect because in the subgame beginning with Little John's choice of w on the left of Fig. 3.1 is not a best response. Nice try, anyway, Little John!

[1]This argument fails if the individuals can condition his behavior in one day on their behavior in previous days (see chapter 9). We assume the players cannot do this.

But what if Little John gets to choose first? Perhaps now Little John can force a better split than getting 1 compared to Big John's 9. This is the extensive form game (Fig. 3.3). We now call Little John player 1 and Big John player 2. Now Big John has four strategies (the strategies that belonged to Little John in the previous version of the game) and Little John only has two (the ones that belonged to Big John before). Little John notices that Big John's best response to w is c, and Big John's best response to c is w. Because Little John gets 4 in the first case and only 1 in the second, Little John chooses w. Big John's best choice is then cw, and the payoffs are (4,4). Note that *by going first, Little John is able to precommit to a strategy that is an incredible threat when going second.*

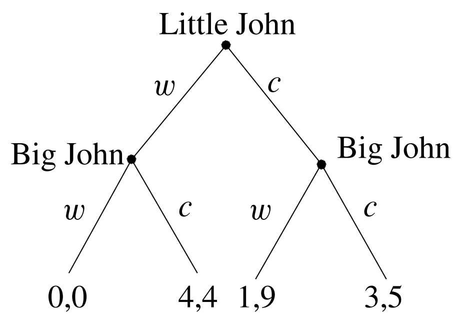

Figure 3.3. Big John and Little John: Little John chooses first.

The normal form for the case when Little John goes first is illustrated in Fig. 3.4. Again we find the two Nash equilibria (w, cc) and (w, cw), and again we find another Nash equilibrium not evident at first glance from the game tree: now it is Big John who has an incredible threat, by playing ww, to which Little John's best response is c.

The final possibility is that the players choose simultaneously or, equivalently, each player chooses an action without seeing what the other player chooses. In this case, each player has two options: climb the tree (c), or wait on the ground (w). We then get the situation in Fig. 3.5. Note the new element in the game tree: the dotted line connecting the two places where Little John chooses. This is called an *information set*. Roughly speaking, an information set is a set of nodes at which (a) the same player chooses, and (b) the player choosing does not know which particular node represents the actual choice node. Note also that we could just as well interchange Big John and Little John in the diagram, reversing their payoffs at the terminal nodes, of course. This illustrates an important point: there may be more than one extensive form game representing the same real strategic situation.

		Big John			
		cc	cw	wc	ww
Little John	w	4,4	4,4	0,0	0,0
	c	3,5	1,9	3,5	1,9

Figure 3.4. Normal form of Big John and Little John game when Little John moves first.

Even though there are fewer strategies in this game, it is hard to see what an equilibrium might be by looking at the game tree. This is because what Little John does cannot depend on which choice Big John makes, because Little John does not see Big John's choice. So let's look at the normal form game, in Fig. 3.6. From this figure, it is easy to see that both (w, c) and (c, w) are Nash equilibria, the first obviously favoring Big John and the second favoring Little John. In fact, there is a third Nash equilibrium that is more difficult to pick out. In this equilibrium Big John randomizes by choosing c and w with probability $1/2$, and Little John does the same. This is called a *mixed strategy Nash equilibrium*; you will learn how to find and analyze it in §3.7. In this equilibrium Big John has payoff 4.5 and Little John has payoff 2. The reason for this meager total payoff is that with probability $1/4$, both wait and get zero reward, and sometimes both climb the tree!

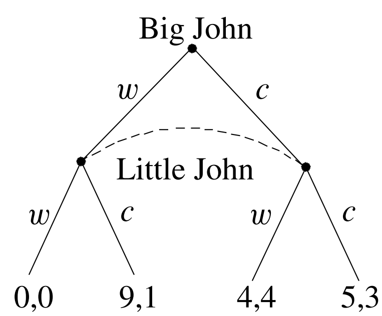

Figure 3.5. Big John and Little John choose simultaneously.

		Little John	
		c	w
Big John	c	5,3	4,4
	w	9,1	0,0

Figure 3.6. Big John and Little John: normal form in the simultaneous move case.

3.2 The Extensive Form

An *extensive form game* $\mathcal{G}$ consists of a number of *players*, a *game tree*, and a set of *payoffs*. A game tree consists of a number of *nodes* connected by *branches*. Each branch connects a *head node* to a distinct *tail node* If b is a branch of the game tree, we denote the head node of b by b^h, and the tail node of b by b^t.

A *path* from node a to node a' in the game tree is a sequence of branches starting at a and ending at a'.[2] If there is a path from node a to a', we say a is an *ancestor* of a', and a' is a *successor* to a. We call k the *length* of the path. If a path from a to a' has length one, we call a the *parent* of a', and a' is a *child* of a.

We require that the game tree have a unique node r, called the *root node*, that has no parent, and a set T of nodes called *terminal nodes* or *leaf nodes*, that have no children. We associate with each terminal node $t \in T$ ($\in$ means "is an element of"), and each player i, a *payoff* $\pi_i(t) \in \mathbf{R}$ ($\mathbf{R}$ is the set of real numbers). We say the game is *finite* if it has a finite number of nodes. We assume all games are finite, unless otherwise stated.

We also require that the graph of $\mathcal{G}$ have the following *tree property*. There must be *exactly one path* from the root node to any given terminal node in the game tree. Equivalently, *every node except the root node has exactly one parent*.

Players relate to the game tree as follows. Each nonterminal node is assigned to a player who moves at that node. Each branch b with head node b^h node represents a particular *action* that the player assigned to b^h that node can take there, and hence determines either a terminal node or the next point of play in the game—the particular child node b^t to be visited next.[3]

If a stochastic event occurs at a node a (for instance, the weather is Good or Bad, or your partner is Nice or Nasty), we assign the fictitious player "Nature" to that node, the actions Nature takes representing the possible outcomes of the stochastic event, and we attach a *probability* to each branch of which a is the head node, representing the probability that Nature chooses that branch (we assume all such probabilities are strictly positive).

[2]Technically, a path is a sequence $b_1, \ldots, b_k$ of branches such that $b_1^h = a$, $b_i^t = b_{i+1}^h$ for $i = 1, \ldots k-1$, and $b_k^t = a'$; i.e., the path starts at a, the tail of each branch is the head of the next branch, and the path ends at a'.

[3]Thus if $\mathbf{p} = (b_1, \ldots, b_k)$ is a path from a to a', then starting from a, if the actions associated with the b_j are taken by the various players, the game moves to a'.

The tree property thus means that there is a *unique* sequence of moves by the players (including Nature) leading from the root node to any specific node of the game tree, and for any two nodes, there is *at most one* sequence of player moves leading from the first to the second.

A player may know the exact node in the game tree when it is his turn to move (e.g., the first two cases in Big John and Little John, above), but he may know only that he is at one of several possible nodes. This is the situation Little John faces in the simultaneous choice case (Fig. 3.6). We call such a collection of nodes an *information set*. For a set of nodes to form an information set, the same player must be assigned to move at each of the nodes in the set and have the same array of possible actions at each node.

We also require that if two nodes a and a' are in the same information set for a player, the moves that player made up to a and a' must be the same. This criterion is called *perfect recall*, because if a player never forgets his moves, he cannot make two different choices that subsequently land him in the same information set.[4]

Suppose each player $i = 1, \ldots, n$ chooses strategy s_i. We call $s = (s_1, \ldots, s_n)$ a *strategy profile* for the game, and we define the *payoff* to player i, given strategy profile s, as follows. If there are no moves by Nature, then s determines a unique path through the game tree, and hence a unique terminal node $t \in T$. The payoff $\pi_i(s)$ to player i under strategy profile s is then defined to be simply $\pi_i(t)$.

Suppose there are moves by Nature, by which we mean that a one or more nodes in the game tree, there is a lottery over the various branches emanating from that node, rather than a player choosing at that node. For every terminal node $t \in T$, there is a unique path $\mathbf{p}_t$ in the game tree from the root node to t. We say $\mathbf{p}_t$ is *compatible* with strategy profile s if, for every branch b on $\mathbf{p}_t$, if player i moves at b^h (the head node of b), then s_i chooses action b at b^h. If $\mathbf{p}_t$ is not compatible with s, we write $p(s,t) = 0$. If $\mathbf{p}_t$ is compatible with s, we define $p(s,t)$ to be the product of all the probabilities associated with the nodes of $\mathbf{p}_t$ at which Nature moves along $\mathbf{p}_t$, or 1 if Nature makes no moves along $\mathbf{p}_t$. We now define the payoff to

[4]Another way to describe perfect recall is to note that the information sets $\mathcal{N}_i$ for player i are the nodes of a graph in which the children of an information set $\nu \in \mathcal{N}_i$ are the $\nu' \in \mathcal{N}_i$ that can be reached by one move of player i, plus some combination of moves of the other players and Nature. Perfect recall means that this graph has the tree property.

player i as

$$\pi_i(s) = \sum_{t \in T} p(s,t)\pi_i(t). \tag{3.1}$$

Note that this is the expected payoff to player i given strategy profile s, assuming that Nature's choices are independent, so that $p(s,t)$ is just the probability that path $\mathbf{p}_t$ is followed, given strategy profile s. We generally assume in game theory that players attempt to maximize their expected payoffs, as defined in (3.1).

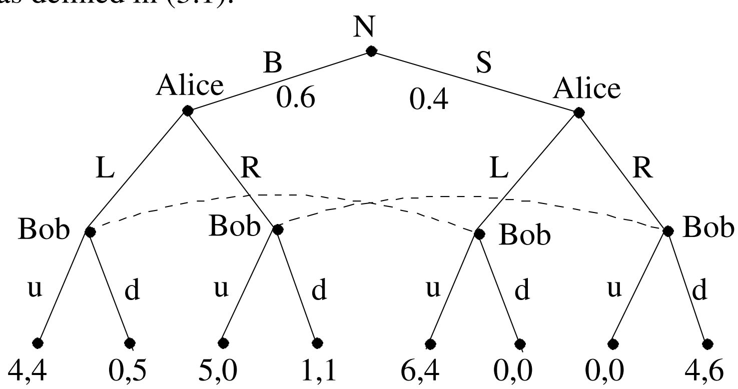

Figure 3.7. Evaluating Payoffs when Nature Moves

For example, consider the game depicted in Figure 3.7. Here, Nature moves first, and with probability $p_l = 0.6$ goes B where the game between Alice and Bob is known as the Prisoner's Dilemma (§3.11), and with probability $p_l = 0.4$ goes S, where the game between Alice and Bob is known as the Battle of the Sexes (§3.9). Note that Alice knows Nature's move, because she has separate information sets on the two branches where Nature moves, but Bob does not, because when he moves, he does not know whether he is on the left or right hand branch. If we write $\pi_A(x,y,z)$ and $\pi_B(x,y,z)$ for the payoffs to Alice and Bob, respectively, when Alice plays $x \in \{L,R\}$, Bob plays $y \in \{u,d\}$, and Nature plays $z \in \{B,S\}$, then (3.1) gives, for instance

$$\pi_A(L,u) = p_u\pi_A(L,u,B) + p_r\pi_A(L,u,S) = 0.6(4) + 0.4(6) = 4.8;$$
$$\pi_B(L,u) = p_u\pi_B(L,u,B) + p_r\pi_B(L,u,S) = 0.6(4) + 0.4(4) = 4.0;$$
$$\pi_A(R,u) = p_u\pi_A(R,u,B) + p_r\pi_A(R,u,S) = 0.6(5) + 0.4(0) = 3.0;$$
$$\pi_B(R,u) = p_u\pi_B(R,u,B) + p_r\pi_B(R,u,S) = 0.6(0) + 0.4(0) = 0;$$

The reader should fill in the payoffs at the remaining nodes.

3.3 The Normal Form

The *strategic form* or *normal form* game consists of a number of players, a set of strategies for each of the players, and a payoff function that associates a payoff to each player with a choice of strategies by each player. More formally, an n-player normal form game consists of

a. A set of *players* $i = 1, \dots, n$.
b. A set S_i of *strategies* for player $i = 1, \dots, n$. We call $s = (s_1, \dots, s_n)$, where $s_i \in S_i$ for $i = 1, \dots, n$, a *strategy profile* for the game.[5]
c. A function $\pi_i : S \to \mathbf{R}$ for player $i = 1, \dots, n$, where S is the set of strategy profiles, so $\pi_i(s)$ is player i's payoff when strategy profile s is chosen.

Two extensive form games are said to be *equivalent* if they correspond to the same normal form game, except perhaps for the labeling of the actions and the naming of the players. But given an extensive form game, how exactly do we form the corresponding normal form game? First, the players in the normal form are the same as the players in the extensive form. Second, for each player i, let S_i be the set of strategies of that player, each strategy consisting of a choice of an action at each information set where i moves. Finally, the payoff functions are given by equation (3.1). If there are only two players and a finite number of strategies, we can write the payoff function in the form of a matrix, as in Fig. 3.2.

As an exercise, you should work out the normal form matrix for the game depicted in Figure 3.7.

3.4 Mixed Strategies

Suppose a player has pure strategies $s_1, \dots, s_k$ in a normal form game. A *mixed strategy* for the player is a probability distribution over $s_1, \dots, s_k$; i.e., a mixed strategy has the form

$$\sigma = p_1 s_1 + \dots + p_k s_k,$$

where $p_1, \dots p_k$ are all nonnegative and $\sum_1^n p_j = 1$. By this we mean that the player chooses s_j with probability p_j, for $j = 1, \dots, k$. We call p_j the *weight* of s_j in σ. If all the p_j's are zero except one, say $p_l = 1$, we say σ

[5]Technically, these are *pure strategies*, because later we will consider *mixed strategies* that are probabilistic combinations of pure strategies.

is a *pure strategy*, and we write $\sigma = s_l$. We say that pure strategy s_j is *used* in mixed strategy σ if $p_j > 0$. We say a strategy is *strictly mixed* if it is not pure, and we say that it is *completely mixed* if it uses all pure strategies. We call the set of pure strategies used in a mixed strategy σ_i the *support* of σ_i.

In an n-player normal form game where, for $i = 1, \ldots, n$, player i has pure strategy set S_i a *mixed strategy profile* $\sigma = (\sigma_1, \ldots, \sigma_n)$ is the choice of a mixed strategy σ_i by each player. We define the *payoffs* to σ as follows. Let $\pi_i(s_1, \ldots, s_n)$ be the payoff to player i when players use the pure strategy profile $(s_1, \ldots, s_n)$, and if s is a pure strategy for player i, let p_s be the weight of s in σ_i. Then we define

$$\pi_i(\sigma) = \sum_{s_1 \in S_1} \ldots \sum_{s_n \in S_n} p_{s_1} p_{s_2} \ldots p_{s_n} \pi_i(s_1, \ldots, s_n).$$

This is a formidable expression, but the idea behind it is simple. We assume the players' choices are made independently, so the probability that the particular pure strategies $s_1 \in S_1, \ldots, s_n \in S_n$ will be used is simply the product $p_{s_1} \ldots p_{s_n}$ of their weights, and the payoff to player i in this case is just $\pi_i(s_1, \ldots, s_n)$. We get the expected payoff by multiplying and adding up over all n-tuples of mixed strategies.

3.5 Nash Equilibrium

The concept of a Nash equilibrium of a game is formulated most easily in terms of the normal form. Suppose the game has n players, with strategy sets S_i and payoff functions $\pi_i : S \rightarrow \mathbf{R}$, for $i = 1, \ldots, n$ where S is the set of strategy profiles. We use the following very useful notation. Let ΔS_i be the set of mixed strategies for player i, and let $\Delta^* S = \Delta S_1 \times \ldots \times \Delta S_n$. If $\sigma \in \Delta^* S$, we write σ_i for the ith component of σ (i.e., σ_i is player i mixed strategy in σ). If $\sigma \in \Delta^* S$, and $\tau_i \in \Delta S_i$, we write

$$(\sigma_{-i}, \tau_i) = (\tau_i, \sigma_{-i}) = \begin{cases} (\tau_1, \sigma_2, \ldots, \sigma_n) & \text{if } i = 1 \\ (\sigma_1, \ldots, \sigma_{i-1}, \tau_i, \sigma_{i+1}, \ldots, \sigma_n) & \text{if } 1 < i < n \\ (\sigma_1, \ldots, \sigma_{n-1}, \tau_n) & \text{if } i = n \end{cases}.$$

In other words, (σ_{-i}, τ_i) is the strategy profile obtained by replacing σ_i with τ_i for player i.

We say a strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \Delta^* S$ is a *Nash equilibrium* if, for every player $i = 1, \ldots, n$, and every $\sigma_i \in \Delta S_i$, we have $\pi_i(\sigma^*) \geq \pi_i(\sigma_{-i}^*, \sigma_i)$; i.e., choosing σ_i^* is at least as good for player i as choosing any other σ_i given that the other players choose σ^*_{-i}. Note that in a Nash equilibrium, the strategy of each player is a *best response* to the strategies chosen by all the other players. Finally, notice that a player could have responses that are *equally good* as the one chosen in the Nash equilibrium—there just cannot be a strategy that is strictly better.

The Nash equilibrium concept is important because in many cases we can accurately (or reasonably accurately) predict how people will play a game by assuming they will choose strategies that implement a Nash equilibrium. It will also turn out that, in dynamic games that model an evolutionary process whereby successful strategies drive out unsuccessful ones over time, stable stationary states are always Nash equilibria. Conversely, we will see that Nash equilibria that seem implausible are actually *unstable* equilibria of an evolutionary process, so we would not expect to see them in the real world. Where people appear to deviate systematically from implementing Nash equilibria, we will sometimes find that they do not understand the game, we have misspecified the game they are playing or the payoffs we attribute to them. But, in important cases, as we shall see in later chapters, people simply do not play Nash equilibria at all.

3.6 The Fundamental Theorems of Game Theory

John Nash showed that every finite game has a Nash equilibrium in mixed strategies (Nash 1950). More concretely, we have

THEOREM 3.1 Nash Existence Theorem. If each player in an n-player game has a finite number of pure strategies, then the game has a (not necessarily unique) Nash equilibrium in (possibly) mixed strategies.

The following Fundamental Theorem of Mixed Strategy Equilibrium develops the principles for finding Nash equilibria. Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a mixed strategy profile for an n-player game. For any player $i = 1, \ldots, n$, let σ_{-i} represent the mixed strategies used by all the players other than player i. The *Fundamental Theorem of Mixed Strategy Nash Equilibrium* says that σ is a Nash equilibrium if and only if, for any player $i = 1, \ldots, n$ with pure strategy set S_i,

a. If $s, s' \in S_i$ occur with positive probability in σ_i, then the payoffs to s and s', when played against σ_{-i}, are equal.

b. If s occurs with positive probability in σ_i and s' occurs with zero probability in σ_i, then the payoff to s' is less than or equal to the payoff to s, when played against σ_{-i}.

The proof of the Fundamental Theorem is straightforward. Suppose σ is the player's mixed strategy in a Nash equilibrium that uses s with probability $p > 0$ and s' with probability $p' > 0$. If s has a higher payoff than s' when played against σ_{-i}, then i's mixed strategy that uses s with probability $(p + p')$, does not use s', and assigns the same probabilities to the other pure strategies as does σ, has a higher payoff than σ, so σ is not a best response to σ_{-i}. This is a contradiction, which proves the assertion. The rest of the proof is similar.

3.7 Solving for Mixed Strategy Nash Equilibria

	L	R
U	a_1, a_2	b_1, b_2
D	c_1, c_2	d_1, d_2

This problem asks you to apply the general method of finding mixed strategy equilibria in normal form games. Consider the game at the right. First, of course, you should check for pure strategy equilibria. To check for a completely mixed strategy equilibrium, we use the Fundamental Theorem (3.6). Suppose the column player uses the strategy $\sigma = \alpha L + (1-\alpha)R$ (i.e., plays L with probability α). Then, if the row player uses both U and D, they must both have the same payoff against σ. The payoff to U against σ is $\alpha a_1 + (1-\alpha)b_1$, and the payoff to D against σ is $\alpha c_1 + (1-\alpha)d_1$. Equating these two, we find

$$\alpha = \frac{d_1 - b_1}{d_1 - b_1 + a_1 - c_1}.$$

For this to make sense, the denominator must be nonzero, and the right hand side must lie between zero and one. Note that *column* player's strategy is determined by the requirement that *row* player's two strategies be equal.

Now suppose the row player uses strategy $\tau = \beta U + (1-\beta)D$ (i.e., plays U with probability β). Then, if the column player uses both L and R, they must both have the same payoff against τ. The payoff to L against τ is $\beta a_2 + (1-\beta)c_2$, and the payoff to R against τ is $\beta b_2 + (1-\beta)d_2$. Equating these two, we find

$$\beta = \frac{d_2 - c_2}{d_2 - c_2 + a_2 - b_2}.$$

Again, for this to make sense, the denominator must be nonzero, and the right-hand side must lie between zero and one. Note that now *row* player's strategy is determined by the requirement that *column* player's two strategies are equal.

a. Suppose the above really is a mixed strategy equilibrium. What are the payoffs to the two players?
b. Note that to solve a 2×2 game, we have checked for five different "configurations" of Nash equilibria—four pure and one mixed. But there are four more possible configurations, in which one player uses a pure strategy and the second player uses a mixed strategy. Show that if there is a Nash equilibrium in which the row player uses a pure strategy (say UU) and the column player uses a completely mixed strategy, then *any* strategy for the column player is a best response to UU.
c. How many different configurations are there to check for in a 2×3 game? In a 3×3 game?
d. Can you generalize to the number of possible configurations of Nash equilibria in an $n \times m$ normal form game?

3.8 Throwing Fingers

	c_1	c_2
c_1	$1,-1$	$-1,1$
c_2	$-1,1$	$1,-1$

Alice and Bob each throws one (c_1) or two (c_2) fingers, simultaneously. If they are the same, Alice wins; otherwise, Bob wins. The winner takes \$1 from the loser. The normal form of this game is depicted to the right. There are no pure strategy equilibria, so suppose player 2 uses the mixed strategy σ that consists of playing c_1 (one finger) with probability α and c_2 (two fingers) with probability $1-\alpha$. We write this as $\sigma = \alpha c_1 + (1-\alpha)c_2$. If Alice uses both c_1 (one finger) and c_2 (two fingers) with positive probability, they both must have the same payoff against σ, or else Alice should drop the lower-payoff strategy and use only the higher-payoff strategy. The payoff to c_1 against σ is $\alpha \cdot 1 + (1-\alpha) \cdot -1 = 2\alpha - 1$, and the payoff to c_2 against σ is $\alpha \cdot -1 + (1-\alpha) \cdot 1 = 1 - 2\alpha$. If these are equal, then $\alpha = 1/2$. A similar reasoning shows that Alice chooses each strategy with probability 1/2. The expected payoff to Alice is then $2\alpha - 1 = 1 - 2\alpha = 0$, and the same is true for Bob.

3.9 Battle of the Sexes

Violetta and Alfredo love each other so much that they would rather be together than apart. But Alfredo wants to go gambling, and Violetta wants to go to the opera. Their payoffs are described to the right. There are two pure strategy equilibria and one mixed strategy equilibrium for this game. We will show that Alfredo and Violetta would be better off if they stuck to either of their pure strategy equilibria.

Alfredo \ Violetta	g	o
g	2,1	0,0
o	0,0	1,2

Let α be the probability of Alfredo going to the opera, and let β be the probability of Violetta going to the opera. Because in a mixed strategy equilibrium, the payoff to gambling and opera must be equal for Alfredo, we must have $\beta = 2(1-\beta)$, which implies $\beta = 2/3$. Because the payoff to gambling and opera must also be equal for Violetta, we must have $2\alpha = 1-\alpha$, so $\alpha = 1/3$. The payoff of the game to each is then

$$\frac{2}{9}(1{,}2) + \frac{5}{9}(0{,}0) + \frac{2}{9}(2{,}1) = \left(\frac{2}{3},\frac{2}{3}\right),$$

because both go gambling $(1/3)(2/3) = 2/9$ of the time, both go to the opera $(1/3)(2/3) = 2/9$ of the time, and otherwise they miss each other.

Both players do better if they can coordinate, because (2,1) and (1,2) are both better than (2/3,2/3).

We get the same answer if we find the Nash equilibrium by finding the intersection of the players' best response function. To see this, note that the payoffs to the two players are

$$\begin{aligned}\pi_A &= \alpha\beta + 2(1-\alpha)(1-\beta) = 3\alpha\beta - 2\alpha - 2\beta + 2\\ \pi_V &= 2\alpha\beta + (1-\alpha)(1-\beta) = 3\alpha\beta - \alpha - \beta + 1.\end{aligned}$$

Thus,

$$\frac{\partial \pi_A}{\partial \alpha} = 3\beta - 2 \begin{cases} > 0 & \text{if } \beta > 2/3 \\ = 0 & \text{if } \beta = 2/3 \\ < 0 & \text{if } \beta < 2/3 \end{cases},$$

so the optimal α is given by

$$\alpha = \begin{cases} 1 & \text{if } \beta > 2/3 \\ [0,1] & \text{if } \beta = 2/3 \\ 0 & \text{if } \beta < 2/3 \end{cases}.$$

Similarly,

$$\frac{\partial \pi_V}{\partial \beta} = 3\alpha - 1 \begin{cases} > 0 & \text{if } \alpha > 1/3 \\ = 0 & \text{if } \alpha = 1/3 \\ < 0 & \text{if } \alpha < 1/3 \end{cases},$$

so the optimal β is given by

$$\beta = \begin{cases} 1 & \text{if } \alpha > 1/3 \\ [0,1] & \text{if } \alpha = 1/3 \\ 0 & \text{if } \alpha < 1/3 \end{cases}.$$

This gives the diagram depicted in Figure 3.8. Note that the three Nash equilibria are the three intersections of the two best response curves, marked by large dots in the figure.

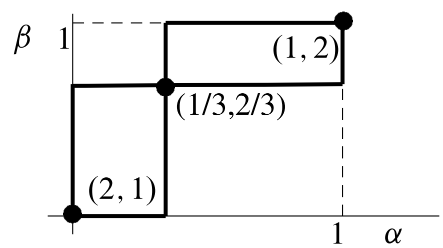

Figure 3.8. Nash equilibria in the Battle of the Sexes

3.10 The Hawk-Dove Game

Consider a population of birds that fight over valuable territory. There are two possible strategies. The "hawk" (H) strategy is to escalate battle until injured or your opponent retreats. The "dove" (D) strategy is to display hostility but retreat before sustaining injury if your opponent escalates. The payoff matrix is given in the figure, where $v > 0$ is the value of territory and $w > v$ is the cost of injury. This figure shows only the payoff to player 1, because the payoff to player 2 is the payoff to player 1 symmetrically across the diagonal. The birds can play mixed strategies, but they cannot condition their play on whether they are player one or player two, and hence both players must used the same strategy.

	H	D
H	$(v-w)/2$	v
D	0	$v/2$

As an exercise, explain the entries in the payoff matrix, show that there are no symmetrical pure strategy Nash equilibria. We will find three Nash equilibria to this game, but two are not symmetric, in the sense that they depend on players making different choices when they are player 1 vs. when they are player 2. There is only one symmetric Nash equilibrium, in which players do not condition their behavior on whether they are player 1 or player 2. This is the game's unique mixed strategy Nash equilibrium.

Let α be the probability of playing Hawk if you are player 1, and let β be the probability of playing Hawk if you are player 2. The payoffs to the two players are

$$\pi_1 = \alpha\beta(v-w)/2 + \alpha(1-\beta)v + (1-\alpha)\beta(0) + (1-\alpha)(1-\beta)v/2$$
$$\pi_2 = \alpha\beta(v-w)/2 + \alpha(1-\beta)(0) + (1-\alpha)\beta v + (1-\alpha)(1-\beta)v/2,$$

which simplifies to

$$\pi_1 = \frac{1}{2}(v(1+\alpha-\beta) - w\alpha\beta)$$
$$\pi_2 = \frac{1}{2}(v(1-\alpha+\beta) - w\alpha\beta)$$

Thus,

$$\frac{\partial\pi_1}{\partial\alpha} = (v-w\beta)/2 \begin{cases} >0 & \text{if } \beta < v/w \\ =0 & \text{if } \beta = v/w \\ <0 & \text{if } \beta > v/w \end{cases},$$

so the optimal α is given by

$$\alpha = \begin{cases} 1 & \text{if } \beta < v/w \\ [0,1] & \text{if } \beta = v/w \\ 0 & \text{if } \beta > v/w \end{cases}$$

Similarly,

$$\frac{\partial\pi_2}{\partial\beta} = (v-w\alpha)/2 \begin{cases} >0 & \text{if } \alpha < v/w \\ =0 & \text{if } \alpha = v/w \\ <1 & \text{if } \alpha > v/w \end{cases},$$

so the optimal β is given by

$$\beta = \begin{cases} 0 & \text{if } \alpha > v/w \\ [0,1] & \text{if } \alpha = v/w \\ 1 & \text{if } \alpha < v/w \end{cases}.$$

This gives the diagram depicted in Figure 3.9. The best response functions intersect in three places, each of which is a Nash equilibrium. However, the only symmetric Nash equilibrium, in which the players cannot condition their move on whether they are player 1 or player 2, is the mixed strategy Nash equilibrium $(v/w, v/w)$.

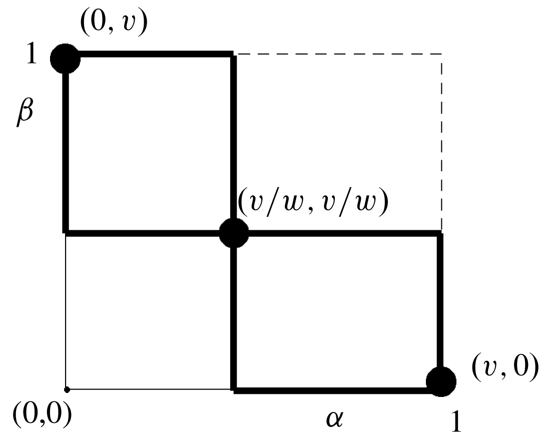

Figure 3.9. Nash equilibria in the Hawk-Dove Game

3.11 The Prisoner's Dilemma

	C	D
C	R,R	S,T
D	T,S	P,P

Alice and Bob can each earn a profit R if they both work hard (pure strategy C). However, either can shirk by working secretly on private jobs (pure strategy D), earning $T > R$, but the other player will earn only $S < R$. If both shirk, however, they will each earn P, where $S < P < R$. Each must decide independently of the other whether to choose C or D. The game tree is depicted in the figure on the right. The payoff T stands for the 'temptation' to defect on a partner, S stands for 'sucker' (for cooperating when your partner defected), P stands for 'punishment' (for both shirking), and R stands for 'reward' (for both cooperating). We usually assume also that $S + T < 2R$, so there is no gain from "taking turns" playing C (Cooperate) and D (Defect).

Let α be the probability of playing C if you are Alice, and let β be the probability of playing C if you are Bob. To simplify the algebra, we assume $P = 1$, $R = 0$, $T = 1+t$ and $S = -s$, where $s, t > 0$. It is easy to see that these assumptions involve no loss of generality, because adding a constant to all payoffs, or multiplying all payoffs by a positive constant does not change the Nash equilibria of the game. The payoffs to Alice and Bob are now

$$\pi_A = \alpha\beta + \alpha(1-\beta)(-s) + (1-\alpha)\beta(1+t) + (1-\alpha)(1-\beta)(0)$$
$$\pi_B = \alpha\beta + \alpha(1-\beta)(1+t) + (1-\alpha)\beta(-s) + (1-\alpha)(1-\beta)(0),$$

which simplify to

$$\pi_A = \beta(1+t) - \alpha(s(1-\beta) + \beta t)$$
$$\pi_B = \alpha(1+t) - \beta(s(1-\alpha) + \alpha t).$$

It is clear from these equations that π_A is maximized by choosing $\alpha = 0$, no matter what Bob does, and similarly π_B is maximized by choosing $\beta = 0$, no matter what Alice does. This is the mutual defect equilibrium.

This is not how many people play this game in the experimental laboratory. Rather, people very often prefer to cooperate, provided their partners cooperate as well (Kiyonari, Tanida and Yamagishi 2000). We can capture this phenomonon by assuming that there is a psychic gain $\lambda_A > 0$ for Alice and $\lambda_B > 0$ for Bob when both players cooperate, above the temptation payoff $T = 1 + t$. If we rewrite the payoffs using this assumption, we get

$$\pi_A = \alpha\beta(1+t+\lambda_A) + \alpha(1-\beta)(-s) + (1-\alpha)\beta(1+t) + (1-\alpha)(1-\beta)(0)$$
$$\pi_B = \alpha\beta(1+t+\lambda_B) + \alpha(1-\beta)(1+t) + (1-\alpha)\beta(-s) + (1-\alpha)(1-\beta)(0),$$

which simplify to

$$\pi_A = \beta(1+t) - \alpha(s - \beta(s+\lambda_A))$$
$$\pi_B = \alpha(1+t) - \beta(s - \alpha(s+\lambda_B)).$$

The first equation shows that if $\beta > s/(s+\lambda_A)$, then Alice plays C, and if $\alpha > s/(s+\lambda_B)$, then Bob plays C. If the opposite equalities hold, then both play D.