

Chapter 1

Backward induction

This chapter deals with situations in which two or more opponents take actions one after the other. If you are involved in such a situation, you can try to think ahead to how your opponent might respond to each of your possible actions, bearing in mind that he is trying to achieve his own objectives, not yours. However, we shall see that it may not be helpful to carry this idea too far.

1.1 Tony's Accident

When one of us (Steve) was a college student, his friend Tony caused a minor traffic accident. We'll let Steve tell the story:

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 the game tree in Figure 1.1. For definiteness, we'll assume that the value to Vic of repairing the damage is \$20.

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.




Figure 1.1. Tony's Accident. Tony's payoff is given first.

- (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 of wine 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 became the one in Figure 1.2.




Figure 1.2. Tony's Accident: second game tree. Tony's payoff is given first.

Most of the new 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

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 also on the choices of others.

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 of them in this case.
- (3) A set B of *actions* or *moves*. In Figure 1.2, the moves are the lines (seven in this case). Each move connects two nodes, one its *start* and the other 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 i th 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.

Here are two consequences of assumption (b):

1. Each node other than the root node is the end of exactly one move. (Proof: Let c be a node that is not the root node. It is the end of *at least* one move, because there is a path from the root node to c . If c were the end of two moves m_1 and m_2 , then there would be two paths from the root node to c : one from the root node to the start of m_1 , followed by m_1 ; the other from the root node to the start of m_2 , followed by m_2 . But this can’t happen because of assumption (b).)

2. *Every* complete path, not just those of finite length, starts at the root node. (If c is any node other than the root node, there is exactly one path p from the root node to c . If a path that contains c is complete, it must contain p .)

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 this book we only discuss finite horizon games. An *infinite horizon game* is one with arbitrarily long paths. We discuss these games in Chapter 6.

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 with 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 in a finite game is always one less than the number of nodes.) Such a game is necessarily a finite horizon game.

Games in extensive form with complete information are good models of situations in which players act one after the other, players understand the situation completely, and nothing depends on chance. In Tony's Accident it was important that Tony knew Vic's payoffs, at least approximately, or he would not have been able to choose what to do.

1.3 Strategies

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 "every situation that the player might encounter" is interpreted to mean every node that is 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 k th 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 k th 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 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, $(\pi_1(s_3, t_2), \pi_2(s_3, t_2)) = (-40, 40)$. Of course, in Figure 1.2, this is the pair of payoffs associated with the terminal node 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

Game theorists often assume that players are *rational*. For a game in extensive form with complete information, rationality is usually considered to imply the following:

- 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. (Game theory discourages wishful thinking.) 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 . *In the rest of this chapter, and until Chapter 6, we shall only deal with situations in which 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 to node c the payoffs to all players that result from this choice, and we “prune the game tree.” This helps us take this choice into account when finding the moves players should 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 move 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 pruned game tree of Figure 1.3.




Figure 1.3. Tony's Accident: pruned game tree.

Now the node 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 node labeled “Tony” and label that node with the payoffs $(-40, 40)$. That is now the only remaining node, 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.

When you are doing problems using backward induction, you may find that recording parts of strategies and then pruning and redrawing game trees is

too slow. Here is another way to do problems. First, find the nodes c such that all moves available at c have ends that are terminal. At each of these nodes, cross out all moves that do not produce the greatest payoff for the player who chooses. If we do this for the game pictured in Figure 1.2, we get Figure 1.4.




Figure 1.4. Tony's Accident: start of backward induction.

Now you can back up a step. In Figure 1.4 we now see that Tony's three possible moves will produce payoffs to him of -80 , -80 , and -40 . Cross out the two moves that produce payoffs of -80 . We obtain Figure 1.5.




Figure 1.5. Tony's Accident: completion of backward induction.

From Figure 1.5 we can read off each player's strategy; for example, we can see what Vic will do at each of the nodes where he chooses, should that node be reached. We can also see how the game will play out if each player uses the strategy we have found.

In more complicated examples, of course, this procedure will have to be continued for more steps.

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. Figure 1.6 shows an example where backward induction fails. At the node where Player 2 chooses, both available moves give him a payoff of 1. Player 2




Figure 1.6. Failure of backward induction. As is standard when the players are numbered 1 and 2, Player 1's payoff is given first.

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 Player 1 chose *b*.

We return to this issue in Chapter 6.

1.5 Big Monkey and Little Monkey 1

Big Monkey and Little Monkey eat coconuts, 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 kilocalories (Kc). If Little Monkey climbs the tree, he incurs a negligible energy cost (because he's 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 what to do first. Payoffs are net gains in kilocalories. The game tree is shown in Figure 1.7. 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.




Figure 1.7. Big Monkey and Little Monkey. Big Monkey's payoff is given first.

1.6 Threats, promises, commitments

The game of Big Monkey and Little Monkey has the following peculiarity. Suppose Little Monkey adopts the strategy wait no matter what Big Monkey does. 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.

A commitment by Little Monkey to wait if Big Monkey waits is called a *threat*. If in fact Little Monkey waits after Big Monkey waits, Big Monkey's payoff is reduced from 9 to 0. Of course, Little Monkey's payoff is also reduced, from 1 to 0. The value of the threat, if it can be made believable, is that it should induce Big Monkey *not* to wait, so that the threat will not have to be carried out.

The ordinary use of the word "threat" includes the idea that the threat, if carried out, would be bad *both* for the opponent *and* for the individual making the threat. Think, for example, of a parent threatening to punish a child, or a country threatening to go to war. If an action would be bad for your opponent and good for you, there is no need to threaten to do it; it is your normal course.

The difficulty with threats is how to make them believable, since if the time comes to carry out the threat, the player making the threat will not want to do it. Some sort of advance commitment is necessary to make the threat believable. Perhaps Little Monkey should break his own leg and show up on crutches!

In this example the threat by Little Monkey works to his advantage. If Little Monkey can somehow convince Big Monkey that he will wait if Big Monkey waits, then from Big Monkey’s point of view, the game tree changes to the one shown in Figure 1.8.




Figure 1.8. Big Monkey and Little Monkey after Little Monkey commits to wait if Big Monkey waits. Big Monkey’s payoff is given first.

If Big Monkey uses backward induction on the new game tree, he will climb!

Closely related to threats are *promises*. In the game of Big Monkey and Little Monkey, Little Monkey could make a promise at the node after Big Monkey climbs. Little Monkey could promise to climb. This would increase Big Monkey’s payoff at that node from 4 to 5, while decreasing Little Monkey’s payoff from 4 to 3. Here, however, even if Big Monkey believes Little Monkey’s promise, it will not affect his action in the larger game. He will still wait, getting a payoff of 9.

The ordinary use of the word “promise” includes the idea that it is both good for the other person and bad for the person making the promise. If an action is also good for you, then there is no need to promise to do it; it is your normal course. Like threats, promises usually require some sort of advance commitment to make them believable.

Let us consider threats and promises more generally. Consider a two-player game in extensive form with complete information G . We first consider a node c such that all moves that start at c have terminal ends. Suppose for simplicity that Player 1 is to move at node c . Suppose Player 1’s “rational” choice at node c , the one she would make if she were using backward induction, is a move m that gives the two players payoffs (π_1, π_2) . Now imagine that Player 1 commits herself to a different move m' at node c , which gives the two players payoffs (π'_1, π'_2) . If m was the unique choice that gave Player 1 her best payoff, we necessarily have $\pi'_1 < \pi_1$, that is, the new move gives Player 1 a lower payoff.

- If $\pi'_2 < \pi_2$ (i.e., if the choice m' reduces Player 2’s payoff as well), Player 1’s commitment to m' at node c is a threat.
- If $\pi'_2 > \pi_2$ (i.e., if the choice m' increases Player 2’s payoff), Player 1’s commitment to m' at node c is a promise.

Now consider any node c where, for simplicity, Player 1 is to move. Suppose Player 1's "rational" choice at node c , the one she would make if she were using backward induction, is a move m . Suppose that if we use backward induction, when we have reduced to a game in which the node c is terminal, the payoffs to the two players at c are (π_1, π_2) . Now imagine that Player 1 commits herself to a different move m' at node c . Remove from the game G all other moves that start at c , and all parts of the tree that are no longer connected to the root node once these moves are removed. Call the new game G' . Suppose that if we use backward induction in G' , when we have reduced to a game in which the node c is terminal, the payoffs to the two players at c are (π'_1, π'_2) . Under the uniqueness assumption we have been using, we necessarily have $\pi'_1 < \pi_1$:

- If $\pi'_2 < \pi_2$, Player 1's commitment to m' at node c is a threat.
- If $\pi'_2 > \pi_2$, Player 1's commitment to m' at node c is a promise.

1.7 Ultimatum Game

Player 1 is given 100 one dollar bills. She must offer some of them (1 to 99) to Player 2. If Player 2 accepts the offer, she keeps the bills she was offered, and Player 1 keeps 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 in Figure 1.9.




Figure 1.9. Ultimatum Game with dollar payoffs. Player 1 offers a number of dollars to Player 2, then Player 2 accepts (a) or rejects (r) the offer.

Backward induction yields:

- 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 way she gets to keep 99!

However, many experiments have shown that people do 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.

A strategy by Player 2 to reject small offers is an implied threat (actually many implied threats, one for each small offer that she would reject). If Player 1 believes this threat—and experimentation has shown that she should—then she should make a fairly large offer. As in the game of Big Monkey and Little Monkey, a threat to make an “irrational” move, if it is believed, can result in a higher payoff than a strategy of always making the “rational” move.

We should also recognize a difficulty in interpreting game theory experiments. The experimenter can set up an experiment with monetary payoffs, but she cannot ensure that those are the only payoffs that are important to the experimental subject.

In fact, experiments suggest that many people prefer that resources not be divided in a grossly unequal manner, which they perceive as unfair; and that most people are especially concerned when it is they themselves who get the short end of the stick. Thus Player 2 may, for example, feel unhappy about accepting an offer x of less than \$50, with the amount of unhappiness equivalent to $4(50 - x)$ dollars (the lower the offer, the greater the unhappiness). Her payoff if she accepts an offer of x dollars is then x if $x > 50$, and $x - 4(50 - x) = 5x - 200$ if $x \leq 50$. In this case she should accept offers of greater than \$40, reject offers below \$40, and be indifferent between accepting and rejecting offers of exactly \$40.

Similarly, Player 1 may have payoffs not provided by the experimenter that lead her to make relatively high offers. She may prefer in general that resources not be divided in a grossly unequal manner, even at a monetary cost to herself. Or she may try to be the sort of person who does not take advantage of others and may experience a negative payoff when she does not live up to her ideals.

The take-home message is that the payoffs assigned to a player must reflect what is actually important to the player.

We have more to say about the Ultimatum Game in Sections 5.6 and 10.12.

1.8 Rosenthal’s Centipede Game

Like the Ultimatum Game, the Centipede Game is a game theory classic.

Mutt and Jeff start with \$2 each. Mutt goes first. On a player’s turn, he has two possible moves:

- (1) Cooperate (*c*): The player does nothing. The game master rewards him with \$1.
- (2) Defect (*d*): 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 1.10.




Figure 1.10. Rosenthal's Centipede Game. Mutt's payoff is given first. The amounts the players have accumulated when a node is reached are shown to the left of the node.

A backward induction analysis begins at the only node both of whose moves end in terminal nodes: 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. So cross out Jeff's last *c* move in Figure 1.10.

Now we back up a step, to Mutt's last node. We see from the figure that if Mutt cooperates, he will end up with \$98, but if he defects, he gets \$101 when the game immediately ends. If Mutt is "rational," he will defect. So cross out Mutt's last *c* move.

If we continue the backward induction procedure, we find that it 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.

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, consult the Wikipedia page for this game, [http://en.wikipedia.org/wiki/Centipede_game_\(game_theory\)](http://en.wikipedia.org/wiki/Centipede_game_(game_theory)).

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

1. *The players care about aspects of the game other than money.* For example, a player may feel better about himself if he cooperates. Alternatively, a player may want to *seem* cooperative, because this normally brings benefits. If a player wants to be, or to seem, cooperative, we should take account of this desire in assigning his payoffs.

2. *The players use a rule of thumb instead of analyzing the game.* 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 time and effort of doing the analysis may be greater than the advantage gained, and (b) if the analysis is complicated enough, you are liable to make a mistake anyway.

3. *The players use a strategy that is correct for a different, more common situation.* We do not typically encounter "games" that we know in advance have exactly or at most n stages, where n is a large number. Instead, we typically encounter games with an unknown number of stages. If the Centipede Game had an unknown number of stages, there would be no place to start a backward induction. In Chapter 6 we will study a class of such games for which it is rational to cooperate as long as your opponent does. When we encounter the unusual situation of a game with at most 196 stages, which is the case with the Centipede Game, perhaps we use a strategy that is correct for the more common situation of a game with an unknown number of stages.

However, the most interesting possibility is that the logical basis for believing that rational players will use long backward inductions is suspect. We address this issue in Section 1.13.

1.9 Continuous games

In the games we have considered so far, when it is a player's turn to move, she has only a finite number of choices. In the remainder of this chapter, we 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 consider games with two players, Player 1 and Player 2. Player 1 goes first. The moves available are all real numbers s in some interval I . Next it is Player 2's turn. The moves available are all real numbers t in some interval J . Player 2 observes Player 1's move s and then chooses her 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 her 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 she is rational, she 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; that is, 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 her best-response strategy, then Player 1 should choose s in I to maximize the function $\pi_1(s, b(s))$. This is again a function of one variable.

1.10 Stackelberg's model of duopoly 1

In a *duopoly*, a certain good is produced by just two firms, which we label 1 and 2. In Stackelberg's model of duopoly (Wikipedia article: <http://>

en.wikipedia.org/wiki/Stackelberg_duopoly), each firm tries to maximize its own profit by choosing an appropriate level of production. Firm 1 chooses its level of production first; then Firm 2 observes this choice and chooses its own level of production. Would you rather run Firm 1 or Firm 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 produced is $q = s + t$. The market price p of the good depends on q : $p = \phi(q)$. At this price, everything produced can be sold.

Suppose Firm 1's cost to produce the quantity s of the good, which we denote $c_1(s)$, is $4s$, and Firm 2's cost to produce the quantity t of the good, which we denote $c_2(t)$, is $4t$. In other words, both Firm 1 and Firm 2 have the same unit cost of production 4.

1.10.1 First model. We assume that price falls linearly as total production of the two firms increases. In particular, we assume

$$p = 20 - 2(s + t). \quad (1.1)$$

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,

$$\pi_1(s, t) = ps - c_1(s) = (20 - 2(s + t))s - 4s = (16 - 2t)s - 2s^2, \quad (1.2)$$

$$\pi_2(s, t) = pt - c_2(t) = (20 - 2(s + t))t - 4t = (16 - 2s)t - 2t^2. \quad (1.3)$$

We regard the profits as payoffs in a game. The players are Firms 1 and 2.

In this subsection we allow the levels of production s and t to be any real numbers, even negative numbers and numbers large enough to make the price negative. This doesn't make sense economically, but it avoids mathematical complications.

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:

$$\frac{\partial \pi_2}{\partial t} = 16 - 2s - 4t = 0.$$

If we solve this equation for t , we will have Firm 2's best-response function:

$$t = b(s) = 4 - \frac{1}{2}s.$$

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. From (1.2), we have

$$\pi_1(s, b(s)) = \pi_1(s, 4 - \frac{1}{2}s) = (16 - 2(4 - \frac{1}{2}s))s - 2s^2 = 8s - s^2.$$

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:

$$\frac{d}{ds}\pi_1(s, b(s)) = 8 - 2s = 0 \quad \Rightarrow \quad s = 4.$$

Thus $\pi_1(s, b(s))$ is maximum at $s^* = 4$. Given this choice of production level for Firm 1, Firm 2 chooses the production level

$$t^* = b(s^*) = 4 - \frac{1}{2}4 = 2.$$

The price is

$$p^* = 20 - 2(s^* + t^*) = 20 - 2(4 + 2) = 8.$$

From (1.2) and (1.3), the profits are

$$\pi_1(s^*, t^*) = (16 - 2 \cdot 2)4 - 2 \cdot 4^2 = 16,$$

$$\pi_2(s^*, t^*) = (16 - 2 \cdot 4)2 - 2 \cdot 2^2 = 8.$$

Firm 1 has twice the level of production and twice the profit of Firm 2. It is better to run the firm that chooses its production level first.

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

First, we only allow nonnegative production levels: $0 \leq s < \infty$ and $0 \leq t < \infty$. Second, we assume that if total production rises above 10, the level at which formula (1.1) for the price gives 0, then the price is 0, not the negative number given by formula (1.1):

$$p = \begin{cases} 20 - 2(s + t) & \text{if } s + t < 10, \\ 0 & \text{if } s + t \geq 10. \end{cases}$$

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:

$$\pi_1(s, t) = ps - c_1(s) = \begin{cases} (20 - 2(s + t))s - 4s & \text{if } 0 \leq s + t < 10, \\ -4s & \text{if } s + t \geq 10, \end{cases} \quad (1.4)$$

$$\pi_2(s, t) = pt - c_2(t) = \begin{cases} (20 - 2(s + t))t - 4t & \text{if } 0 \leq s + t < 10, \\ -4t & \text{if } s + t \geq 10. \end{cases} \quad (1.5)$$

We again begin our analysis by finding Firm 2's best response $t = b(s)$.

Unit cost of production is 4. If Firm 1 produces so much that all by itself it drives the price down to 4 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 by itself drives the price p down to 4 when $20 - 2s = 4$, that is, when its level of production is $s = 8$. We conclude that if Firm 1's level of production s is 8 or higher, Firm 2's best response is 0.

In contrast, if Firm 1 produces $s < 8$, it leaves the price above 4 and gives Firm 2 an opportunity to make a positive profit. In this case Firm 2's profit is given by

$$\pi_2(s, t) = \begin{cases} (20 - 2(s + t))t - 4t = (16 - 2s)t - 2t^2 & \text{if } 0 \leq t < 10 - s, \\ -4t & \text{if } t \geq 10 - s. \end{cases}$$

See Figure 1.11.




Figure 1.11. Graph of $\pi_2(s, t)$ for fixed $s < 8$ in the numerical example.

From the figure, the function $\pi_2(s, t)$ with s fixed is maximum where $(\partial\pi_2/\partial t)(s, t) = 0$, which occurs at $t = 4 - \frac{1}{2}s$.

Thus Firm 2's best-response function is

$$b(s) = \begin{cases} 4 - \frac{1}{2}s & \text{if } 0 \leq s < 8, \\ 0 & \text{if } s \geq 8. \end{cases}$$

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 $0 \leq s < 8$, we have

$$s + b(s) = s + 4 - \frac{1}{2}s = 4 + \frac{1}{2}s < 4 + \frac{1}{2}8 = 8 < 10.$$

Therefore, for $0 \leq s < 8$, we use the first line of (1.4) to calculate $\pi_1(s, b(s))$:

$$\pi_1(s, b(s)) = \pi_1(s, 4 - \frac{1}{2}s) = (16 - 2(4 - \frac{1}{2}s))s - 2s^2 = 8s - s^2.$$

Firm 1 will not choose an $s \geq 8$, since, as we have seen, that would force the price down to the cost of production 4 or lower. Therefore we will not bother to calculate $\pi_1(s, b(s))$ for $s \geq 8$.

The function $\pi_1(s, b(s))$ on the interval $0 \leq s \leq 8$ is maximum at $s^* = 4$, where the derivative of $8s - s^2$ is 0, just as in our first model. The value of $t^* = b(s^*)$ is also the same, as are the price and profits.

1.11 Stackelberg's model of duopoly 2

In this section we give a more general treatment of Stackelberg's model of duopoly.

1.11.1 First model. In this subsection we make the following assumptions, which generalize those used in Subsection 1.10.1.

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

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

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

Firm 1 chooses its level of production s first. Then Firm 2 observes s and chooses t . We ask the question, what will be the production level and profit of each firm?

The payoffs are

$$\begin{aligned}\pi_1(s, t) &= ps - c_1(s) = (\alpha - \beta(s + t))s - cs = (\alpha - \beta t - c)s - \beta s^2, \\ \pi_2(s, t) &= pt - c_2(t) = (\alpha - \beta(s + t))t - ct = (\alpha - \beta s - c)t - \beta t^2.\end{aligned}$$

We find Firm 2's best response $t = b(s)$ by finding 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 an upside-down parabola, we just take the derivative with respect to t and set it equal to 0:

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

We solve this equation for t to find Firm 2's best-response function:

$$t = b(s) = \frac{\alpha - c}{2\beta} - \frac{1}{2}s.$$

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

$$\begin{aligned}\pi_1(s, b(s)) &= \pi_1\left(s, \frac{\alpha - c}{2\beta} - \frac{1}{2}s\right) = \left(\alpha - \beta\left(\frac{\alpha - c}{2\beta} - \frac{1}{2}s\right) - c\right)s - \beta s^2 \\ &= \frac{\alpha - c}{2}s - \frac{\beta}{2}s^2.\end{aligned}$$

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:

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

Thus $\pi_1(s, b(s))$ is maximum at $s^* = (\alpha - c)/2\beta$. 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, which makes sense. The price is

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

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

The profits are

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

As in our numerical example, Firm 1 has twice the level of production and twice the profit of Firm 2.

1.11.2 Second model. As in Subsection 1.10.2, we now complicate the model to prevent the levels of production s and t and the price p from taking negative values. We replace assumption (1) in Subsection 1.11.1 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) and (3) remain unchanged. We replace assumption (4) with:

- (4') The production levels s and t must be nonnegative: $0 \leq s < \infty$ and $0 \leq t < \infty$.

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

The payoffs in the general case are

$$\pi_1(s, t) = ps - c_1(s) = \begin{cases} (\alpha - \beta(s + t))s - cs & \text{if } 0 \leq s + t < \frac{\alpha}{\beta}, \\ -cs & \text{if } s + t \geq \frac{\alpha}{\beta}, \end{cases} \quad (1.6)$$

$$\pi_2(s, t) = pt - c_2(t) = \begin{cases} (\alpha - \beta(s + t))t - ct & \text{if } 0 \leq s + t < \frac{\alpha}{\beta}, \\ -ct & \text{if } s + t \geq \frac{\alpha}{\beta}. \end{cases} \quad (1.7)$$

As usual we begin our analysis by finding Firm 2's best response $t = b(s)$.

If Firm 1 produces so much that by itself it drives the price down to the unit cost of production c or lower, then Firm 2 cannot make a positive profit. In this case Firm 2's best response is to produce nothing. Firm 1 by itself drives the price p down to c when $\alpha - \beta s = c$, that is, when $s = (\alpha - c)/\beta$. We conclude that if $s \geq (\alpha - c)/\beta$, Firm 2's best response is 0.

In contrast, if Firm 1 produces $s < (\alpha - c)/\beta$, it leaves the price above c and gives Firm 2 an opportunity to make a positive profit. In this case Firm 2's profit is given by

$$\pi_2(s, t) = \begin{cases} (\alpha - \beta(s + t))t - ct = (\alpha - \beta s - c)t - \beta t^2 & \text{if } 0 \leq t < \frac{\alpha}{\beta} - s, \\ -ct & \text{if } t \geq \frac{\alpha}{\beta} - s. \end{cases}$$

See Figure 1.12. From the figure, the function $\pi_2(s, t)$ with s fixed is maximum where $(\partial\pi_2/\partial t)(s, t) = 0$, which occurs at

$$t = \frac{\alpha - c}{2\beta} - \frac{1}{2}s.$$

Thus Firm 2's best-response function is:

$$b(s) = \begin{cases} \frac{\alpha - c}{2\beta} - \frac{1}{2}s & \text{if } 0 \leq s < \frac{\alpha - c}{\beta}, \\ 0 & \text{if } s \geq \frac{\alpha - c}{\beta}. \end{cases}$$




Figure 1.12. Graph of $\pi_2(s, t)$ for fixed $s < (\alpha - c)/\beta$.

We now calculate $\pi_1(s, b(s))$, the payoff that Firm 1 can expect from each choice s , assuming that Firm 2 uses its best-response strategy. Notice that for $0 \leq s < (\alpha - c)/\beta$, we have

$$s + b(s) = s + \frac{\alpha - c}{2\beta} - \frac{1}{2}s = \frac{\alpha - c}{2\beta} + \frac{1}{2}s < \frac{\alpha - c}{2\beta} + \frac{\alpha - c}{2\beta} = \frac{\alpha - c}{\beta} < \frac{\alpha}{\beta}.$$

Therefore, for $0 \leq s < (\alpha - c)/\beta$, we use the first line of formula (1.6) to calculate $\pi_1(s, b(s))$:

$$\begin{aligned} \pi_1(s, b(s)) &= \pi_1\left(s, \frac{\alpha - c}{2\beta} - \frac{1}{2}s\right) = \left(\alpha - \beta\left(s + \frac{\alpha - c}{2\beta} - \frac{1}{2}s\right)\right)s - cs \\ &= \frac{\alpha - c}{2}s - \frac{\beta}{2}s^2. \end{aligned}$$

Firm 1 will not choose an $s \geq (\alpha - c)/\beta$, 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 $s \geq (\alpha - c)/\beta$.

The function $\pi_1(s, b(s))$ on the interval $0 \leq s \leq (\alpha - c)/\beta$ is maximum at $s^* = (\alpha - c)/2\beta$, where the derivative of $\frac{1}{2}(\alpha - c)s - \frac{1}{2}\beta s^2$ is 0, just as in our first model. The value of $t^* = b(s^*)$ is also the same, as are the price and profits.

1.12 Backward induction for finite horizon games

Backward induction as defined 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 with complete information. It is actually this variant that we have been using since Section 1.9.

Let us describe this variant of backward induction in general. The idea is 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 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 C , identify the player $i(c)$ who is to choose at node c . Among all the moves available to her 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 C , 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 one of the nodes in C . The nodes c in C are now terminal nodes. Assign to each node c in C the payoffs that were previously assigned to the node at the end of the move $m(c)$.
- (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.

1.13 Critique of backward induction

The basic insight of backward induction is that you should think ahead to how your opponent, acting in his own interest, is liable to react to what you do, and act accordingly to maximize your chance of success. This idea clearly makes sense even in situations that are not as completely defined as the games we analyze. For example, the mixed martial arts trainer Greg Jackson has analyzed countless fight videos and used them to make game trees showing what moves lead to what responses. From these game trees he can figure out which moves in various situations will increase the likelihood of a win. As another example, consider the game of chess. Because of the rule that a draw results when a position is repeated three times, the game tree for chess is finite. Unfortunately, it has 10^{123} nodes and hence is far too big for a computer to analyze. (The number of atoms in the observable universe is estimated to be around 10^{80} .) Thus computer chess programs cannot use backward induction from the terminal nodes. Instead they investigate paths through the game tree from a given position to a given depth and assign values to the end nodes based on estimates of the probability of winning from that position. They then use backward induction from those nodes.

Despite successes like these, it is not clear that backward induction is always a good guide to choosing a move.

Let's first consider Tony's Accident. To justify using backward induction at all, Tony has to assume that Vic will always choose his own best move in response to Tony's move. In addition, Tony should know Vic's payoffs, or at least he should know the order in which Vic values the different outcomes, so that he will know which of Vic's available moves Vic will choose in response to Tony's move. If Tony does not know the order in which Vic values the outcomes, he can still use backward induction based on his belief about Vic's order. This is what Tony did. The success of the procedure then depends on the correctness of Tony's beliefs about Vic.

In Chapter 6 we consider a game, the Samaritan's Dilemma, that raises an additional issue. In that game, Daughter wants to decide how much to save from her earnings this year toward her college expenses next year. Father will then observe how much she saves and chip in some of his own earnings. To figure out how much to give, he will balance his desire to keep his earnings to spend on himself against his desire to help his daughter. To justify her use of backward induction in this situation, Daughter has to assume that she knows Father's desires well enough to figure out his best response, from his own point of view, to each of her possible saving levels. She also has to assume that Father will actually make his best response. Here this second assumption becomes hard to justify. To justify it, she needs to assume both that Father is able to figure out his best response, and that he is willing to do

so. Recall from our discussion of the Centipede Game that it may not even be rational for Father to use a complicated rational analysis to figure out what to do.

Finally, let's consider the Centipede Game. Would a rational player in the Centipede Game (Section 1.8) really defect at his first opportunity, as is required by backward induction? We examine this question under the assumption that the payoffs in the Centipede Game are exactly as given in Figure 1.10, that both players know these payoffs, and that both players are rational. The assumption that players know the payoffs and are rational motivates backward induction. The issue now is whether the assumption that players know the payoffs and are rational *requires* them to use the moves recommended by backward induction.

By a rational player, we mean one whose preferences are consistent enough to be represented by a payoff function; who attempts to discern the facts about the world; who forms beliefs about the world consistent with the facts he has discerned; and who acts on the basis of his beliefs to best achieve his preferred outcomes.

With this "definition" of a rational player in mind, let us consider the first few steps of backward induction in the Centipede Game.

1. If the node labeled (100, 99) in Figure 1.10 is reached, Jeff will see that if he defects, his payoff is 101, and if he cooperates, his payoff is 100. *Since Jeff is rational*, he defects.

2. If the node labeled (99, 99) in Figure 1.10 is reached, Mutt will see that if he defects, his payoff is 101. If he cooperates, the node labeled (100, 99) is reached. *If Mutt believes that Jeff is rational*, then he sees that Jeff will defect at that node, leaving Mutt with a payoff of only 98. *Since Mutt is rational*, he defects.

3. If the node labeled (99, 98) in Figure 1.10 is reached, Jeff will see that if he defects, his payoff is 100. If he cooperates, the node labeled (99, 99) is reached. *If Jeff believes that Mutt believes that Jeff is rational*, and *if Jeff believes that Mutt is rational*, then Jeff concludes that Mutt will act as described in step 2. This would leave Jeff with a payoff of 97. *Since Jeff is rational*, he defects.

4. You probably see that this is getting complicated fast, but let's do one more step. If the node labeled (98, 98) (not shown in Figure 1.10) is reached, Mutt will see that if he defects, his payoff is 100. If he cooperates, the node labeled (99, 98) is reached. *If Mutt believes that Jeff believes that Mutt believes that Jeff is rational*, and *if Mutt believes that Jeff believes that Mutt is rational*, and *if Mutt believes that Jeff is rational*, then Mutt concludes that Jeff will act as described in step 3. This would leave Mutt with a payoff of 97. *Since Mutt is rational*, he defects.

You can see that by the time we get back to the root node in Figure 1.10, at step 196 in this process, Mutt must hold many complicated beliefs to justify using backward induction to choose his move! (At the k th step in this process, the player who chooses must hold $k - 1$ separate beliefs about the other player, the most complicated of which requires $k - 1$ uses of word “believes” to state.) The question is whether a rational player, “who attempts to discern the facts about the world, who forms beliefs about the world consistent with the facts he has discerned,” is *required* to hold such complicated beliefs.

Now suppose that Mutt decides, for whatever reason, to cooperate at his first turn. Then Jeff can conclude that Mutt did not hold the complicated beliefs that would induce him to defect at his first turn. If Jeff believes that at Mutt’s second turn, he will hold the (slightly less complicated) beliefs that would induce him to defect then, then Jeff should defect at his own first turn. Clearly, rationality does not require Jeff to believe this. Thus a rational Jeff may well decide to cooperate at his first turn.

In this way, rational players may well cooperate through many stages of the Centipede Game. More generally, *rationality does not require players to follow the strategy dictated by a long backward induction.*

Backward induction is required by an assumption called Common Knowledge of Rationality: the players are rational; each player believes that the other players are rational; each player believes that the other players believe that the other players are rational; and so on, for as many steps as are required by the game under discussion. In any particular situation, this assumption may hold, but the assumption that players are rational does not by itself imply Common Knowledge of Rationality.

We shall return to the question of players’ beliefs in Section 8.2.

1.14 Problems

1.14.1 Congress vs. the President. Congress is working on a homeland security spending bill. Congress wants the bill to include \$10 million for each member’s district for “important projects.” The President wants the bill to include \$100 million to upgrade her airplane, Air Force One, to the latest model, Air Force One Extreme. The President can sign or veto any bill that Congress passes.

Congress’s payoffs are 1 if it passes a bill (voters like to see Congress take action), 2 if the President signs a bill that contains money for important projects, and -1 if the President signs a bill that contains money for Air Force One Extreme. Payoffs are added; for example, if the President signs a bill that contains money for both, Congress’s total payoff is $1 + 2 + (-1) = 2$.

The President's payoffs are -1 if she signs a bill that contains money for important projects, and 2 if she signs a bill that contains money for Air Force One Extreme. Her payoffs are also added.

The game tree in Figure 1.13 illustrates the situation. In the tree, C = Congress, P = President; n = no bill passed, i = important projects passed, a = Air Force One Extreme passed, b = both passed; s = President signs the bill, v = President vetoes the bill. The first payoff is Congress's, the second is the President's.




Figure 1.13. Congress vs. the President.

- (1) Use backward induction to predict what will happen.
- (2) Suppose the Constitution were changed so that the President could veto parts of bills she doesn't like but still sign the rest. Draw the new tree and use backward induction to predict what will happen. (The only change in the game is that if Congress passes a bill containing both important projects and Air Force One Extreme, the President will have four choices: sign, veto, sign important projects but veto Air Force One Extreme, sign Air Force One Extreme but veto important projects.)

1.14.2 Battle of the Chain Stores. Sub Station has the only sub restaurant in Town A and the only sub restaurant in Town B. The sub market in each town yields a profit of \$100K per year. Rival Sub Machine is considering opening a restaurant in Town A in year 1. If it does, the two stores will split the profit from the sub market there. However, Sub Machine will have to pay setup costs for its new store. These costs are \$25K in a store's first year.

Sub Station fears that if Sub Machine is able to make a profit in Town A, it will open a store in Town B the following year. Sub Station is considering a price war: if Sub Machine opens a store in either town, it will lower prices in that town, forcing Sub Machine to do the same, to the point where profits from the sub market in that town drop to 0.

The game tree in Figure 1.14 is one way to represent the situation. It takes into account net profits from Towns A and B in years 1 and 2, and it assumes that if Sub Machine loses money in A, it will not open a store in B.

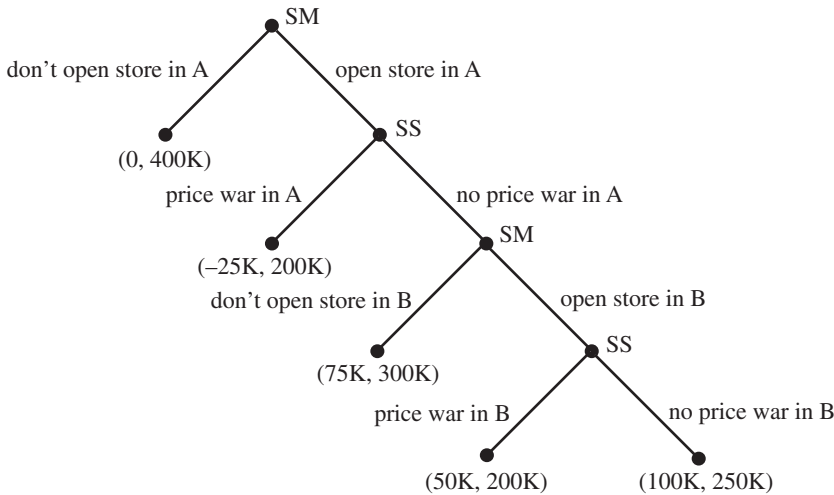


Figure 1.14. SM is Sub Machine, SS is Sub Station. Sub Machine's profits are listed first, Sub Station's profits are second.

The entry (100K, 250K) in the game tree comes about as follows. If there are no price wars, Sub Machine makes net profits of \$25K in Town A in year 1, \$50K in Town A in year 2, and \$25K in Town B in year 2, for a total of \$100K. Sub Station makes profits of \$50K in Town A in year 1, \$50K in Town A in year 2, \$100K in Town B in year 1, and \$50K in Town B in year 2, for a total of \$250K.

- (1) Explain the entry (50K, 200K) in the game tree.
- (2) Use backward induction to figure out what Sub Machine and Sub Station should do.
- (3) How might Sub Station try to obtain a better outcome by using a threat?

1.14.3 Kidnapping. A criminal kidnaps a child and demands a ransom $r > 0$. The value to the parents of the child's return is $v > 0$. If the kidnapper frees the child, he incurs a cost $f > 0$. If he kills the child, he incurs a cost $k > 0$. These costs include the feelings of the kidnapper in each case, the likelihood of being caught in each case, the severity of punishment in each case, and whatever else is relevant. We assume $v > r$. The parents choose

first whether to pay the ransom or not, then the kidnapper decides whether to free the child or kill the child. The game tree is shown in Figure 1.15.

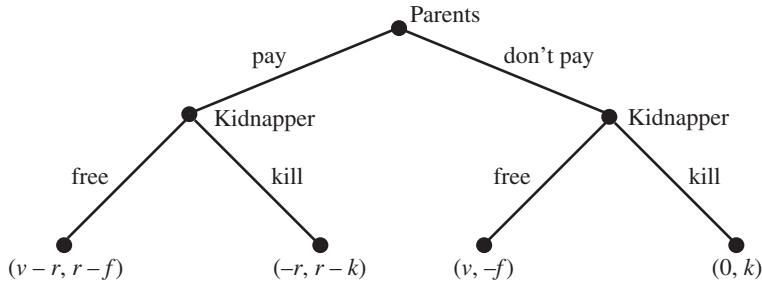


Figure 1.15. Kidnapping. The parents' payoff is given first.

- (1) Explain the payoffs.
- (2) Suppose $k > f$. Use backward induction to show that the parents should not pay the ransom.
- (3) Suppose $f > k$. Use backward induction to show that the parents should not pay the ransom.
- (4) Suppose $k > f$. Find a threat the kidnapper can make to try to get the parents to pay the ransom.
- (5) Suppose $f > k$. Find a promise the kidnapper can make to try to get the parents to pay the ransom.
- (6) Suppose $f > k$, and the parents will incur a guilt cost g if they don't pay the ransom and the child is killed. How large must g be to get the parents to pay the ransom?

1.14.4 The White House Tapes. On March 1, 1974, a grand jury indicted seven former aides to U.S. President Richard Nixon for attempting to cover up White House involvement in a burglary of Democratic National Committee headquarters at the Watergate complex in Washington. On April 18, the judge in the case, John Sirica, issued a subpoena for tapes of President Nixon's conversations with the defendants. The President's attorney, James St. Clair, attempted to delay responding to the subpoena. The prosecutor, Leon Jaworski, then used an unusual procedure to appeal directly to the Supreme Court and request that the Court order the President to supply the tapes. The Court heard oral arguments on July 8, and the justices met on July 9 to decide the case.

One justice, William Rehnquist, withdrew from the case, probably because he had worked in President Nixon's Justice Department. Of the remaining eight justices, six quickly agreed to uphold the prosecutor's request. Two

justices, Warren Burger and Harry Blackmun, were reluctant to uphold the prosecutor's request, because they thought his direct appeal to the Supreme Court was improper.

Also on July 9, President Nixon's attorney said that the President had "not yet decided" whether he would supply the tapes if the Supreme Court ordered him to. This statement was probably intended to pressure the Court into backing down. At minimum, it was probably intended to encourage some justices to vote against upholding the prosecutor's request. If the vote was split, the President could argue that it was not sufficiently definitive for a matter of this magnitude. Jaworski believed that in the event of a split vote, the President would "go on television and tell the people that the presidency should not be impaired by a divided Court."

We model this situation as a two-player game. Player 1 is Justices Burger and Blackmun, whom we assume will vote together; we therefore treat them as one player. Player 2 is President Nixon.

First, Justices Burger and Blackmun decide how to vote. If they vote to uphold the prosecutor's request, the result is an 8-0 Supreme Court decision in favor of the prosecutor. If they vote to reject the prosecutor's request, the result is a 6-2 Supreme Court decision in favor of the prosecutor.

After the Supreme Court has rendered its decision, President Nixon decides whether to comply by supplying the tapes or to defy the decision. President Nixon's preferences are as follows:

- Best outcome (payoff 4): 6-2 decision, President defies the decision.
- Second-best outcome (payoff 3): 6-2 decision, President supplies the tapes.
- Third-best outcome (payoff 2): 8-0 decision, President supplies the tapes.
- Worst outcome (payoff 1): 8-0 decision, President defies the decision.

Explanation: The President's best outcome is a divided decision that he can defy while claiming the decision is not really definitive. His worst outcome is an 8-0 decision that he then defies; this would probably result in immediate impeachment. As for the two intermediate outcomes, the President is better off with the weaker vote, which should give him some wiggle room.

Justices Burger and Blackmun's preferences are as follows:

- Best outcome (payoff 4): 6-2 decision, President supplies the tapes.
- Second-best outcome (payoff 3): 8-0 decision, President supplies the tapes.
- Third-best outcome (payoff 2): 8-0 decision, President defies the decision.
- Worst outcome (payoff 1): 6-2 decision, President defies the decision.

Explanation: In their best outcome, Burger and Blackmun get to vote their honest legal opinion that the prosecutor's direct appeal to the Court was wrong, but a Constitutional crisis is averted, because the President complies

anyway. In their second-best outcome, they vote dishonestly, but they succeed in averting a major Constitutional crisis. In their third-best outcome, the crisis occurs, but because of the strong 8-0 vote, it will probably quickly end. In the worst outcome, the crisis occurs, and because of the weak vote, it may drag out. In addition, in the last outcome, the President may succeed in establishing the principle that a 6-2 Court decision need not be followed, which no member of the Court wants.

- (1) Draw a game tree and use backward induction to predict what happened.
- (2) Can you think of a plausible way that President Nixon might have gotten a better outcome?

(What actually happened: the Court ruled 8-0 in favor of the prosecutor on July 24. On July 30, President Nixon surrendered the tapes. In early August, a previously unknown tape recorded a few days after the break-in was released. It documented President Nixon and his aide Robert Haldeman formulating a plan to block investigations by having the CIA claim to the FBI (falsely) that national security was involved. On August 9, President Nixon resigned.)

1.14.5 Three Pirates. Three pirates must divide 100 gold doubloons. The doubloons cannot be cut into pieces. Pirate A is the strongest, followed by Pirate B, followed by Pirate C. Because of ancient pirate tradition, the coins are divided in the following manner. First Pirate A proposes a division of the coins. The three pirates then vote on whether to accept the proposed division. If the proposal gets a majority vote, it is accepted, and the game is over. If the proposal fails to get a majority vote, Pirate A is executed.

It is then Pirate B's turn to propose a division of the coins between the two remaining pirates. The same rules apply, with one exception: if the vote is a tie (which can happen this time, because the number of pirates is now even), Pirate B, being the strongest remaining pirate, gets an additional vote to break the tie.

Use backward induction to figure out what Pirate A should propose. You don't have to draw a game tree if you don't want to, but you do have to explain your thinking so that your instructor can follow it.

1.14.6 Grocery Store and Gas Station 1. In a certain town, there are two stores, a grocery store and a gas station. The grocery store charges p_1 dollars per pound for food, and the gas station charges p_2 dollars per gallon for gas. The grocery store sells q_1 pounds of food per week, and the gas station sells q_2 gallons of gas per week. The quantities q_1 and q_2 are related to the prices

p_1 and p_2 as follows:

$$\begin{aligned}q_1 &= 10 - 2p_1 - p_2, \\q_2 &= 10 - p_1 - 2p_2.\end{aligned}$$

Thus, if the price of food or gas rises, less of *both* is sold.

Let π_1 be the revenue of the grocery store and π_2 be the revenue of the gas station. Both depend on the two stores' choices of p_1 and p_2 :

$$\begin{aligned}\pi_1(p_1, p_2) &= q_1 p_1 = (10 - 2p_1 - p_2)p_1 = 10p_1 - 2p_1^2 - p_1 p_2, \\ \pi_2(p_1, p_2) &= q_2 p_2 = (10 - p_1 - 2p_2)p_2 = 10p_2 - p_1 p_2 - 2p_2^2.\end{aligned}$$

We interpret this as a game with two players, the grocery store (Player 1) and the gas station (Player 2). The payoff to each player is its revenue.

Allow p_1 and p_2 to be any real numbers, even negative numbers or numbers that produce negative values for q_1 and q_2 .

Suppose the grocery store chooses its price p_1 first, and then the gas station, knowing p_1 , chooses its price p_2 .

- (1) If the grocery store uses backward induction to choose p_1 , what price will it choose?
- (2) What will be the corresponding p_2 , and what will be the revenue of each store?

Partial answer to help keep you on the right track: $p_1 = 2\frac{1}{7}$.

1.14.7 Grocery Store and Gas Station 2. We now change the problem a little. First we change the formulas for q_1 and q_2 to the following more reasonable formulas, which say that when prices get too high, the quantities sold become 0, not negative numbers:

$$\begin{aligned}q_1 &= \begin{cases} 10 - 2p_1 - p_2 & \text{if } 2p_1 + p_2 < 10, \\ 0 & \text{if } 2p_1 + p_2 \geq 10, \end{cases} \\ q_2 &= \begin{cases} 10 - p_1 - 2p_2 & \text{if } p_1 + 2p_2 < 10, \\ 0 & \text{if } p_1 + 2p_2 \geq 10. \end{cases}\end{aligned}$$

Next we make some reasonable restrictions on the prices. First we assume

$$p_1 \geq 0 \quad \text{and} \quad p_2 \geq 0.$$

Now we note that if $p_1 \geq 0$ and $p_2 > 5$, then q_2 becomes 0. The gas station wouldn't want this, so we assume

$$p_2 \leq 5.$$

Finally, we note that if $p_2 \geq 0$ and $p_1 > 5$, then q_1 becomes 0. The grocery store wouldn't want this, so we assume

$$p_1 \leq 5.$$

The formulas for π_1 and π_2 are now a little different, because the formulas for q_1 and q_2 have changed. The payoffs are now only defined for $0 \leq p_1 \leq 5$ and $0 \leq p_2 \leq 5$, and are given by

$$\pi_1(p_1, p_2) = q_1 p_1 = \begin{cases} 10p_1 - 2p_1^2 - p_1 p_2 & \text{if } 2p_1 + p_2 < 10, \\ 0 & \text{if } 2p_1 + p_2 \geq 10, \end{cases}$$

$$\pi_2(p_1, p_2) = q_2 p_2 = \begin{cases} 10p_2 - p_1 p_2 - 2p_2^2 & \text{if } p_1 + 2p_2 < 10, \\ 0 & \text{if } p_1 + 2p_2 \geq 10. \end{cases}$$

We still assume that the grocery store chooses its price p_1 first, and then the gas station, knowing p_1 , chooses its price p_2 .

- (1) Explain the formulas above for the payoffs.
- (2) For a fixed p_1 between 0 and 5, graph the function $\pi_2(p_1, p_2)$, which is a function of p_2 defined for $0 \leq p_2 \leq 5$. Answer: should be partly an upside-down parabola and partly a horizontal line. Make sure you give a formula for the point at which the graph changes from one to the other.
- (3) By referring to the graph you just drew and using calculus, find the gas station's best-response function $p_2 = b(p_1)$, which should be defined for $0 \leq p_1 \leq 5$. Answer: $b(p_1) = (10 - p_1)/4$.
- (4) You are now ready to find p_1 by backward induction. From your formula for π_1 you should be able to see that

$$\pi_1(p_1, b(p_1)) = \begin{cases} (10 - 2p_1 - b(p_1))p_1 & \text{if } 2p_1 + b(p_1) < 10, \\ 0 & \text{if } 2p_1 + b(p_1) \geq 10. \end{cases}$$

Use the formula for $b(p_1)$ from part (3) to show that $2p_1 + b(p_1) < 10$ if $0 \leq p_1 < 4\frac{2}{7}$, and $2p_1 + b(p_1) \geq 10$ if $4\frac{2}{7} \leq p_1 \leq 5$.

- (5) Graph the function $\pi_1(p_1, b(p_1))$, $0 \leq p_1 \leq 5$. (Again, it is partly an upside-down parabola and partly a horizontal line.)
- (6) By referring to the graph you just drew and using calculus, find where $\pi_1(p_1, b(p_1))$ is maximum.

1.14.8 Tax Rate. Government chooses a tax rate x , $0 < x < 1$. Citizens then choose a level of effort y to devote to making money. The resulting level of economic activity a , in trillions of dollars, is

$$a = 4y - 4xy - y^2.$$

Explanation: The term $4y$ expresses the idea that economic activity should be proportional to effort. There are two correction terms. The term $-4xy$ expresses the idea that when the tax rate is high, much effort goes into avoiding taxes, not into economic activity. The term $-y^2$ expresses the idea that too much effort is counterproductive.

The government's payoff is the taxes it collects, computed as tax rate x times level of economic activity a :

$$\pi_1(x, y) = xa = x(4y - 4xy - y^2).$$

The citizens' payoff is the part of economic activity that they keep after taxes:

$$\pi_2(x, y) = a - xa = (1 - x)a = (1 - x)(4y - 4xy - y^2).$$

We regard this as a two-player game. Player 1 is the government. Player 2 is the citizens. Player 1 chooses the tax rate x , then Player 2 observes x and chooses y . The payoffs are given above.

Use backward induction to find the tax rate x that maximizes the government's payoff.

1.14.9 Continuous Ultimatum Game with Inequality Aversion. Players 1 and 2 must divide some Good Stuff. Player 1 offers Player 2 a fraction y , $0 \leq y \leq 1$, of the Good Stuff. If Player 2 accepts the offer, she gets the fraction y of Good Stuff, and Player 1 gets the remaining fraction $x = 1 - y$. If Player 2 rejects the offer, both players get nothing.

In this game, Player 1 has an interval of possible strategies. We can describe this interval as $0 \leq x \leq 1$, where x is the fraction Player 1 keeps, or as $0 \leq y \leq 1$, where y is the fraction Player 1 offers to Player 2. In contrast, Player 2's strategy is a plan, for each y that she might be offered, whether to accept or reject.

The players' payoffs are partly objective (their fraction of the Good Stuff) and partly subjective. Both players are *inequality averse*: they don't like an unequal division. However, each player feels that an inequality favoring the other player is worse than inequality favoring herself. Their payoffs are as follows:

$$u_1(x, y) = x - \begin{cases} \alpha_1(y - x) & \text{if } y \geq x, \\ \beta_1(x - y) & \text{if } y < x, \end{cases}$$

$$u_2(x, y) = y - \begin{cases} \alpha_2(x - y) & \text{if } x \geq y, \\ \beta_2(y - x) & \text{if } x < y, \end{cases}$$

with $0 < \beta_1 < \alpha_1$ and $0 < \beta_2 < \alpha_2$. Thus Player 1's payoff is the fraction x of Good Stuff that she gets, minus a correction for inequality. The correction

is proportional to the difference between the two allocations. If Player 2 gets more ($y > x$), the difference is multiplied by the bigger number α_1 ; if Player 2 gets less ($y < x$), the difference is multiplied by the smaller number β_1 . Player 2's payoff function is similar.

If Player 1's offer is rejected, both players receive 0 payoff.

We assume that Player 2 will accept offers that make her payoff positive or zero. Player 1 wants to make an offer that (a) Player 2 will accept, and (b) that leaves Player 1 with the largest possible payoff.

Since $y = 1 - x$ and $x = 1 - y$, we simplify the payoff functions as follows:

$$u_1(x) = x - \begin{cases} \alpha_1(1 - 2x) & \text{if } x \leq \frac{1}{2}, \\ \beta_1(2x - 1) & \text{if } x > \frac{1}{2}, \end{cases}$$
$$u_2(y) = y - \begin{cases} \alpha_2(1 - 2y) & \text{if } y \leq \frac{1}{2}, \\ \beta_2(2y - 1) & \text{if } y > \frac{1}{2}. \end{cases}$$

- (1) Graph Player 1's payoff function on the interval $0 \leq x \leq 1$ assuming $\beta_1 > \frac{1}{2}$. The graph consists of two line segments that meet at the point $(\frac{1}{2}, \frac{1}{2})$. Your graph should clearly indicate the points with x -coordinates 0 and 1.
- (2) Explain briefly: If $\beta_1 > \frac{1}{2}$, then Player 1 offers Player 2 half of the Good Stuff, and Player 2 accepts.
- (3) Graph Player 1's payoff function on the interval $0 \leq x \leq 1$ assuming $\beta_1 < \frac{1}{2}$.
- (4) Graph Player 2's payoff function on the interval $0 \leq y \leq 1$ assuming $\beta_2 < 1$. Use your graph to explain the following statement: When $\beta_2 \leq 1$, Player 2 will accept any offer $y \geq y^*$ with $y^* = \alpha_2 / (1 + 2\alpha_2)$.
- (5) If $\beta_1 < \frac{1}{2}$ and $\beta_2 \leq 1$, what fraction of the Good Stuff should Player 1 offer to Player 2?
- (6) If $\beta_1 < \frac{1}{2}$ and $\beta_2 > 1$, what fraction of the Good Stuff should Player 1 offer to Player 2?