Notes on Game Theory

Steve Schecter

Department of Mathematics

North Carolina State University

Contents

Preface

These notes are intended to accompany the book of Herbert Gintis, *Game Theory Evolving*, second edition (Princeton University Press, 2009).

Game theory deals with situations in which your payoff depends not only on your own choices but on the choices of others. How are you supposed to decide what to do, since you cannot control what others will do?

In calculus you learn to maximize and minimize functions, for example to find the cheapest way to build something. This field of mathematics is called optimization. Game theory differs from optimization in that in optimization problems, your payoff depends only on your own choices.

Like the field of optimization, game theory is defined by the problems it deals with, not by the mathematical techniques that are used to deal with them. The techniques are whatever works best.

Also, like the field of optimization, the problems of game theory come from many different areas of study. It is nevertheless helpful to treat game theory as a single mathematical field, since then techniques developed for problems in one area, for example evolutionary biology, become available to another, for example economics.

Game theory has three uses:

(1) Understand the world. For example, game theory helps understand why animals sometimes fight over territory and sometimes don't.
(2) Respond to the world. For example, game theory has been used to develop strategies to win money at poker.
(3) Change the world. Often the world is the way it is because people are responding to the rules of a game. Changing the game can change how they act. For example, rules on using energy can be designed to encourage conservation and innovation.

April 27, 2009

CHAPTER 1

Backward Induction

1.1. Tony's accident

When I was a college student, my friend Tony caused a minor traffic accident. The car of the victim, whom I'll call Vic, was slightly scraped. Tony didn't want to tell his insurance company. The next morning, Tony and I went with Vic to visit some body shops. The upshot was that the repair would cost \$80.

Tony and I had lunch with a bottle of wine, and thought over the situation. Vic's car was far from new and had accumulated many scrapes. Repairing the few that Tony had caused would improve the car's appearance only a little. We figured that if Tony sent Vic a check for \$80, Vic would probably just pocket it.

Perhaps, we thought, Tony should ask to see a receipt showing that the repairs had actually been performed before he sent Vic the \$80.

A game theorist would represent this situation by a game tree. For definiteness, we'll assume that the value to Vic of repairing the damage is \$20.

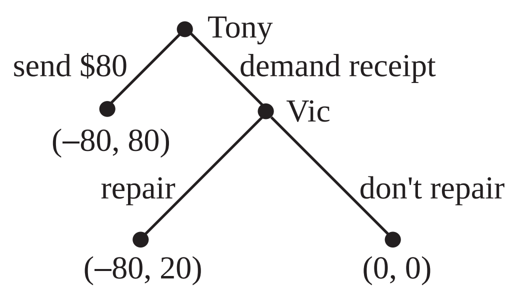

Figure 1.1. Tony's accident.

Explanation of the game tree:

1. Tony goes first. He has a choice of two actions: send Vic a check for \$80, or demand a receipt proving that the work has been done.
2. If Tony sends a check, the game ends. Tony is out \$80; Vic will no doubt keep the money, so he has gained \$80. We represent these payoffs by the ordered pair $(-80, 80)$; the first number is Tony's payoff, the second is Vic's.

(3) If Tony demands a receipt, Vic has a choice of two actions: repair the car and send Tony the receipt, or just forget the whole thing.
(4) If Vic repairs the car and sends Tony the receipt, the game ends. Tony sends Vic a check for $80, so he is out $80; Vic uses the check to pay for the repair, so his gain is $20, the value of the repair.
(5) If Vic decides to forget the whole thing, he and Tony each end up with a gain of 0.

Assuming that we have correctly sized up the situation, we see that if Tony demands a receipt, Vic will have to decide between two actions, one that gives him a payoff of $20 and one that gives him a payoff of 0. Vic will presumably choose to repair the car, which gives him a better payoff. Tony will then be out $80.

Our conclusion was that Tony was out $80 whatever he did. We did not like this game.

When the bottle was nearly finished, we thought of a third course of action that Tony could take: send Vic a check for $40, and tell Vic that he would send the rest when Vic provided a receipt showing that the work had actually been done. The game tree now looked like this:

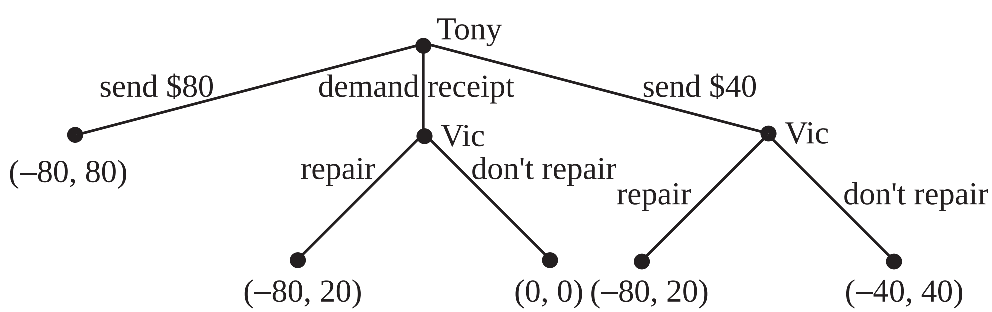

Figure 1.2. Tony's accident: second game tree.

Most of the game tree looks like the first one. However:

(1) If Tony takes his new action, sending Vic a check for $40 and asking for a receipt, Vic will have a choice of two actions: repair the car, or don't.
(2) If Vic repairs the car, the game ends. Vic will send Tony a receipt, and Tony will send Vic a second check for $40. Tony will be out $80. Vic will use both checks to pay for the repair, so he will have a net gain of $20, the value of the repair.
(3) If Vic does not repair the car, and just pockets the the $40, the game ends. Tony is out $40, and Vic has gained $40.

Again assuming that we have correctly sized up the situation, we see that if Tony sends Vic a check for $40 and asks for a receipt, Vic's best course of action is to keep the money and not make the repair. Thus Tony is out only $40.

Tony sent Vic a check for $40, told him he'd send the rest when he saw a receipt, and never heard from Vic again.

1.2. Games in extensive form with complete information

This section is related to Gintis, Sec. 3.2.

Tony's accident is the kind of situation that is studied in game theory, because:

(1) It involves more than one individual.
(2) Each individual has several possible actions.
(3) Once each individual has chosen his actions, payoffs to all individuals are determined.
(4) Each individual is trying to maximize his own payoff.

The key point is that the payoff to an individual depends not only on his own choices, but on the choices of others as well.

We gave two models for Tony's accident, which differed in the sets of actions available to Tony and Vic. Each model was a *game in extensive form with complete information.*

A game in extensive form with complete information consists, to begin with, of the following:

(1) A set P of *players*. In Figure 1.2, the players are Tony and Vic.
(2) A set N of *nodes*. In Figure 1.2, the nodes are the little black circles. There are eight.
(3) A set B of *actions* or *moves*. In Figure 1.2, the moves are the lines. There are seven. Each move connects two nodes, one its *start* and one its *end*. In Figure 1.2, the start of a move is the node at the top of the move, and the end of a move is the node at the bottom of the move.

A *root node* is a node that is not the end of any move. In Figure 1.2, the top node is the only root node.

A *terminal node* is a node that is not the start of any move. In Figure 1.2 there are five terminal nodes.

A *path* is sequence of moves such that the end node of any move in the sequence is the start node of the next move in the sequence. A path is *complete* if it is not part of any longer path. Paths are sometimes called *histories*, and complete paths are called *complete histories*. If a complete path has finite length, it must start at a root node and end at a terminal node.

A game in extensive form with complete information also has:

(4) A function from the set of nonterminal nodes to the set of players. This function, called a *labeling* of the set of nonterminal nodes, tells us which player chooses a move at that node. In Figure 1.2, there are three nonterminal nodes. One is labeled "Tony" and two are labeled "Vic."
(5) For each player, a *payoff function* from the set of complete paths into the real numbers. Usually the players are numbered from 1 to n, and the ith player's payoff function is denoted π_i.

A game in extensive form with complete information is required to satisfy the following conditions:

(a) There is exactly one root node.
(b) If c is any node other than the root node, there is exactly one path from the root node to c.

One way of thinking of (b) is that if you know the node you are at, you know exactly how you got there.

These assumptions imply that *every* complete path, not just those of finite length, starts at a root node. (If c is a node in a complete path p, the path from the root node to c must be part of p; otherwise p would not be complete.)

A *finite horizon game* is one in which there is a number K such that every complete path has length at most K. In chapters 1 to 5 of these notes, we will only discuss finite horizon games.

In a finite horizon game, the complete paths are in one-to-one correspondence with the terminal nodes. Therefore, in a finite horizon game we can define a player's payoff function by assigning a number to each terminal node.

In Figure 1.2, Tony is player 1 and Vic is player 2. Thus each terminal node e has associated to it two numbers, Tony's payoff $\pi_1(e)$ and Vic's payoff $\pi_2(e)$. In Figure 1.2 we have labeled each terminal node with the ordered pair of payoffs $(\pi_1(e), \pi_2(e))$.

A game in extensive form with complete information is *finite* if the number of nodes is finite. (It follows that the number of moves is finite. In fact, the number of moves is always one less than the number of nodes.) Such a game is necessarily a finite horizon game.

1.3. Strategies

This section is also related to Gintis, Sec. 3.2.

In game theory, a player's *strategy* is a plan for what action to take in every situation that the player might encounter. For a game in extensive form with complete information, the phrase “situations that the player might encounter” is interpreted to mean all the nodes that are labeled with his name.

In Figure 1.2, only one node, the root, is labeled “Tony.” Tony has three possible strategies, corresponding to the three actions he could choose at the start of the game. We will call Tony’s strategies s_1 (send \$80), s_2 (demand a receipt before sending anything), and s_3 (send \$40).

In Figure 1.2, there are two nodes labeled “Vic.” Vic has four possible strategies, which we label $t_1, \dots, t_4$:

Vic’s strategy	If Tony demands receipt	If Tony sends \$40
t_1	repair	repair
t_2	repair	don’t repair
t_3	don’t repair	repair
t_4	don’t repair	don’t repair

In general, suppose there are k nodes labeled with a player’s name, and there are n_1 possible moves at the first node, n_2 possible moves at the second node, ..., and n_k possible moves at the kth node. A strategy for that player consists of a choice of one of his n_1 moves at the first node, one of his n_2 moves at the second node, ..., and one of his n_k moves at the kth node. Thus the number of strategies available to the player is the product $n_1 n_2 \cdots n_k$.

If we know each player’s strategy, then we know the complete path through the game tree, so we know both players’ payoffs. With some abuse of notation, we will denote the payoffs to players 1 and 2 when player 1 uses the strategy s_i and player 2 uses the strategy t_j by $\pi_1(s_i, t_j)$ and $\pi_2(s_i, t_j)$. For example, $(\pi_1(s_3, t_2), \pi_2(s_3, t_2)) = (-80, 20)$. Of course, in Figure 1.2, this is the pair of payoffs associated with the terminal vertex on the corresponding path through the game tree.

Recall that if you know the node you are at, you know how you got there. Thus a strategy can be thought of as a plan for how to act after each course the game might take (that ends at a node where it is your turn to act).

1.4. Backward induction

This section is related to Gintis, Sec. 4.2.

Game theorists often assume that players are *rational*. One meaning of rationality for a game in extensive form with complete information is:

- Suppose a player has a choice that includes two moves m and m', and m yields a higher payoff to that player than m'. Then the player will not choose m'.

Thus, if you assume that your opponent is rational in this sense, you must assume that whatever you do, your opponent will respond by doing what is best for him, not what you might want him to do. (Wishful thinking on your part is not allowed.) Your opponent's response will affect your own payoff. You should therefore take your opponent's likely response into account in deciding on your own action. This is exactly what Tony did when he decided to send Vic a check for $40.

The assumption of rationality motivates the following procedure for selecting strategies for all players in a finite game in extensive form with complete information. This procedure is called *backward induction* or *pruning the game tree*.

(1) Select a node c such that all the moves available at c have ends that are terminal. (Since the game is finite, there must be such a node.)
(2) Suppose player i is to choose at node c. Among all the moves available to him at that node, find the move m whose end e gives the greatest payoff to player i. *We assume that this move is unique.*
(3) Assume that at node c, player i will choose the move m. Record this choice as part of player i's strategy.
(4) Delete from the game tree all moves that start at c. The node c is now a terminal node. Assign to it the payoffs that were previously assigned to the node e.
(5) The game tree now has fewer nodes. If it has just one node, stop. If it has more than one node, return to step 1.

In step 2 we find the move that player i presumably will make should the course of the game arrive at node c. In step 3 we assume that player i will in fact make this move, and record this choice as part of player i's strategy. In step 4 we assign the payoffs to all players that result from this choice to the node c and "prune the game tree." This helps us take this choice into account in finding the moves players will presumably make at earlier nodes.

In Figure 1.2, there are two nodes for which all available moves have terminal ends: the two where Vic is to choose. At the first of these nodes, Vic's best move is repair, which gives payoffs of $(-80, 20)$. At the second, Vic's best more is don't repair, which gives payoffs of $(-40, 40)$. Thus after two steps of the backward induction procedure, we have recorded the strategy t_2 for Vic, and we arrive at the following pruned game tree:

Now the vertex labeled "Tony" has all its ends terminal. Tony's best move is to send $40, which gives him a payoff of -40. Thus Tony's strategy is s_3. We delete all moves that start at the vertex labeled "Tony," and label that vertex with the payoffs $(-40, 40)$. That is now the only remaining vertex, so we stop.

Thus the backward induction procedure selects strategy s_3 for Tony and strategy t_2 for Vic, and predicts that the game will end with the payoffs $(-40, 40)$. This is how the game ended in reality.

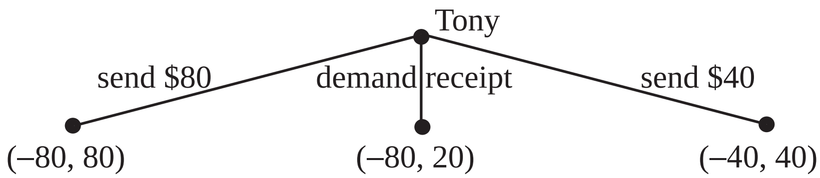

FIGURE 1.3. Tony's accident: pruned game tree.

The backward induction procedure can fail if, at any point, step 2 produces two moves that give the same highest payoff to the player who is to choose. Here is an example where backward induction fails:

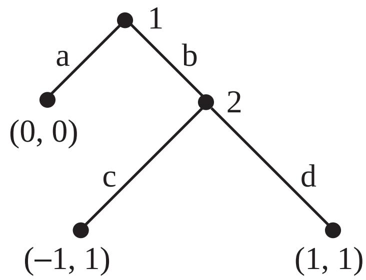

FIGURE 1.4. Failure of backward induction.

At the node where player 2 chooses, both available moves give him a payoff of 1. Player 2 is indifferent between these moves. Hence player 1 does not know which move player 2 will choose if player 1 chooses b. Now player 1 cannot choose between his moves a and b, since which is better for him depends on which choice player 2 would make if he chose b.

We will return to this issue in Chapter 6

1.5. Big Monkey and Little Monkey 1

This section is related to Gintis, Sec. 3.1.

Big Monkey and Little Monkey eat coconut, which dangle from a branch of the coconut palm. One of them (at least) must climb the tree and shake down the fruit. Then both can eat it. The monkey that doesn't climb will have a head start eating the fruit.

If Big Monkey climbs the tree, he incurs an energy cost of 2 Kc. If Little Monkey climbs the tree, he incurs a negligible energy cost (being so little).

A coconut can supply the monkeys with 10 Kc of energy. It will be divided between the monkeys as follows:

	Big Monkey eats	Little Monkey eats
If Big Monkey climbs	6 Kc	4 Kc
If both monkeys climb	7 Kc	3 Kc
If Little Monkey climbs	9 Kc	1 Kc

Let's assume that Big Monkey must decide first. Payoffs are net gains in kilocalories. The game tree is as follows:

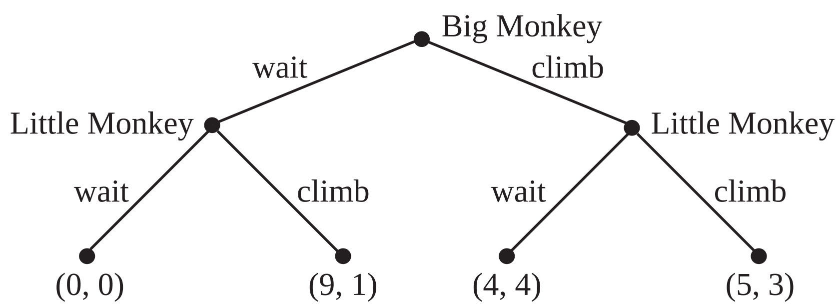

FIGURE 1.5. Big Monkey and Little Monkey.

Backward induction produces the following strategies:

(1) Little Monkey: If Big Monkey waits, climb. If Big Monkey climbs, wait.
(2) Big Monkey: Wait.

Thus Big Monkey waits. Little Monkey, having no better option at this point, climbs the tree and shakes down the fruit. He scampers quickly down, but to no avail: Big Monkey has gobbled most of the fruit. Big Monkey has a net gain of 9 Kc, Little Monkey 1 Kc.

This game has the following peculiarity: Suppose Little Monkey adopts the strategy, no matter what Big Monkey does, wait. If Big Monkey is convinced that this is in fact Little Monkey's strategy, he sees that his own payoff will be 0 if he waits and 4 if he climbs. His best option is therefore to climb. The payoffs are 4 Kc to each monkey.

Little Monkey's strategy of waiting no matter what Big Monkey does is not "rational" in the sense of the last section, since it involves taking an inferior action should Big Monkey wait. Nevertheless it produces a better outcome for Little Monkey than his "rational" strategy.

In game theory terms, a strategy for Little Monkey that includes waiting in the event that Big Monkey waits is a strategy that includes an *incredible threat*. Choosing to wait after Big Monkey waits reduces Big Monkey's payoff (that's why it's a threat) at the price of also reducing Little Monkey's payoff (that's why it's not credible that Little Monkey would do it).

How can such a threat be made credible? Perhaps Little Monkey could arrange to break a leg!

If Little Monkey can somehow convince Big Monkey that he will wait no matter what, then from Big Monkey's point of view, the game tree changes to the one shown in Figure 1.6.

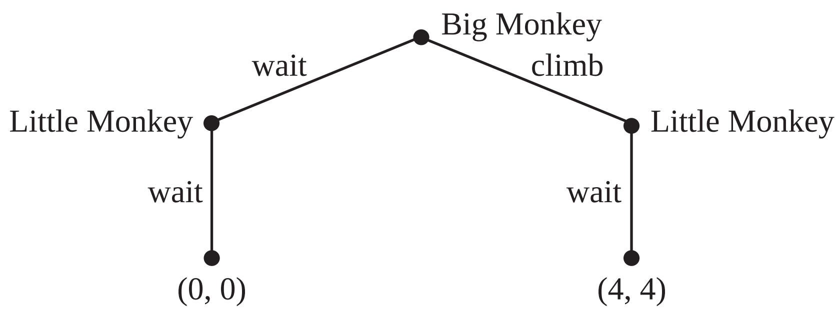

FIGURE 1.6. Big Monkey and Little Monkey if Little Monkey commits to wait.

Now if Big Monkey uses backward induction, he will climb!

1.6. Rosenthal's Centipede Game

Mutt and Jeff start with \$2 each. Mutt goes first.

On a player's turn, he has two possible moves:

(1) Cooperate: The player does nothing. The game master rewards him with \$1.
(2) Defect: The player steals \$2 from the other player.

The game ends when either (1) one of the players defects, or (2) both players have at least \$100.

Payoffs are dollars gained in the game. The game tree is shown in Figure 1.7.

A backward induction analysis begins at the only node both of whose moves end in terminal vertices: Jeff's node at which Mutt has accumulated \$100 and Jeff has accumulated \$99. If Jeff cooperates, he receives \$1 from the game master, and the game ends with Jeff having \$100. If he defects by stealing \$2 from Mutt, the game ends with Jeff having \$101. Assuming Jeff is "rational," he will defect.

In fact, the backward induction procedure yields the following strategy for each player: whenever it is your turn, defect.

Hence Mutt steals \$2 from Jeff at his first turn, and the game ends with Mutt having \$4 and Jeff having nothing.

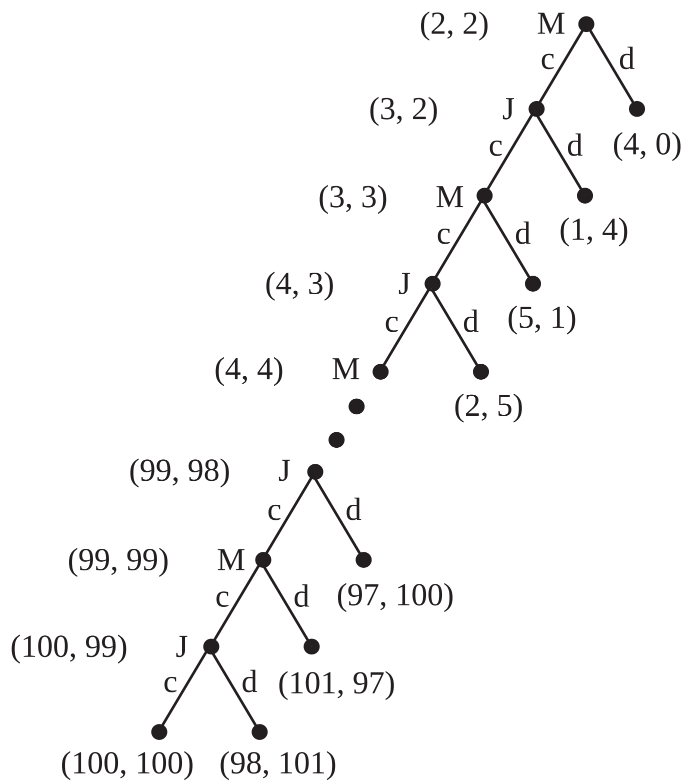

FIGURE 1.7. Rosenthal's Centipede Game. Mutt is Player 1, Jeff is Player 2. The amounts the players have accumulated when a node is reached are shown to the left of the node.

This is a disconcerting conclusion. If you were given the opportunity to play this game, don't you think you could come away with more than $4?

In fact, in experiments, people typically do not defect on the first move. For more information, you may consult the Wikipedia page devoted to this game (http://en.wikipedia.org/wiki/Centipede_game).

What's wrong with our analysis? Here are a few possibilities.

- The players may feel better about themselves if they cooperate. If this is the case, we should take account of it in the players' payoff functions. Even if *you* only care about money, your opponent may have a desire to be cooperative, and you should take this into account in assigning his payoff function.
- People do not typically make decisions on the basis of a complicated rational analysis. Instead they follow rules of thumb, such as be cooperative and don't steal. In fact, it may not be rational to make most decisions on the basis of a complicated rational analysis, because (a) the cost in terms

of time and effort of doing the analysis may be so great as to undo the advantage gained, and (b) if the analysis is complicated enough, you are liable to make a mistake anyway.

- We do not typically encounter "games" that we know in advance will be repeated exactly n times, where n is a large number. Instead, we typically encounter games that will be repeated an unknown number of times. In such a situation, it is not clear that backward induction is relevant. When we encounter the unusual situation of a game that is to be repeated 100 times, we use a rule of thumb for the more usual situation of a game that is to be repeated an unknown number of times. We will study such games in Chapter **??**. The bottom line is that for such games, cooperation may well be rational.

1.7. Ultimatum Game

Player 1 is given 100 one dollar bills. He must offer some of them (one to 99) to Player 2. If Player 2 accepts the offer, he gets to keep the bills he was offered, and Player 1 gets to keep the rest. If Player 2 rejects the offer, neither player gets to keep anything.

Let's assume payoffs are dollars gained in the game. Then the game tree is shown below.

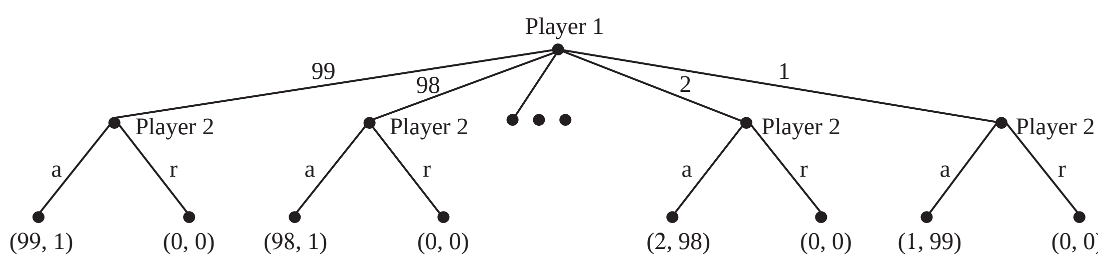

FIGURE 1.8. Ultimatum Game with dollar payoffs. Player 1 offers a number of dollars to Player 2, then Player 2 accepts or rejects the offer.

Backward induction shows:

- Whatever offer Player 1 makes, Player 2 should accept it, since a gain of even one dollar is better than a gain of nothing.
- Therefore Player 1 should only offer one dollar. That we he gets to keep 99!

However, many experimenters have shown that people do no not actually play the Ultimatum Game in accord with this analysis; see the Wikipedia page for this game (http://en.wikipedia.org/wiki/Ultimatum_game). Offers of less than about $40 are typically rejected.

Perhaps we have not correctly represented the players' payoffs. For example, most people don't like to feel that they have been suckers. Perhaps Player 2 will feel like a sucker if he accepts less than $40, and feeling like a sucker is equivalent to a payoff of −$40. How does this change the game tree and the outcome?

1.8. Continuous games

In the games we have considered so far, when it is a player's turn to move, he has only a finite number of choices. In the remainder of this chapter, we will consider some games in which each player may choose an action from an interval of real numbers. For example, if a firm must choose the price to charge for an item, we can imagine that the price could be any nonnegative real number. This allows us to use the power of calculus to find which price produces the best payoff to the firm.

More precisely, we will consider games with two players, player 1 and player 2. Player 1 goes first. The moves available to him are all real numbers s in some interval I. Next it is player 2's turn. The moves available to him are all real numbers t in some interval J. Player 2 observes player 1's move s and then chooses his move t. The game is now over, and payoffs $\pi_1(s,t)$ and $\pi_2(s,t)$ are calculated.

Does such a game satisfy the definition that we gave in Section 1.2 of a game in extensive form with complete information? Yes, it does. In the previous paragraph, to describe the type of game we want to consider, we only described the moves, not the nodes. However, the nodes are still there. There is a root node at which player 1 must choose his move s. Each move s ends at a new node, at which player 2 must choose t. Each move t ends at a terminal node. The set of all complete paths is the set of all pairs (s,t) with s in I and t in J. Since we described the game in terms of moves, not nodes, it was easier to describe the payoff functions as assigning numbers to complete paths, not as assigning numbers to terminal nodes. That is what we did: $\pi_1(s,t)$ and $\pi_2(s,t)$ assign numbers to each complete path.

Such a game is not finite, but it is a finite horizon game: the length of the longest path is 2.

Let us find strategies for players 1 and 2 using the *idea* of backward induction. Backward induction as we described it in Section 1.4 cannot be used because the game is not finite.

We begin with the last move, which is player 2's. Assuming he is rational, he will observe player 1's move s and then choose t in J to maximize the function $\pi_2(s,t)$ with s fixed. For fixed s, $\pi_2(s,t)$ is a function of one variable t. Suppose it takes on its maximum value in J at a unique value of t. This number t is player 2's *best response* to player 1's move s. Normally the best response t will depend on s, so we write $t = b(s)$. The function $t = b(s)$ gives a strategy for player 2, i.e., it gives player 2 a choice of action for every possible choice s in I that player 1 might make.

Player 1 should choose s taking into account player 2's strategy. If player 1 assumes that player 2 is rational and hence will use his best response strategy, then player 1 should choose s in I to maximize the function $\pi_1(s, b(s))$. This is again of function of one variable.

1.9. Stackelberg's model of duopoly

In a *duopoly*, a certain good is produced by just two firms, which we label 1 and 2. Let s be the quantity produced by firm 1 and let t be the quantity produced by firm 2. Then the total quantity of the good that is produced is $q = s + t$. In Stackelberg's model of duopoly (Wikipedia article: http://en.wikipedia.org/wiki/Stackelberg_duopoly), the market price p of the good depends on q: $p = \phi(q)$. At this price, everything that is produced can be sold.

Suppose firm 1's cost to produce the quantity s of the good is $c_1(s)$, and firm 2's cost to produce the quantity t of the good is $c_2(t)$. We denote the profits of the two firms by π_1 and π_2. Now profit is revenue minus cost, and revenue is price times quantity sold. Since the price depends on $q = s + t$, each firm's profit depends in part on how much is produced by the other firm. More precisely,

$$\pi_1(s, t) = \phi(s + t)s - c_1(s), \quad \pi_2(s, t) = \phi(s + t)t - c_2(t).$$

In Stackelberg's model of duopoly, each firm tries to maximize its own profit by choosing an appropriate level of production.

1.9.1. First model.

Let us begin by making the following assumptions:

(1) Price falls linearly with total production. In other words, there are numbers α and β such that the formula for the price is

$$p = \alpha - \beta(s + t),$$

and $\beta > 0$.

(2) Each firm has the same unit cost of production $c > 0$. Thus $c_1(s) = cs$ and $c_2(t) = ct$.

(3) $\alpha > c$. In other words, the price of the good when very little is produced is greater than the unit cost of production. If this assumption is violated, the good will not be produced.

(4) Firm 1 chooses its level of production s first. Then firm 2 observes s and chooses t.

(5) The production levels s and t can be any real numbers.

We ask the question, what will be the production level and profit of each firm?

The payoff in this game is the profit:

$$\pi_1(s,t) = \phi(s+t)s - cs = (\alpha - \beta(s+t) - c)s = (\alpha - \beta t - c)s - \beta s^2,$$
$$\pi_2(s,t) = \phi(s+t)t - ct = (\alpha - \beta(s+t) - c)t = (\alpha - \beta s - c)t - \beta t^2.$$

Since firm 1 chooses s first, we begin our analysis by finding firm 2's best response $t = b(s)$. To do this we must find where the function $\pi_2(s,t)$, with s fixed, has its maximum. Since $\pi_2(s,t)$ with s fixed has a graph that is just an upside down parabola, we can do this by taking the derivative with respect to t and setting it equal to 0:

$$\frac{\partial \pi_2}{\partial t} = \alpha - \beta s - c - 2\beta t = 0 \quad \Rightarrow \quad t = \frac{\alpha - \beta s - c}{2\beta}.$$

Therefore firm 2's best response function is

$$b(s) = \frac{\alpha - \beta s - c}{2\beta}.$$

Finally we must maximize $\pi_1(s, b(s))$, the payoff that firm 1 can expect from each choice s assuming that firm 2 uses its best response strategy. We have

$$\pi_1(s,b(s)) = \pi_1(s, \frac{\alpha - \beta s - c}{2\beta}) = (\alpha - \beta(s + \frac{\alpha - \beta s - c}{2\beta}) - c)s = \frac{\alpha - c}{2}s - \frac{\beta}{2}s^2.$$

Again this function has a graph that is an upside down parabola, so we can find where it is maximum by taking the derivative and setting it equal to 0:

$$\frac{d}{ds}\pi_1(s,b(s)) = \frac{\alpha - c}{2} - \beta s = 0 \quad \Rightarrow \quad s = \frac{\alpha - c}{2\beta}.$$

We see from this calculation that $\pi_1(s,b(s))$ is maximum at $s^* = \frac{\alpha - c}{2\beta}$. Given this choice of production level for firm 1, firm 2 chooses the production level

$$t^* = b(s^*) = \frac{\alpha - c}{4\beta}.$$

Since we assumed $\alpha > c$, the production levels s^* and t^* are positive. This is reassuring. The price is

$$p^* = \alpha - \beta(s^* + t^*) = \alpha - \beta(\frac{\alpha - c}{2\beta} + \frac{\alpha - c}{4\beta}) = \frac{1}{4}\alpha + \frac{3}{4}c = c + \frac{1}{4}(\alpha - c).$$

Since $\alpha > c$, this price is greater than the cost of production c, which is also reassuring.

The profits are

$$\pi_1(s^*,t^*) = \frac{(\alpha - c)^2}{8\beta}, \quad \pi_2(s^*,t^*) = \frac{(\alpha - c)^2}{16\beta}.$$

Firm 1 has twice the level of production and twice the profit of firm 2. In this model, it is better to be the firm that chooses its price first.

1.9.2. Second model. The model in the previous subsection has a disconcerting aspect: the levels of production s and t, and the price p, are all allowed in the model to be negative. We will now complicate the model to deal with this objection.

We replace assumption (1) with the following:

(1) Price falls linearly with total production until it reaches 0; for higher total production, the price remains 0. In other words, there are positive numbers α and β such that the formula for the price is

$$p = \begin{cases} \alpha - \beta(s+t) & \text{if } s+t < \frac{\alpha}{\beta}, \\ 0 & \text{if } s+t \geq \frac{\alpha}{\beta}. \end{cases}$$

Assumptions (2), (3), and (4) remain unchanged. We replace assumption (5) with:

(5) The production levels s and t must be nonnegative.

We again ask the question, what will be the production level and profit of each firm?

The payoff is again the profit, but the formulas are different:

$$\pi_1(s,t) = \phi(s+t)s - cs = \begin{cases} (\alpha - \beta(s+t) - c)s & \text{if } 0 \leq s+t < \frac{\alpha}{\beta}, \\ -cs & \text{if } s+t \geq \frac{\alpha}{\beta}, \end{cases}$$

$$\pi_2(s,t) = \phi(s+t)t - ct = \begin{cases} (\alpha - \beta(s+t) - c)t & \text{if } 0 \leq s+t < \frac{\alpha}{\beta}, \\ -ct & \text{if } s+t \geq \frac{\alpha}{\beta}. \end{cases}$$

The possible values of s and t are now $0 \leq s < \infty$ and $0 \leq t < \infty$.

We again begin our analysis by finding firm 2's best response $t = b(s)$. It is easier to do this if you think like a businessman rather than like a mathematician.

Unit cost of production is c. If firm 1 produces so much that all by itself it drives the price down to c or lower, there is no way for firm 2 to make a positive profit. In this case firm 2's best response is to produce nothing: that way its profit is 0, which is better than losing money.

Firm 1 drives the price p down to c when its level of production s satisfies the equation

$$c = \alpha - \beta s.$$

The solution of this equation is $s = \frac{\alpha - c}{\beta}$. We conclude that if $s \geq \frac{\alpha - c}{\beta}$, firm 2's best response is 0.

On the other hand, if firm 1 produces $s < \frac{\alpha - c}{\beta}$, it leaves the price above c, and gives firm 2 an opportunity to make a positive profit. Firm 2 certainly does not want to produce so much as to drive the price down to c or lower. Firm 2 drives the price p down to c when its level of production t satisfies the equation

$$c = \alpha - \beta(s+t).$$

The solution of this equation is $t = \frac{\alpha-\beta s-c}{\beta}$. Therefore firm 2 wants to choose a level of production t between 0 and $\frac{\alpha-\beta s-c}{\beta}$.

On the interval $0 \le t \le \frac{\alpha-\beta s-c}{\beta}$, firm 2's profit is given by

$$\pi_2(s,t) = (\alpha - \beta(s+t) - c)t = (\alpha - \beta s - c)t - \beta t^2,$$

because

$$t \le \frac{\alpha-\beta s-c}{\beta} \quad \Rightarrow \quad s+t \le s + \frac{\alpha-\beta s-c}{\beta} = \frac{\alpha-c}{\beta} < \frac{\alpha}{\beta}.$$

Thus firm 2 wants to maximize the function $(\alpha - \beta s - c)t - \beta t^2$ on the interval $0 \le t \le \frac{\alpha-\beta s-c}{\beta}$.

The function $(\alpha - \beta s - c)t - \beta t^2$ equals 0 at the endpoints of our interval, and it has a graph that is an upside-down parabola. Therefore it is maximum where $\frac{\partial \pi_2}{\partial t}(s,t) = 0$, which occurs at $t = \frac{\alpha-\beta s-c}{2\beta}$. See Figure 1.9.

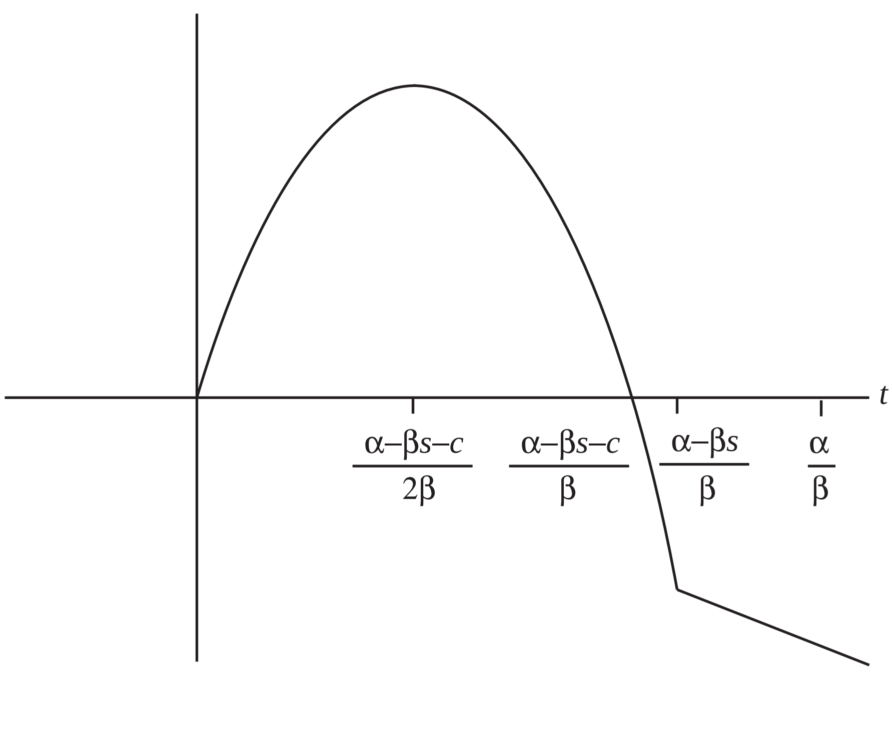

Figure 1.9. Graph of $\pi_2(s,t)$ for fixed $s < \frac{\alpha-c}{\beta}$.

Thus firm 2's best response function is:

$$b(s) = \begin{cases} \frac{\alpha-\beta s-c}{2\beta} & \text{if } 0 \leq s < \frac{\alpha-c}{\beta}, \\ 0 & \text{if } s \geq \frac{\alpha-c}{\beta}. \end{cases}$$

We now turn to calculating $\pi_1(s, b(s))$, the payoff that firm 1 can expect from each choice s assuming that firm 2 uses its best response strategy.

Notice that for $0 \leq s < \frac{\alpha-c}{\beta}$, we have

$$s + b(s) = s + \frac{\alpha - \beta s - c}{2\beta} = \frac{\alpha + \beta s - c}{2\beta} < \frac{\alpha - c}{\beta} < \frac{\alpha}{\beta}.$$

Therefore, for $0 \leq s < \frac{\alpha-c}{\beta}$,

$$\pi_1(s, b(s)) = \pi_1(s, \frac{\alpha - \beta s - c}{2\beta}) = (\alpha - \beta(s + \frac{\alpha - \beta s - c}{2\beta}) - c)s = \frac{\alpha - c}{2}s - \frac{\beta}{2}s^2.$$

Firm 2 will not choose an $s \geq \frac{\alpha-c}{\beta}$, since, as we have seen, that would force the price down to c or lower. Therefore we will not bother to calculate $\pi_1(s, b(s))$ for $s \geq \frac{\alpha-c}{\beta}$.

The function $\pi_1(s, b(s))$ on the interval $0 \leq s \leq \frac{\alpha-c}{\beta}$ is maximum at $s^* = \frac{\alpha-c}{2\beta}$, where the derivative of $\frac{\alpha-c}{2}s - \frac{\beta}{2}s^2$ is 0, just as in our first model. The value of $t^* = b(s^*)$ is also the same, as are the price and profits.

1.10. Economics and calculus background

In this section we give some background that will be useful for the next two examples, as well as later in the course.

1.10.1. Utility functions. A salary increase from \$20,000 to \$30,000 and a salary increase from \$220,000 to \$230,000 are not equivalent in their effect on your happiness. This is true even if you don't have to pay taxes!

Let s be your salary and $u(s)$ the "utility" of your salary to you. Two commonly assumed properties of $u(s)$ are:

(1) $u'(s) > 0$ for all s ("strictly increasing utility function"). In other words, more is better!
(2) $u''(s) < 0$ ("strictly concave utility function"). In other words, $u'(s)$ decreases as s increases.

1.10.2. Discount factor. Happiness now is different from happiness in the future.

Suppose your boss proposes to you a salary of s this year and t next year. The total utility to you *today* of this offer is $U(s, t) = u(s) + \delta u(t)$, where δ is a "discount factor." Typically, $0 < \delta < 1$. The closer δ is to 1, the more important the future is to you.

Which would you prefer, a salary of s this year and s next year, or a salary of $s - a$ this year and $s + a$ next year? Assume $0 < a < s$, $u' > 0$, and $u'' < 0$. Then

$$\begin{aligned} U(s,s) - U(s-a, s+a) &= u(s) + \delta u(s) - (u(s-a) + \delta u(s+a)) \\ &= u(s) - u(s-a) - \delta(u(s+a) - u(s)) \\ &= \int_{s-a}^{s} u'(t)\,dt - \delta \int_{s}^{s+a} u'(t)\,dt > 0. \end{aligned}$$

Hence you prefer s each year.

Do you see why the last line is positive? Part of the reason is that $u'(s)$ decreases as s increases, so $\int_{s-a}^{s} u'(t)\,dt > \int_{s}^{s+a} u'(t)\,dt$.

1.10.3. Maximum value of a function. Suppose f is a differentiable function on an interval $a \le x \le b$. From calculus we know:

(1) f attains a maximum value somewhere on the interval.
(2) The maximum value of f occurs at a point where $f' = 0$ or at an endpoint of the interval.
(3) If $f'(a) > 0$, the maximum does not occur at a.
(4) If $f'(b) < 0$, the maximum does not occur at b.

Suppose that $f'' < 0$ everywhere in the interval $a \le x \le b$. Then we know a few additional things:

(1) f attains attains its maximum value at *unique* point c in $[a, b]$.
(2) Suppose $f'(x_0) > 0$ at some point $x_0 < b$. Then $x_0 < c$.

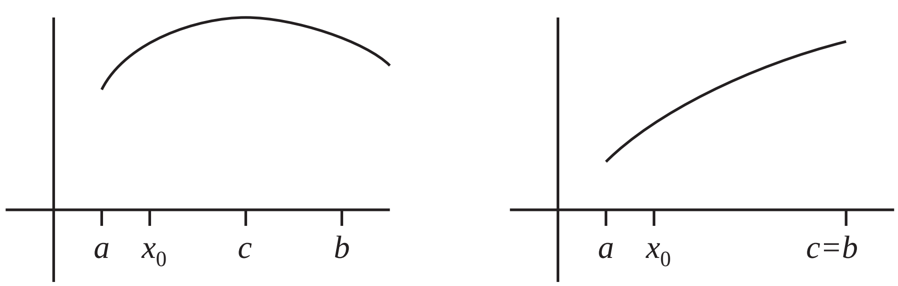

FIGURE 1.10. Two functions on $[a, b]$ with negative second derivative everywhere and positive first derivative at one point $x_0 < b$. Such functions always attain their maximum at a point c to the right of x_0.

(3) Suppose $f'(x_1) < 0$ at some point $x_1 > a$. Then $c < x_1$.

1.11. The Samaritan's Dilemma

There is someone you want to help should she need it. However, you are worried that the very fact that you are willing to help may lead her to do less for herself than she otherwise would. This is the Samaritan's Dilemma.

The Samaritan's Dilemma is an example of *moral hazard*. Moral hazard is "the prospect that a party insulated from risk may behave differently from the way it would behave if it were fully exposed to the risk." There is a Wikipedia article on moral hazard: `http://en.wikipedia.org/wiki/Moral_hazard`.

Here is an example of the Samaritan's Dilemma analyzed by James Buchanan (Nobel Prize in Economics, 1986). A young woman plans to go to college next year. This year she is working and saving for college. If she needs additional help, her father will give her some of the money he earns this year.

Notation and assumptions regarding income and savings:

(1) Father's income this year is $y > 0$, which is known. Of this he will give $0 \leq t \leq y$ to his daughter next year.
(2) Daughter's income this year is $z > 0$, which is also known. Of this she saves $0 \leq s \leq z$ to spend on college next year.
(3) Daughter chooses the amount s of her income to save for college. Father then observes s and chooses the amount t to give to his daughter.

The important point is (3): after Daughter is done saving, Father will choose an amount to give to her. Thus the daughter, who goes first in this game, can use backward induction to figure out how much to save. In other words, she can take into account that different savings rates will result in different levels of support from Father.

Utility functions:

(1) Daughter's utility function $\pi_1(s, t)$, which is her payoff in this game, is the sum of
    (a) her first-year utility v_1, a function of the amount she has to spend in the first year, which is $z - s$; and
    (b) her second-year utility v_2, a function of the amount she has to spend in the second year, which is $s + t$. Second-year utility is multiplied by a discount factor $\delta > 0$.

    Thus we have

$$\pi_1(s, t) = v_1(z - s) + \delta v_2(s + t).$$

(2) Father's utility function $\pi_2(s, t)$, which is his payoff in this game, is the sum of

(a) his personal utility u, a function of the amount he has to spend in the first year, which is $y - t$; and

(b) his daughter's utility π_1, multiplied by a "coefficient of altruism" $\alpha > 0$.

Thus we have

$$\pi_2(s,t) = u(y-t) + \alpha\pi_1(s,t) = u(y-t) + \alpha(v_1(z-s) + \delta v_2(s+t)).$$

Notice that a component of Father's utility is Daughter's utility. The Samaritan's Dilemma arises when the welfare of someone else is important to us.

We assume:

(A1) The functions v_1, v_2, and u have positive first derivative and negative second derivative.

Let's first gather some facts that we will use in our analysis

(1) Formulas we will need for partial derivatives:

$$\frac{\partial \pi_1}{\partial s}(s,t) = -v_1'(z-s) + \delta v_2'(s+t),$$
$$\frac{\partial \pi_2}{\partial t}(s,t) = -u'(y-t) + \alpha\delta v_2'(s+t).$$

(2) Formulas we will need for second partial derivatives:

$$\frac{\partial^2 \pi_1}{\partial s^2}(s,t) = v_1''(z-s) + \delta v_2''(s+t),$$
$$\frac{\partial^2 \pi_2}{\partial s \partial t}(s,t) = \alpha\delta v_2''(s+t),$$
$$\frac{\partial^2 \pi_2}{\partial t^2}(s,t) = u''(y-t) + \alpha\delta v_2''(s+t).$$

All three of these are negative everywhere.

To figure out Daughter's savings rate using backward induction, we must first maximize $\pi_2(s,t)$ with s fixed and $0 \le t \le y$. Let's keep things simple by arranging that for s fixed, $\pi_2(s,t)$ will attain its maximum at some t strictly between 0 and y. This is guaranteed to happen if $\frac{\partial \pi_2}{\partial t}(s,0) > 0$ and $\frac{\partial \pi_2}{\partial t}(s,y) < 0$.

For $0 \le s \le z$, we have

$$\frac{\partial \pi_2}{\partial t}(s,0) = -u'(y) + \alpha\delta v_2'(s) \ge -u'(y) + \alpha\delta v_2'(z)$$

and

$$\frac{\partial \pi_2}{\partial t}(s,y) = -u'(0) + \alpha\delta v_2'(s+y) \le -u'(0) + \alpha\delta v_2'(y).$$

We therefore make two more assumptions:

(A2) $\alpha\delta v_2'(z) > u'(y)$. This assumption is reasonable. We expect Daughter's income z to be much less than Father's income y. Since, as we have discussed, each dollar of added income is less important when income is higher, we expect $v_2'(z)$ to be much greater than $u'(y)$. If the product $\alpha\delta$ is not too small (meaning that Father cares quite a bit about Daughter, and Daughter cares quite a bit about the future), we get our assumption.

(A3) $u'(0) > \alpha\delta v_2'(y)$. This assumption is reasonable because $u'(0)$ should be large and $v_2'(y)$ should be small.

With these assumptions, we have

$$\frac{\partial \pi_2}{\partial t}(s,0) > 0 \quad \text{and} \quad \frac{\partial \pi_2}{\partial t}(s,y) < 0 \quad \text{for all } 0 \le s \le z.$$

Since $\frac{\partial^2 \pi_2}{\partial t^2}$ is always negative, there is a single value of t where $\pi_2(s,t)$, s fixed, attains its maximum value; moreover, $0 < t < y$, so, $\frac{\partial \pi_2}{\partial t}(s,t) = 0$ at this value of t. We denote this value of t by $t = b(s)$. This is Father's best response strategy, the amount Father will give to Daughter if the amount Daughter saves is s.

The daughter now chooses her saving rate $s = s^*$ to maximize the function $\pi_1(s, b(s))$, which we shall denote $V(s)$:

$$V(s) = \pi_1(s, b(s)) = v_1(z - s) + \delta v_2(s + b(s)).$$

Father then contributes $t^* = b(s^*)$.

Here is the punchline: suppose it turns out that $0 < s^* < z$, i.e., Daughter saves some of her income but not all. (This is the usual case.) Then, had Father simply committed himself in advance to providing t^* in support to his daughter no matter how much she saved, Daughter would have chosen a savings rate s greater than s^*. *Both* Daughter and Father would have ended up with higher utility.

We can see this in a series of steps.

(1) In order to maximize $V(s)$, we calculate

$$V'(s) = -v_1'(z - s) + \delta v_2'(s + b(s))(1 + b'(s)).$$

(2) If $V(s)$ is maximum at $s = s^*$ with $0 < s^* < z$, we must have $V'(s^*) = 0$, i.e.,

$$0 = -v_1'(z - s^*) + \delta v_2'(s^* + t^*)(1 + b'(s^*)).$$

(3) We have

$$\frac{\partial \pi_1}{\partial s}(s^*, t^*) = -v_1'(z - s^*) + \delta v_2'(s^* + t^*).$$

(4) Subtracting (2) from (3), we obtain

$$\frac{\partial \pi_1}{\partial s}(s^*, t^*) = -\delta v_2'(s^* + t^*) b'(s^*).$$

(5) We expect that $b'(s) < 0$; this simply says that if Daughter saves more, Father will contribute less. To check this, we note that

$$\frac{\partial \pi_2}{\partial t}(s, b(s)) = 0 \text{ for all } s.$$

Differentiating both sides of this equation with respect to s, we get

$$\frac{\partial^2 \pi_2}{\partial s \partial t}(s, b(s)) + \frac{\partial^2 \pi_2}{\partial t^2}(s, b(s)) b'(s) = 0.$$

Since $\frac{\partial^2 \pi_2}{\partial s \partial t}$ and $\frac{\partial^2 \pi_2}{\partial t^2}$ are always negative, we must have $b'(s) < 0$.

(6) From (4), since v_2' is always positive and $b'(s)$ is always negative, we see that $\frac{\partial \pi_1}{\partial s}(s^*, t^*)$ is positive.

(7) From (6) and the fact that $\frac{\partial^2 \pi_1}{\partial s^2}(s, t)$ is always negative, we see that $\pi_1(s, t^*)$ is maximum at a value $s = s^\sharp$ greater than s^*. (See Subsection 1.10.3.)

(8) We of course have $\pi_1(s^\sharp, t^*) > \pi_1(s^*, t^*)$, so Daughter's utility is higher. Since Daugher's utility is higher, we see from the formula for π_2 that $\pi_2(s^\sharp, t^*) > \pi_2(s^*, t^*)$, so Father's utility is also higher.

This problem has implications for government social policy. It suggests that social programs be made available to everyone rather than on an if-needed basis.

1.12. The Rotten Kid Theorem

A rotten son manages a family business. The amount of effort the son puts into the business affects both his income and his mother's. The son, being rotten, cares only about his own income, not his mother's. To make matters worse, Mother dearly loves her son. If the son's income is low, Mother will give part of her own income to her son so that he will not suffer. In this situation, can the son be expected to do what is best for the family?

We shall give the analysis of Gary Becker (Nobel Prize in Economics, 1992; Wikipedia article http://en.wikipedia.org/wiki/Gary_Becker).

We denote the mother's annual income by y and the son's by z. The amount of effort that the son devotes to the family business is denoted by a. His choice of a will affect both his income and his mother's, so we regard both y and z as functions of a: $y = y(a)$ and $z = z(a)$.

After mother observes a, and hence observes her own income $y(a)$ and her son's income $z(a)$, she chooses an amount t, $0 \leq t \leq y(a)$, to give to her son.

The mother and son have personal utility functions u and v respectively. Each is a function of the amount they have to spend.

The son chooses his effort a to maximize his own utility v, without regard for his mother's utility u. Mother, however, chooses the amount t to transfer to her son to maximize $u(y-t)+\alpha v(z+t)$, where α is her coefficient of altruism. Thus the payoff functions for this game are

$$\pi_1(a,t) = v(z(a)+t),$$
$$\pi_2(a,t) = u(y(a)-t)+\alpha v(z(a)+t).$$

Since the son chooses first, he can use backward induction to decide how much effort to put into the family business. In other words, he can take into account that even if he doesn't put in much effort, and so doesn't produce much income for either himself or his mother, his mother will help him out.

Assumptions:

(1) The functions u and v have positive first derivative and negative second derivative.
(2) The son's level of effort is chosen from an interval $I = [a_1, a_2]$.
(3) For all a in I, $\alpha v'(z(a)) > u'(y(a))$. This assumption expresses two ideas: (1) Mother dearly loves her son, so α is not small; and (2) no matter how little or how much the son works, Mother's income $y(a)$ is much larger than son's income $z(a)$. (Recall that the derivative of a utility function gets smaller as the income gets larger.) This makes sense if the income generated by the family business is small compared to Mother's overall income
(4) For all a in I, $u'(0) > \alpha v'(z(a)+y(a))$. This assumption is reasonable because $u'(0)$ should be large and $v'(z(a)+y(a))$ should be small.
(5) Let $T(a) = y(a)+z(a)$ denote total family income. Then $T'(a) = 0$ at a unique point $a^\sharp$, $a_1 < a^\sharp < a_2$, and $T(a)$ attains its maximum value at this point. This assumption expresses the idea that if the son works too hard, he will do more harm than good. As they say in the software industry, if you stay at work too late, you're just adding bugs.

To find the son's level of effort using backward induction, we must first maximize $\pi_2(a,t)$ with a fixed and $0 \le t \le y(a)$. We calculate

$$\frac{\partial \pi_2}{\partial t}(a,t) = -u'(y(a)-t)+\alpha v'(z(a)+t),$$
$$\frac{\partial \pi_2}{\partial t}(a,0) = -u'(y(a))+\alpha v'(z(a)) > 0,$$
$$\frac{\partial \pi_2}{\partial t}(a,y(a)) = -u'(0)+\alpha v'(z(a)+y(a)) < 0,$$
$$\frac{\partial^2 \pi_2}{\partial t^2}(a,t) = u''(y(a)-t)+\alpha v''(z(a)+t) < 0.$$

Then there is a single value of t where $\pi_2(a,t)$, a fixed, attains its maximum; moreover, $0 < t < y(a)$, so $\frac{\partial \pi_2}{\partial t}(a,t) = 0$. (See Subsection 1.10.3.) We denote this value of t by $t = b(a)$. This is Mother's strategy, the amount Mother will give to her son if his level of effort in the family business is a.

The son now chooses his level of effort $a = a^*$ to maximize the function $\pi_1(a, b(a))$, which we shall denote $V(a)$:

$$V(a) = \pi_1(a, b(a)) = v(z(a) + b(a)).$$

Mother then contributes $t^* = b(a^*)$.

So what? Here is Becker's point.

Suppose $a_1 < a^* < a_2$ (the usual case). Then $V'(a^*) = 0$, i.e.,

$$v'(z(a^*) + b(a^*))(z'(a^*) + b'(a^*)) = 0.$$

Since v' is positive everywhere, we have

$$z'(a^*) + b'(a^*) = 0. \tag{1.1}$$

Now $-u'(y(a) - b(a)) + \alpha v'(z(a) + b(a)) = 0$ for all a. Differentiating this equation with respect to a, we find that, for all a,

$$-u''(y(a) - b(a))(y'(a) - b'(a)) + \alpha v''(z(a) + b(a))(z'(a) + b'(a)) = 0.$$

In particular, for $a = a^*$,

$$-u''(y(a^*) - b(a^*))(y'(a^*) - b'(a^*)) + \alpha v''(z(a^*) + b(a^*))(z'(a^*) + b'(a^*)) = 0.$$

This equation and (1.1) imply that

$$y'(a^*) - b'(a^*) = 0.$$

Adding this equation to (1.1), we obtain

$$y'(a^*) + z'(a^*) = 0.$$

Therefore $T'(a^*) = 0$. But then, by our last assumption, $a^* = a^\sharp$, the level of effort that maximizes total family income.

Thus, if the son had not been rotten, and instead had been trying to maximize total family income $y(a) + z(a)$, he would have chosen the same level of effort a^*.

1.13. Backward induction for finite horizon games

Backward induction as we defined it in Section 1.4 does not apply to any game that is not finite. However, a variant of backward induction can be used on any finite horizon game of complete information. It is actually this variant that we have been using since Section 1.8.

Let us describe this variant of backward induction in general. The idea that, in a game that is not finite, we cannot remove nodes one-by-one, because we will never finish. Instead we must remove big collections of nodes at each step.

(1) Let $k \geq 1$ be the length of the longest path in the game. (This number is finite since we are dealing with a finite horizon game.) Consider the collection $\mathcal{C}$ of all nodes c such that every move that starts at c is the last move in a path of length k. Each such move has an end that is terminal.
(2) For each node c in $\mathcal{C}$, identify the player $i(c)$ who is to choose at node c. Among all the moves available to him at that node, find the move $m(c)$ whose end gives the greatest payoff to player $i(c)$. We assume that this move is unique.
(3) Assume that at each node c in $\mathcal{C}$, the player $i(c)$ will choose the move $m(c)$. Record this choice as part of player $i(c)$'s strategy.
(4) Delete from the game tree all moves that start at all of the nodes in $\mathcal{C}$. The nodes c in $\mathcal{C}$ are now terminal nodes. Assign to them the payoffs that were previously assigned to the nodes $m(c)$.
(5) In the new game tree, the length of the longest path is now $k - 1$. If $k - 1 = 0$, stop. Otherwise, return to step 1.

CHAPTER 2

Eliminating Dominated Strategies

2.1. Prisoner's Dilemma

Two corporate executives are accused of preparing false financial statements. The prosecutor has enough evidence to send both to jail for one year. However, if one confesses and tells the prosecutors what he knows, the prosecutor will be able to send the other to jail for 10 years. In exchange for the help, the prosecutor will let the executive who confesses go free.

If both confess, both will go to jail for 6 years.

The executives are held in separate cells and cannot communicate. Each must decide individually whether to talk or refuse.

Since each executive decides what to do without knowing what the other has decided, it is not natural or helpful to draw a game tree. Nevertheless, we can still identify the key elements of a game: players, strategies, and payoffs.

The players are the two executives. Each has the same two strategies: talk or refuse. The payoffs to each player are the number of years in jail (preceded by a minus sign, since we want higher payoffs to be more desirable). The payoff to each executive depends on the strategy choices of both executives.

In this two-player game, we can indicate how the strategies determine the payoffs by a matrix.

		Executive 2	
		talk	**refuse**
Executive 1	**talk**	$(-6,-6)$	$(0,-10)$
	refuse	$(-10,0)$	$(-1,-1)$

The rows of the matrix represent player 1's strategies. The columns represent player 2's strategies. Each entry of the matrix is an ordered pair of numbers that gives the payoffs to the two players if the corresponding strategies are used. Player 1's payoff is given first.

Notice:

(1) If player 2 talks, player 1 gets a better payoff by talking than by refusing.

(2) if player 2 refuses to talk, player 1 still gets a better payoff by talking than by refusing.

Thus, no matter what player 2 does, player 1 gets a better payoff by talking than by refusing. Player 1's strategy of talking *strictly dominates* his strategy of refusing: it gives a better payoff to player 1 no matter what player 2 does.

Of course, player 2's situation is identical: his strategy of talking gives him a better payoff no matter what player 1 does.

Thus we expect both executives to talk.

Unfortunately for them, the result is that they both go to jail for 6 years. Had they both refused to talk, they would have gone to jail for only one year.

Prosecutors like playing this game. Defendants don't like it much. Hence there have been attempts over the years by defendants' attorneys and friends to change the game.

For example, if the Mafia were involved with the financial manipulations that are under investigation, it might have told the two executives in advance: "If you talk, something bad could happen to your child." Suppose each executive believes this warning and considers something bad happening to his child to be equivalent to 6 years in prison. The payoffs in the game are changed as follows:

		Executive 2	
		talk	**refuse**
Executive 1	**talk**	$(-12, -12)$	$(-6, -10)$
	refuse	$(-10, -6)$	$(-1, -1)$

Now, for both executives, the strategy of refusing to talk dominates the strategy of talking. Thus we expect both executives to refuse to talk, so both go to jail for only one year.

The Mafia's threat sounds cruel. In this instance, however, it helped the two executives achieve a better outcome for themselves than they could achieve on their own.

Prosecutors don't like the second version of the game. One mechanism they have of returning to the first version is to offer "witness protection" to prisoners who talk. In a witness protection program, the witness and his family are given new identities in a new town. If the prisoner believes that the Mafia is thereby prevented from carrying out its threat, the payoffs return to something close to those of the original game.

Another way to change the game of Prisoner's Dilemma is by additional rewards. For example, the Mafia might work hard to create a culture in which prisoners who don't talk are honored by their friends, and their families are taken care of. If the two executives buy into this system, and consider the rewards of not talking to be worth 5 years in prison, the payoffs become the following:

		Executive 2	
		talk	**refuse**
Executive 1	**talk**	$(-6, -6)$	$(0, -5)$
	refuse	$(-5, 0)$	$(4, 4)$

Once again refusing to talk has become a dominant strategy.

The Prisoner's Dilemma is the best-known and most-studied model in game theory. It models many common situations. We illustrate this point in Sections 2.5 and 2.4, where we discuss the Israeli-Palestian conflict and the problem of global warming. You may also want to look at the Wikipedia page on the Prisoners' Dilemma (`http://en.wikipedia.org/wiki/Prisoners_dilemma`).

2.2. Games in normal form

This section is related to Gintis, Sec. 3.3.

A game in normal form consists of:

(1) A finite set P of *players*. We will usually take $P = \{1, \dots, n\}$.
(2) For each player i, a set S_i of available strategies.

Let $S = S_1 \times \cdots \times S_n$. An element of S is an n-tuple $(s_1, \dots, s_n)$ where each s_i is a strategy chosen from the set S_i. Such an n-tuple $(s_1, \dots, s_n)$ is called a *strategy profile*. It represents a choice of strategy by each of the n players.

(3) For each player i, a *payoff function* $\pi_i : S \to \mathbb{R}$.

In the Prisoner's Dilemma, $P = \{1, 2\}$, $S_1 = \{\text{talk, refuse}\}$, $S_2 = \{\text{talk, refuse}\}$, and S is a set of four ordered pairs, namely (talk, talk), (talk, refuse), (refuse, talk), and (refuse, refuse). As to the payoff functions, we have, for example, π_1(refuse, talk) $= -10$ and π_2(refuse, talk) $= 0$.

If there are two players, player 1 has m strategies, and player 2 has n strategies, then a game in normal form can be represented by an $m \times n$ matrix of ordered pairs of numbers, as in the previous section. We will refer to such a game as an $m \times n$ game.

2.3. Dominated strategies

This section is related to Gintis, Sec. 4.1.

For a game in normal form, let s_i and s_i' be two of player i's strategies.

- We say that s_i *strictly dominates* s_i' if, *for every choice of strategies by the other players*, the payoff to player i from using s_i is greater than the payoff to player i from using s_i'.
- We say that s_i *weakly dominates* s_i' if, *for every choice of strategies by the other players*, the payoff to player i from using s_i is at least as great as the payoff to player i from using s_i'; and, *for some choice of strategies by the other players*, the payoff to player i from using s_i is greater than the payoff to player i from using s_i'.

As mentioned in Section 1.4, game theorists often assume that players are rational. One meaning of rationality for a game in normal form is:

- Suppose one of player i's strategies s_i weakly dominates another of his strategies s_i'. Then player i will not use the strategy s_i'.

This is the assumption we used to analyze the Prisoner's Dilemma. Actually, in that case, we only needed to eliminate strictly dominated strategies.

A Prisoner's Dilemma occurs when (1) each player has a strategy that strictly dominates all his other strategies, but (2) each player has another strategy such that, if all players were to use this alternative, all players would receive higher payoffs than those they receive when they all use their dominant strategies.

2.4. Israelis and Palestinians

Henry Kissinger was National Security Advisor and later Secretary of State during the administrations of Richard Nixon of Gerald Ford. Previously he was a professor of international relations at Harvard. In his view, the most important contribution of the game theory point of view in international relations was that it forced you to make a very explicit model of the situation you wanted to understand.

Let's look at the Israeli-Palestinian conflict with this opinion of Kissinger's in mind.

The roots of the conflict go back to Roman times, when most Jews were expelled from the land of Israel, their homeland, after a series of failed revolts against the Romans. A small number of Jews continued to live there. Jewish immigration to the region increased during the 19th century, largely because of persecution in Eastern Europe. In the late 19th century, nationalism became an increasingly powerful force in Central and Eastern Europe, as various nationalities in the multiethnic Austro-Hungarian, Russian, and Ottoman Empires, such as the Poles, Czechs, Slovaks, Bulgarians, Croats, and many others, began to demand their own countries. The Jewish national movement that developed at this time, Zionism, eventually focused on the ancient Jewish homeland, then the Ottoman region of Palestine, as the future site for a Jewish nation. From the late 19th century on, Jewish immigration to Palestine had as its goal the eventual formation of a Jewish nation. This goal was in obvious conflict with the aspirations of the existing Arab population.

In the 1930's and 40's Jewish immigration to Palestine from Europe increased further due to Nazi persecution. In 1948 the United Nations divided Palestine, which had been run by the British since 1918, into two countries, one for the Jews and one for the Arabs. The Arabs rejected the division. In the ensuing fighting, the Jews took over more of Palestine than they had been awarded by the U.N. The Jewish-occupied areas became the state of Israel; the remaining areas, in Gaza and on the West Bank of the Jordan River, did not become a Palestinian Arab state as the U.N. had envisioned, but were instead administered by Egypt and Jordan respectively. Many Arab residents of the Jewish-occupied areas ended up in refugee camps in Gaza, the West Bank, and nearby Arab countries.

In a war between Israel and the neighboring Arab countries in 1967, the Israeli army occupied both Gaza and the West Bank, as well as other territories. The West Bank especially was seen by many Israelis as being a natural part of the state of Israel for both religious reasons (the Jewish heartland in Biblical times was in what is now the West Bank) and military reasons. Considerable Jewish settlement took place in the West Bank, and to a lesser extent in Gaza, after 1968, with the goal of retaining at least part of these territories in an eventual resolution of the conflict.

In 2000 negotiations between Israeli Prime Minister Ehud Barak and Palestinian leader Yasser Arafat, with the mediation of U.S. President Bill Clinton, perhaps came close to resolving the conflict. Barak offered to remove most of the Israeli settlements and allow establishment of a Palestinian state. Arafat rejected the offer. The level of conflict between the two sides increased greatly and has remained high. In 2005 the Israelis abandoned their settlements in Gaza and ended their occupation of that region.

Discussion of this conflict usually focuses on two issues: control of the West Bank and terrorism. Most proposals for resolving the conflict envision a trade-off in which the Israelis would end their occupation of the West Bank and the Palestinians would stop terrorism, the means by which they carry on their conflict with Israel.

In a simple model of the conflict, the Israelis have two possible strategies: continue to occupy the West Bank or end the occupation. The Palestinians also have two possible strategies: continue terrorism or end terrorism. What are the payoffs?

The Israelis certainly value both keeping the West Bank and an end to Palestinian terrorism. The Palestinians certainly value ending the Israeli occupation of the West Bank. I will assume that the Palestinians also value retaining their freedom to continue terrorism. The reason is that for the Palestinians, giving up terrorism essentially means giving up hope of regaining the pre-1967 territory of Israel, which was the home of many Palestinians, and which many Palestinians feel is rightfully their territory.

Let's consider two ways to assign payoffs that are consistent with this analysis.

1. At the time of the negotiations between Barak and Arafat, Barak apparently considered an end to terrorism to be of greater value than continued occupation of the West Bank, since he was willing to give up the latter in exchange for the former. Therefore we will assign the Israelis 2 points if terrorism ends, and 1 point if they continue to occupy the West Bank.

Arafat apparently considered retaining the freedom to engage in terrorism to be of greater value than ending the Israeli occupation of the West Bank, since he was not willing to give up the former to achieve the latter. Therefore we assign the Palestinians 2 points if they keep the freedom to engage in terrorism, and 1 point if the Israelis end their occupation of the West Bank. We get the following game in normal form.

		Palestinians	
		terrorism	**end terrorism**
Israelis	**occupy West Bank**	$(1,2)$	$(3,0)$
	end occupation	$(0,3)$	$(2,1)$

The table shows that the Israelis have a dominant strategy, occupy the West Bank, and the Palestinians have a dominant strategy, continue terrorism. These strategies yield the actual outcome of the negotiations.

This game is not a Prisoner's Dilemma. In the Prisoner's Dilemma, each player has a dominant strategy, but the use of the dominated strategies by each player would result in a higher payoff to both. Here, if each player uses his dominated strategy, the Israeli outcome improves, but the Palestinian outcome is worse.

2. The previous assignment of payoffs was appropriate for Israeli "moderates" and Palestinian "radicals." We will now assign payoffs on the assumption that both governments are "moderate." The Israeli payoffs are unchanged. The Palestinians are now assumed to value ending the occupation of the West Bank above keeping the freedom to engage in terrorism. We therefore assign the Palestinians 1 point if they keep the freedom to engage in terrorism, and 2 points if the Israelis end their occupation of the West Bank. We get the following game in normal form.

		Palestinians	
		terrorism	**end terrorism**
Israelis	**occupy West Bank**	$(1,1)$	$(3,0)$
	end occupation	$(0,3)$	$(2,2)$

The Israelis still have the dominant strategy, occupy the West Bank, and the Palestinians still have the dominant strategy, continue terrorism. This indicates that even "moderate" governments on both sides will have difficulty resolving the conflict.

The game is now a Prisoner's Dilemma: if both sides use their dominated strategies, there will be a better outcome for both, namely an end to both the Israeli occupation of the West Bank and to Palestinian terrorism. As in the original Prisoner's Dilemma, this outcome is not easy to achieve. As in the original Prisoner's Dilemma, one solution is for an outside player to change the payoffs by supplying punishments or rewards, as the Mafia could there. In the context of the Israeli-Palestinian conflict, the most plausible such outside player is the United States.

Of course, once one considers whether the United States wants to become involved, one has a three-person game. If one considers subgroups within the Israelis and Palestinians (for example, the radical Palestinian group Hamas), the game becomes even more complicated.

ADD REFERENCE TO GUY AT NYU.

2.5. Global warming

Ten countries are considering fighting global warming. Each country must choose to spend an amount x_i to reduce its carbon emissions, where $0 \leq x_i \leq 1$. The total benefits produced by these expenditures equal twice the total expenditures: $2(x_1 + \cdots + x_{10})$. Each country receives $\frac{1}{10}$ of the benefits.

This game has ten players, the ten countries. The set of strategies available to country i is just the closed interval $0 \leq x_i \leq 1$. A strategy profile is therefore a 10-tuple $(x_1, \ldots, x_{10})$, where $0 \leq x_i \leq 1$ for each i. The ith country's payoff function is its benefits minus its expenditures:

$$\pi_i(x_1, \ldots, x_{10}) = \frac{1}{10} \cdot 2(x_1 + \cdots + x_{10}) - x_i = \frac{1}{5}(x_1 + \cdots + x_{10}) - x_i.$$

We will show that for each country, the strategy $x_i = 0$ (spend nothing to fight global warming) *strictly dominates* all its other strategies.

We will just show this for country 1, since the arguments for the other countries are the same. Let $x_1 > 0$ be a different strategy for country 1. Let $x_2, \ldots, x_n$ be *any* strategies for the other countries. Then

$$\pi_1(0, x_2, \ldots, x_{10}) - \pi_1(x_1, x_2, \ldots, x_{10}) =$$
$$(\frac{1}{5}(0 + x_2 + \cdots + x_{10}) - 0) - (\frac{1}{5}(x_1 + x_2 + \cdots + x_{10}) - x_1) = -\frac{1}{5}x_1 + x_1 = \frac{4}{5}x_1 > 0.$$

Thus we expect each country to spend nothing to fight global warming, and each country to get a payoff of 0.

If all countries could somehow agree to spend 1 each to fight global warming, each country's payoff would be $\frac{1}{5}(1+\cdots+1)-1=2-1=1$, and each country would be better off. In fact, each country would receive benefits of 2 in return for expenditures of 1, an excellent deal.

Nevertheless, each country would be constantly tempted to cheat. A reduction in country i's expenditures by y_i dollars reduces total benefits to all countries by $2y_i$ dollars, but only reduces benefits to country i by $\frac{1}{5}y_i$ dollars.

This example suggests that the problem of global warming is a type of prisoner's dilemma.

Of course, one can try to change the game by changing the payoffs with punishments or rewards. For example, one might try to raise the environmental consciousness of people around the world by a publicity campaign. Then perhaps governments that fight global warming would get the approval of their own people and the approval of others around the world, which they might see as a reward. In addition, governmental leaders might get subjective rewards by doing what they feel is the right thing.

2.6. Hagar's Battles

This section is related to Gintis, Sec. 4.9.

There are ten villages with values $a_1 < a_2 < \cdots < a_{10}$. There are two players. Player 1 has n_1 soldiers, and player 2 has n_2 soldiers, with $0 < n_1 < 10$ and $0 < n_2 < 10$. Each player independently decides which villages to send his soldiers to. A player is not allowed to send more than one soldier to a village.

A player wins a village if he sends a soldier there but his opponent does not.

A player's score is the sum of the values of the villages he wins. The winner of the game is the player with the higher score.

Where should you send your soldiers?

Since each player decides where to send his soldiers without knowledge of the other player's decision, we will model this game as a game in normal form. To do that, we must describe precisely the *players*, the *strategies*, and the *payoff functions*.

- Players. There are two.
- Strategies. The villages are numbered from 1 to 10. A strategy for player i is just a set of n_i numbers between 1 and 10. The numbers represent the n_i different villages to which he sends his soldiers. Thus if S_i is the set of all of player i's strategies, an element s_i of S_i is simply a set of n_i numbers between 1 and 10.

- Payoff functions. A player's payoff in this game is his score minus his opponent's score. If this number is positive, he wins; if it is negative, he loses.

A neat way to analyze this game is to find a nice formula for the payoff function. Lets look at an example. Suppose $n_1 = n_2 = 3$, $s_1 = \{6, 8, 10\}$, and $s_2 = \{7, 9, 10\}$. Player 1 wins villages 6 and 8, and player 2 wins villages 7 and 9. Thus player 1's payoff is $(a_6 + a_8) - (a_7 + a_9)$, and player 2's payoff is $(a_7 + a_9) - (a_6 + a_8)$. Since $a_6 < a_7$ and $a_8 < a_9$, player 2 wins.

We could also calculate player i's payoff by adding the values of *all* the villages to which he sends his soldiers, and subtracting the values of all the villages to which his opponent sends his soldiers. Then we would have

- Player 1's payoff $= (a_6 + a_8 + a_{10}) - (a_7 + a_9 + a_{10}) = (a_6 + a_8) - (a_7 + a_9)$.
- Player 2's payoff $= (a_7 + a_9 + a_{10}) - (a_6 + a_8 + a_{10}) = (a_7 + a_9) - (a_6 + a_8)$.

Clearly this always works. Thus we have the following formulas for the payoff functions.

$$\pi_1(s_1, s_2) = \sum_{j \in s_1} a_j - \sum_{j \in s_2} a_j,$$
$$\pi_2(s_1, s_2) = \sum_{j \in s_2} a_j - \sum_{j \in s_1} a_j.$$

We claim that for each player, the strategy of sending his n_i soldiers to the n_i villages of highest values *strictly dominates* all his other strategies.

We will just show that for player 1, the strategy of sending his n_1 soldiers to the n_1 villages of highest values strictly dominates all his other strategies. The argument for player 2 is the same.

Let s_1 be the set of the n_1 highest numbers between 1 and 10. (For example, if $n_1 = 3$, $s_1 = \{8, 9, 10\}$). Let s_1' a different strategy for player 1, i.e., a different set of n_1 numbers between 1 and 10. Let s_2 be *any* strategy for player 2, i.e., any set of n_2 numbers between 1 and 10. We must show that

$$\pi_1(s_1, s_2) > \pi_1(s_1', s_2).$$

We have

$$\pi_1(s_1, s_2) = \sum_{j \in s_1} a_j - \sum_{j \in s_2} a_j,$$
$$\pi_1(s_1', s_2) = \sum_{j \in s_1'} a_j - \sum_{j \in s_2} a_j.$$

Therefore

$$\pi_1(s_1, s_2) - \pi_1(s_1', s_2) = \sum_{j \in s_1} a_j - \sum_{j \in s_1'} a_j.$$

This is clearly positive: the sum of the n_1 biggest numbers between 1 and 10 is greater than the sum of some other n_1 numbers between 1 and 10.

2.7. Second-price auctions

This section is related to Gintis, Sec. 4.6.

An item is to be sold at auction. Each bidder submits a sealed bid. All the bids are opened. The object is sold to the highest bidder, but *the price is the bid of the second-highest bidder*.

(If two or more bidders submit equal highest bids, that is the price, and one of those bidders is chosen by chance to buy the object. However, we will ignore this possibility in our analysis.)

If you are a bidder at such an auction, how much should you bid?

Clearly the outcome of the auction depends not only on what you do, but on what the other bidders do. Thus we can think of the auction as a game. Since the bidders bid independently, without knowledge of the other bids, we will try to model this auction as a game in normal form. We must describe precisely the players, the strategies, and the payoff functions.

- Players. We'll suppose there are n bidders.
- Strategies. The ith player's strategy is simply his bid, which we will denote b_i. At this point we must decide whether to allow just integer bids, arbitrary real number bids, or something else. Let's try allowing the bids b_i to be any nonnegative real number. The ith player's set of available bids is then $S_i = [0, \infty)$. Thus each player has an infinite number of possible strategies, indeed an interval. The definition of a game in normal form allows this.
- Payoff functions. A reasonable idea for the payoff function is that the payoff to player i is 0 unless he wins the auction, in which case the payoff to player i is the value of the object to player i minus the price he has to pay for it. Thus the payoff to player i depends on
  (1) the value of the object to player i, which we will denote v_i;
  (2) the bid of player i, b_i; and
  (3) the highest bid of the other players, which we denote $h_i = \max\{b_j : j \neq i\}$.

  The formula is

$$\pi_i(b_1, \dots, b_n) = \begin{cases} 0 & \text{if } b_i < h_i, \\ v_i - h_i & \text{if } h_i < b_i. \end{cases}$$

(Recall that we are ignoring the possibility that two bidders submit equal highest bids, i.e., we ignore the possibility that $b_i = h_i$.)

We claim that for player i, the strategy v_i weakly dominates every other strategy. In other words, you should bid exactly what the object is worth to you. (This is the great thing about second-price auctions.)

To show this, we will just show that for player 1, the strategy v_1 weakly dominates every other strategy. The argument for any other player is the same.

Let $b_1 \neq v_1$ be another possible bid by player 1. We must show two things:

(1) If $b_2, \ldots, b_n$ are *any* bids by the other players, then

$$\pi_1(v_1, b_2, \ldots, b_n) \geq \pi_1(b_1, b_2, \ldots, b_n).$$

(2) There are *some* bids $b_2, \ldots, b_n$ by the other players such that

$$\pi_1(v_1, b_2, \ldots, b_n) > \pi_1(b_1, b_2, \ldots, b_n).$$

To show (1), let $h_1 = \max(b_2, \ldots, b_n)$. We show all the possibilities in a table.

Relation of v_1 to h_1	Relation of b_1 to h_1	$\pi_1(v_1, b_2, \ldots, b_n)$	$\pi_1(b_1, b_2, \ldots, b_n)$
$v_1 < h_1$	$b_1 < h_1$	0	0
$v_1 < h_1$	$h_1 < b_1$	0	$v_1 - h_1 < 0$
$h_1 < v_1$	$b_1 < h_1$	$v_1 - h_1 > 0$	0
$h_1 < v_1$	$h_1 < b_1$	$v_1 - h_1 > 0$	$v_1 - h_1 > 0$

In every case, $\pi_1(v_1, b_2, \ldots, b_n) \geq \pi_1(b_1, b_2, \ldots, b_n)$. This shows (1). The second and third lines of the table, in which $\pi_1(v_1, b_2, \ldots, b_n) > \pi_1(b_1, b_2, \ldots, b_n)$, show (2).

There is a Wikipedia page about second-price auctions:
http://en.wikipedia.org/wiki/Sealed_second-price_auction.

2.8. Iterated elimination of dominated strategies

This section is related to Gintis, Sec. 4.1.

With games in extensive form, a simple notion of rationality was not to choose a move if one that yielded a higher payoff was available. This notion inspired the idea of repeatedly eliminating such moves, thereby repeatedly simplifying the game, a procedure we called backward induction.

With games in normal form, the corresponding simple notion of rationality is not to use a dominated strategy. If we remove a dominated strategy from a game in normal form, we obtain a game in normal form with one less strategy. If the smaller game has a dominated strategy, it can then be removed. This procedure, known as *iterated elimination of dominated strategies*, can be repeated until no dominated strategies remain. The result is a smaller game to analyze.

If the smaller game includes only one strategy s_i^* for player i, s_i^* is called a *dominant strategy* for player i. If the smaller game includes only one strategy s_i^* for *every* player, the strategy profile $(s_1^*, \ldots, s_n^*)$ is called a *dominant strategy equilibrium*.

Iterated elimination of strictly dominated strategies produces the same reduced game in whatever order it is done. However, we shall see that iterated elimination of weakly dominated strategies can produce different reduced games when done in different orders.

2.9. The Battle of the Bismarck Sea

The following description of the Battle of the Bismarck Sea is drastically simplified. For a fuller story, see the Wikipedia page (http://en.wikipedia.org/wiki/Battle_of_the_Bismarck_Sea).

In 1943, during the Second World War, a Japanese admiral was ordered to reinforce a base on the island of New Guinea. The supply convoy could take either a rainy northern route or a sunny southern route. The Americans knew the day the convoy would sail and wanted to bomb it. They only had enough reconnaissance aircraft to search one route per day. The northern route was too rainy for bombing one day in three, although it could still be searched by the reconnaissance aircraft. The sailing time was three days.

The Japanese admiral, who was aware that the Americans knew when the convoy would sail, had to decide which route to send it by. The Americans had to decide which route to search on that day.

The payoff to the Americans is the number of days they are able to bomb the convoy. The payoff to the Japanese is minus this number.

The payoff matrix is shown below.

		Japanese	
		sail north	**sail south**
Americans	**search north**	$(1\frac{2}{3}, -1\frac{2}{3})$	$(2, -2)$
	search south	$(1\frac{1}{3}, -1\frac{1}{3})$	$(2\frac{1}{2}, -2\frac{1}{2})$

Explanation:

- If the Americans search the northern route while the Japanese sail south, they will spend a day searching and finding nothing. They will then know that the Japanese sailed south, and will be able to bomb for the remaining two days of sailing.
- If the Americans search the southern route while the Japanese sail south, they will on average find the convoy toward midday, and will be able to bomb for $2\frac{1}{2}$ days.

- Similar arguments give the payoffs when the Japanese sail north, except that one must take into account that it is clear enough to bomb only two-thirds of the time.

Neither of the American strategies dominates the other. For the Japanese, however, sailing north strictly dominates sailing south. We therefore eliminate the Japanese strategy sail south. The resulting game has two American strategies but only one Japanese strategy. In this smaller game, the American strategy search north dominates search south. We therefore eliminate search south. All that remains is sail north for the Japanese, and search north for the Americans. This is in fact what happened.

2.10. Normal form of a game in extensive form with complete information

This section is related to Gintis, Sec. 3.1.

Recall that for a game in extensive form, a player's strategy is a plan for what action to take in every situation that the player might encounter. We can convert a game in extensive form to one in normal form by simply listing the possible strategies for each of the n players, then associating to each strategy profile the resulting payoffs.

For example, consider the game of Big Monkey and Little Monkey described in Section 1.5. Big Monkey has two strategies, wait (w) and climb (c). Little Monkey has four strategies:

- ww: if Big Monkey waits, wait; if Big monkey climbs, wait.
- wc: if Big Monkey waits, wait; if Big monkey climbs, climb.
- cw: if Big Monkey waits, climb; if Big monkey climbs, wait.
- cc: if Big Monkey waits, climb; if Big monkey climbs, climb.

The normal form of this game has the following payoff matrix.

		Little Monkey			
		ww	**wc**	**cw**	**cc**
Big Monkey	**w**	$(0,0)$	$(0,0)$	$(9,1)$	$(9,1)$
	c	$(4,4)$	$(5,3)$	$(4,4)$	$(5,3)$

2.11. Big Monkey and Little Monkey 2

The game of Big Monkey and Little Monkey with Big Monkey going first provides an example showing that iterated elimination of weakly dominated strategies can produce different reduced games when done in different orders.

Here is one way of doing iterated elimination of weakly dominated strategies in the game.

(1) Eliminate Little Monkey's strategy ww because it is weakly dominated by cw, and eliminate Little Monkey's strategy wc because it is weakly dominated by cc.
(2) Eliminate Little Monkey's strategy cc because in the reduced 2×2 game it is weakly dominated by cw.
(3) Eliminate Big Monkey's strategy c because in the reduced 2×1 game it is dominated by w.
(4) What remains is the 1×1 game in which Big Monkey's strategy is w and Little Monkey's strategy is cw. Thus each is a dominant strategy, and (w, cw) is a dominant strategy equilibrium.

These are the strategies we found for Big Monkey and Little Monkey by backward induction.

However, here is another way of doing iterated elimination of weakly dominated strategies in this game.

(1) As before, begin by eliminating Little Monkey's strategy ww because it is weakly dominated by cw, and eliminate Little Monkey's strategy wc because it is weakly dominated by cc.
(2) Eliminate Big Monkey's strategy c because in the reduced 2×2 game it is dominated by w.
(3) We are left with a 2×1 game in which no more strategies can be eliminated. The remaining strategy profiles are (w, cw) (found before) and (w, cc). This way of doing iterated elimination shows that w is a dominant strategy for Big Monkey, but does not show that cw is a dominant strategy for player 2.

Can you find other ways of doing iterated elimination of weakly dominated strategies in this game?

2.12. Backward induction and iterated elimination of dominated strategies

For a game in extensive form, each way of going through the backward induction procedure is equivalent to a corresponding way of performing iterated elimination of *weakly dominated* strategies in the normal form of the same game.

We will now show that in the game of Big Monkey and Little Monkey with Big Monkey going first, one way of doing backward induction corresponds to the first way of doing iterated elimination of weakly dominated strategies that was described in the previous section.

(1) Suppose you begin backward induction by noting that if Big Monkey waits, Little Monkey should climb, and reduce the game tree accordingly. In iterated elimination of weakly dominated strategies, this corresponds to eliminating Little Monkey's strategy ww because it is weakly dominated by cw, and eliminating Little Monkey's strategy wc because it is weakly dominated by cc. The payoff matrix of the reduced game has two rows and just two columns (cw and cc).

(2) The second step in backward induction is to note that if Big Monkey climbs, Little Monkey should wait, and reduce the game tree accordingly. In iterated elimination of weakly dominated strategies, this corresponds to eliminating Little Monkey's strategy cc in the 2×2 game because it is weakly dominated by cw. The payoff matrix of the reduced game has two rows and just the cw column.

(3) The last step in backward induction is to use the reduced game tree to decide that Big Monkey should wait. In iterated elimination of weakly dominated strategies, this corresponds to eliminating Big Monkey's strategy c because, in the reduced game with only the cw column, it is dominated by w.

We will now describe the correspondence between backward induction and iterated elimination of weakly dominated strategies for games in extensive form with *two* players. The general situation just requires more notation.

Consider a game in extensive form with two players.

- Player 1 moves at nodes c_i, $1 \leq i \leq p$. At each node c_i he has available a set M_i of moves.
- Player 2 moves at nodes d_j, $1 \leq j \leq q$. At each node d_j he has available a set N_j of moves.
- A strategy for player 1 is an assignment to each of his nodes of one of the moves available at that node. Thus player 1's strategy set is $M = M_1 \times \cdots \times M_p$, and a strategy for player 1 is an ordered p-tuple $(m_1, \ldots, m_p)$ with each $m_i \in M_i$.
- Similarly, player 2's strategy set is $N = N_1 \times \cdots \times N_q$. A strategy for player 2 is an ordered q-tuple $(n_1, \ldots, n_q)$ with each $n_j \in N_j$.
- The normal form of the game associates to each pair $(m, n) \in M \times N$ payoffs $\pi_1(m, n)$ and $\pi_2(m, n)$. To determine these payoffs, the game in extensive form is played with the strategies m and n. It ends in a uniquely defined terminal node, whose payoffs are then used.

We consider a backward induction procedure in the extensive form of the game. Assume the players' nodes are numbered so that each player's nodes are reached in the order last to first.

1. Without loss of generality, suppose the first node treated is player 1's node c_p. Each move in M_p ends in a terminal vertex, which has assigned to it payoffs to

players 1 and 2. Let m_p^* be the move in M_p that gives the greatest payoff to player 1. We assume m_p^* is unique.

Backward induction records the fact that at node c_p, player 1 will choose m_p^*; deletes from the game tree all moves that start at c_p; and assigns to the now terminal node c_p the payoffs previously assigned to the end of m_p^*.

The corresponding step in iterated elimination of weakly dominated strategies is to remove all of player 1's strategies $(m_1, \ldots, m_{p-1}, m_p)$ with $m_p \neq m_p^*$. The reason is that each such strategy is weakly dominated by $(m_1, \ldots, m_{p-1}, m_p^*)$. Against any strategy of player 2, the latter gives a better payoff if play reaches node c_p, and the same payoff if it does not.

2. Assume now that backward induction has reached player 1's nodes $c_{k+1}, \ldots, c_p$ and player 2's nodes $d_{l+1}, \ldots, d_q$. Without loss of generality, suppose the next node treated is player 1's node c_k. At this point in backward induction, each move in M_k ends in a terminal vertex, which has assigned to it payoffs to players 1 and 2. Let m_k^* be the move in M_k that gives the greatest payoff to player 1. We assume m_k^* is unique.

Backward induction records the fact that at node c_k, player 1 will choose m_k^*; deletes from the game tree all moves that start at c_k; and assigns to the now terminal node c_k the payoffs previously assigned to the end of m_k^*.

At the corresponding step in iterated elimination of weakly dominated strategies, the remaining strategies of player 1 are those of the form $(m_1, \ldots, m_k, m_{k+1}^*, \ldots, m_p^*)$, and the remaining strategies of player 2 are those of the form $(n_1, \ldots, n_l, n_{l+1}^*, \ldots, n_q^*)$. We now remove all of player 1's strategies $(m_1, \ldots, m_{k-1}, m_k, m_{k+1}^*, \ldots, m_p^*)$ with $m_k \neq m_k^*$. The reason is that each such strategy is weakly dominated by $(m_1, \ldots, m_{k-1}, m_k^*, m_{k+1}^*, \ldots, m_p^*)$. Against any of player 2's *remaining strategies*, the latter gives a better payoff if play reaches node c_k, and the same payoff if it does not.

3. Backward induction or iterated elimination of weakly dominated strategies eventually produces unique strategies $(m_1^*, \ldots, m_p^*)$ for player 1 and $(n_1^*, \ldots, n_q^*)$ for player 2.

CHAPTER 3

Nash equilibria

3.1. Big Monkey and Little Monkey 3 and the definition of Nash equilibria

This section is related to Gintis, Sec. 3.5.

There are many games in normal form that cannot be analyzed by elimination of dominated strategies. For example, in the encounter between Big Monkey and Little Monkey (see Section 1.5), suppose Big Monkey and Little Monkey decide *simultaneously* whether to wait or climb. Then we get the following payoff matrix.

		Little Monkey	
		wait	**climb**
Big Monkey	**wait**	$(0,0)$	$(9,1)$
	climb	$(4,4)$	$(5,3)$

There are no dominated strategies.

Consider a game in normal form with n players, strategy sets $S_1, \ldots, S_n$, and payoff functions $\pi_1, \ldots, \pi_n$. A *Nash equilibrium* is a strategy profile $(s_1^*, \ldots, s_n^*)$ with the following property: if any *single* player, say the ith, changes his strategy, his own payoff will not increase.

In other words, $(s_1^*, \ldots, s_n^*)$ is a Nash equilibrium provided

- For every $s_1 \in S_1$, $\pi_1(s_1^*, s_2^*, \ldots, s_n^*) \geq \pi_1(s_1, s_2^*, \ldots, s_n^*)$.
- For every $s_2 \in S_2$, $\pi_2(s_1^*, s_2^*, s_3^*, \ldots, s_n^*) \geq \pi_2(s_1^*, s_2, s_3^*, \ldots, s_n^*)$.

  $\vdots$
- For every $s_n \in S_n$, $\pi_n(s_1^*, \ldots, s_{n-1}^*, s_n^*) \geq \pi_n(s_1^*, \ldots, s_{n-1}^*, s_n)$.

A *strict* Nash equilibrium is a strategy profile $(s_1^*, \ldots, s_n^*)$ with the property: if any single player, say the ith, changes his strategy, his own payoff will decrease.

In other words, $(s_1^*, \ldots, s_n^*)$ is a strict Nash equilibrium provided

- For every $s_1 \neq s_1^*$ in S_1, $\pi_1(s_1^*, s_2^*, \ldots, s_n^*) > \pi_1(s_1, s_2^*, \ldots, s_n^*)$.
- For every $s_2 \neq s_2^*$ in S_2, $\pi_2(s_1^*, s_2^*, s_3^*, \ldots, s_n^*) > \pi_2(s_1^*, s_2, s_3^*, \ldots, s_n^*)$.

  $\vdots$
- For every $s_n \neq s_n^*$ in S_n, $\pi_n(s_1^*, \ldots, s_{n-1}^*, s_n^*) > \pi_n(s_1^*, \ldots, s_{n-1}^*, s_n)$.

In the game of Big Monkey and Little Monkey described above, there are two strict Nash equilibria, (wait, climb) and (climb, wait):

- The strategy profile (wait, climb) is a Nash equilibrium. It produces the payoffs $(9, 1)$. If Big Monkey changes to climb, his payoff decreases from 9 to 5. If Little Monkey changes to wait, his payoff decreases from 1 to 0.
- The strategy profile (climb, wait) is also a Nash equilibrium. It produces the payoffs $(4, 4)$. If Big Monkey changes to wait, his payoff decreases from 4 to 0. If Little Monkey changes to climb, his payoff decreases from 4 to 3.

Big Monkey prefers the first of these Nash equilibria, Little Monkey the second.

Note that in this game, the strategy profiles (wait, wait) and (climb, climb) are *not* Nash equilibria. In fact, for these strategy profiles, *either* monkey could improve his payoff by changing his strategy.

Game theorists often use the following notation when discussing Nash equilibria. Let $s = (s_1, \dots, s_n)$ denote a strategy profile. Suppose in s we replace the ith player's strategy s_i by another of his strategies, say s_i'. The resulting strategy profile is then denoted (s_i', s_{-i}).

In this notation, a strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is a Nash equilibrium if, for each $i = 1, \dots, n$, $\pi_i(s^*) \geq \pi_i(s_i, s_{-i}^*)$ for every $s_i \in S_i$. The strategy profile s^* is a strict Nash equilibrium if, for each $i = 1, \dots, n$, $\pi_i(s^*) > \pi_i(s_i, s_{-i}^*)$ for every $s_i \neq s_i^*$ in S_i.

The notion of Nash equilibrium is the most important idea in game theory.

We will consider three ways of finding Nash equilibria:

- Inspection (Sections 3.2–3.4).
- Iterated elimination of dominated strategies (Sections 3.5–3.6).
- Using best response (Sections 3.7–3.10).

3.2. Finding Nash equilibria by inspection

One way to find Nash equilibria is by inspection. This is how we found the Nash equilibria in the game of Big Monkey and Little Monkey in the previous section. Here are two more simple examples.

3.2.1. Prisoner's Dilemma. Recall the Prisoner's Dilemma from Section 2.1:

		Executive 2	
		talk	**refuse**
Executive 1	**talk**	$(-6, -6)$	$(0, -10)$
	refuse	$(-10, 0)$	$(-1, -1)$

Let's inspect all four strategy profiles

- The strategy profile (talk,talk) is a strict Nash equilibrium. If either executive alone changes his strategy to refuse, his payoff falls from -6 to -10.
- The strategy profile (refuse, refuse) is not a Nash equilibrium. If either executive alone changes his strategy to talk, his payoff increases from -1 to 0.
- The strategy profile (talk, refuse) is not a Nash equilibrium. If executive 2 changes his strategy to talk, his payoff increases from -10 to -6.
- The strategy profile (refuse, talk) is not a Nash equilibrium. If executive 1 changes his strategy to talk, his payoff increases from -10 to -6.

The Prisoner's Dilemma illustrates an important fact about Nash equilibria: they are not necessarily good for the players! Instead they are strategy profiles where a game can get stuck, for better or worse.

3.2.2. Stag hunt. Two hunters on horseback are pursuing a stag. If both work together, they will succeed. However, the hunters notice that they are passing some hare. If either hunter leaves the pursuit of the stag to pursue a hare, he will succeed, but the stag will escape.

Let's model this situation as a 2-player game in which the players decide their moves simultaneously. The players are the hunters. The possible strategies for each are pursue the stag and pursue a hare. Let's suppose that the payoff for catching a hare is 1 to the hunter who caught it, and the payoff for catching the stag is 2 to each hunter. The payoff matrix is

		Hunter 2	
		stag	**hare**
Hunter 1	**stag**	$(2,2)$	$(0,1)$
	hare	$(1,0)$	$(1,1)$

There are no dominated strategies. If we inspect all four strategy profiles, we find that there are two strict Nash equilibria, (stag, stag) and (hare, hare). Both hunters prefer (stag, stag) to (hare, hare).

The Wikipedia page for this game is http://en.wikipedia.org/wiki/Stag_hunt. Games like Stag Hunt, in which there are several Nash equilibria, one of which is best for *all* players, are sometimes called *pure coordination games*. The players should somehow coordinate their actions so that they are in the best of the Nash equilibria. For more information, see the Wikipedia page http://en.wikipedia.org/wiki/Coordination_game.

3.3. Water Pollution 1

Three firms use water from a lake. When a firm returns the water to the lake, it can purify it or fail to purify it (and thereby pollute the lake).

The cost of purifying the used water before returning it to the lake is 1. If two or more firms fail to purify the water before returning it to the lake, all three firms incur a cost of 3 to treat the water before they can use it.

The payoffs are therefore as follows:

- If all three firms purify: -1 to each firm.
- If two firms purify and one pollutes: -1 to each firm that purifies, 0 to the polluter.
- If one firm purifies and two pollute: -4 to the firm that purifies, -3 to each polluter.
- If all three firms pollute: -3 to each firm.

We inspect these possibilities to see if any are Nash equilibria.

- Suppose all three firms purify. If one switches to polluting, his payoff increases from -1 to 0. This is not a Nash equilibrium.
- Suppose two firms purify and one pollutes. If a purifier switches to polluting, his payoff decreases from -1 to -3. If the polluter switches to purifying, his payoff decreases from 0 to -1. Thus there are three Nash equilibria in which two firms purify and one pollutes.
- Suppose one firm purifies and two pollute. If the purifier switches to polluting, his payoff increases from -4 to -3. This is not a Nash equilibrium.
- Suppose all three firms pollute. If one switches to purifying, his payoff decreases from -3 to -4. This is a Nash equilibrium.

We see here an example of the *free rider problem* (Wikipedia page: http://en.wikipedia.org/wiki/Free_rider_problem). Each firm wants to be the one that gets the advantage of the other firms' efforts without making any effort itself.

The free rider problem arises in negotiating treaties to deal with climate change. For example, the U.S. objected to the 1997 Kyoto protocol because it did not require action by developing countries such as China and India (Wikipedia page: http://en.wikipedia.org/wiki/Kyoto_treaty).

The free rider problem also arises in connection with the Troubled Asset Relief Program (TARP), under which the U.S. Treasury invested several hundred billion dollars in U.S. banks in late 2008 and early 2009. (Wikipedia page: http://en.wikipedia.org/wiki/Troubled_Asset_Relief_Program.) You may have read that some banks want to pay back this investment. Here is the difficulty:

- The problem is that the banks have made many loans to people and corporations who now appear unable to repay them. If too many of these loans are not repaid, the banks will not have enough money to repay their depositors, rendering these banks bankrupt.
- The banks therefore need to conserve what cash they have, so they are unwilling to make new loans.
- If the banks are unwilling to lend, the economy will slow, and more people and corporations will go bankrupt, making it even less likely that that the original problem loans will be repaid.
- The government therefore invests in ("injects capital into") the banks. The banks now have more money, so they will, the government hopes, lend more. Then the economy will pick up, some of the problem loans will be repaid, and the banks will be okay.
- Unfortunately for the banks, the government's investment is accompanied by annoying requirements, such as limitations on executive pay.
- If a few banks can repay the government's investment, they can avoid the annoying requirements, but still benefit from the economic boost and loan repayments due to the em other banks' increased lending. The banks that repay the government become free riders.

3.4. Tobacco market

This section is related to Gintis, Sec. 5.4.

At a certain warehouse, the supply of tobacco in pounds, q, is related to the price of tobacco per pound in dollars, p, by the formula

$$q = 100,000(10 - p).$$

However, a price support program ensures that the price never falls below \$.25 per pound. In other words, if the supply is so high that the price would be below \$.25 per pound, the price is set at $p = .25$, and a government agency purchases whatever cannot be sold at that price.

One day three farmers are the only ones bringing their tobacco to this warehouse. Each has harvested 600,000 pounds and can bring as much of his harvest as he wants. Whatever is not brought must be discarded.

There are three players, the farmers. The ith farmer's strategy is simply the amount of tobacco he brings to the warehouse, and hence is a number q_i, $0 \leq q_i \leq 600,000$. The payoff to farmer i is $\pi_i(q_1, q_2, q_3) = pq_i$, where p is found by solving the equation

$$q_1 + q_2 + q_3 = 100,000(10 - p),$$

provided the solution p is at least .25. We find

$$p = \begin{cases} 10 - \frac{q_1+q_2+q_3}{100{,}000} & \text{if } q_1+q_2+q_3 \leq 975,000, \\ .25 & \text{if } q_1+q_2+q_3 > 975,000. \end{cases}$$

We will find the Nash equilibria by inspecting all strategy profiles (q_1, q_2, q_3), $0 \leq q_i \leq 600,000$.

1. Suppose some q_i=0. Then farmer i's payoff is 0, which he could increase by bringing some of his tobacco to market. This is not a Nash equilibrium.

2. Suppose $q_1 + q_2 + q_3 \geq 975,000$. The price is then \$.25, and will stay the same if any farmer brings more of his tobacco to market. Thus if any q_i is less than 600,000, that farmer could increase his own payoff by bringing all 600,000 pounds of his tobacco to market. Thus the only possible Nash equilibrium in this region is (600,000, 600,000, 600,000). It really is one: if any farmer alone brings less to market, the price will not rise, so his payoff will certainly decrease.

3. Suppose $q_1+q_2+q_3 < 975,000$ and $0 < q_i < 600,000$ for all i. In this region the payoff functions π_i are given by

$$\pi_i(q_1, q_2, q_3) = (10 - \frac{q_1 + q_2 + q_3}{100,000})q_i.$$

Suppose (q_1, q_2, q_3) is a Nash equilibrium in this region. Let us consider first farmer 1. The maximum value of $\pi_1(q_1, q_2, q_3)$, with q_2 and q_3 fixed at their Nash equilibrium values, must occur at the Nash equilibrium. Since q_1 is not an endpoint of the interval $0 \leq q_1 \leq 600,000$, we must have $\frac{\partial \pi_1}{\partial q_1} = 0$ at the Nash equilibrium. By considering farmers 2 and 3, we get the additional equations $\frac{\partial \pi_2}{\partial q_2} = 0$ and $\frac{\partial \pi_3}{\partial q_3} = 0$.

Our system of equations is therefore

$$10 - \frac{q_1 + q_2 + q_3}{100,000} - \frac{q_1}{100,000} = 0,$$
$$10 - \frac{q_1 + q_2 + q_3}{100,000} - \frac{q_2}{100,000} = 0,$$
$$10 - \frac{q_1 + q_2 + q_3}{100,000} - \frac{q_3}{100,000} = 0.$$

Hence $q_1 = q_2 = q_3$. Then the first equation implies that

$$10 - \frac{3q_1}{100,000} - \frac{q_1}{100,000} = 0,$$

so $q_1 = 250,000$. Hence the only possible Nash equilibrium in this region is (250,000, 250,000, 250,000).

To check that this really is a Nash equilibrium, let us consider farmer 1. (The others are similar). For $(q_2, q_3) = (250{,}000,\ 250{,}000)$, farmer 1's payoff function is

$$\pi_1(q_1,\ 250{,}000,\ 250{,}000) = \begin{cases} (10 - \frac{q_1+500{,}000}{100{,}000})q_1 & \text{if } 0 \leq q_1 \leq 475,000, \\ .25q_1 & \text{if } 475,000 < q_1 \leq 600,000. \end{cases}$$

The quadratic function $(10 - \frac{q_1+500{,}000}{100{,}000})q_1$, $0 \leq q_1 \leq 475,000$, is maximum at the point we have found, $q_1 = 250,000$, where $\pi_1 = 2.50 \times 250,000 = 625,000$. Moreover, for $475,000 < q_1 \leq 600,000$, $\pi_1 = .25q_1$ is at most $150,000$. Therefore farmer 1 cannot improve his payoff by changing q_1. The same is true for farmers 2 and 3, so (250,000, 250,000, 250,000) is indeed a Nash equilibrium.

4. There is one case we have not yet considered: $q_1 + q_2 + q_3 < 975,000$, $0 < q_i < 600,000$ for two i, and $q_i = 600,000$ for one i. It turns out that there are no Nash equilibria in this case. The analysis is left as homework.

Thus there are two Nash equilibria, one of which is preferred to the others by all three farmers. This is a pure coordination game, like Stag Hunt (Subsection 3.2.2). If the farmers can agree among themselves to each bring 250,000 pounds of tobacco to market and discard 350,000 pounds, none will have an incentive to cheat. On the other hand, the tobacco buyers would prefer that they each bring 600,000 pounds to market. If all farmers do that, none can improve his own payoff be bringing less.

3.5. Finding Nash equilibria by iterated elimination of dominated strategies

This section is related to Gintis, Sec. 2.2.

The relation between iterated elimination of dominated strategies, which we discussed in Chapter 2, and Nash equilibria is summarized in the following theorems.

Theorem 3.1. *Suppose we do iterated elimination of* weakly *dominated strategies on a game G in normal form. Let H be the reduced game that results. Then:*

(1) *Each Nash equilibrium of H is also a Nash equilibrium of G.*
(2) *In particular, if H has only one strategy s_i^* for each player, then the strategy profile $(s_1^*, \ldots, s_n^*)$ is a Nash equilibrium of G.*

The last conclusion of Theorem 3.1 just says that every dominant strategy equilibrium is a Nash equilibrium.

When one does iterated elimination of *strictly* dominated strategies, one can say more.

Theorem 3.2. *Suppose we do iterated elimination of* strictly *dominated strategies on a game G in normal form. Then:*

(1) *Any order yields the same reduced game H.*
(2) *Each strategy that is eliminated is not part of any Nash equilibrium of G.*
(3) *Each Nash equilibrium of H is also a Nash equilibrium of G.*
(4) *If H has only one strategy s_i^* for each player, $(s_1^*, \dots, s_n^*)$ is a strict Nash equilibrium. Moreover, there are no other Nash equilibria.*

Theorem 3.2 justifies using iterated elimination of *strictly* dominated strategies to reduce the size of the game to be analyzed. It says in part that we do not miss any Nash equilibria by doing the reduction.

On the other hand, we certainly *can* miss Nash equilibria by using iterated elimination of *weakly* dominated strategies to reduce the size of the game. This is the second problem with iterated elimination of weakly dominated strategies that we have identified; the first was that the resulting smaller game can depend on the order in which the elimination is done.

We shall prove statements (2), (3), and (4) of Theorem 3.2, and we shall indicate how our proof of statement (3) of Theorem 3.2 can be modified to give a proof of statement (1) of Theorem 3.1. Of course, statement (2) of Theorem 3.1 follows from statement (1) of Theorem 3.1.

Proof. We consider iterated elimination of strictly dominated strategies on a game G in normal form.

To prove statement (2) of Theorem 3.2, we will prove by induction the statement: The kth strategy that is eliminated is not part of any Nash equilibrium.

$k = 1$: Let s_i, a strategy of player i, be the first strategy that is eliminated. It was eliminated because the ith player has a strategy t_i that strictly dominates it. Suppose s_i is part of a strategy profile. Replacing s_i by t_i, and leaving all other players' strategies the same, increases the payoff to player i. Therefore this strategy profile is not a Nash equilibrium.

Assume the statement is true for $k = 1, \dots, l$. Let s_i, a strategy of player i, be the $l + 1$st strategy that is eliminated. It was eliminated because the ith player has a strategy t_i that strictly dominates it, assuming no player uses any previously eliminated strategy. Suppose s_i is part of a strategy profile. If any of the previously eliminated strategies is used in this strategy profile, then by assumption it is not a Nash equilibrium. If none of the previously eliminated strategies is used, then replacing s_i by t_i, and leaving all other players' strategies the same, increases the payoff to player i. Therefore this strategy profile is not a Nash equilibrium.

This completes the proof of statement (2) of Theorem 3.2.

To prove statement (3) of Theorem 3.2, let $(s_1^*, s_2^*, \dots, s_n^*)$ be a Nash equilibrium of H, and let S_1 be player 1's strategy set for G. We must show that for every

$s_1 \neq s_1^*$ in S_1,

$$\pi_1(s_1^*, s_2^*, \dots, s_n^*) \geq \pi_1(s_1, s_2^*, \dots, s_n^*). \tag{3.1}$$

(Of course, we must also prove an analogous statement for the other players, but the argument would be the same.)

If s_1 is a strategy of player 1 that remains in the reduced game H, then of course (3.1) follows from the fact that $(s_1^*, s_2^*, \dots, s_n^*)$ is a Nash equilibrium of H. To complete the proof we will show that if s_1 is any strategy of player 1 that was eliminated in the course of iterated elimination of strictly dominated strategies, then

$$\pi_1(s_1^*, s_2^*, \dots, s_n^*) > \pi_1(s_1, s_2^*, \dots, s_n^*). \tag{3.2}$$

Note the strict inequality.

In fact, we will prove by reverse induction the statement: If s_1 is the kth strategy of player 1 to be eliminated, then (3.2) holds.

Let s_1 be the *last* strategy of player 1 to be eliminated. It was eliminated because one of the remaining strategy of player 1, say t_1, when used against any remaining strategies $s_2, \dots, s_n$ of the other players, satisfied $\pi_1(t_1, s_2, \dots, s_n) > \pi_1(s_1, s_2, \dots, s_n)$. In particular, since $s_2^*, \dots, s_n^*$ were among the remaining strategies, we get

$$\pi_1(t_1, s_2^*, \dots, s_n^*) > \pi_1(s_1, s_2^*, \dots, s_n^*). \tag{3.3}$$

Since $(s_1^*, s_2^*, \dots, s_n^*)$ is a Nash equilibrium of H and t_1 is a strategy available to player 1 in H (it was never eliminated),

$$\pi_1(s_1^*, s_2^*, \dots, s_n^*) \geq \pi_1(t_1, s_2^*, \dots, s_n^*). \tag{3.4}$$

Combining (3.4) and (3.3), we get (3.2).

Assume the statement is true for all $k \geq l$. Let s_1 be the $(\ell - 1)$st strategy of player 1 to be eliminated. It was eliminated because one of the other remaining strategies of player 1, say t_1, when used against any remaining strategies $s_2, \dots, s_n$ of the other players, satisfied $\pi_1(t_1, s_2, \dots, s_n) > \pi_1(s_1, s_2, \dots, s_n)$. In particular, since $s_2^*, \dots, s_n^*$ were among the remaining strategies, we get (3.3). If t_1, is never eliminated, then it is still available to player 1 in H, so we have (3.4). Combining (3.4) and (3.3), we again get (3.2). If t_1 *is* one of the strategies of player 1 that is eliminated after s_1, then, by the induction hypothesis,

$$\pi_1(s_1^*, s_2^*, \dots, s_n^*) > \pi_1(t_1, s_2^*, \dots, s_n^*). \tag{3.5}$$

Combining (3.5) and (3.3), we get (3.2).

This completes the proof of statement (3) of Theorem 3.2.

To prove statement (4) of Theorem 3.2, we must show that if H has only one strategy s_i^* for each player, then for every $s_1 \neq s_1^*$ in S_1,

$$\pi_1(s_1^*, s_2^*, \dots, s_n^*) > \pi_1(s_1, s_2^*, \dots, s_n^*). \tag{3.6}$$

In this case, every $s_1 \neq s_1^*$ in S_1 is eliminated in the course of iterated elimination of strictly dominated strategies, so the proof we gave of statement (3) actually yields the conclusion.

The proof of statement (1) of Theorem 3.1 is essentially the same as the proof of statement (3) of Theorem 3.2, except that no inequalities are strict. □

3.6. Big Monkey and Little Monkey 4

Let us consider again the game of Big Monkey and Little Monkey (Sec. 1.5) with Big Monkey going first. The normal form of this game is repeated below.

		Little Monkey			
		ww	**wc**	**cw**	**cc**
Big Monkey	**w**	$(0,0)$	$(0,0)$	$(9,1)$	$(9,1)$
	c	$(4,4)$	$(5,3)$	$(4,4)$	$(5,3)$

We have seen two ways of doing iterated elimination of weakly dominated strategies for this game.

One, which corresponds to a way of doing backward induction in the extensive form of the game, led to the 1×1 reduced game consisting of the strategy profile (w, cw). This strategy profile is therefore a dominant strategy equilibrium, and hence a Nash equilibrium. However, since it was found by iterated elimination of *weakly* dominated strategies, it is not guaranteed to be the *only* Nash equilibrium.

The second way we did iterated elimination of weakly dominated strategies led to a 2×1 reduced game. Both remaining strategy profiles (w, cw) (found before) and (w, cc) are Nash equilibria.

One can check that (c, ww) is also a Nash equilibrium.

Both of the Nash equilibria (w, cc) and (c, ww) use strategies that were eliminated in one way of doing iterated elimination of weakly dominated strategies. If one returns from the normal form of the game to the extensive form, both appear peculiar.

In the Nash equilibrium (c, ww), Little Monkey's strategy ww includes an incredible threat, the threat that if Big Monkey waits, he will wait also. When the time comes to carry out this threat, it would not be profitable to Little Monkey to do it.

In the Nash equilibrium (w, cc), Little Monkey's plan to climb if Big Monkey climbs is an example of a *promise*: it would help Big Monkey, at the cost of hurting Little Monkey. In this particular game, the promise does not affect Big Monkey's behavior. Little Monkey is promising that if Big Monkey climbs, he will get a payoff of 5, rather than the payoff of 4 he would normally expect. Big Monkey ignores this promise because by waiting, he gets an even bigger payoff, namely 9.

3.7. Finding Nash equilibria using best response

Consider a game in normal form with n players, strategy sets $S_1, \ldots, S_n$, and payoff functions $\pi_1, \ldots, \pi_n$. Let $s_2, \ldots, s_n$ be fixed strategies for players 2, ..., n. Suppose s_1^* is a strategy for player 1 with the property that

$$\pi_1(s_1^*, s_2, \ldots, s_n) \geq \pi_1(s_1, s_2, \ldots, s_n) \text{ for all } s_1 \in S_1. \tag{3.7}$$

Then s_1^* is a *best response* of player 1 to the strategy choices $s_2, \ldots, s_n$ of the other players. Of course, player 1 may have more than one such best response.

For each choices $s_2, \ldots, s_n$ of strategies by the other players, let $B_1(s_2, \ldots, s_n)$ denote the set best responses of player 1. In other words,

$$s_1^* \in B_1(s_2, \ldots, s_n) \text{ if and only if } \pi_1(s_1^*, s_2, \ldots, s_n) \geq \pi_1(s_1, s_2, \ldots, s_n) \text{ for all } s_1 \in S_1.$$

The mapping that associates to each $(s_2, \ldots, s_n) \in S_2 \times \cdots \times S_n$ the corresponding set $B_1(s_2, \ldots, s_n)$, a subset of S_1, is called player 1's *best response correspondence*. If each set $B_1(s_2, \ldots, s_n)$ consists of a single point, we have player 1's *best response function* $b_1(s_2, \ldots, s_n)$.

Best response correspondences for the other players are defined analogously.

At a Nash equilibrium, each player's strategy is a best response to the other players' strategies. In other words, the strategy profile $(s_1^*, \ldots, s_n^*)$ is a Nash equilibrium if and only if

- $s_1^* \in B_1(s_2^*, \ldots, s_n^*)$.
- $s_2^* \in B_2(s_1^*, s_3^* \ldots, s_n^*)$.

  $\vdots$
- $s_n^* \in B_n(s_1^*, \ldots, s_{n-1}^*)$.

This property of Nash equilibria can be used to find them. Just graph all players' best response correspondences in one copy of strategy profile space and find where they intersect! Alternatively, describe each best response correspondence by an equation, and solve the equations simultaneously.

3.8. Big Monkey and Little Monkey 5

Once again we consider the game of Big Monkey and Little Monkey (Sec. 1.5) with Big Monkey going first. The normal form of this game with both players' best response correspondences graphed is shown below.

		Little Monkey			
		ww	**wc**	**cw**	**cc**
Big Monkey	**w**	$(0,0)$	$(0,0)$	$(\boxed{9},\boxed{1})$	$(\boxed{9},\boxed{1})$
	c	$(\boxed{4},\boxed{4})$	$(\boxed{5},3)$	$(4,\boxed{4})$	$(5,3)$

Explanation:

- Big Monkey's best response correspondence is actually a function:
  - If Little Monkey does ww, do c.
  - If Little Monkey does wc, do c.
  - If Little Monkey does cw, do w.
  - If Little Monkey does cc, do w.

  This correspondence is indicated in the payoff matrix by drawing a box around the associated payoffs to Big Monkey. In other words, in each of the four columns of the matrix, I drew a box around Big Monkey's highest payoff.
- Little Monkey's best response correspondence is not a function:
  - If Big Monkey does w, do cw or cc.
  - If Big Monkey does c, do ww or cw.

  This correspondence is indicated in the payoff matrix by drawing a box around the associated payoffs to Little Monkey. In other words, in each of the two rows of the matrix, I drew a box around Little Monkey's highest payoffs.

Notice that three of the entries in the matrix have both payoffs boxed. These entries correspond to intersections of the graphs of the two players' best response correspondences, and hence to Nash equilibria.

For a two-player game in normal form where each player has only a finite number of strategies, graphing the best response correspondences as was done in this example is the best way to find Nash equilibria.

3.9. Water Pollution 2

The payoffs in the Water Polution game (Sec. 3.3) can be represented by two 2×2 matrices as follows.

Firm 3 purifies

		Firm 2 purify	pollute
Firm 1	**purify**	$(-1,-1,-1)$	$(-1,0,-1)$
	pollute	$(0,-1,-1)$	$(-3,-3,-4)$

Firm 3 pollutes

		Firm 2 purify	pollute
Firm 1	**purify**	$(-1,-1,0)$	$(-4,-3,-3)$
	pollute	$(-3,-4,-3)$	$(-3,-3,-3)$

These two matrices should be thought of as stacked one above the other. One then indicates

- player 1's best response to each choice of strategies by the other players by boxing the highest first entry in each column;
- player 2's best response to each choice of strategies by the other players by boxing the highest second entry in each row;
- player 3's best response to each choice of strategies by the other players by boxing the highest third entry in each "stack".

Firm 3 purifies

		Firm 2	
		purify	**pollute**
Firm 1	**purify**	$(-1,-1,-1)$	$(\boxed{-1},\boxed{0},\boxed{-1})$
	pollute	$(\boxed{0},\boxed{-1},\boxed{-1})$	$(-3,-3,-4)$

Firm 3 pollutes

		Firm 2	
		purify	**pollute**
Firm 1	**purify**	$(\boxed{-1},\boxed{-1},\boxed{0})$	$(-4,-3,-3)$
	pollute	$(-3,-4,-3)$	$(\boxed{-3},\boxed{-3},\boxed{-3})$

As before, we have found four Nash equilibria. In three of them, two firms purify and one pollutes. In the fourth, all firms pollute.

3.10. Cournot's model of duopoly

Cournot's model of duopoly (Wikipedia article: http://en.wikipedia.org/wiki/Cournot_duopoly) is the same as Stackelberg's, except that that the players choose their production levels simultaneously. This is a game in normal form with two players, strategy sets $0 \le s < \infty$ and $0 \le t < \infty$, and payoff functions

$$\pi_1(s,t) = p(s+t)s - cs = \begin{cases} (\alpha - \beta(s+t) - c)s & \text{if } s+t < \frac{\alpha}{\beta}, \\ -cs & \text{if } s+t \ge \frac{\alpha}{\beta}, \end{cases}$$

$$\pi_2(s,t) = p(s+t)t - ct = \begin{cases} (\alpha - \beta(s+t) - c)t & \text{if } s+t < \frac{\alpha}{\beta}, \\ -ct & \text{if } s+t \ge \frac{\alpha}{\beta}. \end{cases}$$

To calculate player 2's best response function, we must maximize $\pi_2(s,t)$, s fixed, on the interval $0 \le t < \infty$. This was done in Section 1.9; the answer is

$$t = b_2(s) = \begin{cases} \frac{\alpha-\beta s-c}{2\beta} & \text{if } s < \frac{\alpha-c}{\beta}, \\ 0 & \text{if } s \ge \frac{\alpha-c}{\beta}. \end{cases}$$

From the symmetry of the problem, player 1's best response function is

$$s = b_1(t) = \begin{cases} \frac{\alpha-\beta t-c}{2\beta} & \text{if } t < \frac{\alpha-c}{\beta}, \\ 0 & \text{if } t \ge \frac{\alpha-c}{\beta}. \end{cases}$$

See Figure 3.1. Notice that in order to make the figure analagous to the table in Section 3.8, which can be interpreted as the graph of a best response correspondence, we have made the s-axis, which represents player 1's strategies, the vertical axis.

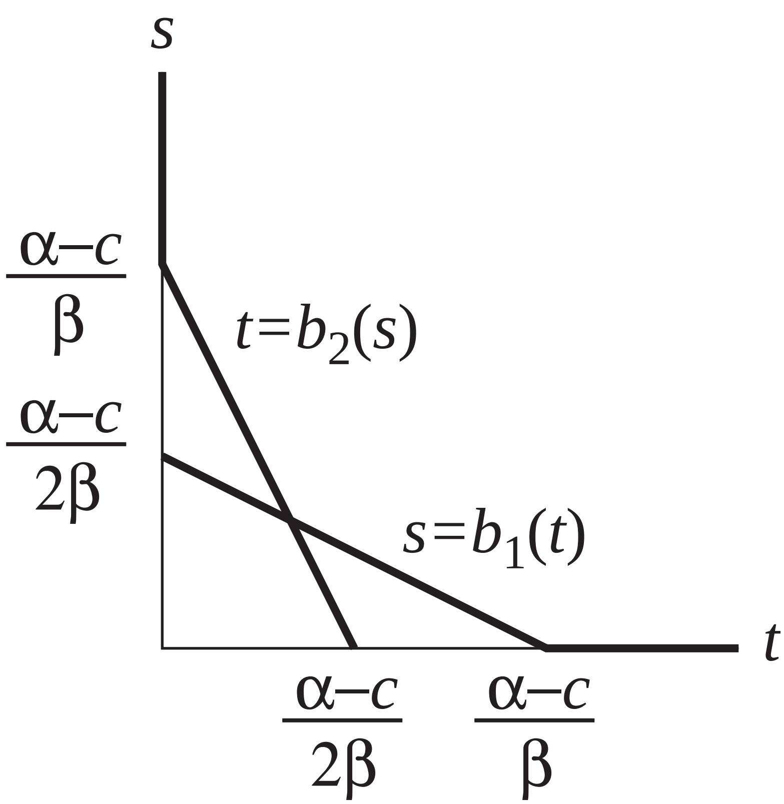

FIGURE 3.1. Best response functions in Cournot's model of duopoly.

There is a Nash equiibrium where the two best response curves intersect. From the figure, we see that to find this point we must solve simultaneously the two equations

$$t = \frac{\alpha - \beta s - c}{2\beta}, \quad s = \frac{\alpha - \beta t - c}{2\beta}.$$

We find that $s = t = \frac{\alpha - c}{3\beta}$.

CHAPTER 4

Games in Extensive Form with Incomplete Information

4.1. Lotteries

A *lottery* has n possible outcomes. The outcome depends on chance. The ith outcome occurs with probability p_i and yields a *payoff* x_i, which is a real number. We have all $p_i \geq 0$ and $p_1 + \ldots + p_n = 1$. The *expected value* of the lottery is

$$E[x] = p_1 x_1 + \cdots + p_n x_n.$$

If the recipient of the payoff has a utility function $u(x)$, its expected value is

$$E[u(x)] = p_1 u(x_1) + \cdots + p_n u(x_n).$$

The *expected utility principle* states that the lottery with the higher expected utility is preferred. For some discussion of the conditions under which this principle is true, see Chapter 2 of Gintis and the Wikipedia page
http://en.wikipedia.org/wiki/Expected_utility_hypothesis.

4.2. Buying fire insurance

You have a warehouse worth \$1.2 million. The probability of fire in any given year is 5%. Fire insurance costs \$100,000 per year. Should you buy it?

To answer the question, we compare two lotteries. Without fire insurance the

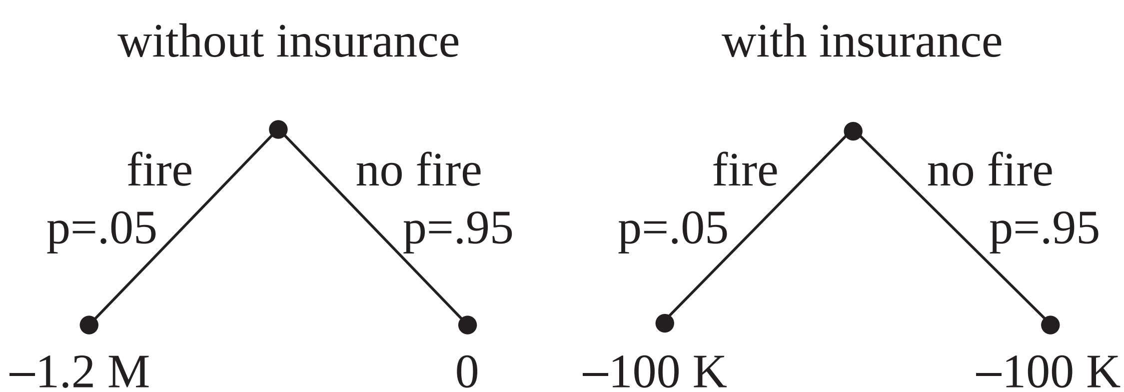

FIGURE 4.1. Should you buy the insurance?

expected outcome is

$$E[x] = .05 \times -1.2M + .95 \times 0 = -60K.$$

With fire insurance the expected outcome is

$$E[x] = .05 \times -100K + .95 \times -100K = -100K.$$

Don't buy the insurance.

However, people typically have concave utility functions. Suppose your utility function is $u(x) = \ln(1.3M + x)$. (This function is just $\ln x$ shifted $1.3M$ to the left, so it is continuous on the interval $-1.3M < x < \infty$.) Now, without fire insurance the expected outcome is

$$E[u(x)] = .05 \ln(.1M) + .95 \ln(1.3M) = 13.95.$$

With fire insurance, the expected outcome is

$$E[u(x)] = .05 \ln(1.2M) + .95 \ln(1.2M) = 14.00.$$

Buy the insurance.

4.3. Games in extensive form with incomplete information

This section is related to Gintis, Sec. 3.2.

In order to treat games in which players have incomplete information, we will add two ingredients to our allowed models of games in extensive form:

- Certain nodes may be assigned not to a player but to Nature. Nature's moves are chosen by chance. Therefore, if c is a node assigned to nature, each move that starts at c will be assigned a probability $0 \leq p \leq 1$, and these probabilities will sum to 1.
- The nodes assigned to a player may be partitioned into *information sets*. If several of a player's nodes are in the same information set, then the player does not know which of these nodes represents the true state of the game. The sets of available moves at the different nodes of an information set must be identical.

A player's strategy is required to assign the same move to every node of an information set.

A strategy profile determines a set of complete paths through the game tree, not necessarily a single complete path. Each of these paths is assigned a probability: probability 1 if the path does not include any of Nature's moves, and the product of the probabilities of Nature's moves along the path if it does. To determine the payoff to a player of a strategy profile, one sums the payoffs of the terminal nodes of the corresponding paths, each multiplied by the probability of the path.

In games in extensive form with incomplete information, a player may have information about Nature's moves that other players lack. For example, in a card game, Nature decides the hand you are dealt. You know it, but other players do not.

Backward induction often does not work in games in extensive form with incomplete information. The difficulty comes when you must decide on a move at an information set that includes more than one node. Usually the payoffs for the available moves depend on which node you are truly at, but you don't know that.

If the only nodes that precede the information set in question are Nature's, then you can calculate the probability that you are at each node, and use expected value to make a choice. In our treatment of the Cuban Missile Crisis later in this chapter, we will encounter information sets of this type. On the other hand, if, preceding the information set in question, there are nodes at which a player made a choice, then you do not know the probability that you are at each node. Games with information sets of this type can be treated by converting them to games in normal form. The game Buying a Used Car in the next section is an example.

4.4. Buying a Used Car

This section is related to Gintis, Sec. 3.21.

A customer is interested in a used car. He asks the salesman if it is worth the price. The salesman wants to sell the car. He also wants a reputation for telling the truth. How does the salesman respond? And should the customer believe his response?

We will assume that for cars of the type being sold, the probability the car is worth the price (i.e., is a good car) is p, so the probability it is not worth the price (i.e., is a bad car) is $1 - p$. These probabilities are known to both salesman and customer. In addition, the salesman knows whether this particular car is good or bad.

The payoffs are:

- The salesman gets 2 points if he sells the car and 1 point if he tells the truth.
- The customer gets 1 point if he correctly figures out whether the car is good or bad.

We model this situation as a game in extensive form with incomplete information. Nature moves first and decides if the car is good or bad. Then the salesman tells the customer whether it is good or bad. Then the customer decides whether it is good or bad, and on that basis decides whether to buy it. See Figure 4.2. Notice that the customer's moves are divided into two information sets, reflecting the fact

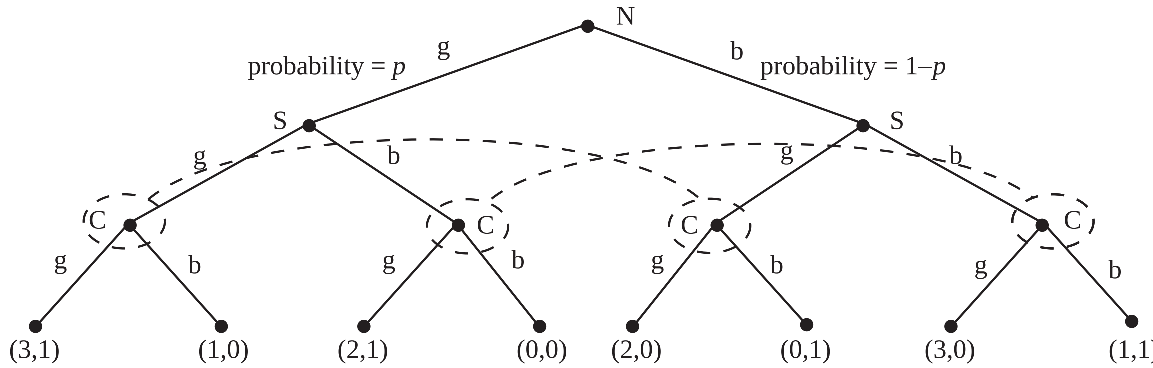

FIGURE 4.2. Buying a used car. The salesman's payoffs are given first.

that when the salesman says the car is good, the customer does not know whether it is good or not.

This game is an example of a game in extensive form with incomplete information in which backward induction cannot be used at all. For example, what should the customer do if the salesman says the car is good? The corresponding information set contains two nodes: at one of them, the car is actually good, but at the other it is bad. The customer's best move depends on which is the case, but he doesn't know which is the case. More importantly, he doesn't know the probabilities, because there are nodes preceding the nodes in this information set at which the salesman made a choice of what to say.

We will therefore analyze this game by converting it to a game in normal form.

The salesman has four strategies:

- *gg*: If the car is good, say it is good; if the car is bad, say it is good. (Always say the car is good.)
- *gb*: If the car is good, say it is good; if the car is bad, say it is bad. (Always tell the truth.)
- *bg*: If the car is good, say it is bad; if the car is bad, say it is good. (Always lie.)
- *bb*: If the car is good, say it is bad; if the car is bad, say it is bad. (Always say the car is bad.)

The customer must make the same move at two nodes that are in the same information set. Thus he only has four strategies:

- *gg*: If the salesman says the car is good, believe it is good; if the salesman says the car is bad, believe it is good. (Always believe the car is good.)
- *gb*: If the salesman says the car is good, believe it is good; if the salesman says the car is bad, believe it is bad. (Always believe the salesman.)

- *bg*: If the salesman says the car is good, believe it is bad; if the salesman says the car is bad, believe it is good. (Never believe the salesman.)
- *bb*: If the salesman says the car is good, believe it is bad; if the salesman says the car is bad, believe it is bad. (Always believe the car is bad.)

We will consider the salesman to be player 1 and the customer to be player 2.

In this game, each strategy profile is associated with *two* paths through the game tree. For example, consider the strategy profile (gg, bg): the salesman always says the car is good; and the customer never believes the salesman. This profile is associated with the two paths *ggb* and *bgb*:

- *ggb*: Nature decides the car is good, the salesman says the car is good, the customer does not believe him and decides it is bad (and hence does not buy the car): Payoffs: 1 to the salesman for telling the truth, 0 to the customer for miscalculating.
- *bgb*: Nature decides the car is bad, the salesman says the car is good, the customer does not believe him and decides it is bad (and hence does not buy the car): Payoffs: 0 to the salesman (he lied and still didn't sell the car), 1 to the customer (for correctly deciding the car was bad).

The first path through the tree is assigned probability p (the probability of a good car), the second is assigned probability $1 - p$ (the probability of a bad car). Thus the ordered pair of payoffs assigned to the strategy profile (gg, bg) is $p(1, 0) + (1 - p)(0, 1) = (p, 1 - p)$.

In this game, a good way to derive the payoff matrix is to separately write down the payoff matrix when Nature chooses a good car and when Nature chooses a bad car. The payoff matrix for the game is then p times the first matrix plus $1-p$ times the second.

If the car is good, the payoffs are:

		Customer			
		gg	**gb**	**bg**	**bb**
	gg	$(3, 1)$	$(3, 1)$	$(1, 0)$	$(1, 0)$
Salesman	**gb**	$(3, 1)$	$(3, 1)$	$(1, 0)$	$(1, 0)$
	bg	$(2, 1)$	$(0, 0)$	$(2, 1)$	$(0, 0)$
	bb	$(2, 1)$	$(0, 0)$	$(2, 1)$	$(0, 0)$

If the car is bad, the payoffs are:

		Customer			
		gg	**gb**	**bg**	**bb**
	gg	$(2,0)$	$(2,0)$	$(0,1)$	$(0,1)$
Salesman	**gb**	$(3,0)$	$(1,1)$	$(3,0)$	$(1,1)$
	bg	$(2,0)$	$(2,0)$	$(0,1)$	$(0,1)$
	bb	$(3,0)$	$(1,1)$	$(3,0)$	$(1,1)$

The payoff matrix for the game is p times the first matrix plus $1-p$ times the second:

		Customer			
		gg	**gb**	**bg**	**bb**
	gg	$(2+p,p)$	$(2+p,p)$	$(p,1-p)$	$(p,1-p)$
Salesman	**gb**	$(3,p)$	$(1+2p,1)$	$(3-2p,0)$	$(1,1-p)$
	bg	$(2,p)$	$(2-2p,0)$	$(2p,1)$	$(0,1-p)$
	bb	$(3-p,p)$	$(1-p,1-p)$	$(3-p,p)$	$(1-p,1-p)$

Let us assume $p > \frac{1}{2}$. We look for Nash equilibria by drawing a box around the highest first entry in each column and the highest second entry in each row.

		Customer			
		gg	**gb**	**bg**	**bb**
	gg	$(2+p,\boxed{p})$	$(\boxed{2+p},\boxed{p})$	$(p,1-p)$	$(p,1-p)$
Salesman	**gb**	$(\boxed{3},p)$	$(1+2p,\boxed{1})$	$(3-2p,0)$	$(\boxed{1},1-p)$
	bg	$(2,p)$	$(2-2p,0)$	$(2p,\boxed{1})$	$(0,1-p)$
	bb	$(3-p,\boxed{p})$	$(1-p,1-p)$	$(\boxed{3-p},\boxed{p})$	$(1-p,1-p)$

There are two Nash equilibria, (gg, gb) and (bb, bg). At the first equilibrium, the salesman always says the car is good, and the customer always believes the salesman. At the second, the salesman always says the car is bad, and the customer always assumes the salesman is lying. The two Nash equilibria give the same payoff to the customer, but the first gives a better payoff to the salesman.

4.5. Cuban Missile Crisis

You may want to compare my account of the Cuban Missile Crisis to the Wikipedia article: http://en.wikipedia.org/wiki/Cuban_missile_crisis.

In late summer of 1962, the Soviet Union began to place about 40 nuclear-armed medium- and intermediate-range ballistic missiles in Cuba. These missiles could target most of the eastern United States. The missile sites were guarded by surface-to-air missiles. There were also bombers.

A U.S. spy plane discovered the missiles on October 14, 1962. In the view of the U.S. government, the missiles posed several dangers: (1) they were a direct military threat to the U.S., and could perhaps be used to compel U.S. withdrawal from contested territories such as Berlin; (2) they promised to deter any possible U.S. attack against Cuba; and (3) their successful placement in Cuba would be seen by the world as a Soviet victory and a U.S. defeat.

President Kennedy and his associates at first considered an air strike against the missiles sites. Military leaders argued for a massive air strike against airfields and other targets as well. The civilian leaders decided this proposal was too risky and settled on a naval blockade.

The blockade went into effect October 24. Several apparently civilian freighters were allowed through with minimal inspection. Other questionable Soviet ships were heading toward Cuba, however, and the President and his associates feared that a confrontation with them could get out of hand. The Soviet premier, Khrushchev, indicated he might be willing to remove the missiles from Cuba if the U.S. removed its own missiles from Turkey. The U.S. had been planning to remove these missiles anyway and to substitute missiles on nuclear submarines, but did not want to appear to be giving in to pressure, or to make the Turks feel that the U.S. would not protect them.

On October 26 the U.S. discovered that the Soviets had also installed tactical nuclear weapons in Cuba that could be used against invading troops. On October 27, a U.S. spy plane was shot down over Cuba, and Cuban antiaircraft defenses fired on other U.S. aircraft. The U.S. Air Force commander, General Curtis Lemay, sent U.S. nuclear-armed bombers toward the Soviet Union, past their normal turnaround points.

On October 28 the crisis suddenly ended: Khrushchev announced that his missiles and other nuclear weapons in Cuba would be dismantled and brought home. Negotiations over the next month resulted in the withdrawal of the Soviet bombers as well, and, in a semisecret agreement, the removal of U.S. missiles from Turkey.

Today more is known about Soviet intentions in the crisis than was known to the U.S. government at the time. On the Soviet side, as on the U.S. side, there was considerable division over what to do. Apparently Khrushchev made the decision on his own to install missiles in Cuba; some of his advisors thought it reckless. He thought Kennedy would accept the missiles as a *fait accompli*, and he planned to issue an ultimatum to resolve the Berlin issue, using the missiles as a threat. The Soviet Presidium apparently decided as early as October 22 that it would back down rather than allow the crisis to lead to war. Two years later it removed Khrushchev from power; one if its main charges against him was the disastrous Cuban adventure.

On both sides it was not certain that decisions taken by leaders would be carried out as they intended. On the U.S. side, civilian leaders proposed methods of carrying out the blockade, but the U.S. Navy mostly followed its standard procedures. General LeMay, the inspiration for General Jack D. Ripper in the 1963 movie *Dr. Strangelove*, acted on his own in sending bombers toward the Soviet Union. He regarded the end of the crisis as a U.S. defeat: “We lost! We ought to just go in there today and knock ’em off.” On the Soviet side, the decision to shoot down a U.S. spy plane on October 27 was taken by the deputy to the Soviet general in charge while the general was away from his desk. The Cubans’ decision to fire on U.S. aircraft was taken by the Cuban president Castro over objections from the Soviet ambassador.

We will consider several models of the Cuban missile crisis beginning at the point where the missiles were discovered. A very simple model captures the essence of what happened: the U.S. can either accept the missiles or threaten war to remove them; if war is threatened, the Soviets can either defy the U.S., which would lead to war, or can back down and remove the missiles. Backward induction should tell us what the two parties will do.

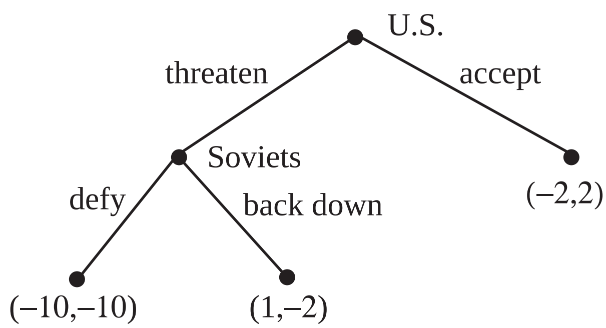

FIGURE 4.3. A simple model of the Cuban Missile Crisis Payoffs to U.S. are given first.

If the U.S. accepts the Soviet missiles in Cuba, we take the payoffs to be -2 to the U.S. and 2 to the Soviets. If the U.S. threatens war and the Soviets back down, we take the payoffs to be 1 to the U.S. and -2 to the Soviets. If the U.S. threatens war and the Soviets do not back down, we take the payoffs to be -10 to both the U.S. and the Soviets.

We conclude that if the U.S. threatens war, the Soviets will back down. Using backward induction, the U.S. decides to threaten war rather than accept the Soviet missiles. The Soviets then back down rather than go to war.

This is in fact what happened. On the other hand, the Soviets could do this analysis, too, so why did they place the missiles in Cuba to begin with?

The payoffs in Figure 4.3 assume a rather reasonable Soviet leadership. If the Soviet leadership regarded backing down in the face of a U.S. threat as totally unacceptable, and was less fearful of nuclear war, we get a situation like that in Figure 4.4.

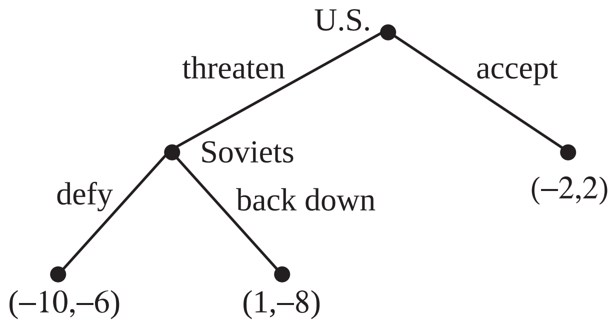

FIGURE 4.4. Cuban Missile Crisis with hard-line Soviets.

In this case, if the U.S. threatens war, the Soviets will not back down. Using backward induction, the U.S. will decide to accept the missiles.

In fact the U.S. government was not sure if the Soviets would turn out to be reasonable or hard-line. The evidence was conflicting; for example, an accommodating letter from Khrushchev on October 24 was followed by the shooting down of a U.S. plane on October 26. In addition, the U.S. understood at least in general that there were conflicting attitudes within the Soviet leadership. Figure 4.5 shows the situation if the U.S. is unsure whether the Soviet leadership is reasonable, as in Figure 4.3, or hard-line, as in Figure 4.4.

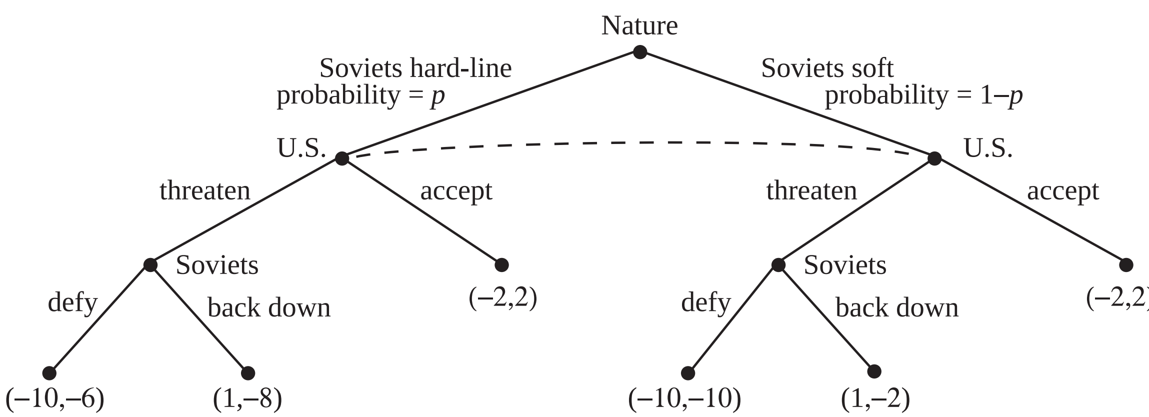

FIGURE 4.5. Cuban Missile Crisis with unknown Soviet leadership.

In the first move of the game, Nature decides, with certain probabilities, whether the Soviets are hard-line or reasonable. The Soviets know which they are, but the U.S. does not. Therefore, when the U.S. makes its move, which is the first move in the game by a player, both its nodes are in the same information set; this is indicated by a dashed line. The U.S. must make the same move (threaten war or accept the missiles) at both of these nodes. However, if the U.S. threatens war, the Soviets, knowing who they are, may reply differently in the two cases.

This game differs from Buying a Used Car in that the nodes toward the bottom of the tree are not in information sets containing more than one node. We can therefore at least start to analyze it by backward induction. In fact, the one information set that contains more than one move is preceded only by a move of Nature's. This kind of information set does not pose a difficulty for backward induction.

We must look first at the two nodes where the Soviets move, following a U.S. threat of war. (These are the only nodes that are followed only by terminal vertices.) If the Soviets are hard-line, they will choose to defy, with payoffs $(-10, -6)$. If the Soviets are reasonable, they will choose to back down, with payoffs $(1, -2)$.

Proceeding by backward induction, we look next at the two nodes where the U.S. moves, which are in the same information set. In this case, the probability that we are at each node in the information set is clear, so we can describe the payoffs of our choices using expected value.

If the U.S. threatens war, the payoffs are

$$p(-10, -6) + (1-p)(1, -2) = (1 - 11p, -2 - 4p).$$

If the U.S. accepts the Soviet missiles, the payoffs are

$$p(-2, 2) + (1-p)(-2, 2) = (-2, 2).$$

The U.S. will threaten war provided $1 - 11p > -2$, i.e., provided $p < \frac{3}{11}$. If $p > \frac{3}{11}$, the U.S. will accept the missiles.

President Kennedy apparently considered the probability of an unreasonable Soviet leadership to be somewhere between $\frac{1}{3}$ and $\frac{1}{2}$; even $\frac{1}{3}$ is greater than $\frac{3}{11}$. Now we have the opposite question: why didn't the U.S. accept the missiles?

In fact, the U.S. did not exactly threaten the Soviets with war if they did not remove their missiles. Instead it took actions, including a naval blockade, increased overflights of Cuba, and other military preparations, that increased the chance of war, even if war was not the U.S. intention. Both sides recognized that commanders on the scene might take actions the leadership did not intend, and events might spiral out of control. As we have seen, in the course of the crisis, dangerous decisions were in fact taken that leaders had trouble interpreting and bringing under control.

Brinkmanship (Wikipedia article: http://en.wikipedia.org/wiki/Brinkmanship) typically refers to the creation of a probabilistic danger in order to win a better outcome. On the one hand, it requires a great enough probability of disaster to persuade a reasonable opponent to concede rather than face the possibility that the dangerous situation will get out of hand. On the other hand, it requires a low probability of the nightmare outcome: the dangerous situation gets out of hand, *and* the opponent turns out to be a hard-line one who will not back down. The term dates to the 1950's. It is generally believed that the U.S. successfully practiced brinkmanship in the Cuban Missile Crisis. Figure 4.6 is a brinkmanship model of the Cuban Missile Crisis.

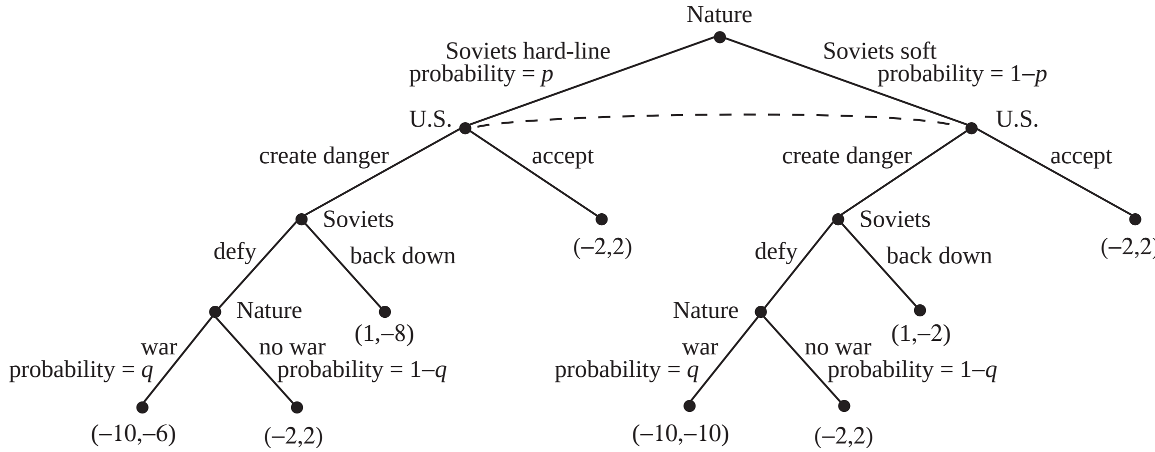

FIGURE 4.6. Cuban Missile Crisis with brinkmanship.

The change from Figure 4.5 is that, instead of threatening war, the U.S. now creates a dangerous situation by its military moves. The Soviets can defy or back down. If they defy, Nature decides whether there is war (with probability q) or no war (with probability $1-q$). If there is war, the payoffs are what they were in Figure 4.5 if the Soviets defied the U.S. threat. If there is no war, then the missiles remain; we take the payoffs to be the same as those when the U.S. accepts the Soviet missiles to begin with.

The game in Figure 4.6 can be analyzed by backward induction. Nature's last move results in the following payoffs if the Soviets defy the U.S.

- If Soviets are hard-line: $q(-10,-6)+(1-q)(-2,2)=(-2-8q,2-8q)$.
- If Soviets are reasonable: $q(-10,-10)+(1-q)(-2,2)=(-2-8q,2-12q)$.

Backing up one step, we find:

- If the Soviets are hard-line and the U.S. creates a dangerous situation, the Soviets will defy if $2-8q>-8$. Since this inequality is true for all q between 0 and 1, the Soviets will certainly defy.
- If the Soviets are reasonable and the U.S. creates a dangerous situation, the Soviets will defy if $2-12q>-2$, i.e., if $q<\frac{1}{3}$. If $q>\frac{1}{3}$, the Soviets will back down.

Now we back up one more step and ask whether the U.S. should create a dangerous situation or accept the Soviet missiles. If the U.S. accepts the Soviet missiles, the payoffs are of course $p(-2,2)+(1-p)(-2,2)=(-2,2)$. If the U.S. creates a dangerous situation, the payoffs depend on the probability of war q associated with the situation that the U.S. creates. The payoffs are:

- If $q<\frac{1}{3}$: $p(-2-8q,2-8q)+(1-p)(-2-8q,2-12q)=(-2-8q,2-12q+4pq)$.

- If $q > \frac{1}{3}$: $p(-2-8q, 2-8q) + (1-p)(1,-2) = (1-3p-8pq, -2+4p-8pq)$.

Let's consider both cases.

- The payoff to the U.S. from creating a dangerous situation with $q < \frac{1}{3}$ is $-2-8q$. The U.S. is better off simply accepting with missiles, which yields a payoff of -2. The probability of war is too low to induce even reasonable Soviets to back down; making such a threat increases the danger to the U.S. without any offsetting benefit.
- The payoff to the U.S. from creating a dangerous situation with $q > \frac{1}{3}$ is $1-3p-8pq$. The U.S. benefits from creating such a situation if $1-3p-8pq > -2$, i.e., if

$$q < \frac{3(1-p)}{8p}.$$

Therefore, if $\frac{1}{3} < \frac{3(1-p)}{8p}$, the U.S. can benefit by creating a dangerous situation in which the probability of war q is any number between $\frac{1}{3}$ and $\frac{3(1-p)}{8p}$.

Figure 4.7 helps in interpreting this result. If $0 < p < \frac{3}{11}$, any q between $\frac{1}{3}$ and 1 gives the U.S. a better result than accepting the missiles. Of course, we already knew that for p in this range, a simple threat of war (equivalent to $q = 1$) would give the U.S. a better result than accepting the missiles. More interesting is the interval $\frac{3}{11} < p < \frac{9}{17}$, which includes the U.S. government's guess as to the true value of p. For each p in this interval there is a corresponding interval $\frac{1}{3} < q < \frac{3(1-p)}{8p}$ that gives the U.S. gets a better result than accepting the missiles.

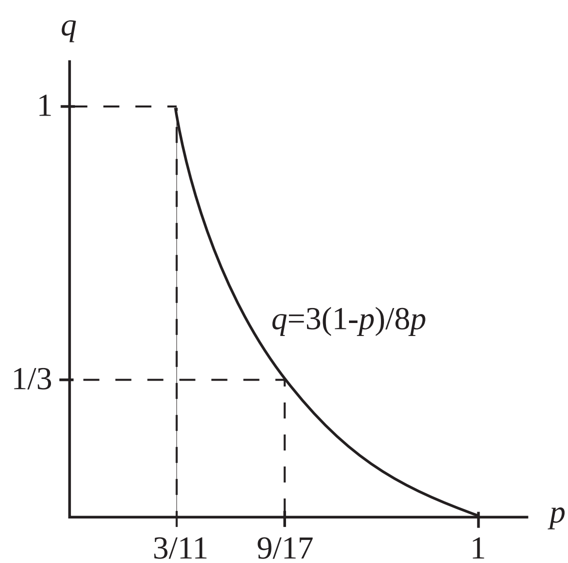

FIGURE 4.7. Where brinkmanship is helpful.

CHAPTER 5

Mixed-Strategy Nash Equilibria

5.1. Mixed-strategy Nash equilibria

This section is related to Gintis, Secs. 3.4–3.5 and 6.4.

Even simple games can fail to have Nash equilibria in the sense we have so far discussed.

Let us consider, for example, two people playing tennis (Gintis, Sec. 6.4). In tennis, the player serving can serve to his opponent's backhand or forehand. The player receiving the serve can anticipate a serve to his backhand or a serve to his forehand.

For a certain pair of tennis players, the probability that the serve is returned is given by the following table.

		Receiver anticipates serve to	
		backhand	**forehand**
Server serves to	**backhand**	.6	.2
	forehand	.3	.9

(These numbers are consistent with Gintis's verbal description; his table has mistakes.) We regard this as a game in normal form. The payoff to the receiver is the fraction of serves he returns; the payoff to the server is the fraction of serves that are not returned. Thus the payoff matrix is

		Receiver anticipates serve to	
		backhand	**forehand**
Server serves to	**backhand**	$(.4, .6)$	$(.8, .2)$
	forehand	$(.7, .3)$	$(.1, .9)$

You can easily check that there are no Nash equilibria in the sense we have so far discussed.

It is easy to understand why this game has no equilibria. If I plan serve to your forehand, your best response is to anticipate a serve to your forehand. But if you anticipate a serve to your forehand, my best response is to serve to your backhand. But if I plan serve to your backhand, your best response is to anticipate a serve to your backhand. But if you anticipate a serve to your backhand, my best response is to serve to your forehand. We are back where we started from, without having found an equilibrium!

Does game theory have any suggestions for these players? Yes: the suggestion to the server is to mix his serves randomly, and the suggestion to the receiver is to mix his expectations randomly. But what fraction of the time should the server serve to the forehand, and what fraction of the time should the receiver anticipate a serve to his forehand?

To answer this question we must develop the idea of a mixed strategy.

Consider a game in normal form with players $1, \dots, n$ and corresponding strategy sets $S_1, \dots, S_n$, all finite. Suppose that for each i, player i's strategy set consists of k_i strategies, which we denote $s_{i1}, \dots, s_{ik_i}$. A *mixed strategy* σ_i for player i consists of using strategy s_{i1} with probability p_{i1}, strategy s_{i2} with probability p_{i2}, ..., strategy s_{ik_i} with probability p_{ik_i}. Of course, each $p_{ij} \geq 0$, and $\sum_{j=1}^{k_i} p_{ij} = 1$. Formally, $\sigma_i = \sum_{j=1}^{k_i} p_{ij} s_{ij}$.

If $p_{ij} > 0$, we say that the pure strategy s_{ij} is *active* in the mixed strategy σ_i.

A mixed strategy σ_i is called *pure* if only one pure strategy is active, i.e., if one p_{ij} is 1 and all the rest are 0. We will usually denote by s_{ij} the pure strategy of player i that uses s_{ij} with probability 1 and his other strategies with probability 0. Up until now we have only discussed pure strategies.

We will try whenever possible to avoid double subscripting. Thus we will often denote a strategy of player i by s_i, and the associated probability by p_{s_i}. This will require summing over all $s_i \in S_i$ instead of summing from $j = 1$ to k_i. Thus a mixed strategy of player i will be written

$$\sigma_i = \sum_{\text{all } s_i \in S_i} p_{s_i} s_i. \tag{5.1}$$

If each player i chooses a mixed strategy σ_i, we get a *mixed strategy profile* $(\sigma_1, \dots, \sigma_n)$.

Recall that if each player i chooses a pure strategy s_i, we get a pure strategy profile $(s_1, \dots, s_n)$. Recall that associated with each pure strategy profile $(s_1, \dots, s_n)$ is a payoff to each player; the payoff to player i is denoted $\pi_i(s_1, \dots, s_n)$.

Suppose

- player 1's mixed strategy σ_1 uses his strategy s_1 with probability p_{s_1},
- player 2's mixed strategy σ_2 uses his strategy s_2 with probability p_{s_2},

  ⋮
- player n's mixed strategy σ_n uses his strategy s_n with probability p_{s_n}.

We assume that the players independently choose pure strategies to use. Then, if the players use the mixed strategy profile $(\sigma_1, \dots, \sigma_n)$, the probability that the pure strategy profile $(s_1, \dots, s_n)$ occurs is $p_{s_1} p_{s_2} \cdots p_{s_n}$. Thus the expected payoff to player i is

$$\pi_i(\sigma_1, \dots, \sigma_n) = \sum_{\text{all } (s_1,\dots,s_n)} p_{s_1} p_{s_2} \dots p_{s_n} \pi_i(s_1, \dots, s_n). \tag{5.2}$$

Let σ denote the mixed strategy profile $(\sigma_1, \dots, \sigma_n)$. Suppose in σ we replace the ith player's mixed strategy σ_i by another of his mixed strategies, say τ_i. We will denote the resulting mixed strategy profile by (τ_i, σ_{-i}). This notation is analogous to that introduced in Subsection 3.1.

If player 1 uses a pure strategy s_1, equation (5.2) simplifies to

$$\pi_i(s_1, \sigma_2 \dots, \sigma_n) = \sum_{\text{all } (s_2,\dots,s_n)} p_{s_2} \dots p_{s_n} \pi_i(s_1, \dots, s_n). \tag{5.3}$$

Equation (5.2) can be rewritten as

$$\pi_i(\sigma_1, \dots, \sigma_n) = \sum_{\text{all } s_1 \in S_1} p_{s_1} \sum_{\text{all } (s_2,\dots,s_n)} p_{s_2} \dots p_{s_n} \pi_i(s_1, \dots, s_n). \tag{5.4}$$

Using (5.3), equation (5.4) can be written

$$\pi_i(\sigma_1, \dots, \sigma_n) = \sum_{\text{all } s_1 \in S_1} p_{s_1} \pi_i(s_1, \sigma_2 \dots, \sigma_n) = \sum_{\text{all } s_1 \in S_1} p_{s_1} \pi_i(s_1, \sigma_{-1}). \tag{5.5}$$

More generally, for any players i and j,

$$\pi_i(\sigma_1, \dots, \sigma_n) = \sum_{\text{all } s_j \in S_j} p_{s_j} \pi_i(s_j, \sigma_{-j}). \tag{5.6}$$

In other words, think of all players' mixed strategies as fixed except player j's. If player j uses one of his pure strategies s_j, player i gets the payoff $\pi_i(s_j, \sigma_{-j})$. If player j uses a mixed strategy, player i's payoff is a weighted average of the payoffs $\pi_i(s_j, \sigma_{-j})$, where the weights are the the probabilities in player j's strategy σ_j.

In particular, for any player i,

$$\pi_i(\sigma_1, \dots, \sigma_n) = \sum_{\text{all } s_i \in S_i} p_{s_i} \pi_i(s_i, \sigma_{-i}). \tag{5.7}$$

In other words, the payoff to player i from using strategy σ_i against the other players' mixed strategies is just a weighted average of his payoffs from using his pure strategies against their mixed strategies, where the weights are the probabilities in his strategy σ_i.

A mixed strategy profile $(\sigma_1^*, \dots, \sigma_n^*)$ is a *mixed strategy Nash equilibrium* if no single player can improve his own payoff by changing his strategy. In other words:

- For every mixed strategy σ_1 of player 1,
$$\pi_1(\sigma_1^*, \sigma_2^*, \ldots, \sigma_n^*) \geq \pi_1(\sigma_1, \sigma_2^*, \ldots, \sigma_n^*).$$
- For every mixed strategy σ_2 of player 2,
$$\pi_2(\sigma_1^*, \sigma_2^*, \sigma_3^*, \ldots, \sigma_n^*) \geq \pi_2(\sigma_1^*, \sigma_2, \sigma_3^*, \ldots, \sigma_n^*).$$

  $\vdots$
- For every mixed strategy σ_n of player n,
$$\pi_n(\sigma_1^*, \ldots, \sigma_{n-1}^*, \sigma_n^*) \geq \pi_n(\sigma_1^*, \ldots, \sigma_{n-1}^*, \sigma_n).$$

More compactly, a mixed strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ is a Nash equilibrium if, for each $i = 1, \ldots, n$, $\pi_i(\sigma^*) \geq \pi_i(\sigma_i, \sigma_{-i}^*)$ for every mixed strategy σ_i of player i.

THEOREM 5.1. *Nash's Existence Theorem. If, in an n-person game in normal form, each player's strategy set is finite, then the game has at least one mixed strategy Nash equilibrium.*

John Nash was awarded the Nobel Prize in Economics in 1994 largely for inventing the notion of Nash equilibrium and discovering this theorem.

The next result gives a characterization of Nash equilibria that is often useful in finding them.

THEOREM 5.2. *Fundamental Theorem of Nash Equilibria. The mixed strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a mixed strategy Nash equilibrium if and only if the following two conditions are satisfied for every $i = 1, \ldots, n$.*

(1) *If the strategies s_i and s_i' are both active in σ_i, then $\pi_i(s_i, \sigma_{-i}) = \pi_i(s_i', \sigma_{-i})$.*
(2) *If the strategy s_i is active in σ_i and the strategy s_i' is not active in σ_i, then $\pi_i(s_i, \sigma_{-i}) \geq \pi_i(s_i', \sigma_{-i})$.*

This theorem is very easy to prove. We will just give the idea of how to prove (1). Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a mixed strategy profile. Suppose that in player i's strategy σ_i, his pure strategies s_i and s_i' are both active, with probabilities $p_i > 0$ and $p_i' > 0$ respectively. Look at equation (5.7). Suppose, for example, that $\pi_i(s_i, \sigma_{-i}) < \pi_i(s_i', \sigma_{-i})$. Then player i can switch to a new strategy τ_i that differs from σ_i only in that the pure strategy s_i is not used at all, but the pure strategy s_i' is used with probability $p_i + p_i'$. This would increase player i's payoff, so $\sigma = (\sigma_1, \ldots, \sigma_n)$ would not be a mixed strategy Nash equilibrium.

Here are two easy consequences of the Fundamental Theorem of Nash Equilibria.

THEOREM 5.3. (1) *If $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ is a mixed strategy Nash equilibrium and the strategy s_i is active in σ_i^*, then $\pi_i(s_i, \sigma_{-i}^*) = \pi_i(\sigma^*)$.*

(2) *If $s^* = (s_1^*, \dots, s_n^*)$ is a profile of pure strategies, then s^* is a mixed strategy Nash equilibrium if and only if s^* is a Nash equilibrium in the sense of Subsection 3.1.*

Proof. (1) Suppose $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a mixed strategy Nash equilibrium and σ_i^* is given by (5.1). Let $I = \{i : p_i > 0\}$. According to the Fundamental Theorem, all the payoffs $\pi_i(s_i, \sigma_{-i}^*)$ with $i \in I$ are equal. Let's say they all equal P. Then from (5.7),

$$\pi_i(\sigma_1^*, \dots, \sigma_n^*) = \sum_{\text{all } s_i \in S_i} p_{s_i} \pi_i(s_i, \sigma_{-i}^*) = \sum_{i \in I} p_{s_i} \pi_i(s_i, \sigma_{-i}^*) = \sum_{i \in I} p_{s_i} P = P.$$

(2) Suppose $s^* = (s_1^*, \dots, s_n^*)$ is a profile of pure strategies that is a mixed strategy Nash equilibrium. Consider player i and one of his strategies s_i other than s_i^*. Since s_i is not active, the second part of the Fundamental Theorem says that $\pi_i(s_i^*, s_{-i}^*) \geq \pi_i(s_i, s_{-i}^*)$. This says that s^* is a Nash equilibrium in the sense of Subsection 3.1.

On the other hand, suppose $s^* = (s_1^*, \dots, s_n^*)$ is a profile of pure strategies that is a Nash equilibrium in the sense of Subsection 3.1. Then condition (1) of the Fundamental Theorem is automatically satisfied since each player has only one active strategy. Condition (2) is an immediate consequence of the definition of Nash equilibrium in Subsection 3.1. Since both conditions hold, s^* is a mixed strategy Nash equilibrium. □

A mixed strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a *strict mixed strategy Nash equilibrium* if, for each $i = 1, \dots, n$, $\pi_i(\sigma^*) > \pi_i(\sigma_i, \sigma_{-i}^*)$ for every mixed strategy $\sigma_i \neq \sigma_i^*$ of player i.

Here is an easy consequence of Theorem 5.3.

Theorem 5.4. *$\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a strict mixed strategy Nash equilibrium if and only if (i) each σ_i^* is a pure strategy, and (ii) σ^* is a strict Nash equilibrium in the sense of Subsection 3.1.*

For most two-person games, at each Nash equilibrium, both players use the same number of active pure strategies. Thus, in two-person games, one can begin by looking for Nash equilibria in which each player uses one active pure strategy, then Nash equilibria in which each player uses two active pure strategies, etc. In most two-person games, this procedure not only finds all mixed strategy Nash equilibria; it also yields as a byproduct a proof that there are no Nash equilibria in which the two players use different numbers of active pure strategies.

We will now look at some examples to see how this works.

5.2. Tennis

This section is related to Gintis, Secs. 6.4.

We recall the game of tennis described in the previous section, with the payoff matrix

		Receiver anticipates serve to	
		backhand	**forehand**
Server serves to	**backhand**	$(.4, .6)$	$(.8, .2)$
	forehand	$(.7, .3)$	$(.1, .9)$

We shall find all mixed strategy Nash equilibria in this game.

Suppose the receiver uses his two strategies with probabilities p and $1-p$, and the server uses his strategies with probabilities q and $1-q$. It is helpful to write these probabilities next to the payoff matrix as follows:

			Receiver anticipates serve to	
			q	$1-q$
			backhand	**forehand**
Server serves to	p	**backhand**	$(.4, .6)$	$(.8, .2)$
	$1-p$	**forehand**	$(.7, .3)$	$(.1, .9)$

We shall look for a mixed strategy Nash equilibria $(pb + (1-p)f, qb + (1-q)f)$.

In accordance with the advice in previous section, we shall first look for equilibria in which both players use one active pure strategy, then look for equilibria in which both players use two active pure strategies.

1. Suppose both players use one active pure strategy. Then we would have a pure strategy Nash equilibrium. You checked in the previous section that there are none.

2. Suppose both players use two active pure strategies. Then $0 < p < 1$ and $0 < q < 1$. Since both of player 2's pure strategies b and f are active, each gives the same payoff to player 2 against player 1's mixed strategy $pb + (1-p)f$:

$$.6p + .3(1-p) = .2p + .9(1-p).$$

Solving this equation for p, we find that $p = .6$.

Similarly, since both of player 1's pure strategies b and f are active, each gives the same payoff to player 1 against player 2's mixed strategy $qb + (1-q)f$:

$$.4q + .8(1-q) = .7q + .1(1-q).$$

Solving this equation for q, we find that $q = .7$.

We conclude that the $(.6b + .4f, .7b + .3f)$ satisfies the equality criterion for a mixed strategy Nash equilibrium. Since there are no unused pure strategies, there is no inequality criterion to check. Therefore we have found a mixed strategy Nash equilibrium in which both players have two active pure strategies.

Note that in the course of finding this Nash equilibrium, we actually did more.

(1) We showed that if both of player 2's pure strategies are active at a Nash equilibrium, then player 1's strategy must be $.6b + .4f$. Hence there are no Nash equilibria in which player 2 uses two pure strategies but player 1 uses only one pure strategy.
(2) Similarly, we showed that if both of player 1's pure strategies are active at a Nash equilibrium, then player 2's strategy must be $7b + .3f$. Hence there are no Nash equilibria in which player 1 uses two pure strategies but player 2 uses only one pure strategy.

This is an example of how, in the course of finding mixed strategy Nash equilibria in which both players use the same number of pure strategies, one usually shows as a byproduct that there are no Nash equilibria in which the two players use different numbers of pure strategies.

5.3. Other ways to find mixed strategy Nash equilibria

Here are two ways to find mixed strategy Nash equilibria in the previous problem without using the Fundamental Theorem. Both may be useful in other problems.

5.3.1. Differentiating the payoff functions. In the tennis problem, there are two payoff functions, π_1 and π_2. Since player 1's strategy is determined by the choice of p and player 2's by the choice of q, we may regard π_1 and π_2 as functions of (p, q) with $0 \le p \le 1$ and $0 \le q \le 1$. From the payoff matrix, we have

$$\pi_1(p, q) = .4pq + .8p(1 - q) + .7(1 - p)q + .1(1 - p)(1 - q),$$
$$\pi_2(p, q) = .6pq + .2p(1 - q) + .3(1 - p)q + .9(1 - p)(1 - q).$$

Suppose (p, q) is a mixed strategy Nash equilibrium with $0 < p < 1$ and $0 < q < 1$. Then the definition of mixed strategy Nash equilibrium implies that

$$\frac{\partial \pi_1}{\partial p}(p, q) = 0 \text{ and } \frac{\partial \pi_2}{\partial q}(p, q) = 0.$$

Therefore

$$\frac{\partial \pi_1}{\partial p}(p, q) = .4q + .8(1 - q) - .7q - .1(1 - q) = .7 - q = 0,$$
$$\frac{\partial \pi_2}{\partial q}(p, q) = .6p - .2p + .3(1 - p) - .9(1 - p) = . - .6 + p = 0.$$

We see that $(p, q) = (.6, .7)$.

5.3.2. Best response correspondences. From the calculation of partial derivatives above,

$$\frac{\partial \pi_1}{\partial p}(p,q) = \begin{cases} + & \text{if } q < .7, \\ 0 & \text{if } q = .7, \\ - & \text{if } q > .7. \end{cases} \qquad \frac{\partial \pi_2}{\partial q}(p,q) = \begin{cases} - & \text{if } p < .6, \\ 0 & \text{if } p = .6, \\ + & \text{if } p > .6, \end{cases}$$

These partial derivatives tell us each player's best response to all strategies of his opponent. For player 1:

- If player 2 chooses q with $0 \le q < .7$, player 1 observes that his own payoff is an increasing function of p. Hence his best response is $p = 1$.
- If player 2 chooses $q = .7$, player 1 observes that his own payoff will be the same whatever p he chooses. Hence he can choose any p between 0 and 1.
- If player 2 chooses q with $.7 < q \le 1$, player 1 observes that his own payoff is a decreasing function of p. Hence his best response is $p = 0$.

Player 1's best response correspondence $B_1(q)$ is graphed below, along with player 2's best response correspondence $B_2(p)$. Note that $B_1(.7)$ is the set $0 \le p \le 1$ and $B_2(.6)$ is the set $0 \le q \le 1$.

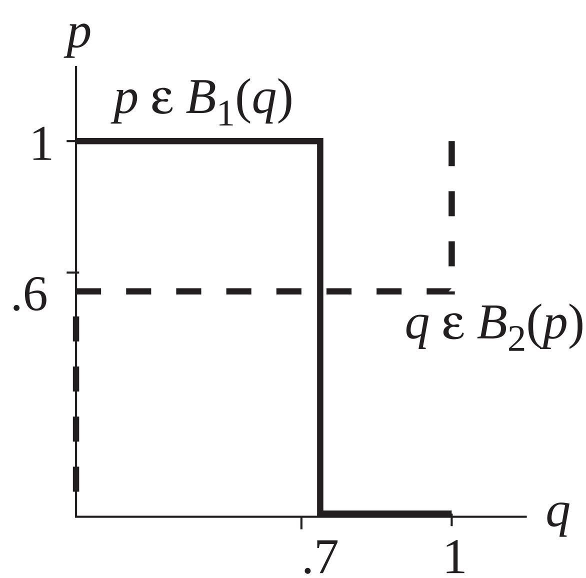

Figure 5.1. Graphs of best response correspondences in the game of tennis. The only point in the intersection of the two graphs is $(p,q) = (.6,.7)$.

5.4. One-card Two-round Poker

This section is related to Gintis, Sec. 6.21.

We will play poker with a deck of two cards, one high (H) and one low (L). There are two players. Play proceeds as follows.

(1) Each player puts \$2 into the pot.
(2) Player 1 is dealt one card, chosen by chance. He looks at it. He either bets \$2 or he folds. If he folds, player 2 gets the pot. If he bets:
(3) Player 2 either bets \$2 or he folds. If he folds, player 1 gets the pot. If he bets:
(4) Player 1 either bets \$2 or he folds. If he folds, player 2 gets the pot. If he bets:
(5) Player 2 either bets \$2 or he folds. If he folds, player 1 gets the pot. If he bets, player 1 shows his card. If it is H, player 1 wins the pot. If it is L, player 2 wins the pot.

The game tree is shown below.

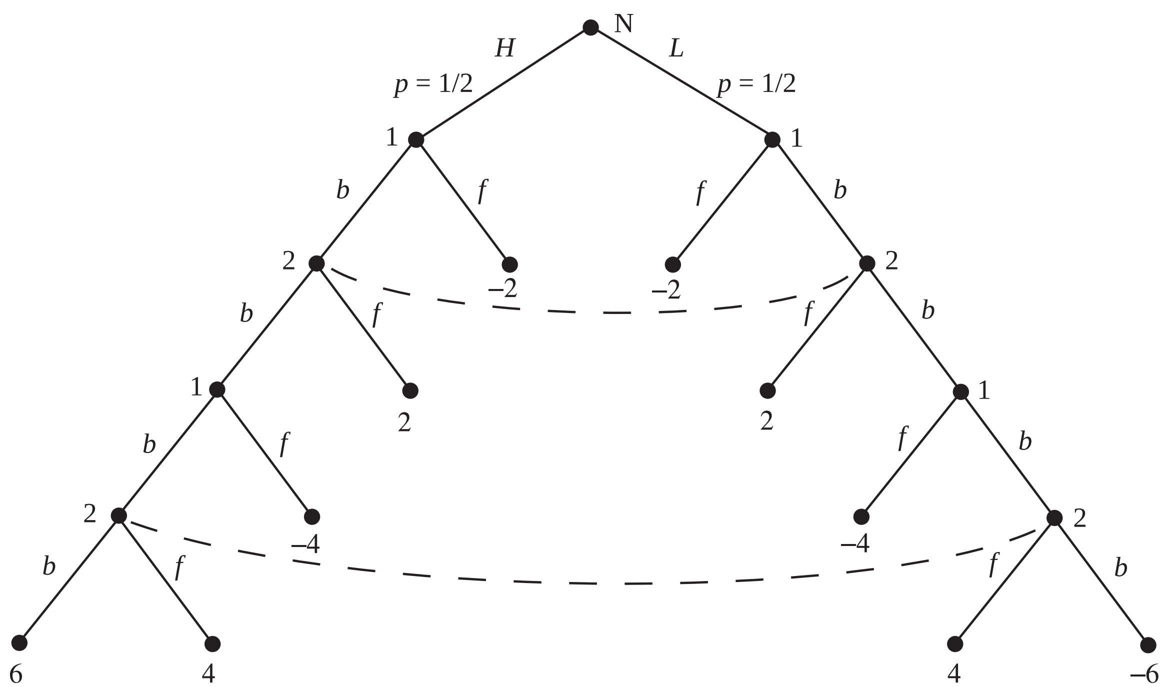

FIGURE 5.2. One-card two-round poker. Nodes in the same information set are linked by a dashed line. Only payoffs to player 1 are shown; payoffs to player 2 are opposite.

Player 2 has three pure strategies:

- *bb*: the first time player 1 bets, respond by betting; the second time player 1 bets, respond by betting.
- *bf*: the first time player 1 bets, respond by betting; the second time player 1 bets, respond by folding.
- *f*: the first time player 1 bets, respond by folding.

To describe player 1's pure strategies, we first note that if player 1 is dealt the high card, he has three options:

- *bb*: bet; if player 2 responds by betting, bet again.
- *bf*: bet; if player 2 responds by betting, fold.
- *f*: fold.

If player 1 is dealt the low card, he has the same three options. Thus player 1 has nine pure strategies: choose one option to use if dealt the high card, and one to option to use if dealt the low card.

The payoff matrix for this game is 9×3. If we draw it, we will quickly see that six of player 1's pure strategies are weakly dominated: every strategy of player 1 that does not use the option *bb* when dealt the high card is weakly dominated by the corresponding strategy that does use this option. This is obviously correct: if player 1 is dealt the high card, he will certainly gain a positive payoff if he continues to bet, and will certainly suffer a negative payoff if he ever folds.

We therefore eliminate six of player 1's strategies and obtain a reduced 3×3 game. In the reduced game, we denote player 1's strategies by *bb*, *bf*, and *f*. The notation represents the option player 1 uses if dealt the low card; if dealt the high card, he uses the option *bb*.

If the card dealt is high, the payoffs are:

		Player 2		
		bb	**bf**	**f**
	bb	$(6, -6)$	$(4, -4)$	$(2, -2)$
Player 1	**bf**	$(6, -6)$	$(4, -4)$	$(2, -2)$
	f	$(6, -6)$	$(4, -4)$	$(2, -2)$

If the card dealt is low, the payoffs are:

		Player 2		
		bb	**bf**	**f**
	bb	$(-6, 6)$	$(4, -4)$	$(2, -2)$
Player 1	**bf**	$(-4, 4)$	$(-4, 4)$	$(2, -2)$
	f	$(-2, 2)$	$(-2, 2)$	$(-2, 2)$

The payoff matrix for the game is $\frac{1}{2}$ times the first matrix plus $\frac{1}{2}$ times the second:

		Player 2		
		bb	**bf**	**f**
	bb	$(0,0)$	$(4,-4)$	$(2,-2)$
Player 1	**bf**	$(1,-1)$	$(0,0)$	$(2,-2)$
	f	$(2,-2)$	$(1,-1)$	$(0,0)$

We shall look for a mixed strategy Nash equilibrium (σ_1, σ_2), with $\sigma_1 = p_1 bb + p_2 bf + p_3 f$ and $\sigma_2 = q_1 bb + q_2 bf + q_3 f$:

			Player 2		
			q_1	q_2	q_3
			bb	**bf**	**f**
	p_1	**bb**	$(0,0)$	$(4,-4)$	$(2,-2)$
Player 1	p_2	**bf**	$(1,-1)$	$(0,0)$	$(2,-2)$
	p_3	**f**	$(2,-2)$	$(1,-1)$	$(0,0)$

We should consider three possibilities:

(1) Both players use a pure strategy.
(2) Both players use exactly two active pure strategies.
(3) Both players use exactly three active pure strategy.

One easily deals with the first case: there are no pure strategy Nash equilibria. The second case divides into nine subcases, since each player has three ways to choose his two active strategies. For now we ignore these possibilities; we will return to them in the next section.

In the third case, we assume that all p_i and q_i are positive. Since all q_i are positive, each of player 2's pure strategies gives the same payoff to player 2 against player 1's mixed strategy σ_1. These three payoffs are:

$$\begin{aligned}
\pi_2(\sigma_1, bb) &= p_1(0) + p_2(-1) + p_3(-2),\\
\pi_2(\sigma_1, bf) &= p_1(-4) + p_2(0) + p_3(-1),\\
\pi_2(\sigma_1, f) &= p_1(-2) + p_2(-2) + p_3(0).
\end{aligned}$$

The fact that these three quantities must be equal yields two independent equations. For example, one can use $\pi_2(\sigma_1, bb) = \pi_2(\sigma_1, f)$ and $\pi_2(\sigma_1, bf) = \pi_2(\sigma_1, f)$:

$$\begin{aligned}
p_1(0) + p_2(-1) + p_3(-2) &= p_1(-2) + p_2(-2) + p_3(0),\\
p_1(-4) + p_2(0) + p_3(-1) &= p_1(-2) + p_2(-2) + p_3(0).
\end{aligned}$$

Simplifying, we have

$$\begin{aligned}
2p_1 + p_2 - 2p_3 &= 0,\\
-2p_1 + 2p_2 - p_3 &= 0.
\end{aligned}$$

A third equation is given by

$$p_1 + p_2 + p_3 = 1.$$

These three equations in the three unknowns (p_1, p_2, p_3) can be solved to yield the solution

$$(p_1, p_2, p_3) = (\frac{1}{5}, \frac{2}{5}, \frac{2}{5}).$$

Had any p_i failed to lie strictly between 0 and 1, we would have had to discard the assumption that there is a Nash equilibrium in which both players use three active pure strategies.

(One way to use the third equation is to use it to substitute $p_3 = 1 - p_1 - p_2$ in the first two equations. This is analogous to how we solved the tennis problem.)

Similarly, since all p_i are positive, each of player 1's three pure strategies gives the same payoff to player 1 against player 2's mixed strategy σ_2. This observation leads to three equations in the three unknowns (q_1, q_2, q_3), which can be solved to yield

$$(q_1, q_2, q_3) = (\frac{8}{15}, \frac{2}{15}, \frac{1}{3}).$$

The mixed strategy Nash equilibrium we have found yields the following prescription for the play of the game.

- Player 1: If dealt the high card, bet at every opportunity. If dealt the low card:
  (1) Bid with probability $\frac{3}{5}$, fold with probability $\frac{2}{5}$.
  (2) If you get to bet a second time, bet with probability $\frac{1}{3}$, fold with probability $\frac{2}{3}$.
- Player 2:
  (1) If player 1 bets, bet with probability $\frac{2}{3}$, fold with probability $\frac{1}{3}$.
  (2) If you get to bet a second time, bet with probability $\frac{4}{5}$, fold with probability $\frac{1}{5}$.

Note that in searching for a Nash equilibria in which both players use all three of their pure strategies, we in fact showed:

(1) If all three of player 2's pure strategies are active at a Nash equilibrium, then player 1's strategy must be $(p_1, p_2, p_3) = (\frac{1}{5}, \frac{2}{5}, \frac{2}{5})$. Hence there are no Nash equilibria in which player 1 uses only one or two pure strategies and player 2 uses three pure strategies.
(2) If all three of player 1's pure strategies are active at a Nash equilibrium, then player 2's strategy must be $(q_1, q_2, q_3) = (\frac{8}{15}, \frac{2}{15}, \frac{1}{3})$. Hence there are no Nash equilibria in which player 1 uses three pure strategies but player 2 uses only one or two pure strategies.

If the two players use these strategies, the expected payoff to each player is given by formula (5.2). For example, the expected payoff to player 1 is

$$\pi_1(\sigma_1, \sigma_2) = p_1 q_1 \cdot 0 + p_1 q_2 \cdot 4 + p_1 q_3 \cdot 2 + p_2 q_1 \cdot 1 + p_2 q_2 \cdot 0 + p_2 q_3 \cdot 2 + p_3 q_1 \cdot 2 + p_3 q_2 \cdot 1 + p_3 q_3 \cdot 0.$$

Substituting the values of the p_i and q_j that we have calculated, we find that the expected payoff to player 1 is $\frac{6}{5}$. Since the payoffs to the two players must add up to 0, the expected payoff to player 2 is $-\frac{6}{5}$. (In particular, we have found an arguably better strategy for player 2 than always folding, which you might have suspected would be his best strategy. If player 2 always folds and player 1 uses his best response, which is to always bid, then player 2's payoff is -2.)

Tom Ferguson is a mathematician at UCLA, where he teaches game theory. His online game theory text is available at http://www.math.ucla.edu/~tom/Game_Theory/Contents.html, and his home page is http://www.math.ucla.edu/~tom. On his home page you will find some recent articles applying game theory to poker that he wrote with his son, Chris Ferguson. Chris is a champion poker player, having won over $7 million. His Wikipedia page is http://en.wikipedia.org/wiki/Chris_Ferguson. A 2009 *New Yorker* article about Chris Ferguson, poker, and game theory is available at http://www.newyorker.com/reporting/2009/03/30/090330fa_fact_wilkinson.

5.5. Two-player zero-sum games

In the previous game, One-card Two-round Poker, we found one mixed strategy Nash equilibrium. Are there others?

One-card Two-round Poker is an example of a *two-player zero-sum game.* "Zero-sum" means that the two players' payoffs always add up to 0; if one player does better, the other must do worse. The Nash equilibria of two-player zero-sum games have two useful properties:

THEOREM 5.5. *Let* (σ_1, σ_2) *and* (τ_1, τ_2) *be two mixed strategy Nash equilibria of a two-player zero-sum game. Then*

(1) *Both give the same payoffs to the two players.*
(2) (σ_1, τ_2) *and* (τ_1, σ_2) *are also Nash equilibria.*

PROOF. Consider the following table of payoffs.

		Player 2	
		σ_2	τ_2
Player 1	σ_1	$(a, -a)$	$(b, -b)$
	τ_1	$(c, -c)$	$(d, -d)$

The table indicates that $\pi_1(\sigma_1, \sigma_2) = a$, $\pi_2(\sigma_1, \sigma_2) = -a$, $\pi_1(\sigma_1, \tau_2) = b$, $\pi_2(\sigma_1, \tau_2) = -b$, etc.

Since (σ_1, σ_2) is a Nash equilibrium,

$$a \geq c \quad \text{and} \quad -a \geq -b.$$

Since (τ_1, τ_2) is a Nash equilibrium,

$$d \geq b \quad \text{and} \quad -d \geq -c.$$

Rewriting these four inequalities in less-than form without minus signs, we have

$$c \leq a, \quad a \leq b, \quad b \leq d, \quad d \leq c.$$

Therefore $a = b = c = d$. Thus all four strategy profiles give the same payoffs to the two players. In particular (σ_1, σ_2) and (τ_1, τ_2) give the same payoffs to the two players, which proves (1).

To prove (2), we will just show that (σ_1, τ_2) is a Nash equilibrium. In fact, we will just show that player 1 cannot improve his own payoff by changing his strategy. Let σ_1' be another mixed strategy for player 1. Since (τ_1, τ_2) is a Nash equilibrium, $\pi_1(\sigma_1', \tau_2) \leq \pi_1(\tau_1, \tau_2)$. By the previous paragraph, $\pi_1(\tau_1, \tau_2) = \pi_1(\sigma_1, \tau_2)$. Therefore $\pi_1(\sigma_1', \tau_2) \leq \pi_1(\sigma_1, \tau_2)$. □

The first part of this theorem says that in a two-person zero-sum game, once you have found one Nash equilibrium, the others won't be any better (or worse).

The second part says that once you have found one Nash equilibrium in a two-person zero-sum game, just go ahead and play the strategy you have found. Even if your opponent finds a different Nash equilibrium and uses the strategy he finds, the profile of your two strategies is still a Nash equilibrium, and gives the same payoffs as those the two of you found.

In fact, in many problems, Theorem 5.5 can be used to rule out the existence of other Nash equilibria once you have found one. In One-card Two-round Poker, for example, we found a Nash equilibrium in which both players had three active pure strategies. We saw that our calculations to find this Nash equilibrium ruled out the existence of Nash equilibria in which player 1 used one or two pure strategies and player 2 used three pure strategies. But then, since One-card Two-round Poker is a two-person zero-sum game, Theorem 5.5 implies that there are no Nash equilibria in which player 1 uses one or two pure strategies and player 2 uses one or two pure strategies. (If there were such a Nash equilibrium, Theorem 5.5 part (2) would give a Nash equilibrium in which player 1 used one or two pure strategies and player 2 used three pure strategies; but there is no such Nash equilibrium.)

You may recall that when we were working on One-card Two-round Poker, we did not investigate the nine types of possible Nash equilibria in which each player used exactly two pure strategies. It turns out that there are no Nash equilibria of these types.

The same ideas work more generally for two-person constant-sum games, in which the two players' payoffs always add up to the same constant. In Tennis, for example, they always add up to 1.

5.6. Colonel Blotto vs. the People's Militia

This section is related to Gintis, Sec. 6.29.

There are two valuable towns. Col. Blotto has four regiments. The People's Militia has three regiments. Each decides how many regiments to send to each town.

If Col. Blotto send m regiments to a town and the People's Militia sends n, Col. Blotto's payoff for that town is

$$\begin{aligned} 1+n &\text{ if } m>n, \\ 0 &\text{ if } m=n, \\ -(1+m) &\text{ if } m<n. \end{aligned}$$

Col. Blotto's total payoff is the sum of his payoffs for each town. The People's Militia's payoff is the opposite of Col. Blotto's.

We consider this to be a game in normal form. Col. Blotto has 5 strategies, which we denote 40, 31, 22, 13, and 04. Strategy 40 is to send 4 regiments to town 1 and 0 town 2, etc. Similarly, the People's Militia has 4 strategies, which we denote 30, 21, 12, and 03.

We shall look for a mixed strategy Nash equilibrium (σ_1, σ_2), with $\sigma_1 = p_1 40 + p_2 31 + p_3 22 + p_4 13 + p_5 04$ and $\sigma_2 = q_1 30 + q_2 21 + q_3 12 + q_4 03$:

			People's Militia			
			q_1 **30**	q_2 **21**	q_3 **12**	q_4 **03**
	p_1	**40**	$(4,-4)$	$(2,-2)$	$(1,-1)$	$(0,0)$
	p_2	**31**	$(1,-1)$	$(3,-3)$	$(0,0)$	$(-1,1)$
Col. Blotto	p_3	**22**	$(-2,2)$	$(2,-2)$	$(2,-2)$	$(-2,2)$
	p_4	**13**	$(-1,1)$	$(0,0)$	$(3,-3)$	$(1,-1)$
	p_5	**04**	$(0,0)$	$(1,-1)$	$(2,-2)$	$(4,-4)$

We should consider the following possibilities.

(1) Both players use a pure strategy.
(2) Both players use exactly two active strategies.
(3) Both players use exactly three active strategies.
(4) Both players use exactly four active strategies.

We will also look briefly at the possibility that both players use all their active strategies (five for Col. Blotto, four for the People's Militia).

1. If both players use a pure strategy, we have a pure strategy Nash equilibrium. One easily checks that there are none.

2. We will not look at any possibilities in which both players use exactly two active strategies.

3. Suppose both player use exactly three active strategies. There are 40 ways this can happen. (Col. Blotto has 10 ways to choose 3 of his 5 strategies; the People's Militia has 4 ways to choose 3 of their 4 strategies; $10 \times 4 = 40$.)

We will consider just one of these possibilities. Suppose Col. Blotto uses only his 40, 22, and 04 strategies, and the People's Militia uses only its 30, 21, and 12 strategies. Thus we look for a Nash equilibrium $\sigma = (\sigma_1, \sigma_2)$, $\sigma_1 = p_1 40 + p_3 22 + p_5 04$, $\sigma_2 = q_1 30 + q_2 21 + q_3 12$.

Each of the People's Militia's pure strategies 30, 21, and 12 must yield the same payoff to the People's Militia against Col. Blotto's mixed strategy σ_1. These three payoffs are

$$\begin{aligned}\pi_2(\sigma_1, 30) &= -4p_1 + 2p_3, \\ \pi_2(\sigma_1, 21) &= -2p_1 - 2p_3 - p_5, \\ \pi_2(\sigma_1, 12) &= -p_1 - 2p_3 - 2p_5.\end{aligned}$$

We obtain three equations in three unknowns by using $\pi_2(\sigma_1, 30) = \pi_2(\sigma_1, 12)$ and $\pi_2(\sigma_1, 21) = \pi_2(\sigma_1, 12)$, together with $p_1 + p_3 + p_5 = 1$:

$$\begin{aligned}-3p_1 + 4p_3 + 2p_5 &= 0, \\ -p_1 + p_5 &= 0, \\ p_1 + p_3 + p_5 &= 1.\end{aligned}$$

The solution is $(p_1, p_3, p_5) = (\frac{4}{9}, \frac{1}{9}, \frac{4}{9})$.

Each of Col. Blotto's pure strategies 40, 22, and 04 must yield the same payoff to Col. Blotto against the People's Militia's mixed strategy σ_2. These three payoffs are

$$\begin{aligned}\pi_1(40, \sigma_2) &= 4q_1 + 2q_2 + q_3, \\ \pi_1(22, \sigma_2) &= -2q_1 + 2q_2 + 2q_3, \\ \pi_1(04, \sigma_2) &= q_2 + 2q_3.\end{aligned}$$

We obtain three equations in three unknowns by using $\pi_1(40, \sigma_2) = \pi_1(04, \sigma_2)$ and $\pi_1(22, \sigma_2) = \pi_1(04, \sigma_2)$, together with $q_1 + q_2 + q_3 = 1$:

$$\begin{aligned}4q_1 + q_2 - q_3 &= 0, \\ -2q_1 + q_2 &= 0, \\ q_1 + q_2 + q_3 &= 1.\end{aligned}$$

The solution is $(q_1, q_2, q_3) = (\frac{1}{9}, \frac{2}{9}, \frac{2}{3})$.

These calculations rule out the existence of Nash equilibria in which the People's Militia's active strategies are 30, 21, 12, and Col. Blotto's active strategies are one or two of his strategies 40, 22, and 04. They also rule out the existence of Nash equilibria in which Col. Blotto's active strategies are 40, 22, and 04, and the People's Militia's active strategies are one or two of their strategies 30, 21, 12.

We have seen that (σ_1, σ_2) with $\sigma_1 = \frac{4}{9}40 + \frac{1}{9}22 + \frac{4}{9}04$ and $\sigma_2 = \frac{1}{9}30 + \frac{2}{9}21 + \frac{2}{3}12$ satisfies the equality conditions for a Nash equilibrium. We now check the inequality conditions.

For Col Blotto:

$$\pi_1(40, \sigma_2) = \pi_1(22, \sigma_2) = \pi_1(04, \sigma_2) = \frac{14}{9},$$
$$\pi_1(31, \sigma_2) = \frac{1}{9}(1) + \frac{2}{9}(3) + \frac{2}{3}(0) = \frac{7}{9},$$
$$\pi_1(13, \sigma_2) = \frac{1}{9}(-1) + \frac{2}{9}(0) + \frac{2}{3}(3) = \frac{17}{9}$$

For the People's Militia:

$$\pi_2(\sigma_1, 30) = \pi_2(\sigma_1, 21) = \pi_2(\sigma_1, 12) = -\frac{14}{9},$$
$$\pi_2(\sigma_1, 03) = -\frac{14}{9}$$

Since $\pi_1(13, \sigma_2) > \frac{14}{9}$, the inequality conditions are not satisfied; (σ_1, σ_2) is not a Nash equilibrium.

Notice, however, that $\pi_2(\sigma_1, 03) = -\frac{14}{9}$, i.e., player 2's strategy 03 does just as well against σ_1 as the strategies that are active in σ_2. *When this happens, it is possible that there is a mixed strategy Nash equilibrium in which the two players use different numbers of active strategies.* In this case, we must check the possibility that there is a Nash equilibrium in which Col. Blotto's active strategies are 40, 22, and 04, and the People's Militia's strategies include 03 in addition to 30, 21, and 12.

3′. Thus we suppose Col. Blotto's active strategies are 40, 22, and 04, and all of the People's Militia's strategies are active. In other words, we look for a Nash equilibrium $\sigma = (\sigma_1, \sigma_2)$, $\sigma_1 = p_1 40 + p_3 22 + p_5 04$, $\sigma_2 = q_1 30 + q_2 21 + q_3 12 + q_4 03$.

Each of the People's Militia's pure strategies must yield the same payoff to the People's Militia against Col. Blotto's mixed strategy σ_1. We obtain four equations in three unknowns by using $\pi_2(\sigma_1, 30) = \pi_2(\sigma_1, 03)$, $\pi_2(\sigma_1, 21) = \pi_2(\sigma_1, 03)$, and $\pi_2(\sigma_1, 12) = \pi_2(\sigma_1, 03)$, together with $p_1 + p_3 + p_5 = 1$. Usually, if there are more equations than unknowns, there are no solutions. In this case, however, there is a solution: $(p_1, p_3, p_5) = (\frac{4}{9}, \frac{1}{9}, \frac{4}{9})$, i.e., the same solution we had before allowing the People's Militia to use its strategy 03. A little thought indicates that this is what will always happen.

Each of Col. Blotto's pure strategies 40, 22, and 04 must yield the same payoff to Col. Blotto against the People's Militia's mixed strategy σ_2. These three payoffs are

$$\pi_1(40,\sigma_2) = 4q_1 + 2q_2 + q_3,$$
$$\pi_1(22,\sigma_2) = -2q_1 + 2q_2 + 2q_3 - 2q_4,$$
$$\pi_1(04,\sigma_2) = q_2 + 2q_3 + 4q_4.$$

We obtain three equations in four unknowns by using $\pi_1(40,\sigma_2) = \pi_1(04,\sigma_2)$ and $\pi_1(22,\sigma_2) = \pi_1(04,\sigma_2)$, together with $q_1 + q_2 + q_3 + q_4 = 1$:

$$4q_1 + q_2 - q_3 - 4q_4 = 0,$$
$$-2q_1 + q_2 - 6q_4 = 0,$$
$$q_1 + q_2 + q_3 + q_4 = 1.$$

As usually happens with fewer equations than unknowns, there are many solutions. One way to list them all is as follows:

$$q_1 = \frac{1}{9} - q_4, \quad q_2 = \frac{2}{9} + 4q_4, \quad q_3 = \frac{2}{3} - 4q_4, \quad q_4 \text{ arbitrary.}$$

In order to keep all the q_i's between 0 and 1, we must restrict q_4 to the interval $0 \le q_4 \le \frac{1}{9}$.

Thus, if $\sigma_1 = \frac{4}{9}40 + \frac{1}{9}22 + \frac{4}{9}04$ and

$$\sigma_2 = (\frac{1}{9} - q_4)30 + (\frac{2}{9} + 4q_4)21 + (\frac{2}{3} - 4q_4)12 + q_4 03, \quad 0 \le q_4 \le \frac{1}{9},$$

then (σ_1, σ_2) satisfies the equality conditions for a Nash equilibrium. We now consider the inequality conditions when $0 < q_4 < \frac{1}{9}$, so that all of the People's Militia's strategies are active. Then for the People's Militia, there is no inequality constraint to check. For Col. Blotto:

$$\pi_1(40,\sigma_2) = \pi_1(22,\sigma_2) = \pi_1(04,\sigma_2) = \frac{14}{9},$$
$$\pi_1(31,\sigma_2) = (\frac{1}{9} - q_4)(1) + (\frac{2}{9} + 4q_4)(3) + (\frac{2}{3} - 4q_4)(0) + q_4(-1) = 10q_4 + \frac{7}{9},$$
$$\pi_1(13,\sigma_2) = (\frac{1}{9} - q_4)(-1) + (\frac{2}{9} + 4q_4)(0) + (\frac{2}{3} - 4q_4)(3) + q_4(1) = -10q_4 + \frac{17}{9}$$

To satisfy the inequality constraints for a Nash equilibrium, we need

$$10q_4 + \frac{7}{9} \le \frac{14}{9} \text{ and } -10q_4 + \frac{17}{9} \le \frac{14}{9}.$$

These inequality conditions are satisfied for $\frac{1}{30} \le q_4 \le \frac{7}{90}$.

We have thus found a one-parameter family of Nash equilibria (σ_1, σ_2): $\sigma_1 = \frac{4}{9}40 + \frac{1}{9}22 + \frac{4}{9}04$ and

$$\sigma_2 = (\frac{1}{9} - q_4)30 + (\frac{2}{9} + 4q_4)21 + (\frac{2}{3} - 4q_4)12 + q_4 03, \quad \frac{1}{30} \le q_4 \le \frac{7}{90},$$

The most attractive of these Nash equilibria occurs for q_4 at the midpoint of its allowed interval of values: $(q_1, q_2, q_3, q_4) = (\frac{1}{18}, \frac{4}{9}, \frac{4}{9}, \frac{1}{18})$. At this Nash equilibrium, the People's Militia uses its 30 and 03 strategies equally, and also uses its 21 and 12 strategies equally. Gintis (p. 350) seems to think that "by symmetry" this *must* be the case at a Nash equilibrium where the People's Militia uses all its strategies, but we have seen that this is not the case. We shall discuss symmetry and Nash equilibria in Chapter 7.

4. We will not look at any possibilities in which both players use four active strategies.

5. Suppose Col. Blotto uses all five of his pure strategies. Then at a Nash equilibrium, each of Col. Blotto's five pure strategies gives the same payoff to him against the People's Militia's mixed strategy σ_2. Therefore we have the following system of 5 equations in the 4 unknowns q_1, q_2, q_3, q_4:

$$\begin{aligned}\pi_1(40, \sigma_2) &= \pi_1(04, \sigma_2),\\ \pi_1(31, \sigma_2) &= \pi_1(04, \sigma_2),\\ \pi_1(22, \sigma_2) &= \pi_1(04, \sigma_2),\\ \pi_1(13, \sigma_2) &= \pi_1(04, \sigma_2),\\ q_1 + q_2 + q_3 + q_4 &= 1.\end{aligned}$$

Typically, when there are more equations than unknowns, there is no solution. One can check that that is the case here.

The game we have discussed is one of a class called Colonel Blotto games. They differ in the number of towns and in the number of regiments available to Col. Blotto and his opponent. There is a Wikipedia page devoted to these games: http://en.wikipedia.org/wiki/Colonel_Blotto. There you will learn that it has been argued that U.S. presidential campaigns should be thought of as Colonel Blotto games, it which the candidates must allocate their resources among the different states.

CHAPTER 6

Subgame perfect Nash equilibria and infinite-horizon games

6.1. Subgame perfect equilibria

Consider a game G in extensive form. We say that a node c' is a *successor* of a node c if there is a path in the game tree from c to c'.

Let h be a node that is not terminal and has no other nodes in its information set. Assume:

- If a node c is a successor of h, then every node in the information set of c is also a successor of h.

In this situation it makes sense to talk about the *subgame* H of G whose root is h. H consists of the node h and all its successors, connected by the same moves that connected them in G, and partitioned into the same information sets as in G. The players and the payoffs at the terminal vertices are also the same as in G.

If G is a game with complete information, then any nonterminal vertex of G is the root of a subgame of G.

Let s_i be one of player i's strategies in the game G. Recall that s_i is just a plan for what move to make at every vertex labeled i in the game G. So of course s_i includes a plan for what move to make at every vertex labeled i in the subgame H. Thus s_i contains within it a strategy that player i can use in the subgame H. We call this strategy the *restriction* of s_i to the subgame H, and label it s_{iH}.

Suppose the game G has n players, and $(s_1, \dots, s_n)$ is a Nash equilibrium for G. It is called a *subgame perfect Nash equilibrium* if, for every subgame H of G, $(s_{1H}, \dots, s_{nH})$ is a Nash equilibrium for H.

6.2. Big Monkey and Little Monkey 6

Recall the game of Big Monkey and Little Monkey from Section 1.5, with Big Monkey going first.

Recall that Big Monkey has two possible strategies in this game, and Little Monkey has four. When we find the payoffs for each choice of strategies, we get a game in normal form.

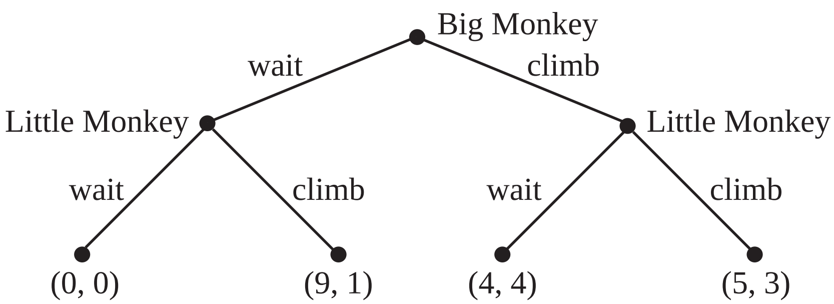

FIGURE 6.1. Big Monkey and Little Monkey.

		Little Monkey			
		ww	**wc**	**cw**	**cc**
Big Monkey	**w**	$(0, 0)$	$(0, 0)$	$(9, 1)$	$(9, 1)$
	c	$(4, 4)$	$(5, 3)$	$(4, 4)$	$(5, 3)$

We have seen (see Section 3.8 that there are three Nash equilibria in this game: (w, cw), (c, ww), and (w, cc).

The equilibrium (w, cw), which is the one found by backward induction, is subgame perfect. The other two are not.

The equilibrium (c, ww) is not subgame perfect because, in the subgame that begins after Big Monkey waits, Little Monkey would wait. By switching to climb, Little Monkey would achieve a better payoff. Little Monkey's plan to wait if Big Monkey waits is what we have called a *incredible threat*: it would hurt Big Monkey, at the cost of hurting Little Monkey

The equilibrium (w, cc) is not subgame perfect because, in the subgame that begins after Big Monkey climbs, Little Monkey would climb. By switching to wait, Little Monkey would achieve a better payoff. Little Monkey's plan to climb if Big Monkey climbs is what we called in Section 3.6 a *promise*: it would help Big Monkey, at the cost of hurting Little Monkey. As we saw there, in this particular game, the promise does not affect Big Monkey's behavior. Little Monkey is promising that if Big Monkey climbs, he will get a payoff of 5, rather than the payoff of 4 he would normally expect. Big Monkey ignores this promise because by waiting, he gets an even bigger payoff, namely 9.

6.3. Subgame perfect equilibria and backward induction

When we find strategies in finite-horizon extensive form games by backward induction, we are finding subgame perfect Nash equilibria. In fact, when we use backward induction, we are essentially considering every subgame in the entire game.

Strategies in a subgame perfect Nash equilibrium make sense no matter where in the game tree you use them. In contrast, at a Nash equilibrium that is not subgame perfect, at least one of the players is using a strategy that at some node tells him to make a move that would not be in his interest to make. For example, at a Nash equilibrium where a player is using a strategy that includes an incredible threat, if the relevant node were reached, the strategy would tell the player to make a move that would hurt him. The success of such a strategy depends on this node not being reached!

There are some finite-horizon games in extensive form for which backward induction does not work. Recall the following game that we discussed in Section 1.4.

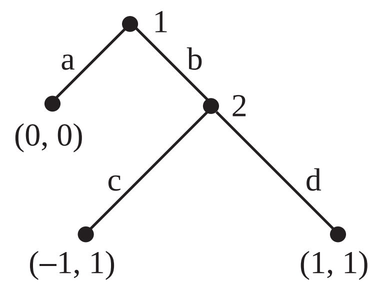

FIGURE 6.2. Failure of backward induction.

The problem with this game is that at the node where player 2 chooses, both available moves give him a payoff of 1. Hence player 1 does not know which move player 2 will choose if that node is reached. However, player 1 certainly wants to know which move player 2 will choose before he decides between a and b!

In this game player 1's strategy set is just $S_1 = \{a, b\}$ and player 2's strategy set is just $S_2 = \{c, d\}$. In normal form, the game is just

		Player 2	
		c	**d**
Player 1	**a**	$(0, 0)$	$(0, 0)$
	b	$(-1, 1)$	$(1, 1)$

This game has two Nash equilibria, (a, c) and (b, d). Both are subgame perfect.

Actually, there is a way to find all subgame perfect Nash equilibria in any finite-horizon game in extensive form with perfect information by a variant of backward induction. Do backward induction as usual. If at any point a player has *several* best choices, record each of them as a possible choice at that point, and *separately* continue the backward induction using each of them. Ultimately you will find all subgame perfect Nash equilibria

For example, in the game we are presently considering, we begin the backward induction at the node where player 2 is to choose, since it is the only node all of whose successors are terminal. Player 2 has *two* best choices, c and d. Continuing the backward induction using c, we find that Player 1 chooses a. Continuing the backward induction using d, we find that Player 1 chooses b. Thus the two strategy profiles we find are (a, c) and (b, d). Both are subgame perfect Nash equilibria.

For a finite-horizon game in extensive form with complete information, this more general backward induction procedure never fails. Therefore *every finite-horizon game in extensive form with complete information has at least one subgame perfect Nash equilibrium.*

For a *two-player zero-sum* finite-horizon game in extensive form with complete information we know more (recall Section 5.5: any subgame perfect Nash equilibrium yields the same payoffs to the two players; and player 1's strategy from one subgame perfect Nash equilibrium, played against player 2's strategy from another subgame perfect Nash equilibrium, also yields those payoffs. Such strategies are "best" for the two players. This applies, for example, to chess: because of the rule that a game is a draw if a position is repeated three times, chess is a finite-horizon game. Actually, it is a finite game: the game tree has about 10^50 nodes. Since the tree is so large, best strategies for white (player 1) and black (player 2) are not known. In particular, it is not known whether the best strategies yield a win for white, a win for black, or a draw.

The notion of subgame-perfect Nash equilibrium is especially valuable for *infinite-horizon* games, for which backward induction cannot be used. There are liable to be many Nash equilibria. Looking for the ones that are subgame perfect is a way of zeroing in on the (perhaps) most plausible ones. The next remainder of this chapter treats examples.

6.4. The Rubinstein bargaining model

This section is related to Gintis, Sec. 5.10, but I have changed the notation.

One dollar is to be split between two players. Player 1 goes first and offers to keep a fraction x_1 of the available money (one dollar). Of course, $0 \leq x_1 \leq 1$, and player 2 would get the fraction $1 - x_1$ of the available money. If player 2 accepts this proposal, the game is over, and the payoffs are x_1 to player 1 and $1 - x_1$ to player 2.

If player 2 rejects the proposal, the money shrinks to δ dollars, $0 < \delta < 1$, and it becomes player 2's turn to make an offer.

Player 2 offers a fraction y_1 of the available money (now δ dollars) to player 1. Of course, $0 \leq y_1 \leq 1$, and player 2 would get the fraction $1 - y_1$ of the available money. If player 1 accepts this proposal, the game is over, and the payoffs are $y_1\delta$ to player 1 and $(1 - y_1)\delta$ to player 2.

If player 1 rejects the proposal, the money shrinks to δ^2, and it becomes player 1's turn to make an offer.

Player 1 offers to keep a fraction x_2 of the available money (now δ^2 dollars) and give the fraction $1 - x_2$ to player 2. ... Well, you probably get the idea. See Figure 6.3.

The payoff to each player is the money he gets. If no proposal is ever accepted, the payoff to each player is 0.

This game models a situation in which it is in everyone's interest to reach an agreement quickly. Think, for example, of labor negotiations during a strike: as the strike goes on, the workers lose pay, and the company loses production. On the other hand, you probably don't want to reach a quick agreement by offering everything to your opponent!

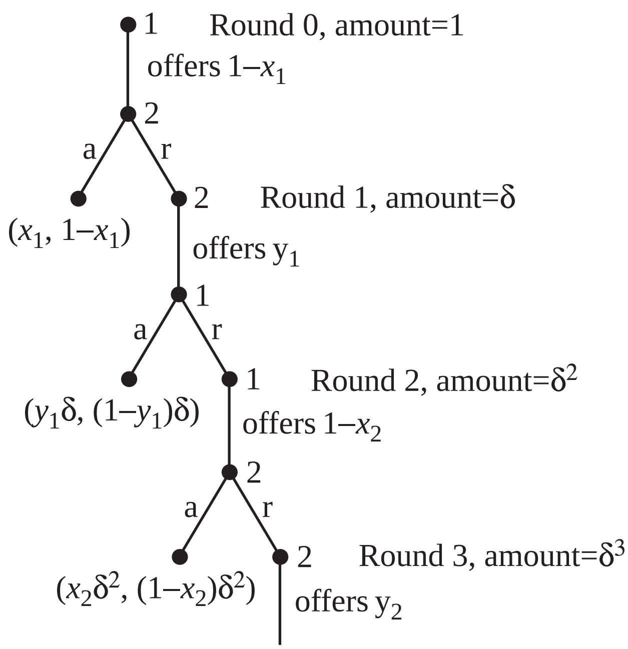

FIGURE 6.3. The Rubinstein bargaining model.

The numbering of the rounds of the game is shown in Figure 6.3.

A strategy for player 1 consists of a plan for each round. For the even rounds, he must plan what offer to make. For the odd rounds, he must plan which offers to accept and which to reject. Of course, his plan can depend on what has happened up to that point.

Player 2's strategies are similar. For the even rounds, he must plan which offers to accept and which to reject. For the odd rounds, he must plan what offer to make.

Notice that at the start of any even round, player 1 faces exactly the same situation that he faced at the start of the game, except that the available money is less. Similarly, at the start of any odd round, player 2 faces exactly the same situation that he faced at the start of round 1, except that the available money is less.

We will make two simplifying assumptions:

(1) Suppose a player has a choice between accepting an offer (thus terminating the game) and rejecting the offer (thus extending the game), and suppose the payoff the player expects from extending the game equals the offer he was just made. Then he will accept the offer, thus terminating the game.
(2) There is a subgame-perfect Nash equilibrium with the following property: if it yields a payoff of x to player 1, then in the subgame that starts at round 2, it yields a payoff of $x\delta^2$ to player 1; in the subgame that starts at round 4, it yields a payoff of $x\delta^4$ to player 1; etc.

With these assumptions, in the subgame perfect equilibrium of assumption (2), if the game were to go to Round 2, the payoffs would be $(x\delta^2, (1-x)\delta^2)$. So the game tree gets pruned to:

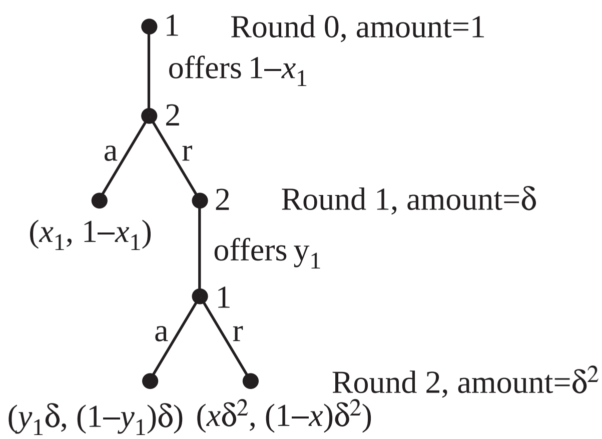

Figure 6.4. Pruned Rubinstein bargaining model.

Let us continue to investigate the subgame perfect equilibrium of assumption (2). Since it is subgame perfect, we should reason backward on the pruned game tree to find the players' remaining moves.

1. At round 1, player 2 must make an offer. If he offers a fraction y_1^* of the available amount δ chosen so that $y_1^*\delta = x\delta^2$, player 1 will be indifferent between accepting the offer and rejecting it. According to assumption 1, he will accept it. Player 2 will get to keep $\delta - y_1^*\delta$. If player 2 offered more than y_1^*, player 1 would accept, but player 2 would get less than $\delta - y_1^*\delta$. If player 2 offered less than y_1^*,

player 1 would not accept, and player 2 would end up with $\delta^2 - x\delta^2$. This is less than $\delta - y_1^*\delta$:

$$\delta^2 - x\delta^2 = \delta^2 - y_1^*\delta < \delta - y_1^*\delta.$$

Thus player 2's best move is to offer the fraction y_1^* of the available amount δ chosen so that $y_1^*\delta = x\delta^2$, i.e., $y_1^* = x\delta$. Since this is his best move, it is the move he makes at round 1 in our subgame perfect equilibrium . Player 1 accepts the offer.

2. At round 0, player 1 must make an offer. If he offers a fraction $1 - x_1^*$ of the available amount 1 chosen so that $1 - x_1^* = (1 - y_1^*)\delta$, player 2 will be indifferent between accepting the offer and rejecting it. According to assumption 1, he will accept it. Player one gets to keep x_1^*. Reasoning as before, we see that if player 1 offer more or less than $1 - x_1^*$, he ends up with less than x_1^*. Thus player 1's best move is to offer the fraction $1 - x_1^*$, player 2 accepts, and player 1's payoff is x_1^*.

3. We conclude that $x_1^* = x$,

4. From the equations $1 - x = 1 - x^* = (1 - y_1^*)\delta$ and $y_1^* = x\delta$, we obtain

$$1 - x = (1 - y_1^*)\delta = (1 - x\delta)\delta = \delta - x\delta^2,$$

so $1 - \delta = x - x\delta^2$, so

$$x = \frac{1-\delta}{1-\delta^2} = \frac{1}{1+\delta}. \tag{6.1}$$

Then

$$y_1^* = x\delta = \frac{\delta}{1+\delta} = 1 - x. \tag{6.2}$$

Player 1's payoff is higher. For δ close to 1 (i.e., when the money they are bargaining about does not shrink very fast), both payoffs are close to $\frac{1}{2}$.

6.5. Repeated games

Let G be a game with players $1, \dots, n$, strategy sets $S_1, \dots, S_n$, and payoff functions $\Pi_i(s_1, \dots, s_n)$.

We define a *repeated game* R with *stage game* G and *discount factor* δ as follows. The stage game G is played at times $j = 0, 1, 2, \dots$. A strategy for player i is just a way of choosing which of his mixed strategies σ_i to use at each time j. His choice of which strategy to use at time j can depend on all the strategies use by all the players at times before j.

There will be a payoff to player i from the stage game at each time $j = 0, 1, 2, \dots$. His payoff in the repeated game R is just

$$\text{his payoff at time } 0 + \delta \cdot \text{his payoff at time } 1 + \delta^2 \cdot \text{his payoff at time } 2 + \dots.$$

A subgame of R is defined by taking the repeated game that starts at some time j, together with the "memory" of the strategies used by all players in the stage games at earlier times.

6.6. Big Fish and Little Fish

Big Fish has parasites that live in its mouth. Little Fish likes to swim around in Big Fish's mouth and eat the parasites. This is nice for Big Fish. Of course, Big Fish would also enjoy taking a bite out of Little Fish. Similarly, Little Fish would enjoy taking a bite out of Big Fish.

The payoffs are given in the falling table.

Payoffs to	LF eats patasites	LF bites BF	BF bites LF
Big Fish	5	-3	3
Little Fish	5	8	-8

If the two fish decide simultaneously what to do, we have the following game in normal form.

		Little Fish	
		eat parasites	**bite Big Fish**
Big Fish	**don't bite Little Fish**	$(5, 5)$	$(-3, 8)$
	bite Little Fish	$(8, -3)$	$(0, 0)$

Big Fish has a strictly dominant strategy: bite Little Fish. Little Fish also has a strictly dominant strategy: bite Big Fish. If both fish use their nasty dominant strategies, each gets a payoff of 0. (They both get a bite, but also get bit.) On the other hand, if both fish use their nice, dominated strategies (Big Fish leaves Little Fish alone, Little Fish just eats the parasites), both get a payoff of 5. This game is a prisoner's dilemma.

To make this game easier to discuss, let's call each fish's nice strategy C for "cooperate", and let's call each fish's nasty strategy D for "defect." Now we have the following table:

		Little Fish	
		C	**D**
Big Fish	**C**	$(5, 5)$	$(-3, 8)$
	D	$(8, -3)$	$(0, 0)$

Both fish have the same strategy set, namely $\{C, D\}$.

We will take this game to be the stage game in a repeated game R.

Recall that a player's strategy in a repeated game is a way of choosing which of his strategies to use at each time j. The choice can depend on what all players have done at times before j.

We consider the *trigger strategy* in this repeated game, which we denote σ and which is defined as follows. Start by using C. Continue to use C as long as your opponent uses C. If your opponent ever uses D, use D at your next turn, and continue to use D forever.

THEOREM 6.1. *If $\delta \geq \frac{3}{8}$, then (σ, σ) is a Nash equilibrium.*

PROOF. If both fish use σ, then both cooperate in every round. Therefore both receive a payoff of 5 in every round. Taking into account the discount factor δ, we have

$$\Pi_1(\sigma,\sigma) = \Pi_2(\sigma,\sigma) = 5(1 + \delta + \delta^2 + \ldots) = 5\frac{1}{1-\delta}.$$

Suppose Little Fish switches to a different strategy σ'.

Case 1. Little Fish still ends up cooperating in every round. Then Big Fish, who is still using σ, will also cooperate in every round. The payoffs are unchanged.

Case 2. Little Fish first defects in round j. Then Big Fish will cooperate through round j, will defect in round $j+1$, and will defect in every round after that. Does using σ' improve Little Fish's payoff?

The payoffs from the strategy profiles (σ, σ) and (σ, σ') are the same through round $j - 1$, so let's just compare their payoffs to Little Fish from round j on.

With (σ, σ), Little Fish's payoff from round j on is

$$5(\delta^j + \delta^{j+1} + \ldots) = 5\delta^j(1 + \delta + \ldots) = 5\delta^j\frac{1}{1-\delta}.$$

With (σ, σ'), Little Fish's payoff in round j is 8: the benefit of biting Big Fish when Big Fish doesn't bite back. From round $j+1$ on, unfortunately, Big Fish will defect (bite Little Fish) in every round. Little Fish's best response to this is to bite back, giving him a payoff of 0. Therefore, taking into account the discount factor, Little Fish's payoff from round j on is at most

$$8\delta^j + 0(\delta^{j+1} + \delta^{j+2} + \ldots) = 8\delta^j.$$

From round j on, Little Fish's payoff from (σ, σ) is greater than or equal to his payoff from (σ, σ') provided

$$5\delta^j\frac{1}{1-\delta} \geq 8\delta^j \quad \text{or} \quad \delta \geq \frac{3}{8}.$$

□

Actually, the proof shows that (σ, σ) is a subgame perfect Nash equilibrium.

CHAPTER 7

Symmetric games

We have seen several examples of games in normal form in which the players are interchangeable:

- Prisoner's Dilemma
- Stag Hunt
- Water Pollution
- Tobacco Market
- Cournot Duopoly

For each of these games, we found at least one pure strategy Nash equilibrium in which every player used the same strategy. However, this is not always the case.

For example, the game of Chicken was (supposedly) played by teenagers in the 1950's. A variant of the game (not the version we will describe) is shown in the James Dean movie "Rebel without a Cause." In this game, two teenagers drive their cars toward each other at high speed. Each has two strategies: drive straight or swerve. The payoffs are as follows.

		Teenager 2	
		straight	**swerve**
Teenager 1	**straight**	$(-2,-2)$	$(1,-1)$
	swerve	$(-1,1)$	$(0,0)$

If one teen drives straight and one swerves, the one who drives straight gains in reputation, and the other loses face. However, if both drive straight, there is a crash, and both are injured. There are two Nash equilibria: (straight, swerve) and (swerve, straight). In both of these Nash equilibria, of course, the players use different strategies.

Let's look for a mixed strategy equilibrium

$$(p \text{ straight} + (1-p) \text{ curve}, q \text{ straight} + (1-q) \text{ curve}),$$

with $0 < p < 1$ and $0 < q < 1$. From the Fundamental Theorem,

$$p(-2) + (1-p)(1) = p(-1) + (1-p)(0),$$
$$q(-2) + (1-q)(1) = q(-1) + (1-q)(0).$$

The two equations are the same; this is a consequence of the symmetry of the game. The solution is $(p, q) = (\frac{1}{2}, \frac{1}{2})$.

A game in normal form in which the players are interchangeable is called *symmetric*. In such a game, all players, of course, have the same strategy set.

We will give a formal definition only in the case of a symmetric two-player game. Let S denote the set of (pure) strategies available to either player. Then we require

- For all s and t in S, $\pi_1(s, t) = \pi_2(t, s)$.

The following result should not be surprising.

THEOREM 7.1. *In a symmetric game in which the (pure) strategy set is finite, there is a mixed strategy Nash equilibrium in which every player uses the same mixed strategy.*

One can take advantage of this theorem by looking for such Nash equilibria instead of more general equilibria.

We note one other fact about symmetric games.

- Suppose a symmetric game has a Nash equilibrium in which it is not the case that all players use the same mixed strategy. Then other Nash equilibria can be found by interchanging the rolls of the players.

We have seen this phenomenon in the games of Chicken and Preservation of Ecology.

7.1. Water Pollution 3

In the Water Pollution (Sec. 3.3), we have already considered pure strategy Nash equilibria. Now we will consider mixed strategy Nash equilibria in which all three players use completely mixed strategies. Let g and b denote the strategies purify and pollute respectively. Then we search for a mixed strategy Nash equilibrium $(\sigma_1, \sigma_2, \sigma_3) = (xg + (1 - x)b, yg + (1 - y)b, zg + (1 - z)b)$. Since the numbers x, y, and z determine the player's strategies, we shall think of the payoff functions π_i as functions of (x, y, z). We will not assume that $x = y = z$. In fact, since we have already found a Nash equilibrium in which all players use the same strategy, it is not guaranteed that there will exist additional mixed strategy Nash equilibria in which all players use the same strategy.

The following chart helps to keep track of the notation:

Firm 3 z **g**

			Firm 2 y **g**	$1-y$ **b**
Firm 1	x	**g**	$(-1,-1,-1)$	$(-1,0,-1)$
	$1-x$	**b**	$(0,-1,-1)$	$(-3,-3,-4)$

Firm 3 $1-z$ **b**

			Firm 2 y **g**	$1-y$ **b**
Firm 1	x	**g**	$(-1,-1,0)$	$(-4,-3,-3)$
	$1-x$	**b**	$(-3,-4,-3)$	$(-3,-3,-3)$

The criteria for a mixed strategy Nash equilibrium in which all three players have two active strategies (i.e., $0 < x < 1$, $0 < y < 1$, $0 < z < 1$) are:

$$\pi_1(1,y,z) = \pi_1(0,y,z), \quad \pi_2(x,1,z) = \pi_2(x,0,z), \quad \pi_3(x,y,1) = \pi_3(x,y,0).$$

The first equation, for example, says that if players 2 and 3 use the mixed strategies $yg + (1-y)b$ and $zg + (1-z)b$ respectively, the payoff to player 1 if he uses g must equal the payoff to him if he uses b.

The third of these equations is

$$-xy - x(1-y) - (1-x)y - 4(1-x)(1-y) = 0xy - 3x(1-y) - 3(1-x)y - 3(1-x)(1-y),$$

which simplifies to

$$\frac{1-3x}{x} = 3\frac{y}{1-3y} \tag{7.1}$$

The other two equations are

$$\frac{1-3x}{x} = 3\frac{z}{1-3z} \quad \text{and} \quad \frac{1-3y}{y} = 3\frac{z}{1-3z}. \tag{7.2}$$

Equations (7.2) yield

$$\frac{1-3x}{x} = \frac{1-3y}{y}, \tag{7.3}$$

which easily yields $x = y$. Similarly, $x = z$, so $x = y = z$. In other words, all players *must* use the same strategy.

Substituting $x = y$ in (7.1), we obtain

$$\frac{1-3y}{y} = 3\frac{y}{1-3y},$$

so $6y^2 - 6y + 1 = 0$. The quadratic formula yields $y = \frac{1}{2} \pm \frac{1}{6}\sqrt{3}$. We have therefore found two mixed strategy Nash equilibria: $x = y = z = \frac{1}{2} + \frac{1}{6}\sqrt{3}$ and $x = y = z = \frac{1}{2} - \frac{1}{6}\sqrt{3}$.

7.2. Reporting a crime

In 1964 a young woman named Kitty Genovese was murdered outside her home in Queens, New York. According to a New York Times article written two weeks later, 38 of her neighbors witnessed the murder, but none of them called the police. While the accuracy of the article has since been called into question (http://en.wikipedia.org/wiki/Kitty_genovese), at the time it horrified the country.

Here is a model that has been proposed for such events. A crime is observed by n people. Each wants the police to be informed but prefers that someone else make the call. We will suppose that each person receives a payoff of v as long as at least one person calls the police; if no one calls the police, each person receives a payoff of 0. Each person who calls the police incurs a cost of c. We assume $0 < c < v$.

We view this as a game with n players. Each has two strategies: call the police (C) or don't call the police (N). The total payoffs are:

- If at least one person calls the police: v to each person who does not call, $v - c$ to each person who calls.
- If no one calls the police: 0 to everyone.

You can easily check that there are exactly n pure strategy Nash equilibria. In each of them, exactly one of the n people calls the police.

Motivated by Theorem 7.1, we shall look for a mixed strategy Nash equilibrium $(\sigma_1, \dots, \sigma_n)$ in which all players use the same strictly mixed strategy $\sigma_i = (1-p)C + pN$, $0 < p < 1$.

Let's consider player 1. By the Fundamental Theorem, each of her pure strategies gives her the same expected payoff when players 2 through n use their mixed strategies:

$$\pi_1(C, \sigma_2, \dots, \sigma_n) = \pi_1(N, \sigma_2, \dots, \sigma_n).$$

Now $\pi_1(C, \sigma_2, \dots, \sigma_n) = v - c$, since the payoff to a player who calls is $v - c$ no matter what the other players do. On the other hand,

$$\pi_1(N, \sigma_2, \dots, \sigma_n) = \begin{cases} 0 \text{ if no one else calls,} \\ v \text{ if at least one other person calls.} \end{cases}$$

The probability that no one else calls is p^{n-1}, so the probability that at least one other person calls is $1 - p^{n-1}$. Therefore

$$\pi_1(N, \sigma_2, \dots, \sigma_n) = 0 \cdot p^{n-1} + v \cdot (1 - p^{n-1}) = v(1 - p^{n-1}).$$

Therefore

$$v - c = v(1 - p^{n-1}),$$

so

$$p = \left(\frac{c}{v}\right)^{\frac{1}{n-1}}.$$

Since $0 < \frac{c}{v} < 1$, p is a number between 0 and 1.

What does this formula mean? Notice first that as $n \to \infty$, $p \to 1$, so $1-p \to 0$. Thus, as the size of the group increases, each individual's probability of calling the police declines toward 0. However, it is more important to look at the probability that at least one person calls the police. This probability is

$$1 - p^n = 1 - \left(\frac{c}{v}\right)^{\frac{n}{n-1}}.$$

As n increases, $\frac{n}{n-1} = 1 + \frac{1}{n-1}$ *decreases* toward 1, so $\left(\frac{c}{v}\right)^{\frac{n}{n-1}}$ *increases* toward $\frac{c}{v}$, so $1 - \left(\frac{c}{v}\right)^{\frac{n}{n-1}}$ *decreases* toward $1 - \frac{c}{v}$. Thus, as the size of the group increases, the probability that the police are called *decreases*.

For a large group, the probability that the police are called is approximately $1-\frac{c}{v}$. Anything that increases c (the perceived cost of calling the police) or decreases v (the value to people of seeing that the police get called) will decrease the likelihood that the police are called.

7.3. Sex ratio

This section is related to Gintis, Sec. 6.26.

Most organisms that employ sexual reproduction come in two types: male and female. In many species, the percentages of male and female offspring that survive to actually reproduce are very different. Nevertheless, in most species, approximately half of all births are male and half are female. What is the reason for this? This puzzle goes back to Darwin.

One can find an answer by focussing on the number of grandchildren of each female.

Suppose a cohort of males and females is about to reproduce. We think of this as a game in which the players are the females, a female's strategy is her fraction of male offspring, and a female's payoff is her number of grandchildren.

7.3.1. Many many females. There are lots of players! For a first pass at analyzing this situation, we will imagine that one female has a fraction u of male offspring, $0 \le u \le 1$, and the females as a group have a fraction v of male offspring, $0 \le v \le 1$. We imagine the group is so large that what our one female does has no appreciable effect on v. For each v we will calculate our female's best response set $B(v)$. Motivated by the notion of Nash equilibrium, we ask, for what values of v does the set $B(v)$ include v?

Notation:

- σ_m = fraction of males that survive to reproduce.
- σ_f = fraction of females that survive to reproduce.
- c = number of offspring per female.
- r = number of offspring per male.
- y = number of females.

Then we have:

	Female 1	**All Females**
Sons	uc	vcy
Daughters	$(1-u)c$	$(1-v)cy$
Surviving sons	$\sigma_m uc$	$\sigma_m vcy$
Surviving daughters	$\sigma_f(1-u)c$	$\sigma_f(1-v)cy$

Let $f(u, v)$ denote the number of grandchildren of female 1. Then we have

$$\begin{aligned} f(u,v) &= \text{surviving sons} \cdot \text{offspring per son} \\ &\quad + \text{ surviving daughters} \cdot \text{offspring per daughter} = \sigma_m uc \cdot r + \sigma_f(1-u)c \cdot c. \end{aligned}$$

We can calculate r as follows. For the population as a whole,

$$\text{surviving sons} \cdot \text{offspring per son} = \text{surviving daughters} \cdot \text{offspring per daughter},$$

i.e.,

$$\sigma_m vcy \cdot r = \sigma_f(1-v)cy \cdot c.$$

Therefore $r = \frac{\sigma_f(1-v)}{\sigma_m v}c$. Substituting this value into our formula for $f(u, v)$, we obtain

$$f(u,v) = \sigma_m uc\frac{\sigma_f(1-v)}{\sigma_m v}c + \sigma_f(1-u)c^2 = \sigma_f c^2\left(1 + u\frac{1-2v}{v}\right), \quad 0 \le u \le 1,\ 0 \le v \le 1.$$

Notice

$$\frac{\partial f}{\partial u}(u,v) = \begin{cases} + & \text{if } 0 < v < \frac{1}{2}, \\ 0 & \text{if } v = \frac{1}{2}, \\ - & \text{if } v > \frac{1}{2}, \end{cases}$$

Therefore female 1's best response to v is $u = 1$ if $0 < v < \frac{1}{2}$; any u if $v = \frac{1}{2}$; and $u = 0$ if $v > \frac{1}{2}$. Only in the case $v = \frac{1}{2}$ does female 1's best response include v.

For Darwin's views on sex ratios, see http://www.ucl.ac.uk/~ucbhdjm/courses/b242/Sex/D71SexRatio.html. For further discussion of sex ratios, see the Wikipedia page http://en.wikipedia.org/wiki/Fisher's_principle.

7.3.2. Not so many females. Suppose there are n females. The ith female's strategy is a number u_i, $0 \le u_i \le 1$, that represents the fraction of male offspring she has. Her payoff is her number of grandchildren. Motivated by our work in the previous subsection, we will now derive a formula for the payoff.

We continue to use the notation:

- σ_m = fraction of males that survive to reproduce.
- σ_f = fraction of females that survive to reproduce.
- c = number of offspring per female.
- r = number of offspring per male.

Let

$$v = \frac{1}{n}(u_1 + \ldots + u_n).$$

Then we have:

	Female i	**All Females**
Sons	$u_i c$	vcn
Daughters	$(1-u_i)c$	$(1-v)cn$
Surviving sons	$\sigma_m u_i c$	$\sigma_m vcn$
Surviving daughters	$\sigma_f(1-u_i)c$	$\sigma_f(1-v)cn$

The formula for r is unchanged. We therefore have

$$\pi_i(u_1, \ldots, u_n) = \sigma_f c^2 \left(1 + u_i \frac{1-2v}{v}\right).$$

We will look for a Nash equilibrium at which $0 < u_i < 1$ for every i. Then for every i, we must have $\frac{\partial \pi_i}{\partial u_i}(u_1, \ldots, u_n) = 0$, i.e.,

$$0 = \frac{\partial \pi_i}{\partial u_i} = \sigma_f c^2 \left(1 \cdot \frac{1-2v}{v} + u_i \frac{\partial}{\partial v}\left(\frac{1-2v}{v}\right)\frac{\partial v}{\partial u_i}\right)$$
$$= \sigma_f c^2 \left(\frac{1-2v}{v} - u_i \frac{1}{v^2}\frac{1}{n}\right) = \sigma_f c^2 \frac{nv - 2nv^2 - u_i}{nv^2}.$$

Hence, for every i, $u_i = nv - 2nv^2$. Therefore all u_i are equal; denote their common value by u. Then $v = u$, so

$$0 = \frac{nu - 2nu^2 - u}{nu^2} = \frac{n - 2nu - 1}{nu}.$$

Therefore $u = \frac{n-1}{2n} = \frac{1}{2} - \frac{1}{2n}$. For n large, u is very close to $\frac{1}{2}$.

CHAPTER 8

Evolutionary stability

8.1. Evolutionary stability

This section is related to Gintis, Secs. 10.1–10.3.

We consider a symmetric two-person game in normal form in which the strategy set S is finite: $S = \{s_1, \ldots, s_n\}$.

Consider a population of many individuals who play this game with each other. An individual uses a mixed strategy $\tau = \sum q_i s_i$ with (of course) all $q_i \geq 0$ and $\sum q_i = 1$. We will say that the individual is of *type* τ. If τ is a pure strategy, say $\tau = s_i$, we will say that the individual is of type i.

The population taken as a whole uses strategy s_1 with probability p_1, ..., strategy s_n with probability p_n; all $p_i \geq 0$ and $\sum p_i = 1$. We say that the *state* of the population is σ and write $\sigma = \sum p_i s_i$.

Let $\pi_{ij} = \pi_1(s_i, s_j)$. From the symmetry of the game,

$$\pi_2(s_i, s_j) = \pi_1(s_j, s_i) = \pi_{ji}.$$

If an individual of type i plays this game against a random individual from a population of type σ, his expected payoff, which we denote $\pi_{i\sigma}$, is

$$\pi_{i\sigma} = \sum_{j=1}^{n} \pi_{ij} p_j.$$

Of course, $\pi_{i\sigma} = \pi_1(s_i, \sigma)$, the expected payoff to an individual who uses the pure strategy s_i against another individual using the mixed strategy σ.

If an individual of type τ plays this game against a random individual from a population of type σ, his expected payoff, which we denote $\pi_{\tau\sigma}$, is

$$\pi_{\tau\sigma} = \sum_{i=1}^{n} q_i \pi_{i\sigma} = \sum_{i,\, j=1}^{n} q_i \pi_{ij} p_j. \tag{8.1}$$

Of course, $\pi_{\tau\sigma} = \Pi_1(\tau, \sigma)$, the expected payoff to an individual who uses the mixed strategy τ against another individual using the mixed strategy σ.

Let σ_1 and σ_2 be population states, $\sigma_1 = \sum p_{1i}s_i$, $\sigma_2 = \sum p_{2i}s_i$, and let $0 < \epsilon < 1$. Then we can define a new population state

$$(1-\epsilon)\sigma_1 + \epsilon\sigma_2 = \sum((1-\epsilon)p_{1i} + \epsilon p_{2i})s_i.$$

If τ is a mixed strategy, one can easily show from (8.1) that

$$\pi_{\tau\ (1-\epsilon)\sigma_1+\epsilon\sigma_2} = (1-\epsilon)\pi_{\tau\sigma_1} + \epsilon\pi_{\tau\sigma_2}. \tag{8.2}$$

Suppose, in a population with state σ, we replace a fraction ϵ of the population, $0 < \epsilon < 1$, with individuals of type τ. The new population state is $(1-\epsilon)\sigma + \epsilon\tau$.

We will say that a population state σ is *evolutionarily stable* if for every $\tau \neq \sigma$ there is a number $\epsilon_0 > 0$ such that if $0 < \epsilon < \epsilon_0$ then

$$\pi_{\sigma\ (1-\epsilon)\sigma+\epsilon\tau} > \pi_{\tau\ (1-\epsilon)\sigma+\epsilon\tau}. \tag{8.3}$$

This inequality says that if a population of type σ is invaded by a small number of individuals of any other type τ, an individual of type σ will have a better expected payoff against a random member of the mixed population than will an individual of type τ.

THEOREM 8.1. *A population state σ is evolutionarily stable if and only if for all $\tau \neq \sigma$,*

(1) $\pi_{\sigma\sigma} \geq \pi_{\tau\sigma}$.
(2) *If* $\pi_{\sigma\sigma} = \pi_{\tau\sigma}$ *then* $\pi_{\sigma\tau} > \pi_{\tau\tau}$.

The first condition says that (σ, σ) is a mixed strategy Nash equilibrium of the symmetric two-player game. Therefore σ is a best response to σ. The second condition says that if τ is another best response to σ, then σ is a better response to τ than τ is to itself.

PROOF. Assume σ is evolutionarily stable. Let $\tau \neq \sigma$. Then for all sufficiently small $\epsilon > 0$, (8.3) holds. Therefore, for all sufficiently small $\epsilon > 0$,

$$(1-\epsilon)\pi_{\sigma\sigma} + \epsilon\pi_{\sigma\tau} > (1-\epsilon)\pi_{\tau\sigma} + \epsilon\pi_{\tau\tau}.$$

Dividing by $1-\epsilon$, a positive number, we obtain

$$\pi_{\sigma\sigma} + \frac{\epsilon}{1-\epsilon}\pi_{\sigma\tau} > \pi_{\tau\sigma} + \frac{\epsilon}{1-\epsilon}\pi_{\tau\tau}. \tag{8.4}$$

Letting $\epsilon \to 0$, we obtain (1). If $\pi_{\sigma\sigma} = \pi_{\tau\sigma}$, (8.4) becomes

$$\frac{\epsilon}{1-\epsilon}\pi_{\sigma\tau} > \frac{\epsilon}{1-\epsilon}\pi_{\tau\tau}.$$

Multiplying by $\frac{1-\epsilon}{\epsilon}$ we obtain (2).

The converse is similar. □

One consequence of Theorem 8.1 is

THEOREM 8.2. *If (σ, σ) is a strict Nash equilibrium of the symmetric two-player game, then σ is an evolutionarily stable state.*

The reason is that for any Nash equilibrium, (1) holds; and for a strict Nash equilibrium, (2) is irrelevant. Of course, strict Nash equilibria use only pure strategies, so such populations are simple: they consist entirely of individual of one pure type i.

Another consequence of Theorem 8.1 is

THEOREM 8.3. *If σ is an evolutionarily stable state in which all pure strategies are active, then for all $\tau \neq \sigma$, (τ, τ) is not a Nash equilibrium of the symmetric two-player game. Hence there are no other evolutionarily stable states.*

PROOF. For such a σ we have that for all i, $\pi_{i\sigma} = \pi_{\sigma\sigma}$. Therefore, for all τ,

$$\pi_{\tau\sigma} = \sum_{1=1}^{n} q_i \pi_{i\sigma} = \sum_{1=1}^{n} q_i \pi_{\sigma\sigma} = \pi_{\sigma\sigma}.$$

Therefore, since σ is evolutionarily stable, (2) implies that for all $\tau \neq \sigma$, $\pi_{\sigma\tau} > \pi_{\tau\tau}$. Therefore (τ, τ) is not a Nash equilibrium of the symmetric two-player game. □

The same argument shows:

THEOREM 8.4. *Let $\sigma = \sum p_i s_i$ be an evolutionarily stable state, and let $I = \{i : p_i > 0\}$. Let $\tau \neq \sigma$, $\tau = \sum q_i s_i$, be a population state for which the the set of i's such that $q_i > 0$ is a subset of I. Then (τ, τ) is not a Nash equilibrium of the symmetric two-player game. Hence there are no other evolutionarily stable states in which the set of active strategies is a subset of the set of active strategies in σ.*

8.2. Evolutionary stability with two pure strategies

8.2.1. Theory. This subsection is related to Gintis, Sec. 10.3.

Consider a symmetric two-person game in normal form with just two pure strategies. The payoff matrix must have the form

		Player 2	
		s_1	s_2
Player 1	s_1	(a, a)	(b, c)
	s_1	(c, b)	(d, d)

THEOREM 8.5. *Suppose $a \neq c$ and $b \neq d$.*

(1) *If $a > c$ and $d < b$, then strategy s_1 strictly dominates strategy s_2. There is one Nash equilibrium: (s_1, s_1). It is symmetric and strict, so the population state s_1 is evolutionarily stable.*

(2) *If $a < c$ and $d > b$, then strategy s_2 strictly dominates strategy s_1. There is one Nash equilibrium: (s_2, s_2). It is symmetric and strict, so the population state s_2 is evolutionarily stable.*

To describe the other two cases, let

$$p = \frac{d-b}{(a-c)+(d-b)}, \quad \text{so} \quad 1-p = \frac{a-c}{(a-c)+(d-b)}. \tag{8.5}$$

(3) *If $a > c$ and $d > b$, there are three Nash equilibria: (s_1, s_1), (s_2, s_2), and (σ, σ) with $\sigma = ps_1 + (1-p)s_2$ and p given by (8.5). The first two are symmetric and strict, so the population states s_1 and s_2 are evolutionarily stable. The population state σ is not evolutionarily stable.*

(4) *If $a < c$ and $d < b$, there are three Nash equilibria: (s_1, s_2), (s_2, s_1), and (σ, σ) with $\sigma = ps_1 + (1-p)s_2$ and p given by (8.5). Only the last is symmetric. The population state σ is evolutionarily stable.*

Case 3 is a pure coordination game if $a \neq d$. Case 4 includes Chicken.

PROOF. You can find the pure strategy Nash equilibria by circling best responses.

Now we look for mixed strategy Nash equilibria (σ, τ) with $\sigma = ps_1 + (1-p)s_2$, $\tau = qs_1 + (1-q)s_2$, and at least one of σ and τ has two active strategies. To find p and q, we first add the probabilities to the payoff matrix:

			Player 2	
			q	$1-q$
			s_1	s_2
Player 1	p	s_1	(a, a)	(b, c)
	$1-p$	s_1	(c, b)	(d, d)

At least one player has two active strategies; suppose it is player 2. Then if player 2 uses either pure strategy s_1 or pure strategy s_2, he gets the same expected payoff when player 1 uses σ. Therefore

$$pa + (1-p)b = pc + (1-p)d \quad \text{so} \quad b - d = \big((a-c) + (d-b)\big)p.$$

Since $b - d \neq 0$ by assumption, we must have $(a-c) + (d-b) \neq 0$ in order to solve for p. Then p is given by (8.5). In cases 1 and 2, this value of p is not between 0 and 1, so it cannot be used. In cases 3 and 4, on the other hand, we see that $0 < p < 1$, so both of player 1's strategies are active. Then we can calculate q the same way. We find that $q = p$.

Now that we have a symmetric Nash equilibrium (σ, σ) in cases 3 and 4, we check whether the corresponding population state σ is evolutionarily stable. Since (σ, σ) is a Nash equilibrium, σ satisfies (1) of Theorem 8.1. Since both pure strategies are active in σ, every τ satisfies $\pi_{\sigma\sigma} = \pi_{\tau\sigma}$, so (2) must be checked for every $\tau \neq \sigma$. For $\tau = qs_1 + (1-q)s_2$, we calculate

$$\begin{aligned}
\pi_{\sigma\tau} - \pi_{\tau\tau} &= paq + pb(1-q) + (1-p)cq + (1-p)d(1-q) \\
&\quad - \big(qaq + qb(1-q) + (1-q)cq + (1-q)d(1-q)\big) \\
&= (p-q)\big(aq + b(1-q) - cq - d(1-q)\big) \\
&= (p-q)\big(b - d + \big((a-c) + (d-b)\big)q\big) \\
&= (p-q)\big((a-c)+(d-b)\big)\left(\frac{b-d}{(a-c)+(d-b)} + q\right) \\
&= -(p-q)\big((a-c)+(d-b)\big)\left(\frac{d-b}{(a-c)+(d-b)} - q\right) \\
&= -(p-q)^2\big((a-c)+(d-b)\big).
\end{aligned}$$

If $\tau \neq \sigma$, then $q \neq p$, so $(p-q)^2 > 0$. Thus we see that in case 3 ($a-c$ and $d-b$ both positive), $\pi_{\sigma\tau} - \pi_{\tau\tau} < 0$ for all $\tau \neq \sigma$, so σ is *not* evolutionarily stable; and in case 4 ($a-c$ and $d-b$ both negative), $\pi_{\sigma\tau} - \pi_{\tau\tau} > 0$ for all $\tau \neq \sigma$, so σ *is* evolutionarily stable. □

In case 3, the population state σ is the opposite of evolutionarily stable: if τ is any invading population type, σ does worse against τ than τ does against itself.

8.2.2. Stag hunt. Let's reconsider the game of Stag Hunt (Subsection 3.2.2). The payoff matrix is reproduced below.

		Hunter 2	
		stag	**hare**
Hunter 1	**stag**	$(2,2)$	$(0,1)$
	hare	$(1,0)$	$(1,1)$

In Theorem 8.5, we are in case 3. There are two symmetric pure-strategy strict Nash equilibria, (stag, stag) and (hare, hare). Both pure populations, all stag hunters and all hare hunters, are evolutionarily stable. There is also a symmetric mixed-strategy Nash equilibrium in which each player uses the strategy stag half the time and the strategy hare half the time. However, the corresponding population state is not evolutionarily stable.

8.2.3. Stag hunt variation. Suppose in the game of Stag Hunt, a hunter who hunts the stag without help from the other hunter has a $\frac{1}{4}$ chance of catching it. (Previously we assumed he had no chance of catching it.) Then the payoff matrix becomes

		Hunter 2	
		stag	**hare**
Hunter 1	**stag**	$(2,2)$	$(1,1)$
	hare	$(1,1)$	$(1,1)$

This game is not covered by Theorem 8.5, because $d = \pi_{22}$ and $b = \pi_{21}$ are equal. There are two symmetric pure-strategy Nash equilibria, (stag, stag) and (hare, hare). However, only (stag, stag) is a strict Nash equilibrium. Indeed, the strategy hare is now weakly dominated by the strategy stag. There are no mixed strategy Nash equilibria.

By Theorem 8.2, the pure population consisting of all stag hunters is evolutionarily stable.

What about the pure population consisting of all hare hunters? Since $\pi_{22} = \pi_{12} = 1$, when we check (2) of Theorem 8.1, among the strategies τ that must be checked is the pure strategy, hunt stags. However, $\pi_{21} = 1$ and $\pi_{11} = 2$, i.e., if some stag hunters invade the population of hare hunters, they do better against themselves than the hare-hunters do against them. Thus a pure population of hare hunters is not evolutionarily stable.

8.2.4. Hawks and Doves. Consider a population of animals that fight over food, territory, or mates. We will consider two possible strategies:

- Hawk (H): fight until either you are injured or your opponent retreats.
- Dove (D): display hostility, but if your opponent won't retreat, you retreat.

Let

- v = value of what you are fighting over.
- w = cost of injury.
- t = cost of protracted display.

We assume v, w, and t are all positive, and $v < w$. The payoff matrix is

		Animal 2	
		H	**D**
Animal 1	**H**	$(\frac{v-w}{2}, \frac{v-w}{2})$	$(v,0)$
	D	$(0,v)$	$(\frac{v}{2}-t, \frac{v}{2}-t)$

In Theorem 8.5, we are in case 4. Thus there are no symmetric pure-strategy Nash equilibria, and there is a mixed-strategy Nash equilibrium (σ,σ), $\sigma = pH + (1-p)D$; you can check that $p = \frac{v+2t}{w+2t}$. The population state σ evolutionarily stable.

For more information about Hawks and Doves, see Gintis, Sec. 3.10, and the Wikipedia page http://en.wikipedia.org/wiki/Hawk-dove_game.

8.3. Sex ratio

Recall the sex ration pseudo-game analyzed in Subsection 7.3.1. There was no actual game. There was, however, a situation very close to that considered in this chapter. The female population as a whole produced a fraction v of male offspring and a fraction $1 - v$ of female offspring. The number v can be regarded as the population state. An individual female produced a fraction u of male offspring and a fraction $1 - u$ of female offspring. The number u can be regarded as the type of an individual. The payoff to this individual $\pi_1(u, v)$ is her number of grandchildren. We derived the formula

$$\pi_1(u,v) = \sigma_f c^2 \left(1 + u\frac{1-2v}{v}\right).$$

In this situation, our analog of a Nash equilibrium was the pair $(\frac{1}{2}, \frac{1}{2})$, in the sense that if the population state was $\frac{1}{2}$ (i.e., females as a whole have $\frac{1}{2}$ male offspring), an individual could do no better than by choosing also to have $\frac{1}{2}$ male offspring.

Is this population state evolutionarily stable? We saw that for any individual type u, we have

$$\pi_1(\frac{1}{2}, \frac{1}{2}) = \pi_1(u, \frac{1}{2}) = \sigma_f c^2.$$

Thus we must check (2) of Theorem 8.1 for every u. We have

$$\begin{aligned}\pi_1(\frac{1}{2}, u) - \pi_1(u,u) &= \sigma_f c^2 \left(1 + \frac{1}{2}\frac{1-2u}{u}\right) - \sigma_f c^2 \left(1 + u\frac{1-2u}{u}\right)\\ &= \sigma_f c^2 (\frac{1}{2} - u)\frac{1-2u}{u} = \frac{2\sigma_f c^2}{u}(\frac{1}{2} - u)^2.\end{aligned}$$

Since this is positive for $u \neq \frac{1}{2}$, the population state $\frac{1}{2}$ is evolutionarily stable.

CHAPTER 9

Dynamical systems and differential equations

CHAPTER 10

Replicator dynamics

10.1. Replicator system

This section is related to Gintis, Secs. 9.2.1, 9.3, 9.5, and 9.6.

As in our study of evolutionary stability, we consider a symmetric two-person game in normal form with finite strategy set $S = \{s_1, \ldots, s_n\}$. There is a population that uses strategy s_1 with probability p_1, ..., strategy s_n with probability p_n; all $p_i \geq 0$ and $\sum p_i = 1$. The population state is $\sigma = \sum p_i s_i$.

When an individual of type i plays the game against a randomly chosen individual from a population with state σ, his expected payoff is $\pi_{i\sigma}$. When two randomly chosen individuals from a population with state σ play the game, the expected payoff to the first is $\pi_{\sigma\sigma}$.

In this chapter we will explicitly regard the population state σ as changing with time. Thus we will write $\sigma(t) = \sum p_i(t) s_i$.

It is reasonable to expect that if $\pi_{i\sigma} > \pi_{\sigma\sigma}$, then individuals using strategy i will in general have an above average number of offspring. Thus $p_i(t)$ should increase. On the other hand, if $\pi_{i\sigma} < \pi_{\sigma\sigma}$, we expect $p_i(t)$ to decrease.

In fact, it is reasonable to suppose that the growth rate of p_i is proportional to $\pi_{i\sigma} - \pi_{\sigma\sigma}$. For simplicity we will assume that the constant of proportionality for each population is one.

With these assumptions we obtain the replicator system:

$$\dot{p}_i = (\pi_{i\sigma} - \pi_{\sigma\sigma})p_i, \quad i = 1, \ldots, n. \tag{10.1}$$

The replicator system is a differential equation on $\mathbb{R}^n$. The physically relevant subset of $\mathbb{R}^n$ is the simplex

$$\Sigma = \{(p_1, \ldots, p_n) : \text{all } p_i \geq 0 \text{ and } \sum p_i = 1\}.$$

Σ can be decomposed as follows. For each nonempty subset I of $\{1, \ldots, n\}$, let

$$\Sigma_I = \{(p_1, \ldots, p_n) : p_i > 0 \text{ if } i \in I, p_i = 0 \text{ if } i \notin I, \text{ and } \sum p_i = 1\}.$$

Then Σ is the disjoint union of the Σ_I, where I ranges over all nonempty subsets of $\{1, \ldots, n\}$.

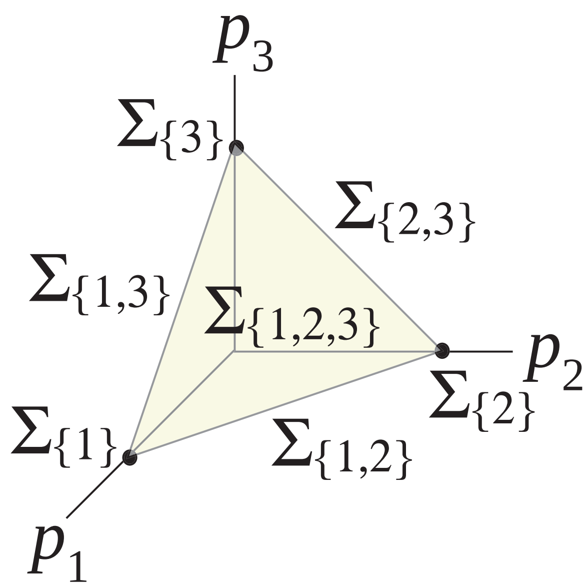

FIGURE 10.1. The set Σ with $n = 3$ and its decomposition.

Let $|I|$ denote the size of the set I. The dimension of the set Σ_I is $|I|-1$. For example, $\Sigma_{\{1,\ldots,n\}}$, the interior of Σ, has dimension $n-1$, and for each $i \in \{1,\ldots,n\}$, $\Sigma_{\{i\}}$ is a point (which has dimension 0).

THEOREM 10.1. *The replicator system has the following properties.*

(1) *If $p_i = 0$ then $\dot{p}_i = 0$.*
(2) *Let $S(p_1,\ldots,p_n) = \sum p_i$. If $S = 1$ then $\dot{S} = 0$.*
(3) *Each set Σ_I is invariant.*

PROOF. (1) follows immediately from (10.1).

To show (2), just note that if $\sum p_i = 1$, then

$$\dot{S} = \sum(\pi_{i\sigma} - \pi_{\sigma\sigma})p_i = \sum p_i\pi_{i\sigma} - \sum p_i\pi_{\sigma\sigma} = \pi_{\sigma\sigma} - \pi_{\sigma\sigma} = 0.$$

To prove (3), let

$$A_I = \{(p_1,\ldots,p_n) : p_i = 0 \text{ if } i \notin I \text{ and } \sum p_i = 1\}.$$

$\Sigma_I \subset A_I$, and from (1) and (2), A_I is invariant. Let $p(t)$ be a solution of the replicator system with $p(0) \in \Sigma_I$. Then $p(t)$ stays in A_I. Thus the only way $p(t)$ can leave Σ_I is if some $p_i(t)$, $i \in I$, becomes 0. If this happens, $p(t)$ enters A_J where J is some proper subset of I. Since A_J is invariant, this is impossible. $\square$

THEOREM 10.2. *Let $p \in \Sigma$. Then p is an equilibrium of the replicator system if and only if (σ, σ) satisfies condition (1) of the Fundamental Theorem of Nash Equilibria (Theorem 5.2).*

In other words, if $p \in \Sigma_I$, then p is an equilibrium of the replicator system if and only if all $\pi_{i\sigma}$ with $i \in I$ are equal.

For example, if only one strategy is active at σ (i.e., one $p_i = 1$ and the others are 0), then p is automatically an equilibrium of the replicator system. Actually, we already knew this: $\Sigma_{\{i\}}$ is a single point and is invariant, so it must be an equilibrium.

Note that if all $p_i > 0$, then condition (2) of Theorem 8.1 for a Nash equilibrium is irrelevant. Hence, if all $p_i > 0$, then p is an equilibrium of the replicator system if and only if (σ, σ) is a Nash equilibrium of the game.

PROOF. Let $p \in \Sigma_I$.

(1) Suppose p is an equilibrium of the replicator system. If $i \in I$, then $p_i > 0$, so we see from (10.1) that $\pi_{i\sigma} = \pi_{\sigma\sigma}$. Hence all $\pi_{i\sigma}$ with $i \in I$ are equal.

(2) Suppose $\pi_{i\sigma} = c$ for all $i \in I$, i.e., for all i such that $p_i > 0$. We have

$$\pi_{\sigma\sigma} = \sum_{i=1}^{n} p_i \pi_{i\sigma} = \sum_{i \in I} p_i \pi_{i\sigma} = \sum_{i \in I} p_i c = c.$$

Hence if $p_i > 0$ then $\pi_{i\sigma} = \pi_{\sigma\sigma}$. Now we see from (10.1) that p is an equilibrium of the replicator system. □

THEOREM 10.3. *Let $p^* \in \Sigma_I$ be an equilibrium of the replicator system. Suppose (σ^*, σ^*) does not satisfy condition (2) of Theorem 8.1, i.e., suppose there is an $i \notin I$ such that $\pi_{i\sigma^*} > \pi_{\sigma^*\sigma^*}$. Let J be a subset of $\{1, \ldots, n\}$ that includes both i and I. Then no solution $p(t)$ of the replicator system that lies in Σ_J approaches p^* as $t \to \infty$.*

PROOF. Let $\pi_{i\sigma^*} - \pi_{\sigma^*\sigma^*} = a > 0$. There is a number $\delta > 0$ such that if the distance from p to p^* is less than δ, then $\pi_{i\sigma} - \pi_{\sigma\sigma} > \frac{a}{2}$. If the distance from p to p^* is less than δ and $p_i > 0$, then $\dot{p}_i = (\pi_{i\sigma} - \pi_{\sigma\sigma}) p_i > \frac{a}{2} p_i > 0$. Therefore, if $p(t)$ is a solution of the replicator system that lies in Σ_J, so $p_i(t) > 0$, and stays within δ of p^*, then $p_i(t)$ is increasing. Since $p_i^* = 0$, clearly $p(t)$ does not approach p^* as t increases. □

THEOREM 10.4. *Let $p^* \in \Sigma$. Suppose σ^* is an evolutionarily stable state. Then p^* is an asymptotically stable equilibrium of the replicator system.*

The proof uses the following fact.

THEOREM 10.5. *Assume*

- $x_i > 0$ *for* $i = 1, \ldots, n$.
- $p_i > 0$ *for* $i = 1, \ldots, n$.
- $\sum p_i = 1$.

Then $\ln(\sum p_i x_i) > \sum p_i \ln x_i$ *unless* $x_1 = \cdots = x_n$.

Given this fact, we shall prove Theorem 10.4 assuming all $p_i^* > 0$. Define a function W with domain $\Sigma_{\{1,\ldots,n\}}$ by $W(p_1, \ldots, p_n) = \sum p_i^* \ln \frac{p_i}{p_i^*}$. Then $W(p^*) = 0$. For $p \neq p^*$,

$$W(p) = \sum p_i^* \ln \frac{p_i}{p_i^*} < \ln(\sum p_i^* \frac{p_i}{p_i^*}) = \ln(\sum p_i) = \ln 1 = 0.$$

Let $V = -W$. Then $V(p^*) = 0$, and, for $p \neq p^*$, $V(p) > 0$. We can write $V(p) = -\sum p_i^*(\ln p_i - \ln p_i^*)$. Then

$$\begin{aligned}\dot{V} &= -\sum p_i^* \frac{1}{p_i} \dot{p}_i = -\sum \frac{p_i^*}{p_i}(\pi_{i\sigma} - \pi_{\sigma\sigma})p_i = -\sum p_i^*(\pi_{i\sigma} - \pi_{\sigma\sigma}) \\ &= -\sum p_i^* \pi_{i\sigma} + \sum p_i^* \pi_{\sigma\sigma} = -\pi_{\sigma^*\sigma} + \pi_{\sigma\sigma} < 0.\end{aligned}$$

Therefore V is a strict Liapunov function, so p^* is asymptotically stable.

10.2. Studying the replicator system in practice

Let

$$D = \{(p_1, \ldots, p_{n-1}) : p_i \geq 0 \text{ for } i = 1, \ldots, n-1, \text{ and } \sum_{i=1}^{n-1} p_i \leq 1\}.$$

Then

$$\Sigma = \{(p_1, \ldots, p_n) : (p_1, \ldots, p_{n-1}) \in D \text{ and } p_n = 1 - \sum_{i=1}^{n-1} p_i\}.$$

Instead of studying the replicator system on Σ, one can instead take the space to be D and use only the first $n-1$ equations of the replicator system. In these equations, one must of course let $p_n = 1 - \sum_{i=1}^{n-1} p_i$.

For $n = 2$, the set D is simply the line segment $0 \leq p_1 \leq 1$. The endpoints $p_1 = 0$ and $p_1 = 1$ are always equilibria. Hence the differential equation always has p_1 and $1 - p_1$ as factors.

For $n = 3$, the set D is the triangle $\{(p_1, p_2) : p_1 \geq 0, p_2 \geq 0, \text{ and } p_1 + p_2 \leq 1\}$. The vertices $(0, 0)$, $(1, 0)$, and $(0, 1)$ are equilibria, and the lines $p_1 = 0$, $p_2 = 0$, and $p_1 + p_2 = 1$ are invariant. Therefore $\dot{p}_1$ has p_1 as a factor, $\dot{p}_2$ has p_2 as a factor, and $\dot{p}_1 + \dot{p}_2$ has $1 - (p_1 + p_2)$ as a factor.

10.3. Microsoft vs. Apple

In the early days of personal computing, people faced the following dilemma. You could buy a computer running Microsoft Windows, or one running the Apple operating system. Either was reasonably satisfactory, although Apple's was better. However, neither type of computer dealt well with files produced by the other. Thus if your coworker used Windows and you used Apple, not much got accomplished.

We model this dilemma as a symmetric two-person game in normal form. The strategies are buy Microsoft (M) or buy Apple (A). The payoffs are given by the following matrix.

		Player 2	
		M	**A**
Player 1	**M**	$(1,1)$	$(0,0)$
	A	$(0,0)$	$(2,2)$

There are two pure-strategy strict Nash equilibria, (M,M) and (A,A). Both M and A are evolutionarily stable states by Theorem 8.2. Hence both correspond to attractors of the replicator system. There is also a symmetric mixed strategy Nash equilibrium (σ^*,σ^*) with $\sigma^* = \frac{2}{3}M + \frac{1}{3}E$.

The Nash equilibria (M,M) and (A,A) are easy to understand intuitively. Clearly if your coworker is using Microsoft, you should use it too. Since, for each player, Microsoft is the best response to Microsoft, (M,M) is a Nash equilibrium. The same reasoning applies to (A,A).

The mixed strategy Nash equilibrium is harder to understand. Certainly one can calculate that if your opponent uses the mixed strategy $\sigma^* = \frac{2}{3}M + \frac{1}{3}E$, any mixed strategy you use will give you the same payoff, so σ^* itself is among the best responses to σ^*. Since this calculation applies to both players, (σ^*,σ^*) is a Nash equilibrium. Nevertheless, one feels intuitively that this Nash equilibrium does not correspond to any behavior one would ever observe. Even if for some reason people picked computers randomly, why would they choose the worse computer with higher probability?

To resolve this mystery, we imagine a large population of people who randomly encounter each other and play this two-player game. People observe which strategy, buy Microsoft or buy Apple, is on average producing higher payoffs. They will tend to use the strategy that they observe produces the higher payoff.

Let a state of the population be $\sigma = p_1M + p_2A$, and let $p = (p_1,p_2)$. It is consistent with our understanding of the situation to assume that $p(t)$ evolves by the replicator system. Then we have

$$\pi_{1\sigma} = p_1, \quad \pi_{2\sigma} = 2p_2, \quad \pi_{\sigma\sigma} = p_1\pi_{1\sigma} + p_2\pi_{2\sigma} = p_1^2 + 2p_2^2,$$

so the replicator system is

$$\begin{aligned}\dot{p}_1 &= (\pi_{1\sigma} - \pi_{\sigma\sigma})p_1 &= (p_1 - (p_1^2 + 2p_2^2))p_1, \\ \dot{p}_2 &= (\pi_{2\sigma} - \pi_{\sigma\sigma})p_2 &= (2p_2 - (p_1^2 + 2p_2^2))p_2.\end{aligned}$$

We only need the first equation, in which we substitute $p_2 = 1 - p_1$:

$$\dot{p}_1 = (p_1 - (p_1^2 + 2(1-p_1)^2))p_1 = (1-p_1)(p_1 - 2(1-p_1)) = (1-p_1)(3p_1-2)p_1.$$

The phase portrait on the interval $0 \le p_1 \le 1$ is shown below.

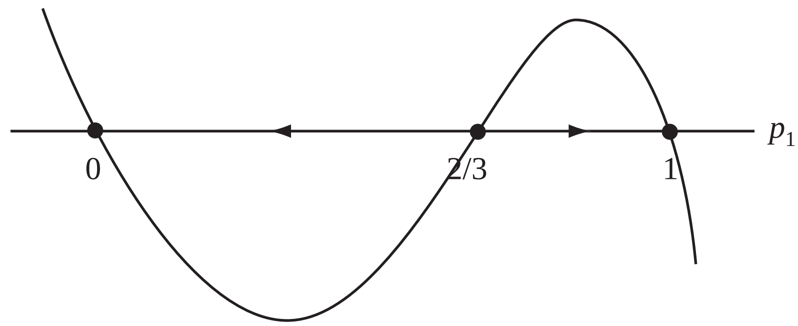

FIGURE 10.2. Graph of $\dot{p}_1 = (1-p_1)(3p_1-2)p_1$ and phase portrait.

We see that there are attracting equilibria at $p_1 = 0$ (everyone uses Apple) and $p_1 = 1$ (everyone uses Microsoft) as expected. The equilibrium at $p_1 = \frac{2}{3}$ is unstable. It separates initial conditions that tend to $p_1 = 0$ and $p_1 = 1$ as $t \to \infty$. The location of this equilibrium now makes intuitive sense: a larger range of initial conditions produces the eventual outcome $p_1 = 0$, in which everyone uses the better computer Apple, than produces the eventual outcome $p_1 = 1$, in which everyone uses the worse computer Microsoft. Nevertheless, if initially more than $\frac{2}{3}$ of the population uses Microsoft, eventually everyone uses Microsoft, even though it is worse.

10.4. Hawks and doves

This section is related to Gintis, Sec. 9.2.2.

We consider again the game of Hawks and Doves from Section 8.2.4. The payoff matrix is reproduced below.

		Animal 2	
		H	**D**
Animal 1	**H**	$(\frac{v-w}{2}, \frac{v-w}{2})$	$(v, 0)$
	D	$(0, v)$	$(\frac{v}{2} - t, \frac{v}{2} - t)$

We recall that there are no pure-strategy Nash equilibria; that if $p^* = \frac{v+2t}{w+2t}$ and $\sigma^* = p^*H + (1-p^*)V$, then (σ^*, σ^*) is a symmetric mixed-strategy Nash equilibrium; and that σ^* is an evolutionarily stable state.

We have

$$\pi_{1\sigma} = p_1\frac{v-w}{2} + p_2 v, \quad \pi_{2\sigma} = p_2(\frac{v}{2} - t),$$

$$\pi_{\sigma\sigma} = p_1\pi_{1\sigma} + p_2\pi_{2\sigma} = p_1^2\frac{v-w}{2} + p_1p_2v + p_2^2(\frac{v}{2} - t).$$

Hence the replicator system is

$$\dot{p}_1 = (\pi_{1\sigma} - \pi_{\sigma\sigma})p_1 = (p_1\frac{v-w}{2} + p_2v - (p_1^2\frac{v-w}{2} + p_1p_2v + p_2^2(\frac{v}{2} - t)))p_1,$$

$$\dot{p}_2 = (\pi_{2\sigma} - \pi_{\sigma\sigma})p_2 = (p_2(\frac{v}{2} - t) - (p_1^2\frac{v-w}{2} + p_1p_2v + p_2^2(\frac{v}{2} - t)))p_2.$$

Again we only need the first equation, in which we substitute $p_2 = 1 - p_1$:

$$\begin{aligned}\dot{p}_1 &= (p_1\frac{v-w}{2} + (1-p_1)v - (p_1^2\frac{v-w}{2} + p_1(1-p_1)v + (1-p_1)^2(\frac{v}{2} - t)))p_1 \\ &= (1-p_1)(p_1\frac{v-w}{2} + (1-p_1)v - (1-p_1)(\frac{v}{2} - t))p_1 \\ &= (1-p_1)(\frac{v}{2} + t - p_1(\frac{w}{2} + t))p_1 = \frac{1}{2}(1-p_1)(v + 2t - p_1(w+2t))p_1\end{aligned}$$

The phase portrait on the interval $0 \le p_1 \le 1$ is shown below. In this case, if the population is initially anything other than $p_1 = 0$ or $p_1 = 1$, the population tends toward the state σ^*.

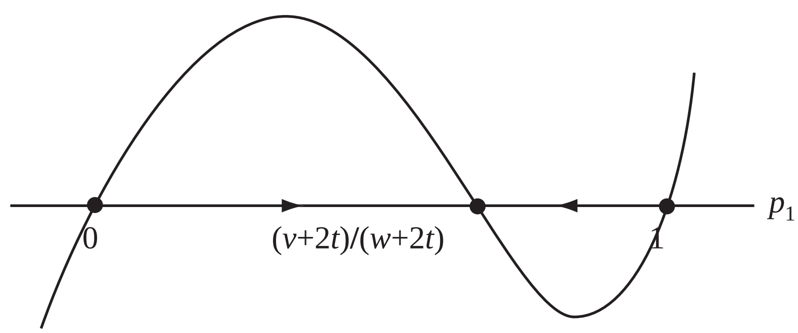

FIGURE 10.3. Graph of $\dot{p}_1 = \frac{1}{2}(1-p_1)(v+2t-p_1(w+2t))p_1$ and phase portrait.

10.5. Orange-throat, blue-throat, and yellow-striped lizards

This section is related to Gintis, Secs. 4.16 and 9.12.

The side-blotched lizard, which lives in California, has three types of males:

- **Orange-throats** are very aggressive, keep large harems of females (up to seven), and defend large territories.
- **Blue-throats** are less aggressive, keep small harems of females (about three), and defend small territories.

- **Yellow-stripes** are docile and resemble females. They do not keep a harem or defend a territory. Instead they sneak into other males' territories and mate with their females.

Field reports indicate that populations of side-blotched lizards cycle: mostly orange-throats one generation, mostly yellow-stripes the next, mostly blue-throats the next, then back to mostly orange-throats.

Let's consider a competition between two different types of male side-blotched lizards:

- Orange-throats vs. Yellow-stripes. The orange-throats are unable to defend their large territories against the sneaky yellow-stripes. The yellow-stripes have the advantage.
- Yellow-stripes vs. blue-throats. The blue-throats are able to defend their small territories against the yellow-stripes. The blue-throats have the advantage.
- Blue-throats vs. orange-throats. Neither type of male bothers the other. The orange-throats, with their larger harems, produce more offspring.

This simple analysis shows why the population should cycle.

To be a little more precise, we'll consider a game in which the players are two male side-blotched lizards. Each has three possible strategies: orange-throat (O), yellow-striped (Y), and blue-throat (B). The payoffs are 0 if both use the same strategies; otherwise we assign a payoff of 1 or -1 to the lizard that does or does not have the advantage according to our analysis. The payoff matrix is therefore:

		Lizard 2		
		O	**Y**	**B**
	O	$(0,0$	$(-1,1)$	$(1,-1)$
Lizard 1	**Y**	$(1,-1)$	$(0,0)$	$(-1,1)$
	B	$(-1,1)$	$(1,-1)$	$(0,0)$

This game is symmetric. You can check that there is no pure-strategy Nash equilibrium. There is one mixed strategy Nash equilibrium (σ^*, σ^*) with $\sigma^* = \frac{1}{3}O + \frac{1}{3}Y + \frac{1}{3}B$. (This game is just rock-paper-scissors in disguise.) However, this Nash equilibrium is not evolutionarily stable.

Let's calculate the replicator system. Let $\sigma = p_1 O + p_2 Y + p_3 B$. Then

$$\pi_{1\sigma} = -p_2 + p_3, \quad \pi_{2\sigma} = p_1 - p_3, \quad \pi_{3\sigma} = -p_1 + p_2.$$

Therefore

$$\pi_{\sigma\sigma} = p_1\pi_{1\sigma} + p_2\pi_{2\sigma} + p_3\pi_{3\sigma} = p_1(-p_2 + p_3) + p_2(p_1 - p_3) + p_3(-p_1 + p_2) = 0.$$

Hence the replicator system is

$$\begin{aligned}
\dot{p}_1 &= (\pi_{1\sigma} - \pi_{\sigma\sigma})p_1 &= (-p_2 + p_3)p_1,\\
\dot{p}_2 &= (\pi_{2\sigma} - \pi_{\sigma\sigma})p_2 &= (p_1 - p_3)p_2,\\
\dot{p}_3 &= (\pi_{3\sigma} - \pi_{\sigma\sigma})p_1 &= (-p_1 + p_2)p_3.
\end{aligned}$$

We only need the first and second equations, in which we substitute $p_3 = 1-(p_1+p_2)$:

$$\dot{p}_1 = (1 - p_1 - 2p_2)p_1, \tag{10.2}$$

$$\dot{p}_2 = (-1 + 2p_1 + p_2)p_2. \tag{10.3}$$

The simplex Σ in $\mathbb{R}^3$ corresponds to the region $p_1 \geq 0$, $p_2 \geq 0$, $p_1 + p_2 \leq 1$ in $\mathbb{R}^2$.

Let's analyze the system (10.2)–(10.3) on Σ.

1. Invariance of the boundary of Σ. This is just a check on our work. Note that if $p_1 = 0$ then $\dot{p}_1 = 0$; if $p_2 = 0$ then $\dot{p}_2 = 0$; and if $p_1 + p_2 = 1$ then

$$\begin{aligned}
\dot{p}_1 + \dot{p}_2 &= (1 - p_1)p_1 - 2p_2p_1 + (-1 + p_2)p_2 + 2p_1p_2\\
&= (1 - p_1)p_1 + (-1 + p_2)p_2 = p_2p_1 - p_1p_2 = 0.
\end{aligned}$$

2. To find all equilibria of the replicator system, we solve simultaneously the pair of equations

$$\dot{p}_1 = (1 - p_1 - 2p_2)p_1 = 0, \quad \dot{p}_2 = (-1 + 2p_1 + p_2)p_2 = 0.$$

We find that the equilibria are $(p_1, p_2) = (0, 0),\ (0, 1),\ (1, 0),$ and $(\frac{1}{3}, \frac{1}{3})$.

3. Nullclines: We have $\dot{p}_1 = 0$ on the lines $1 - p_1 - 2p_2 = 0$ and $p_1 = 0$, and we have $\dot{p}_2 = 0$ on the lines $-1 + 2p_1 + p_2 = 0$ and $p_2 = 0$. See Figure 10.4

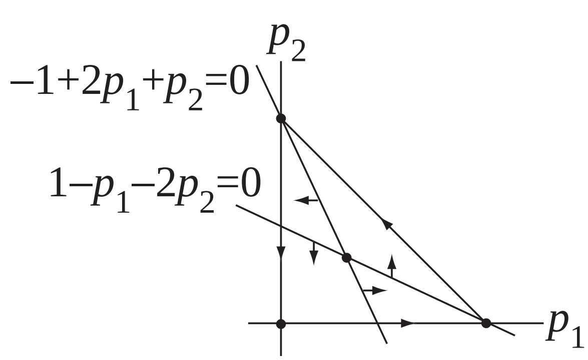

FIGURE 10.4. Vector field for the system (10.2)–(10.3) on Σ.

4. From the figure it appears that solutions circle around the equilibrium $(\frac{1}{3}, \frac{1}{3})$. We cannot, however, tell from the figure if solutions spiral toward the equilibrium, spiral away from the equilibrium, or form closed curves.

5. We can try to get more information by linearizing the system (10.2)–(10.3) at the equilibrium $(\frac{1}{3}, \frac{1}{3})$. The linearization of (10.2)–(10.3) has the matrix

$$\begin{pmatrix} 1 - 2p_1 - 2p_2 & -2p_1 \\ 2p_2 & -1 + 2p_1 + 2p_2 \end{pmatrix}.$$

At $(p_1, p_2) = (\frac{1}{3}, \frac{1}{3})$, the matrix is

$$\begin{pmatrix} -\frac{1}{3} & -\frac{2}{3} \\ \frac{2}{3} & \frac{1}{3} \end{pmatrix}.$$

The characteristic equation is $\lambda^2 + \frac{1}{3} = 0$, so the eigenvalues are $\pm\frac{1}{\sqrt{3}}i$. Since they are pure imaginary, we still don't know what the solutions do.

6. Fortunately, the system (10.2)–(10.3) has the first integral $V(p_1, p_2) = \ln p_1 + \ln p_2 + \ln(1 - p_1 - p_2)$. Check:

$$\begin{aligned} \dot{V} &= \frac{\dot{p}_1}{p_1} + \frac{\dot{p}_2}{p_2} + \frac{-\dot{p}_1 - \dot{p}_2}{1 - p_1 - p_2} = \frac{\dot{p}_1}{p_1} + \frac{\dot{p}_2}{p_2} - \frac{\dot{p}_1 + \dot{p}_2}{1 - p_1 - p_2} \\ &= \frac{(1 - p_1 - 2p_2)p_1}{p_1} + \frac{(-1 + 2p_1 + p_2)p_2}{p_2} - \frac{(1 - p_1)p_1 + (-1 + p_2)p_2}{1 - p_1 - p_2} \\ &= (1 - p_1 - 2p_2) + (-1 + 2p_1 + p_2) - \frac{(1 - p_1 - p_2)(p_1 - p_2)}{1 - p_1 - p_2} \\ &= (1 - p_1 - 2p_2) + (-1 + 2p_1 + p_2) - (p_1 - p_2) = 0. \end{aligned}$$

The level curves of $V(p_1, p_2)$ surround the point $(\frac{1}{3}, \frac{1}{3})$. Thus solutions in the interior of Σ form closed curves around the equilibrium $(\frac{1}{3}, \frac{1}{3})$.

We conclude that we expect the populations of the different types of side-blotched lizard to oscillate.

10.6. Dominated strategies and the replicator system

This section is related to Section 9.5 in Gintis. We prove two results relating iterated elimination of dominated strategies to the replicator system. Since the games we consider have two players and are symmetric, whenever we eliminate a strategy, we shall eliminate it for both players.

THEOREM 10.6. *In a two-player symmetric game, suppose strategy s_i is strictly dominated by strategy s_j. Let I be a subset of $\{1, \ldots, n\}$ that contains both i and j, and let $p(t)$ be a solution of the replicator system in Σ_I. Then $p_i(t) \to 0$ as $t \to \infty$.*

PROOF. Since strategy s_i is strictly dominated by strategy s_j, we have that for every pure strategy s_k, $\pi_{ik} < \pi_{jk}$. Then for any population state $\sigma = \sum p_k s_k$, we have

$$\pi_{i\sigma} = \pi_{i,\sum p_k s_k} = \sum p_k \pi_{ik} < \sum p_k \pi_{jk} = \pi_{j,\sum p_k s_k} = \pi_{j\sigma}.$$

Therefore, for each $p \in \Sigma$, if σ is the corresponding population state, then $\pi_{i\sigma} - \pi_{j\sigma} < 0$.

Now $\pi_{i\sigma} - \pi_{j\sigma}$ depends continuously on p, and Σ is a compact set (closed and bounded). Therefore there is a number $\epsilon > 0$ such that $\pi_{i\sigma} - \pi_{j\sigma} \leq -\epsilon$ for every $p \in \Sigma$.

Let $p(t)$ be a solution of the replicator system in Σ_I. Then $p_i(t) > 0$ and $p_j(t) > 0$ for all t. Therefore we can define the function $V(t) = \ln p_i(t) - \ln p_j(t)$. We have

$$\dot{V} = \frac{1}{p_i}\dot{p}_i - \frac{1}{p_j}\dot{p}_j = \frac{1}{p_i}(\pi_{i\sigma} - \pi_{\sigma\sigma})p_i - \frac{1}{p_j}(\pi_{j\sigma} - \pi_{\sigma\sigma})p_j = \pi_{i\sigma} - \pi_{j\sigma} \leq -\epsilon.$$

Then for $t > 0$,

$$V(t) - V(0) = \int_0^t \dot{V}\,dt \leq \int_0^t -\epsilon\,dt = -\epsilon t.$$

Therefore $V(t) \to -\infty$ as $t \to \infty$.

But $0 < p_j(t) < 1$, so $\ln p_j(t) > 0$, so $-\ln p_j(t) < 0$. Since $V(t) = \ln p_i(t) - \ln p_j(t)$ approaches $-\infty$ and $-\ln p_j(t)$ is positive, it must be that $\ln p_i(t)$ approaches $-\infty$. But then $p_i(t)$ approaches 0. □

Theorem 10.7. *In a two-player symmetric game, suppose that when we do iterated elimination of strictly dominated strategies, the strategy s_k is eliminated at some point. Let $p(t)$ be a solution of the replicator system in $\Sigma_{\{1,\dots,n\}}$. Then $p_k(t) \to 0$ as $t \to \infty$.*

Proof. For simplicity, we will assume that only one strategy is eliminated before s_k. Let that strategy be s_i, eliminated because it is strictly dominated by a strategy s_j. Then s_k is strictly dominated by some strategy s_l *once s_i is eliminated*. This means that $\pi_{km} < \pi_{lm}$ for every m *other than i*.

Let $\tilde{\Sigma}$ denote the subset of Σ on which $p_i = 0$. Then for every $p \in \tilde{\Sigma}$, $\pi_{k\sigma} < \pi_{l\sigma}$. Since $\tilde{\Sigma}$ is compact, there is a number $\epsilon > 0$ such that $\pi_{k\sigma} - \pi_{l\sigma} \leq -\epsilon$ for all $p \in \tilde{\Sigma}$.

By continuity, there is a number $\delta > 0$ such that if $p \in \Sigma$ and $0 \leq p_i < \delta$, then $\pi_{k\sigma} - \pi_{l\sigma} \leq -\frac{\epsilon}{2}$.

Let $p(t)$ be a solution of the replicator system in $\Sigma_{\{1,\dots,n\}}$. By the previous theorem, $p_i(t) \to 0$ as $t \to \infty$. Therefore, for t greater than or equal to some t_0, $0 < p_i(t) < \delta$. Let $V(t) = \ln p_k(t) - \ln p_l(t)$. Then for $t \geq t_0$,

$$\dot{V} = \frac{1}{p_k}\dot{p}_k - \frac{1}{p_l}\dot{p}_l = \frac{1}{p_k}(\pi_{k\sigma} - \pi_{\sigma\sigma})p_k - \frac{1}{p_l}(\pi_{l\sigma} - \pi_{\sigma\sigma})p_l = \pi_{k\sigma} - \pi_{l\sigma} \leq -\frac{\epsilon}{2}.$$

Therefore, for $t > t_0$,

$$V(t) - V(t_0) = \int_{t_0}^t \dot{V}\,dt \leq \int_{t_0}^t -\frac{\epsilon}{2}\,dt = -\frac{\epsilon}{2}t.$$

Therefore $V(t) \to -\infty$ as $t \to \infty$.

As in the proof of the previous theorem, we can conclude that $p_k(t) \to 0$ as $t \to \infty$. □

10.7. Asymmetric evolutionary games

This section is related to Gintis, Sec. 9.15.

Consider an asymmetric two-person game G in normal form. Player 1 has the finite strategy set $S = \{s_1, \ldots, s_n\}$. Player 2 has the finite strategy set $T = \{t_1, \ldots, t_m\}$. If player 1 uses the pure strategy s_i and player 2 uses the pure strategy t_j, the payoff to player 1 is α_{ij}, and the payoff to player 2 is β_{ij}.

Suppose there are two populations, one consisting of individuals like player 1, the other consisting of individuals like player 2. When an individual from the first population encounters an individual from the second population, they play the game G.

Taken as a whole, the first population uses strategy s_1 with probability p_1, $\ldots$, strategy s_n with probability p_n; all $p_i \geq 0$ and $\sum p_i = 1$. Let $p = (p_1, \ldots, p_n)$, the state of the first population. Similarly, taken as a whole, the second population uses strategy t_1 with probability q_1, $\ldots$, strategy t_m with probability q_m; all $q_j \geq 0$ and $\sum q_j = 1$. Let $q = (q_1, \ldots, q_m)$, the state of the second population.

When an individual of type i from the first population plays the game against a randomly chosen individual from the second population, whose state is q, his expected payoff is

$$\alpha_i(q) = \sum_{j=1}^{m} q_j \alpha_{ij}.$$

Similarly, when an individual of type j from the second population plays the game against a randomly chosen individual from the first population, whose state is p, his expected payoff is

$$\beta_j(p) = \sum_{i=1}^{n} p_i \beta_{ij}.$$

When two randomly chosen individuals from the two populations play the game, the expected payoff to the first is

$$\alpha(p, q) = \sum_{i=1}^{n} p_i \alpha_i(q) = \sum_{i=1}^{n} \sum_{j=1}^{m} p_i q_j \alpha_{ij}.$$

Similarly, the expected payoff to the second is

$$\beta(p, q) = \sum_{j=1}^{m} q_j \beta_j(p) = \sum_{j=1}^{m} \sum_{i=1}^{n} q_j p_i \beta_{ij} = \sum_{i=1}^{n} \sum_{j=1}^{m} p_i q_j \beta_{ij}.$$

We combine the two population states p and q into a total population state (p, q), and we will regard (p, q) as changing with time. Reasoning as in Section 10.1, we obtain the replicator system:

$$\dot{p}_i = (\alpha_i(q) - \alpha(p, q))p_i, \quad i = 1, \dots, n; \tag{10.4}$$

$$\dot{q}_j = (\beta_j(p) - \beta(p, q))q_j, \quad j = 1, \dots, m. \tag{10.5}$$

Let

$$\Sigma_n = \{(p_1, \dots, p_n) : \text{all } p_i \geq 0 \text{ and } \sum p_i = 1\},$$
$$\Sigma_m = \{(q_1, \dots, q_m) : \text{all } q_j \geq 0 \text{ and } \sum q_j = 1\},$$

The system (10.4)–(10.5) should be considered on $\Sigma_n \times \Sigma_m$.

Let

$$D_{n-1} = \{(p_1, \dots, p_{n-1}) : p_i \geq 0 \text{ for } i = 1, \dots, n-1, \text{ and } \sum_{i=1}^{n-1} p_i \leq 1\},$$

$$D_{m-1} = \{(q_1, \dots, q_{m-1}) : q_j \geq 0 \text{ for } j = 1, \dots, m-1, \text{ and } \sum_{j=1}^{m-1} q_j \leq 1\}.$$

Instead of studying an asymmetric replicator system on $\Sigma_n \times \Sigma_m$, one can instead take the space to be $D_{n-1} \times D_{m-1}$, and use only the differential equations for $\dot{p}_1, \dots, \dot{p}_{n-1}$ and $\dot{q}_1, \dots, \dot{q}_{m-1}$. In these equations, one must of course let $p_n = 1 - \sum_{i=1}^{n-1} p_i$ and $q_m = 1 - \sum_{j=1}^{m-1} q_j$

A total population state (p, q) can be regarded as a pair of mixed strategies for players 1 and 2. Then the expected payoffs for players 1 and 2 are $\alpha(p, q)$ and $\beta(p, q)$ respectively. Thus a population state (p^*, q^*) is a Nash equilibrium provided

$$\alpha(p^*, q^*) \geq \alpha(p, q^*) \text{ for all } p \in \Sigma_n,$$
$$\beta(p^*, q^*) \geq \beta(p^*, q) \text{ for all } q \in \Sigma_m.$$

However, the notion of evolutionarily stable state for symmetric games does not have an analogue for asymmetric games, since individuals from the same population cannot play the game against each other.

Many results about the replicator system for symmetric games also hold for the replicator system for asymmetric games:

(1) A population state (p, q) is an equilibrium of the replicator system if and only if it satisfies the equality conditions for a Nash equilibrium .
(2) A point on the boundary of $\Sigma_n \times \Sigma_m$ that satisfies the equality conditions for a Nash equilibrium, but does not satisfy one of the inequality conditions, attracts no solution in in which the strategy corresponding to the unsatisfied inequality condition is present.

(3) If a strategy is eliminated in the course of iterated elimination of strictly dominated strategies, then for any solution in the interior of $\Sigma_n \times \Sigma_m$, that strategy dies out.

An important difference, however, is that for asymmetric replicator systems, it is known that equilibria in the interior of $\Sigma_n \times \Sigma_m$ are never asymptotically stable.

10.8. Big Monkey and Little Monkey 7

As in Subsection **??**, suppose Big Monkey and Little Monkey decide simultaneously whether to wait or climb. We have a game in normal form with the following payoff matrix, repeated from Subsection **??**, except that I have changed the order of climb and wait:

			Little Monkey	
			q_1 **climb**	q_2 **wait**
Big Monkey	p_1	**climb**	$(5,3)$	$(4,4)$
	p_2	**wait**	$(9,1)$	$(0,0)$

We now imagine a population of Big Monkeys and a population of Little Monkeys; (p_1,p_2) and (q_1,q_2) represent the two population states, so $p_2 = 1-p_1$ and $q_2 = 1-q_1$. The monkeys randomly encounter a monkey of the other type and play the game.

We could write differential equations for $\dot{p}_1$, $\dot{p}_2$, $\dot{q}_1$, and $\dot{q}_2$, but we only need those for $\dot{p}_1$ and $\dot{q}_1$, so I will omit the other two. Using $p_2 = 1-p_1$ and $q_2 = 1-q_1$, we obtain

$$\begin{aligned}
\dot{p}_1 &= (\alpha_1(q) - \alpha(p,q))p_1 \\
&= ((5q_1 + 4(1-q_1)) - (5p_1q_1 + 4p_1(1-q_1) + 9(1-p_1)q_1))p_1 \\
&= p_1(1-p_1)(5q_1 + 4(1-q_1) - 9q_1) = p_1(1-p_1)(4-8q_1) \\
\dot{q}_1 &= (\beta_1(p) - \beta(p,q))q_1 \\
&= ((3p_1 + 1(1-p_1)) - (3p_1q_1 + 4p_1(1-q_1) + 1(1-p_1)q_1))q_1 \\
&= q_1(1-q_1)(3p_1 + (1-p_1) - 4p_1) = q_1(1-q_1)(1-2p_1)
\end{aligned}$$

We consider this system on

$$D_1 \times D_1 = \{(p_1,q_1) : 0 \le p_1 \le 1 \text{ and } 0 \le q_1 \le 1\}.$$

See Figure 10.5

1. Invariance of the boundary of $D_1 \times D_1$. This is just a check on our work. Note that if $p_1 = 0$ or $p_1 = 1$, then $\dot{p}_1 = 0$; and if $q_1 = 0$ or $q_1 = 1$, then $\dot{q}_1 = 0$.

2. To find all equilibria of the replicator system, we solve simultaneously the pair of equations $\dot{p}_1 = 0$ and $\dot{q}_1 = 0$. We find that the equilibria are $(p_1,q_1) = (0,0),\ (0,1),\ (1,0),\ (1,0),$ and $(\frac{1}{2},\frac{1}{2})$.

3. Nullclines: We have $\dot{p}_1 = 0$ on the lines $p_1 = 0$, $p_1 = 1$, and $q_1 = \frac{1}{2}$; and we have $\dot{q}_1 = 0$ on the lines $q_1 = 0$, $q_1 = 1$, and $p_1 = \frac{1}{2}$ See Figure 10.5.

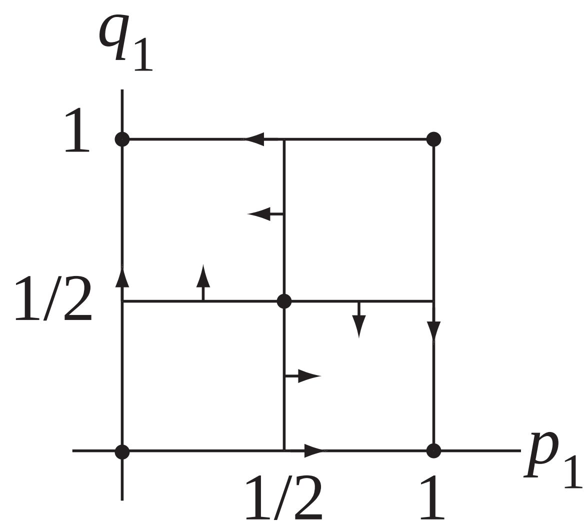

FIGURE 10.5. Vector field for the evolving monkeys.

4. It appears that the corner equilibria are attractors or repellers and the interior equilibrium is a saddle. This is correct and can be checked by linearization. The phase portrait is given in Figure 10.6.

The stable manifold of the saddle divides its complement into two sets. Initial conditions in one are attracted to the point $(0,1)$, where Big Monkey waits and Little Monkey climbs. Initial conditions in the other are attracted to the point $(1,0)$, where Big Monkey climbs and Little Monkey waits. Thus we expect to observe the population in one or the other of these pure states, but we can't guess which without knowing where the population started.

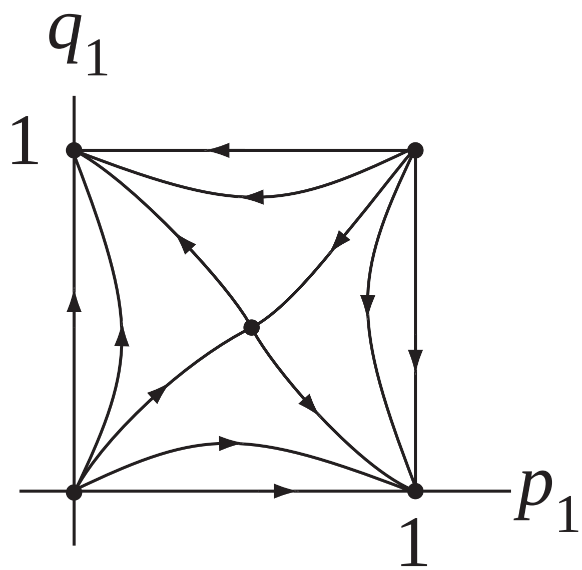

FIGURE 10.6. Phase portrait for the evolving monkeys.

CHAPTER 11

Trust, reciprocity, and altruism