
Game Theory: Basic Concepts

High-rationality solution concepts in game theory can emerge in a world populated by low-rationality agents.

Young (1998)

The philosophers kick up the dust and then complain that they cannot see.

Bishop Berkeley

2.1 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 connected sequence of branches starting at a and ending at a' .¹ If there is a path from node a to node a' , we say a is an *ancestor* of a' , and a' is a *successor* to a . We call the number of branches between a and a' the *length* of the path. If a path from a to a' has length 1, 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*.

¹Technically, a path is a sequence b_1, \dots, b_k of branches such that $b_1^h = a$, $b_i^t = b_{i+1}^h$ for $i = 1, \dots, k-1$, 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' . The length of the path is k .

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 represents a particular *action* that the player assigned to 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.²

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).

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, or he may know only that he is at one of several possible nodes. 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.³

A *strategy* s_i for player i is a choice of an action at every information set assigned to i . Suppose each player $i = 1, \dots, n$ chooses strategy s_i . We call $s = (s_1, \dots, s_n)$ 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

²Thus, if $\mathbf{p} = (b_1, \dots, b_k)$ is a path from a to a' , then starting from a , if the actions associated with b_j are taken by the various players, the game moves to a' .

³Another 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 $v \in \mathcal{N}_i$ are the $v' \in \mathcal{N}_i$ 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.

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 at 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 player i as

$$\pi_i(s) = \sum_{t \in T} p(s, t) \pi_i(t). \tag{2.1}$$

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 (2.1).

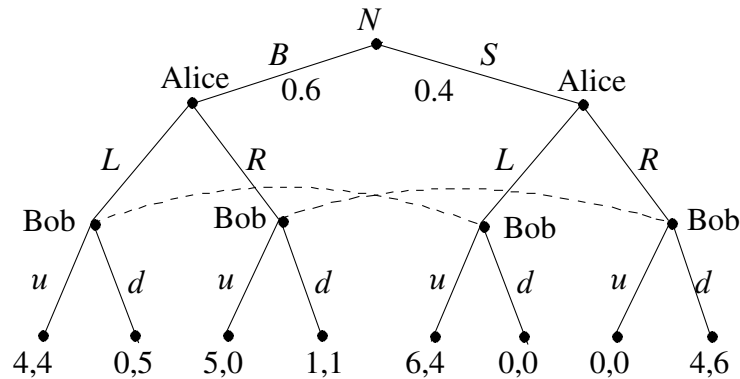


Figure 2.1. Evaluating payoffs when Nature moves

For example, consider the game depicted in figure 2.1. Here, Nature moves first, and with probability $p_l = 0.6$ chooses B where the game between Alice and Bob is known as the Prisoner's Dilemma (§2.10), and with probability $p_l = 0.4$ chooses S , where the game between Alice and Bob is known as the Battle of the Sexes (§2.8). 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.

Alice's strategies can be written LL, LR, RL, and RR, where LL means choose L whatever Nature chooses, RR means choose R whatever Nature chooses, LR means chose L when Nature chooses B, and choose R when Nature chooses S, and finally, RL means choose R when Nature chooses B and choose L when Nature chooses S. Similarly we can write Bob's choices as uu, ud, du, and dd, where uu means choose u whatever Alice chooses, dd means choose d whatever Alice chooses, ud means chose u when Alice chooses L, and choose d when Alice chooses R, and finally, and du means choose d when Alice chooses L and choose u when Alice chooses R.

Let us write, $\pi_A(x, y, z)$ and $\pi_B(x, y, z)$ for the payoffs to Alice and Bob, respectively, when Alice plays $x \in \{LL, LR, RL, RR\}$, Bob plays $y \in \{uu, ud, du, dd\}$ and Nature plays $z \in \{B, S\}$. Then, using the above parameter values, (2.1) gives the following equations.

$$\begin{aligned}\pi_A(LL, uu) &= p_u \pi_A(LL, uu, B) + p_r \pi_A(LL, uu, S) \\ &= 0.6(4) + 0.4(6) = 4.8; \\ \pi_B(LL, uu) &= p_u \pi_B(LL, uu, B) + p_r \pi_B(LL, uu, S) \\ &= 0.6(4) + 0.4(4) = 4.0;\end{aligned}$$

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

2.2 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.⁴
- c. A function $\pi_i : S \rightarrow \mathbf{R}$ 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.

⁴Technically, these are *pure strategies* because in §2.3 we will consider *mixed strategies* that are probabilistic combinations of pure strategies.

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 for that player, each strategy consisting of the choice of an action at each information set where i moves. Finally, the payoff functions are given by equation (2.1). 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 an exercise, you should work out the normal form matrix for the game depicted in figure 2.1.

2.3 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

$$\sigma = p_1s_1 + \dots + p_k s_k,$$

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 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 all pure strategies are used in it. 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 player i has pure-strategy set S_i for $i = 1, \dots, n$, a *mixed-strategy profile* $\sigma = (\sigma_1, \dots, \sigma_n)$ is the choice of a mixed strategy σ_i by each player. We define the *payoffs* to σ as follows. Let $\pi_i(s_1, \dots, s_n)$ be the payoff to player i when players use the pure strategy profile (s_1, \dots, 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

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

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 $s_1 \in S_1, \dots, s_n \in S_n$ will be used is simply the product $p_{s_1} \cdots p_{s_n}$ of their weights, and the payoff to player i in this case is just $\pi_i(s_1, \dots, s_n)$. We get the expected payoff by multiplying and adding up over all n -tuples of mixed strategies.

2.4 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 $i = 1, \dots, n$, 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 $\Delta^* S = \prod_{i=1}^n \Delta S_i$, be the mixed-strategy profiles for the game. If $\sigma \in \Delta^* S$, we write σ_i for the i th component of σ (i.e., σ_i is player i 's mixed strategy in σ). If $\sigma \in \Delta^* S$ and $\tau_i \in \Delta S_i$, we write

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

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 $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Delta^* S$ is a *Nash equilibrium* if, for every player $i = 1, \dots, n$ 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. 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, Nash equilibria that seem implausible are often *unstable* equilibria of an evolutionary process, so we would not expect to see them in the real world (Gintis 2009). Where people appear to deviate systematically from implementing Nash equilibria, we sometimes

find that they do not understand the game, or that 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.

2.5 The Fundamental Theorem of Game Theory

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

THEOREM 2.1 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 $\sigma = (\sigma_1, \dots, \sigma_n)$ be a mixed-strategy profile for an n -player game. For any player $i = 1, \dots, n$, 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 $i = 1, \dots, n$ 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.

2.6 Solving for Mixed-Strategy Nash Equilibria

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 (§2.5). 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

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

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

For this to make sense, the denominator must be nonzero and the right-hand side must lie between zero and one. Note that the *column* player's strategy is determined by the requirement that the *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 $\beta = (d_2 - c_2)/(d_2 - c_2 + a_2 - b_2)$.

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 the *row* player's strategy is determined by the requirement that the *column* player's two strategies be equal.

- Suppose the above really is a mixed-strategy equilibrium. What are the payoffs to the two players?
- 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? In an $n \times m$ game?

2.7 Throwing Fingers

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 with probability α and c_2 with probability $1 - \alpha$. We write this as $\sigma = \alpha c_1 + (1 - \alpha)c_2$. If Alice uses both c_1 and c_2 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 $\alpha \cdot 1 + (1 - \alpha)(-1) = 2\alpha - 1$, and the payoff to c_2 against σ is $\alpha(-1) + (1 - \alpha)1 = 1 - 2\alpha$. If these are equal, then $\alpha = 1/2$. 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.

		Bob	
	Alice	c_1	c_2
c_1		1, -1	-1, 1
c_2		-1, 1	1, -1

2.8 The 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.

		Violetta	
	Alfredo	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 from gambling and from going to the opera must be equal for Alfredo, we must have $\beta = 2(1 - \beta)$, which implies $\beta = 2/3$. Because the payoff from gambling and from going to the 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 2.2. Note that the three Nash equilibria are the three intersections of the two best response curves, marked by large dots in the figure.




Figure 2.2. Nash equilibria in the Battle of the Sexes

2.9 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 1 or player 2, and hence both players must use the same mixed strategy.

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

As an exercise, explain the entries in the payoff matrix and show that there are no symmetric pure-strategy Nash equilibria. We will find three Nash equilibria for 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 behaviors 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 2.3. 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 moves on whether they are player 1 or player 2, is the mixed-strategy Nash equilibrium $(v/w, v/w)$.

2.10 The Prisoner's Dilemma

Alice and Bob can each earn a profit R if they both cooperate (C). However, either can defect by working secretly on private jobs (D), earning $T > R$, but the other player will earn only $S < R$. If both defect, however, they will each earn P , where $S < P < R$.

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

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 “temptation” (to defect on a partner), S stands for “sucker” (for cooperating when your partner defected), P stands for “punishment” (for both




Figure 2.3. Nash equilibria in the Hawk-Dove Game

defecting), 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 and D .

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

$$\begin{aligned} \pi_A &= \alpha\beta + \alpha(1 - \beta)(-s) + (1 - \alpha)\beta(1 + t) + (1 - \alpha)(1 - \beta)(0), \\ \pi_B &= \alpha\beta + \alpha(1 - \beta)(1 + t) + (1 - \alpha)\beta(-s) + (1 - \alpha)(1 - \beta)(0), \end{aligned}$$

which simplify to

$$\begin{aligned} \pi_A &= \beta(1 + t) - \alpha(s(1 - \beta) + \beta t), \\ \pi_B &= \alpha(1 + t) - \beta(s(1 - \alpha) + \alpha t). \end{aligned}$$

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.

As we shall see in §3.5, this is not the way many people play this game in the experimental laboratory. Rather, people very often prefer to cooperate, provided their partners cooperate as well. We can capture this phenomenon 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

$$\begin{aligned}\pi_A &= \alpha\beta(1 + t + \lambda_A) + \alpha(1 - \beta)(-s) \\ &\quad + (1 - \alpha)\beta(1 + t) + (1 - \alpha)(1 - \beta)(0) \\ \pi_B &= \alpha\beta(1 + t + \lambda_B) + \alpha(1 - \beta)(1 + t) \\ &\quad + (1 - \alpha)\beta(-s) + (1 - \alpha)(1 - \beta)(0),\end{aligned}$$

which simplify to

$$\begin{aligned}\pi_A &= \beta(1 + t) - \alpha(s - \beta(s + \lambda_A)) \\ \pi_B &= \alpha(1 + t) - \beta(s - \alpha(s + \lambda_B)).\end{aligned}$$

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 .

2.11 Alice, Bob, and the Choreographer

Consider the game played by Alice and Bob, with the normal form matrix shown to the right. There are two Pareto-efficient pure-strategy equilibria: $(1,5)$ and $(5,1)$. There is also a mixed-strategy equilibrium with payoffs $(2.5,2.5)$, in which Alice plays u with probability 0.5 and Bob plays l with probability 0.5.

	Bob	
Alice	l	r
u	5,1	0,0
d	4,4	1,5

If the players can jointly observe a choreographer who signals ul and dr , each with probability $1/2$, Alice and Bob can then achieve the payoff $(3,3)$ by obeying the choreographer; i.e. by playing (u, l) if they see ul and playing (d, r) if they see dr . Note that this is a Nash equilibrium of a larger game in which the choreographer moves first and acts as Nature. This is a Nash equilibrium because each player chooses a best response to the move of the other, assuming the other carries out the choreographer's directive. This situation is termed a *correlated equilibrium* of the original game (Aumann 1974). The commonly observable event on which their behavior is conditioned is called a *correlating device*.

A more general correlated equilibrium for this game can be constructed as follows. Consider a choreographer who would like to direct Alice to play d and Bob to play l so the joint payoff $(4, 4)$ could be realized. The problem is that if Alice obeys the choreographer, then Bob has an incentive to

choose r , giving him a payoff of 5 instead of 4. Similarly, if Bob obeys the choreographer, then Alice has an incentive to choose u , giving her a payoff of 5 instead of 4. The choreographer must therefore be more sophisticated.

Suppose the choreographer has three states. In ω_1 , which occurs with probability α_1 , he advises Alice to play u and Bob to play l . In ω_2 , which occurs with probability α_2 , the choreographer advises Alice to play d and Bob to play l . In ω_3 , which occurs with probability α_3 , the choreographer advises Alice to play d and Bob to play r . We assume Alice and Bob know α_1 , α_2 , and $\alpha_3 = 1 - \alpha_1 - \alpha_2$, and it is common knowledge that both have a *normative predisposition* (see chapter 7) to obey the choreographer unless they can do better by deviating. However, neither Alice nor Bob can observe the state ω of the choreographer, and each hears only what the choreographer tells them, not what the choreographer tells the other player. We will find the values of α_1 , α_2 , and α_3 for which the resulting game has a Pareto-efficient correlated equilibrium.

Note that Alice has knowledge partition $[\{\omega_1\}, \{\omega_2, \omega_3\}]$ (§1.5), meaning that she knows when ω_1 occurs but cannot tell whether the state is ω_2 or ω_3 . This is because she is told to move u only in state ω_1 but to move d in both states ω_2 and ω_3 . The conditional probability of ω_2 for Alice given $\{\omega_2, \omega_3\}$ is $p_A(\omega_2) = \alpha_2/(\alpha_2 + \alpha_3)$, and similarly $p_A(\omega_3) = \alpha_3/(\alpha_2 + \alpha_3)$. Note also that Bob has knowledge partition $[\{\omega_3\}, \{\omega_1, \omega_2\}]$ because he is told to move r only at ω_3 but to move l at both ω_1 and ω_2 . The conditional probability of ω_1 for Bob given $\{\omega_1, \omega_2\}$ is $p_B(\omega_1) = \alpha_1/(\alpha_1 + \alpha_2)$, and similarly $p_B(\omega_2) = \alpha_2/(\alpha_1 + \alpha_2)$.

When ω_1 occurs, Alice knows that Bob plays l , to which Alice's best response is u . When ω_2 or ω_3 occurs, Alice knows that Bob is told l by the choreographer with probability $p_A(\omega_2)$ and is told r with probability $p_A(\omega_3)$. Thus, despite the fact that Bob plays only pure strategies, Alice knows she effectively faces the mixed strategy l played with probability $\alpha_2/(\alpha_2 + \alpha_3)$ and r played with probability $\alpha_3/(\alpha_2 + \alpha_3)$. The payoff to u in this case is $5\alpha_2/(\alpha_2 + \alpha_3)$, and the payoff to d is $4\alpha_2/(\alpha_2 + \alpha_3) + \alpha_3/(\alpha_2 + \alpha_3)$. If d is to be a best response, we must thus have $\alpha_1 + 2\alpha_2 \leq 1$.

Turning to the conditions for Bob, when ω_3 occurs, Alice plays d so Bob's best response is r . When ω_1 or ω_2 occurs, Alice plays u with probability $p_B(\omega_1)$ and d with probability $p_B(\omega_2)$. Bob chooses l when $\alpha_1 + 4\alpha_2 \geq 5\alpha_2$. Thus, any α_1 and α_2 that satisfy $1 \geq \alpha_1 + 2\alpha_2$ and $\alpha_1 \geq \alpha_2$ permit a correlated equilibrium. Another characterization is $1 - 2\alpha_2 \geq \alpha_1 \geq \alpha_2 \geq 0$.

What are the Pareto-optimal choices of α_1 and α_2 ? Because the correlated equilibrium involves $\omega_1 \rightarrow (u, l)$, $\omega_2 \rightarrow (d, l)$, and $\omega_3 \rightarrow (d, r)$, the payoffs to (α_1, α_2) are $\alpha_1(5, 1) + \alpha_2(4, 4) + (1 - \alpha_1 - \alpha_2)(1, 5)$, which simplifies to $(1 + 4\alpha_1 + 3\alpha_2, 5 - 4\alpha_1 - \alpha_2)$, where $1 - 2\alpha_2 \geq \alpha_1 \geq \alpha_2 \geq 0$. This is a linear programming problem. It is easy to see that either $\alpha_1 = 1 - 2\alpha_2$ or $\alpha_1 = \alpha_2$ and $0 \leq \alpha_2 \leq 1/3$. The solution is shown in figure 2.4.




Figure 2.4. Alice, Bob, and the choreographer

The pair of straight lines connecting $(1, 5)$ to $(10/3, 10/3)$ to $(5, 1)$ is the set of Pareto-optimal points. Note that the symmetric point $(10/3, 10/3)$ corresponds to $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$.

2.12 An Efficiency-Enhancing Choreographer

Consider an n -player game in which each player can choose an integer in the range $k = 1, \dots, 10$. Nature chooses an integer k in this range, and if all n players also choose k , each has payoff 1. Otherwise, each has payoff 0. Nature also supplies to any agent who inquires (one sample per agent) a noisy signal that equals the true value with probability $p > 0.10$. A best response for each player is to sample the signal and choose a number equal to the signal received. The payoff is p^n . For a correlated equilibrium, suppose there is a social rule that obligates the youngest player to reveal his choice. There is then a correlated equilibrium in which each player follows the youngest player's choice, and the payoff to each player is now p . For instance, when $p = 90\%$ and $n = 25$, the Nash equilibrium payoff is 0.071, which is only 8% of the value of the correlated equilibrium payoff.

This example shows that there may be huge gains to groups that manage to find an appropriate correlating device.

2.13 The Correlated Equilibrium Solution Concept

The correlated equilibrium concept will be studied further in chapter 7. This solution concept has been neglected in classical game theory, although we will show that it is a more natural solution concept than the Nash equilibrium. This is because the correlated equilibrium directly addresses the central weaknesses of the Nash equilibrium concept: its lack of a mechanism for choosing among various equally plausible alternatives, for coordinating the behaviors of players who are indifferent among several pure strategies, and for providing incentives for players to follow the suggested strategy even when they may have private payoffs that would lead self-regarding agents to do otherwise (§6.3, 6.4).

Game theorists have not embraced the correlated equilibrium concept because it appears to require an active social agency, in the form of the choreographer, that cannot be accounted for within game theory. We will argue that therein lies the true power of the correlated equilibrium concept: it points away from game theory to a larger, complementary, social epistemology, that will be explored in chapters 7 and 12.