

TEXTBOOK

UNIT 9

UNIT 09

GAME THEORY

TEXTBOOK

UNIT OBJECTIVES

- Game theory is the mathematical study of social interactions as games that have payoffs for the players.
- Traditional game theory assumes that players are rational actors, always acting in ways that maximize their benefits. Real people are not necessarily like this.
- The payoff matrix is the essential way to express a game mathematically.
- In zero-sum games, a winner's gains come at a loser's expense.
- Non-zero-sum games can include both win-win and lose-lose situations.
- Strategies are the actions that players take in a game.
- Payoffs are often frequency dependent; that is, they depend on how many people are playing a particular strategy as well as how often a game is played.
- Equilibrium is reached when each player has no incentive to play differently.
- Game analysis changes when games are played repeatedly. This gives rise to mixed strategies.
- Game theory can provide insight into many situations, phenomena, and subjects, including biology, sociology, and linguistics.

“”

What accounts for TIT FOR TAT's robust success is its combination of being nice, retaliatory, forgiving, and clear. Its niceness prevents it from getting into unnecessary trouble. Its retaliation discourages the other side from persisting when defection is tried. Its forgiveness helps restore mutual cooperation. And its clarity makes it intelligible to the other player, thereby eliciting long-term cooperation.

ROBERT AXELROD

SECTION 9.1

INTRODUCTION

Mathematics has an on-again-off-again relationship with the real world. There are fields of mathematics that exist, more or less, solely “for themselves.” Researchers in these fields are primarily motivated by the most abstract of “what-if” conjectures. The field of topology is a prime example of this, as it is known primarily for the beauty of the thinking behind its results rather than for its connections to reality. Much of mathematics, however, has its foundation in the happenings of the world around us. The field of game theory, in some sense, represents the pinnacle of this type of mathematics. It can be thought of as the mathematical study of our human interactions.

Mathematics started off as a way to apply strict, rigorous thinking to the wildly complicated world around us. Throughout its evolution, mathematical thought has oscillated between two alternative paths of thinking. On the one hand, it has often been at the vanguard of quantitative science, helping to shed new light on previously incomprehensible phenomena through the power of logical thinking. On the other hand, it has sometimes been taken to levels of abstraction far beyond what any pragmatic scientist would be interested in. Even in its more applied incarnations, mathematics involves a great many simplifying assumptions. To get anywhere, we must reduce incredibly complicated situations from our everyday reality into problems that can be strictly defined and rigorously analyzed. The subject of game theory provides a great example of this mathematical reduction in action.

Games, to a mathematician, are simplified versions of situations that arise in the course of interactions among people, or any kinds of agents. A game represents an idealized situation that can be analyzed mathematically to shed light on how or why certain outcomes are reached. For example, the Ultimatum Game is a two-player situation in which one player is given an amount of money to share with another in any proportion desired. The receiving player then can choose either to accept or to decline the giver’s offer. If the offer is accepted, both players get their share of the money, as determined by the giver. If the receiver declines, no one gets any money. This represents a vastly simplified model of a business negotiation. As we will see later in this unit, this simple game has far-reaching and sometimes surprising implications for how humans judge the importance of justice and equality.

Most games have many simplifying assumptions, as is evident throughout this discussion, but the one that provides the foundation for all others is the assumption that players act rationally. The classic assumption is that a rational

SECTION 9.1

INTRODUCTION
CONTINUED

player will always act in a way that seems to maximize personal benefit. This assumption of rationality is what allows the mathematical analysis of games to work. If the people playing a game do not behave rationally, then the results of any game-theory-supported analysis may be less relevant. Nonetheless, the conclusions that are obtained, even under these idealized assumptions, can be useful. Interestingly, an important use of game theory has been to probe the limits of rationality. As we shall see in our study of the Ultimatum Game, there are multiple concepts of what is rational, in addition to the one based on maximizing profit.

One of the most interesting conclusions reached in game theory is that rational actions by both players can result in situations in which both players are worse off. The primary example of this is the Prisoner's Dilemma, which we will see in more detail later on in this unit.

In this unit we will examine the mathematical analysis of games, beginning with a bit of the history and some of the motivations behind the field's development. From there we will examine a few simple games in order to illustrate the basic terminology and concepts. We will then move on to more substantial games, such as the Prisoner's Dilemma and the Hawks and Doves game. We will examine the various types of solutions and equilibria that exist for these games, and we will see how these elements change, depending on whether a game is played once or more than once. In iterated and multi-player games, we will see how the payoff a player can expect from using a particular strategy depends on the strategies that others employ and how frequently they do so. Along the way, we will see how game theory can help us understand the processes of evolution, including the evolution of language. Finally, we will see how analyzing abstract games applies to real-life situations, such as business transactions, language development, and avoiding nuclear war.

SECTION 9.2

ORIGINS OF GAME THEORY

- Games People Play
- Zero-Sum
- Minimize the Worst-Case Scenario

GAMES PEOPLE PLAY

- Mathematicians attempt to analyze types of interactions by treating them as games in which players use strategies to obtain payoffs.
- Games like checkers and chess are games of perfect information; the board shows both players all the information needed to make the right decision.
- Games like poker are games of imperfect information; no player has enough information (the other players' cards are hidden) to make a guaranteed right decision.

Games have appeared throughout history, and in many different cultures, as a form of entertainment and education. The first mathematical analyses of certain games occurred in the 1700s, centered around the card game Le Her, but it was not until the mid-1920s that the field of game theory was properly founded by the multi-talented Hungarian mathematician and physicist, John von Neumann.

Von Neumann was an exceptional character and ranks among the brightest minds of the first half of the twentieth century. He was involved in many fields, including quantum physics, topology, computer science, and economics. He also worked on the Manhattan Project, the top-secret effort that built the first atomic bomb. In his leisure time, he was an avid poker player, and his interest in games and how they “work” is said to have stemmed partly from this fascination. Poker is a game in which mathematics and human nature square off. A player must not only be good at calculating odds, but must also be able to read the other players to tell if they are bluffing or not.

Poker players are able to bluff one another because poker is a game of imperfect information. As a player, you know your own cards and you know what the other players have wagered, but you do not know what cards they hold. It is this missing information that brings excitement to the game and makes it challenging. Many of the games we will explore in this discussion present situations similar to those that arise in poker, in that players do not know all of the information available and yet still must make decisions.

SECTION 9.2

ORIGINS OF GAME
THEORY
CONTINUED

Certain games of imperfect information can lead to interesting scenarios in which players unwittingly act against their own best interests.

By contrast, chess is a game with perfect information. In a game of chess, knowledge of the positions (i.e., the pieces and board set-up) and all historical information regarding the moves are available to both players. Players must use this information to formulate and modify strategies and tactics during a match. The challenge of chess lies not in the ambiguity of not knowing what your opponent has to work with, but rather in the extreme complexity of possible scenarios, each requiring a good deal of analysis. Despite their differences, poker and chess have one very important characteristic in common: both are zero-sum games.

ZERO SUM

- Zero-sum games are games in which a winner's payoff is equal to a loser's loss.
- Utility is an imprecise concept, but biological payoffs are measurable.

Zero-sum games are games in which one player's loss is exactly equal to another player's gain. If you win a hand at poker, your winnings will add up to the sum of your opponents' losses (unless you play at a casino, where the house takes a cut). At the end of the day, poker players (as a group) are no wealthier or poorer than when they started. Likewise, in chess, one person's victory is balanced by the opponent's loss. Even when a chess match ends in a stalemate, then both players are no better off than when they started—neither has gained or lost anything, except perhaps time, but that is not taken into account.

Not all situations in life, or game theory, are zero-sum situations, however. Two notable examples are business transactions and arms races. In an arms race, two or more nations compete to build the most—and the most destructive—weapons possible. This is a situation in which there are many outcomes that leave all “players” worse off than when they started. If they invest heavily in weapons, but then never use them, they have squandered a good deal of their resources and are poorer for it. Also, if the nations use their weapons and go to war, there is much squandering of resources, and all are worse off than before.

In business transactions, both parties exchange something for something that they find more valuable. If you trade ten dollars for a hat, you have obviously decided that the hat is more valuable to you than your ten dollars

SECTION 9.2

ORIGINS OF GAME
THEORY
CONTINUED

(or, at least, equally valuable). Likewise, your ten dollars is more valuable than the hat to the hat-seller. This idea of relative value is what economists call “utility.” It is through increased or decreased utility that non-zero-sum arrangements are possible.

Utility, however, has proved to be a rather controversial concept. When game theory is applied to biology, or evolution, the concept of utility comes through in the term “biological fitness.” In this context, game theory measures payoffs in terms of reproductive success—the objective number of offspring that a “player” leaves behind to carry on a genetic legacy.

MINIMIZE THE WORST-CASE SCENARIO

- The minimax theorem, attributable to Von Neumann and Morgenstern, states that players seek the strategy that minimizes their maximum loss.

Von Neumann focused the bulk of his research on zero-sum games. He and economist Oskar Morgenstern established the field of game theory with the publication of their book “The Theory of Games and Economic Behavior,” which was primarily an analysis of the zero-sum situation. Von Neumann and Morgenstern worked from the fundamental assumption that players always act in a way that increases their utility—that is, they always implement the strategy that they perceive will bring them the greatest reward. This is the classic definition of “rational,” in the context of game theory. Furthermore, Von Neumann showed that in any two-player, zero-sum, game, there will always be a best, or optimal, strategy for each player. This will be the strategy that maximizes the minimum possible gain, or utility, dependent on what the other player does.

This theorem, known as the “minimax theorem,” set the foundation for the mathematical study of games. It states that there is always a strategy that a player can choose that will lead to their most-favorable worst possible outcome. Depending on the specific rules of the game, that outcome might not be better than that of another player, but it will be better than the alternatives, provided the other player is playing in a similar fashion. For instance, even in a game as complex as chess, there exists an outcome that should happen every time, provided both players play perfectly. This means that, theoretically, one of the three possible outcomes of chess—a white win, a black win, or a draw—should be the “right” outcome if both players play perfectly. The game is so complex,

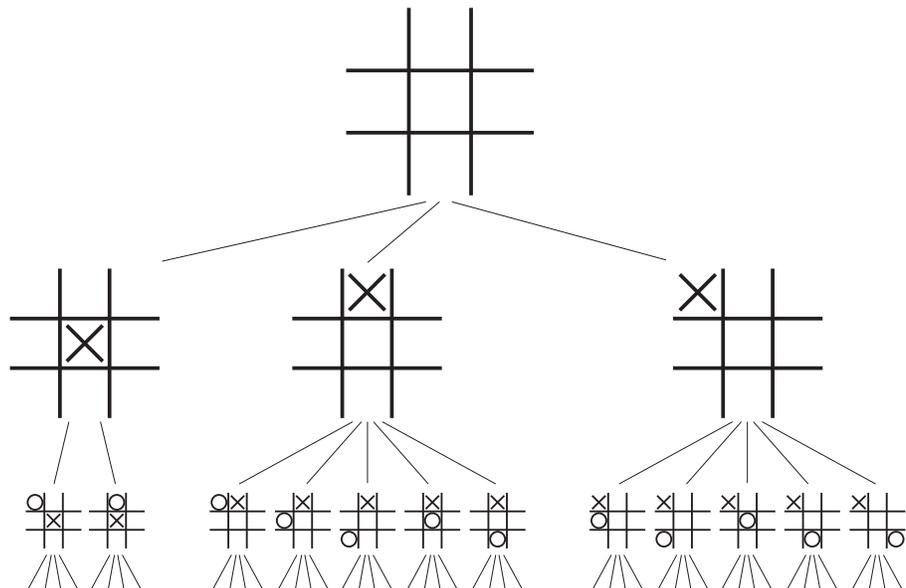
SECTION 9.2

ORIGINS OF GAME THEORY CONTINUED

however, that the strategies that each person should employ to produce this ideal outcome are unknown.

An easier example to consider is the game of tic-tac-toe (TTT). A good TTT player can never be beaten. A good player, given the first move, will always take a corner square, and any rational opponent will then take the center square. Play continues from this point, with the players making the moves that ensure that they will not lose, yet that also leave their best winning options open. Because of the layout and the rules of TTT, this means that when two good players play each other, the result will always be a tie. This is the “right” outcome, provided no one makes a mistake.

One way to analyze the game progression is through the use of a “game tree.” A game tree provides a systematic way to lay out all sequences of moves in a game in visual format. Each move is represented by a node, whose branches represent all the possible countermoves. Even for a game as simple as TTT, the game tree gets very large, as there are over 300,000 (9!) possible sequences of moves and countermoves that must be taken into account.



Now that we’ve been introduced to some of the general concepts and terms that will be used in our discussion of game theory, let’s turn our attention to a few specific games and analyses.

SECTION 9.3

SIMPLE GAMES

- Piece of Cake
- A Penny Saved...

PIECE OF CAKE

- Cake division is a very simple zero-sum game, modeled with a 2×2 payoff matrix.
- If each player is greedy, we assume that each will choose the strategy with the best worst-case scenario.
- Equilibrium is reached when both players have no incentive to change their strategies; these can be considered the “best” strategies.

Let’s start by looking at a simple situation that can be modeled as a game. Suppose that two children at a birthday party both want to have the last piece of cake. If one child gets it, the other will be resentful, and even if an adult intervenes to split the cake in two, one child will inevitably complain that the other’s share is larger. This conundrum can be avoided by letting one child cut the cake and letting the other child have first choice of the pieces. This seems to be an intuitively fair way to solve the problem. We can put this intuition on firmer footing, however, using the techniques of game theory.

To model this situation as a game, we have to make a few simplifying assumptions. First, the child who is to cut the cake, who we’ll call the “cutter,” has a variety of choices of how to make the cut, but we can simplify things by recognizing that the real decision is simply whether or not to attempt to cut the cake fairly. In this model, we reduce the cutter’s choices to just two: namely, cut evenly or cut unevenly. The chooser has only two possible actions, of course: choose the piece perceived to be larger or the one that seems smaller. Finally, we must assume that both children are completely selfish, or rational. That is, they always act in a way that gives them as much cake as possible.

We can organize this information into a matrix that enables us to see and analyze the various possible outcomes.

NOTE: The first value in each cell is the cutter’s payoff, and the second is the chooser’s payoff. We will follow this convention of listing the row player’s payoff first throughout the unit.

SECTION 9.3

PIECE OF CAKE CONTINUED



The last piece of cake

		Chooser	
		Choose larger	Choose smaller
Cutter	Cut evenly	()	()
	Cut unevenly	()	()

In this scheme, the cutter chooses the row of the outcome and the chooser chooses the column. For example, if the cutter chooses to cut evenly and the chooser chooses the larger piece, then the cutter will get “half minus a crumb” and the chooser will get “half plus a crumb,” which is the outcome represented in the upper left cell of the table. Allowing for the difference of a “crumb” is simply a way to acknowledge that actually cutting a cake evenly is extremely difficult.

Now, being selfish, the cutter will choose the action that promises to bring him the most cake regardless of what the chooser chooses. Cutting the cake unevenly creates the possibility of getting the larger piece, but it also opens the door for the chooser to thwart this effort. In other words, the cutter’s maximum payoff in choosing to cut unevenly is the large piece, and his minimum payoff in this case is the smaller piece.

If the cutter chooses to cut evenly, however, his maximum payoff is about half of the cake, and his minimum payoff is also about half of the cake. So, of the cutter’s two choices, the one that has the least downside-or, in other words, the “maximum minimum”-is the one he should choose. Consequently, in this situation, he should choose to cut the piece of cake evenly.

The chooser seeks to do the same thing, make the choice that maximizes her benefit. In this case, if the cutter cuts evenly, the chooser’s best option is to pick the “half plus a crumb.” Notice that even though both children implement their best strategy, one still comes out slightly advantaged over the other. This common feature of games is summed up in the following statement: “You know, the best you can expect is to avoid the worst.”¹ Games do not have to be fair.

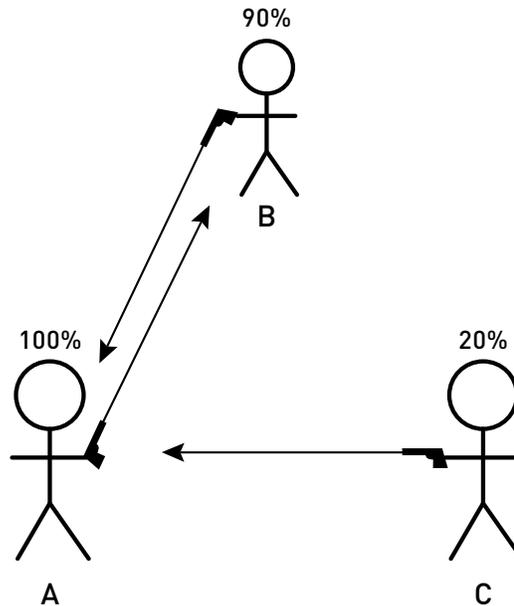
The choices that the cutter and chooser had to make in the above example are known as “pure strategies.” This simply means that the players play the game using the same strategy every time; deviating from the strategy gains

SECTION 9.3

PIECE OF CAKE CONTINUED

them nothing and could potentially end up harming them. This idea of a “best” strategy—in the sense that deviation from it increases the chance of reaching a less-desirable result—is known as an “equilibrium.”

A nice example of equilibrium is the case of a three-way duel (sometimes called a “truel”). Let’s say that person A is an excellent shot, able to hit the target 100% of the time; person B is a great shot, with a 90% success rate; and person C is a terrible shot, striking the target a mere 20% of the time.



Assuming that each person can shoot only once, the best strategy for each in this case is basically: “shoot the person most dangerous to you.” If each player adopts this strategy, then A should shoot at B and B should shoot at A, as each of them is the other’s most imminent threat. Person C should shoot at A, because there is a tiny chance that B will miss, but there is no chance that A will miss. The outcome, if all “players” implement their equilibrium strategies, is that

SECTION 9.3

PIECE OF CAKE CONTINUED

Person C will be the winner of the contest, the one left standing. This example demonstrates how the conclusions of game theory can sometimes be counter-intuitive.

A PENNY SAVED...

- Matching pennies, in its one-off form, is an example of a game in which there is no clear “best” strategy.
- Playing multiple rounds of matching pennies does have a clear equilibrium.

To get a better sense of how equilibrium works, let’s look at a game that has no clear “best” strategy: matching pennies. In this game, two players, called “Mixed” and “Matched,” simultaneously place one penny each on a table, either heads up or heads down. If the two coins are matching, then Matched gets to keep them both; if the two coins are not matching, then Mixed gets them both. We can summarize the situation with the following payoff matrix:

		Matched	
		Play heads	Play tails
Mixed	Play heads	$(-1, 1)$	$(1, -1)$
	Play tails	$(1, -1)$	$(-1, 1)$

We can see from this table that if Mixed plays heads and Matched also plays heads, then Mixed loses a penny and Matched gains a penny. Remember, in this game both players put their coins down at the same time, so neither has an advantage in knowing what to pick. Notice also that this is indeed a zero-sum situation—whatever Mixed loses is gained by Matched, and vice versa. If this game were to be played just once, neither Mixed nor Matched would have any clue as to what the other was going to play, so the choice between heads and tails would be completely random. Therefore, unlike the cake game discussed earlier in which each player had a definite “best” strategy, there is no one strategy that beats all others for a single round of matching pennies.

The plot thickens, so to speak, when multiple rounds are played; this is what game theorists call an “iterated game.” If the two players were to play multiple rounds, then playing a pure strategy of heads every time or of tails every time would definitely put a player at a disadvantage. For instance, if Matched noticed that Mixed always plays heads, then she should play heads as well and win every round.

SECTION 9.3

PIECE OF CAKE CONTINUED

A pure strategy is not the best bet for either side in this situation. Ideally, each player would like to keep the other player guessing as to the next play. The most intuitive, and best, way to accomplish this is to play randomly. Random play is an example of a mixed strategy. Let's look at a modified table that includes this new strategy option. Note that the payoff values in the table also must change a bit in meaning. Whereas previously we were concerned with the payoff of just a single round of play, we are now considering iterated games and mixed strategies, and the payoffs must represent averages per round. Playing randomly results in an average payoff of "zero" per round. Note that, because this is still a zero-sum situation, if one player gets zero, so must the other.

		Matched		
		Play heads	Play tails	Play randomly
Mixed	Play heads	(-1, 1)	(1, -1)	(0, 0)
	Play tails	(1, -1)	(-1, 1)	(0, 0)
	Play randomly	(0, 0)	(0, 0)	(0, 0)

Now, each player has a choice of how to play this iterated game—pure heads, pure tails, or randomly. Using the logic developed in the preceding section, Mixed should choose the strategy that ensures the maximum minimum, and Matched, from his perspective, should do the same. This means that Mixed should choose to play randomly, and so should Matched, and both of them should expect to make nothing from the game. Playing randomly in this case means playing heads and tails with equal probability. Doing this means that the game is at equilibrium: neither player has anything to gain by deviating from the chosen strategy if the other player does not deviate. Not every equilibrium in an iterated game must be composed of equal probabilities, however. The precise probabilities depend on the specific payoffs of the game.

In this analysis, we assume that playing randomly means that the odds of a player playing heads or playing tails are $\frac{50}{50}$. If this were not true, then the opposing player could statistically recognize a bias towards either heads or tails and adjust her play accordingly to take advantage of this. In a scenario such as this, the payoffs for playing randomly would no longer be (0,0), but rather the product of the pure strategy payoffs, (-1, 1) for example, and the proportion of heads or tails played. For example, if out of 100 games, Mixed plays 60% heads, then Matched should also play 60% heads and expect to have an average payoff of 0.1 per round as opposed to the zero that would be expected if both players play heads and tails with equal probability. Conversely, Mixed should expect to lose 0.1 per round, on average.

SECTION 9.3

PIECE OF CAKE CONTINUED

We have until now been concerned with zero-sum games; whatever the winner wins, the loser has to lose. However, many situations in life, and, hence, the games that model these situations, are not zero-sum. These are situations in which the combined outcome can be greater than or less than zero. In other words, some situations are win-win, and some situations are lose-lose. One of the most famous non-zero-sum games is the Prisoner's Dilemma.

SECTION 9.4

PRISONER'S DILEMMA

- Trust No One
- Déjà Vu

TRUST NO ONE

- The Prisoner's Dilemma is a classic example of a non-zero-sum game.
- The equilibrium in Prisoner's Dilemma is not the optimum solution.

The RAND Corporation, located in Santa Monica, California, is the original "think-tank." It was founded after World War II to be a center of national security and global policy ideas and analysis. Whereas today it advises many nations on a variety of issues, its initial focus was national defense. Game theory was one of the early pursuits of RAND thinkers, and in 1950 two RAND scientists, Merrill Flood and Melvin Dresher, framed what would become one of the most fascinating games of all time, the Prisoner's Dilemma.

The basic game is set up like this: imagine that you and your friend are caught robbing a bank. Upon being apprehended you are immediately separated so that you do not have time to communicate with each other. Each of you is taken to a separate cell for interrogation. If you and your buddy cooperate (C) with each other—that is, say nothing to the cops—each of you will get only a year in jail, known as the "reward" payoff, R.

		Your Buddy	
		C	D
You	C	(R, R)	
	D		

If you both rat on each other, or "defect" (D), to use the game theorist's terminology, you will both get three years in prison, known as the "punishment" payoff, P.

SECTION 9.4

PRISONER'S DILEMMA CONTINUED

		Your Buddy	
		C	D
You	C	(R, R)	
	D		(P, P)

If one of you cooperates and the other defects, the cooperator will get five years, known as the “sucker’s” payoff, S, and the defector will get off with no jail time, known as the “temptation to defect” payoff, T.

		Your Buddy	
		C	D
You	C	(R, R)	(S, T)
	D	(T, S)	(P, P)

This matrix concisely expresses the game as we have described it, where $T = 0$, $R = 1$, $P = 3$, and $S = 5$. Note that $T > R > P > S$. (It might be useful here to interpret the “is greater than” sign to mean “is better than,” because the values actually represent negatives-years spent in jail.)

First let’s consider why this is not a zero-sum game. Looking at each cell, we can tell that none of the payoffs for you and your buddy sum to zero. In fact, all of them result in some net jail time for one or both of you, although some outcomes are more favorable than others. For instance, if both of you cooperate, the total time served by the two of you will be two years, which is as close to win-win as this situation can get (after all, you did just rob a bank). If both of you defect, then the total jail time for the two of you will be six years, a lose-lose scenario that is a good deal worse than the best-case scenario. The other two scenarios result in a total of five years of jail time served between the two of you. So, if you could only agree with your buddy that both of you will keep quiet, as a team you’ll be better off. The dilemma comes from the fact that neither you nor your buddy has any incentive to do this.

You have no idea whether or not your buddy is going to cooperate. Even if you have discussed a situation like this with him beforehand, you cannot be sure that he won’t betray you. As a rational being, you are going to make the decision that minimizes your potential downside, or your personal maximum penalty. If you choose to cooperate with your buddy, the maximum penalty you could receive is five years, and your best-case scenario is a one-year prison sentence. However, if you choose to defect, your maximum penalty would be three years, and there

SECTION 9.4

PRISONER'S DILEMMA
CONTINUED

is a chance that you could get off with no jail time. Your rational buddy is faced with the same set of options and the same reasoning. As a rational being, you will choose to defect and so will your buddy, and these actions result in the lose-lose scenario.

What is so interesting in the Prisoner's Dilemma is that it is an example in which the equilibrium solution is not the same as the optimal solution. The equilibrium solution, remember, is the state in which neither player has anything to gain by switching strategies as long as the other player also doesn't switch. The optimal solution is the scenario in which the greatest good, or utility, is realized. In the Prisoner's Dilemma the greatest good, on the whole, comes about when both players cooperate. This scenario is unstable, however, because both players have an incentive to switch strategy. On the other hand, if both players defect, neither has anything to gain by changing strategy if the other doesn't, so the defect-defect solution is stable. Game theorists would say that the defect strategy is strictly dominant over the cooperate strategy as long as $T > R > P > S$.

When versions of the Prisoner's Dilemma are posed to actual people, the results do not always match the mathematical predictions. Real people do not always act rationally, and even if they did, it is very rare that a game is ever played just once in real life. As an example, let's say that you decide to cooperate, but your buddy decides to defect. After you serve your sentence, your buddy offers to rob another bank with you to help you get back on your feet (with friends like this, who needs enemies?!). You agree, and both of you get caught again. This situation is not exactly like the first time you got caught, however, because now each of you has a reputation, a track record. Your buddy might realize that you have already cooperated once and that if you cooperate again, and he chooses to cooperate this time also, then both of you will be better off. On the other hand, you might have revenge on your mind and decide that because your buddy burned you the last time, you will retaliate this time. These kinds of considerations make the Iterated Prisoner's Dilemma more complicated than the one-shot version.

SECTION 9.4

PRISONER'S DILEMMA CONTINUED

DÉJÀ VU

- The Iterated Prisoner's Dilemma admits a wide variety of equilibrium outcomes, depending on the mix of strategies adopted by the players.
- In computer tournaments, strategies that are neither always generous nor always punitive tend to fare the best.

If the Prisoner's Dilemma is to be played over and over again, it is best that the number of times that it is to be played is not pre-determined; otherwise, everyone should just defect, as in the one-round version. The reasoning goes like this: you should always defect in your last game because there is no chance for retaliation. Knowing that your buddy will also think of this strategy, you should always defect in your second-to-last game as well. This thinking naturally extends all the way back to the first move, so everyone should just always defect. If, however, players play without knowledge of when the game will end, strategies other than "always defect" become viable, even dominant. One such alternative strategy is the random strategy, in which a player randomly cooperates or defects, with no consideration given to what has happened in previous rounds. Another strategy is retaliation: always do to your opponent what she did to you the last time.

There are many strategies, some of which are clearly better than others, and others of which are rather obscure in their efficacy. To put all of these strategies to the test, Robert Axelrod of the University of Michigan organized a tournament in the mid-1980s in which different Iterated Prisoner's Dilemma strategies competed against each other over the course of many rounds. The winning strategy was to be the one with the lowest accumulated jail time in the end.

One might suspect that "always-defecting" would still be the best strategy in such a tournament. If two players played five rounds of the always-defecting strategy, their individual scores at the end of five rounds would be 15 years. (In our score-keeping, lower scores are better).

Results of "Pure Defect" vs. "Pure Defect"

P1 STRATEGY DDDDD

P2 STRATEGY DDDDD

P1 PAYOFF = PPPPP = 3 + 3 + 3 + 3 + 3 = 15 years

P2 PAYOFF = PPPPP = 3 + 3 + 3 + 3 + 3 = 15 years

SECTION 9.4

PRISONER'S DILEMMA CONTINUED

On the other hand, if two players who were “always-cooperating” played each other for five rounds, each player’s accumulated score would be five years.

Results of “Pure Cooperate” vs. “Pure Cooperate”

P1 STRATEGY CCCCC

P2 STRATEGY CCCCC

P1 PAYOFF = RRRRR = $1 + 1 + 1 + 1 + 1 = 5$ years

P2 PAYOFF = RRRRR = $1 + 1 + 1 + 1 + 1 = 5$ years

So, there is clearly something to be gained by not defecting all the time if you can get into a repeated mutual cooperation situation, such as that shown above. The question is: how can you get into such a situation, especially when a “Pure Defect” strategy will dominate a “Pure Cooperate” strategy?

Results of “Pure Defect” vs. “Pure Cooperate”

P1 STRATEGY DDDDD

P2 STRATEGY CCCCC

P1 PAYOFF = TTTTT = $0 + 0 + 0 + 0 + 0 = 0$ years

P2 PAYOFF = SSSSS = $5 + 5 + 5 + 5 + 5 = 25$ years

Analysis of the strategies that fared best in Axelrod’s tournament indeed provided answers to this question. Some were very complicated, based on analyzing specific sequences of prior moves to prescribe the next sequence of moves. Others were very simple, such as the Tit-For-Tat (TFT) strategy. As its name implies, TFT relies simply on doing to your opponent what he last did to you. So, if your opponent cooperates on the first turn, then you should cooperate on the second turn. This can lead to the nice “always-cooperating” cycle if two TFT players start off cooperating, while protecting the player from getting too many sucker’s payoffs. However, TFT can also lead to the “always-defecting” situation, if two TFT players start off by defecting.

The best strategies were variants of Tit-For-Tat with Forgiveness (TFTWF). This strategy is basically the same as regular TFT except that some small percentage of the time, you forgive your opponent’s prior defection and do not mimic it. This provides a mechanism for breaking the “always-defecting” trap.

SECTION 9.4

Results of “Tit-For-Tat-with-Forgiveness” vs. “Tit-For-Tat”:

PRISONER’S DILEMMA
CONTINUED

P1 (TFTWF) DDCCC

P2 (TFT) DDDCC

P1 PAYOFF = PPSRR = $3 + 3 + 5 + 1 + 1 = 13$ yearsP2 PAYOFF = PPTRR = $3 + 3 + 0 + 1 + 1 = 8$ years

Note that in this particular match, TFTWF loses to TFT. Remember, however, that the tournament consists of many matches against many different strategies. TFT will invariably get caught in “always defecting” cycles, whereas TFTWF will be able to escape these, providing an advantage over TFT in the long run.

Most of the successful strategies in Axelrod’s tournament were based on some amount of altruistic behavior. It was a stunning mathematical indication that aggression and vindictiveness do not always prevail. It seems that a truly selfish strategy, in the sense that it is designed to maximize one’s own benefit, must include some element of forgiveness. In fact, Axelrod found that successful strategies had four common traits, which he described anthropomorphically in this way:

- First, the strategy should be “nice.” This means that it will not defect unless its opponent defects first.
- Second, the strategy should retaliate against defectors to avoid being exploited by “always defectors.”
- Third, the strategy should be forgiving. After retaliating against a defection, it should begin to cooperate again as soon as its opponent cooperates.
- Fourth, the strategy should not try to score more than its opponent—it should be non-envious. This stems from the fact that the strength of cooperation lies in the reality that both parties benefit equally from it.

SECTION 9.4

PRISONER'S DILEMMA CONTINUED

These traits are fascinating if not heartwarming, showing us that cooperation and altruism really do have a place in a world as starkly defined as that of the Iterated Prisoner's Dilemma. This suggests that studying games can help us to understand some of the behavioral aspects of our natural world, such as why certain types of animals live in cooperative societies and others live as solitary aggressors. This world of conflicting living strategies is characterized by the game called the "Hawks and Doves," and it is to this game that we will next turn our attention.

SECTION 9.5

HAWKS AND DOVES

- A Lover, Not a Fighter
- Might Might Not Always Make Right
- The Nitty Gritty

A LOVER, NOT A FIGHTER

- Hawks and Doves is a rudimentary model of game theory in a biological context in which creatures compete for resources by being either aggressive or passive.
- Aggressive players beat passive players, but they incur costs when they compete against other aggressive players.

In the Prisoner's Dilemma, the players have a choice of whether or not to cooperate with one another or to defect. We saw that what the players should do—that is, their best strategies—depend on whether they will be playing just once or many times. If they are to play only once, they should both defect, even though it would be better overall for them to cooperate. If they are to play many times, however, it behooves them to try different mixes of cooperation and defection. Axelrod's tournament study demonstrated that, over the course of numerous games, players who play pure strategies can be beaten by players who choose mixed strategies.

We can extend this thinking to the natural world if we imagine the competition for survival to be a tournament. The many rounds of this tournament correspond to the daily struggles that certain species face for survival. In the natural world, all species compete for resources using wildly varying strategies. The strategies that are most successful in the long run are the ones that survive to be passed along to offspring.

Strategies for survival in the natural world do vary, but broad trends are discernible. For our purposes, we will simplify things greatly by limiting the options to only two types of behavior, aggressive and passive, and we'll call the actors of these behaviors "Hawks" and "Doves" respectively. Aggressive animals will always fight over resources, whereas passive animals will not. This is the basis for a famous game often known as "Hawks and Doves," which was first proposed by John Maynard Smith and George Price in a 1973 paper.

SECTION 9.5

HAWKS AND DOVES CONTINUED

The assumptions behind the game are pretty straightforward. Imagine a field strewn with piles of food. This field is populated with animals that can behave either passively or aggressively toward one another. The animals, or players, compete with one another for the resource piles. For simplicity's sake, we determine that all competitions occur between individuals, i.e., one-on-one. The standard scenario is that one animal approaches a pile of food, and then another animal presents a challenge for it. Furthermore, neither animal knows the other's behavioral identity until the challenge has begun. The possible interactions are then Hawk-Hawk, Hawk-Dove, Dove-Hawk, and Dove-Dove.

Player 1	Player 2	
	Hawk	Dove
	Hawk	
	Dove	

Whenever a Hawk fights another Hawk, one of them wins the entire food pile, thus getting a benefit, B . The loser gets injured, incurring a cost, C . Both B and C can be thought of as food calories. The resource calories gained by the winner are counted as positive, but the calories that the loser must devote to healing are counted as negative. For the purposes of this game, we assume that all Hawks are equal and win half of all their battles with other Hawks. Also, as with the Prisoner's Dilemma game, we are concerned only with the tendencies established through iterative scenarios. Consequently, on average, a Hawk will gain $\frac{B}{2}$ calories and lose $\frac{C}{2}$ calories in a Hawk-Hawk interaction. This average energy accounting can be simplified to $\frac{B-C}{2}$.

Player 1	Player 2	
	Hawk	Dove
	Hawk	$\frac{B-C}{2}, \frac{B-C}{2}$
	Dove	

When a Hawk challenges a Dove, the Dove does not fight but simply walks away. This means that the Hawk gets the entire benefit, with no cost of fighting. The Dove gets nothing, but also loses nothing. In Hawk-Dove and Dove-Hawk interactions, the Hawk always gets the entire benefit, B , and the Dove always gets nothing and loses nothing. Therefore, the Hawk's average payoff is B , and the Dove's is zero.

SECTION 9.5

HAWKS AND DOVES CONTINUED

		Player 2	
		Hawk	Dove
Player 1	Hawk	$\frac{(B-C)}{2}, \frac{(B-C)}{2}$	(B, 0)
	Dove	(0, B)	

Finally, when a Dove challenges a Dove, they do not fight but, rather, split the resource evenly. Each player gets $\frac{B}{2}$ without any cost to anyone.

		Player 2	
		Hawk	Dove
Player 1	Hawk	$\frac{(B-C)}{2}, \frac{(B-C)}{2}$	(B, 0)
	Dove	(0, B)	$(\frac{B}{2}, \frac{B}{2})$

Notice that this is not a zero-sum situation, because not all of the cells add up to the same value; the Hawk-Hawk interaction yields less in total benefit than the other three scenarios.

MIGHT MIGHT NOT ALWAYS MAKE RIGHT

- Pure strategies can be a bad idea.

Now that we have a grasp of the basic circumstances of the game, let's think about whether it's better to be a Hawk or a Dove. It might seem, at first glance, that being a Hawk is always the best idea. If we imagine that the population of the field is nearly 100% Hawk, it's hard to see how a Dove could ever survive for very long, as it would get to eat only upon encountering another Dove. On the other hand, if the field is nearly 100% Dove, then a single Hawk is going to have it incredibly easy. This would lead us to think, if we had to choose between playing Hawk and playing Dove, that we should always choose Hawk. After all, a Hawk in an all-Dove world is going to do well, whereas a Dove in an all-Hawk world is going to starve.

That's the standard intuition, but let's consider the situation of the lone Dove a little more carefully. He never loses calories, and while the Hawks are gaining calories, they are also losing them in their fights. If these costs end up being more than the resource benefits, then each Hawk will experience an overall calorie loss as time goes on, while the Dove holds steady (this assumes, of

SECTION 9.5

HAWKS AND DOVES
CONTINUED

course, that there is no cost for simply waiting around while everybody else fights amongst themselves). After a while, the lone Dove will be doing much better than the always-fighting Hawks. This suggests that if costs are more than benefits, one might do well to be a Dove in an all-Hawk world.

If Doves do better in an all-Hawk environment when costs outweigh benefits, then over time the population should shift toward all Doves. This is based on the assumption that the most fit, the ones with the highest net calories, survive to reproduce more often than the less fit.

One might then think that whenever costs outweigh benefits, the population will tend to evolve into all Doves. However, a Hawk in an all-Dove environment will do extremely well relative to the Doves, even if costs outweigh benefits. This is because the cost becomes irrelevant if there are no other Hawks around to inflict injuries. We would then be led to believe that the all-Dove scenario is not stable, even when costs outweigh benefits.

The idea of a pure strategy's stability is an important one. In our analysis, we saw that neither the all-Hawk nor the all-Dove strategy is stable when the costs outweigh the benefits. This means that either situation can be infiltrated by the opposing strategy. Note that this is not true when the benefits outweigh the costs. Such a world would be driven towards the all-Hawk state, as a lone Dove would gain nothing while the Hawks gained something from each fight. This suggests to us that the relationship between costs and benefits has something to do with which state will be stable. Furthermore, we can conclude that because neither the all-Hawk nor the all-Dove state is stable, if there is to be a stable state, it must lie somewhere between the pure states. This means that if one has a choice as to whether to be a Hawk or a Dove, it would be best to adopt a mixture of the strategies-but what mixture? Remember that on the level of each individual confrontation, you have to choose your identity, whether to be a Hawk or a Dove, before you know the identity of your opponent. What percentage of the time should you be a Hawk and what percentage of the time should you be a Dove? With just a little algebra, we can find these percentages:

SECTION 9.5

HAWKS AND DOVES CONTINUED

THE NITTY GRITTY

- The optimum mix of passive and aggressive behavior depends on the exact values of the costs and benefits.

To start, let's represent the pure Hawk strategy as H, the pure dove strategy as D, and the Mixed strategy as S. The payoffs for these would be as follows:

$E(H,S)$ = payoff of pure Hawk versus the Mixed strategy

$E(D,S)$ = payoff of pure Dove versus the Mixed strategy

Let's define p as the probability that the Mixed player plays Hawk in a given interaction; then the expression $1-p$ represents the probability that the Mixed player will play Dove. The expected average payoff of H vs. S, $E(H,S)$, will be composed of part of the Hawk-Hawk and part of the Hawk-Dove payoffs.

$E(H,S)$ = (probability that S plays Hawk) x (payoff of Hawk-Hawk) + (probability that S plays Dove) x (payoff of Hawk-Dove)

$$E(H,S) = \frac{p(B-C)}{2} + (1-p)B$$

The expected average payoff for D vs. S, $E(D,S)$, can be found in a similar manner:

$E(D,S)$ = (probability that S plays Hawk) x (payoff of Dove-Hawk) + (probability that S plays Dove) x (payoff of Dove-Dove)

$$E(D,S) = p(0) + (1-p)\left(\frac{B}{2}\right)$$

S's optimum mix will be when both H and D do equally well against it. This means that S has nothing to gain by skewing the mix towards more Hawk or more Dove than prescribed by p and $(1-p)$ respectively. In other words, the optimal mix will be the value of p when $E(H,S) = E(D,S)$.

$$E(H,S) = E(D,S)$$

$$\frac{p(B-C)}{2} + (1-p)B = p(0) + (1-p)\left(\frac{B}{2}\right)$$

Solving this for p yields the percentage of time that S should play Hawk, which turns out to be $\frac{B}{C}$. Note that this percentage is entirely dependent on the benefit-to-cost ratio.

SECTION 9.5

HAWKS AND DOVES
CONTINUED

All of this means that were we to study the population of our field for a long time, we would find that the ratio of benefits given by food piles to costs incurred by fighting would determine the percentage of time that a Mixed animal should play Hawk or Dove. If, for some reason, the system falls out of balance, as when a group of players decides to play Hawk more often than they should, then there will be a clear advantage for others to play Dove more than they should. These counteracting forces would then drive the system back to the appropriate average ratio of Hawks to Doves.

The evolutionary progress of our field, like the process of Axelrod's tournament, shows that pure strategies are neither always stable, nor always optimal. The most successful strategies are usually mixed strategies. In terms of human behavior, this suggests that to be successful, we should not be too quarrelsome, nor should we be pushovers. Additionally, we should be forgiving at times, and at other times we should not hesitate to retaliate against wrongdoers. These conclusions are all well and good in theory, but how do they play out in real life with actual human beings? In our next section, we will examine what happens when game theory's predictions are put to the test in different human cultures.

SECTION 9.6

**FAIRNESS IN
DIFFERENT CULTURES**

- No One Said Life Was Fair
- Get Real

NO ONE EVER SAID LIFE WAS FAIR

- Games can be used as a sociologist's measuring stick to quantify notions of fairness in human cultures.
- The Ultimatum Game gives one player a sum of resources to be shared with another player, who can accept or reject the offer.

What does "fair" really mean? Does it mean the same thing to everybody? Sociologists have been able to explore these questions using the techniques of game theory. Games can serve as one of the essential tools of the sociologist, much as litmus paper serves as a tool for the chemist or a telescope serves as a tool for the astronomer.

First, let's clarify the difference between the terms "rational" and "fair." A rational action, as we have defined it, is one in which a player chooses the strategy with the best chance of producing the most personal benefit, without regard to what happens to the other player. Being fair, on the other hand, takes into account a whole host of other factors, including cultural norms, experience in market transactions, and experience with cooperation. In work done at the turn of the twenty-first century, researchers found that the concept of what is "fair" ranges widely, depending on who's playing. They reached this conclusion after watching how people from 17 different small-scale societies, ranging from hunter-gatherers to nomadic herders to sedentary farmers, played in a variety of cooperative games, such as the Ultimatum Game and the Public Goods Game.

In the Public Goods Game, players are asked to contribute some amount of money to a communal pot, which will be subsequently increased, based on how much everyone gives. In the Ultimatum Game, one player is given a sum of money or other valuable resource and is instructed to share it with another player. The first player decides how much to offer and the second player decides whether or not to accept the offer. If the second player rejects the offer, neither player gets any reward, or benefit.

SECTION 9.6

FAIRNESS IN DIFFERENT CULTURES CONTINUED

Let's examine the Ultimatum Game in a bit more detail. Player 1, the Offerer, can offer any amount that he or she chooses. For the sake of simplicity, let's say that the Offerer can choose to offer a high amount (H) or a low amount (L). If he offers H, then he will be left with L if Player 2 accepts the offer, and vice versa. Player 2, the Receiver, always has the choice of accepting or rejecting the offer. With these simplified assumptions we can create a matrix:

Offerer	Receiver	
	Accept Offer	Reject Offer
Offer High	(L, H)	(0, 0)
Offer Low	(H, L)	(0, 0)

It should be evident from this matrix that a rational Receiver will never reject an offer. From the rational Receiver's point of view, receiving L, even if L is of very low value, is better than getting nothing. A rational Offerer will pick the strategy corresponding to the row with the largest minimum payoff. Both rows in this case have the same minimum, 0, so the Offerer should then choose the strategy with the best potential payoff, which will be to Offer Low. In fact, the rational Offerer should offer the smallest amount possible, because the rational Receiver accepts any offer.

GET REAL

- The notion of what is fair depends on cultural norms.

When actual people play this game, however, the results vary widely and are never in line with the rational model. The study found that average offers across all societies range from 25% of the total to more than 50%. Furthermore, many real players will reject offers, even offers of more than 50%. What is perhaps more illuminating is how offers and acceptances depend on the society in which the players live.

Certain groups of people who are very economically independent, at least at the family level, had the lowest average offers. Other groups of people who depend on communal cooperation to gain food, such as in a whale hunt, had mean offers very close to 50%. Still others, in societies in which gift-giving is an act of status, had average offers above 50%. Quite surprisingly, some of these high-offer societies exhibited high rejection rates as well.

SECTION 9.6

FAIRNESS IN
DIFFERENT CULTURES
CONTINUED

Why would someone reject an offer? The answer relates to the psychology inherent in reiterative games. The researchers surmised that people reject offers that are too low because if they accepted such offers, they would develop a reputation for accepting low offers and, consequently, no one would give them higher offers in the future. Also, rejecting an offer turns the tables of power in the Receiver's favor. The Receiver can punish the low Offerer, who has much more to lose in a rejection than the Receiver does. From the Receiver's point of view, it might be worth incurring the cost of losing the low offer if it discourages the Offerer from being so stingy with future offers.

Why would anyone reject a high offer? In certain cultures, gift giving obligates the receiver to return the favor; receivers who do not wish to be obligated to someone else would then reject any offer that seemed to be too big a burden to pay back. These cultural norms were thought to manifest themselves in how people played the Ultimatum Game, as the participants sought to contextualize their experience of the game. In other words, they often asked themselves, "What does this game remind me of?" and then they adjusted their strategy to align with their perception of the situation.

The Public Goods Game and the Ultimatum Game show that what people perceive as being fair depends heavily on their cultural context. In these cases, games served as tools for measuring and quantifying cultural values in the real world. We see that the concept of fairness develops in human societies in relation to their specific needs and values. Game theory can also be used to examine another very human concept, that of language. We will now turn our attention to how ideas from game theory can contribute to the explanation of how language can arise and develop within a group.

SECTION 9.7

LANGUAGE

- The Name Game
- Pass It On

THE NAME GAME

- Language development can be modeled as a game in which players receive a payoff when they are in agreement about a particular word referring to a particular object.
- The payoff depends on how many players are in agreement.

Human language seems to be a large part of what makes us unique beings. We said in the introduction to this unit that game theory can be thought of as the mathematical study of human interactions. Language is arguably the most fundamental of these interactions. It is, in fact, hard to imagine interactions without language in the first place. Yet, how does a group of people agree on which words to use for particular objects? How is this set of agreements passed along to offspring? In 1999, in a paper written by Nowak, Plotkin, and Krakauer at the Princeton Institute for Advanced Study, principles of game theory were used to illuminate how language can develop.

Imagine a group of human ancestors, a troop of hominids, if you like. Suppose that this group is just starting to communicate about specific objects in their environment. Perhaps they have become concerned with communicating specific threats or other important information more efficiently—knowing the location of a hungry leopard or of a grove of fruit-laden trees is often critical to survival. Each individual in the group develops an internal list of verbal signals, or words, that are associated with objects such as leopards or fruit trees.

In order for communication to take place, a speaker makes an association between an object and a word. A listener either has the same association or they do not. Suppose that the speaker, upon seeing a leopard, says “leopard” to the listener. If the listener has the same association, then they will think “leopard” and act accordingly. If the listener does not have the same association, then they will not understand, and there could be some negative consequence.

SECTION 9.7

LANGUAGE
CONTINUED

This implies that if the speaker and the listener have the same association, then there is some sort of payoff for both of them. That payoff could be that the listener avoids danger, or perhaps learns the location of some food. The payoff for the speaker will be the same, if we imagine that both individuals use the same word in the future. In this discussion, we will assume that, as with the Hawks and Doves game, this language game is played more than once.

When the speaker wishes to alert the listener to the presence of the leopard, there is a certain probability that the speaker will use a given word. Likewise, there is a certain probability that the listener will associate the speaker's word with the concept "leopard." The maximum payoff for these two will increase as the probability that each player uses the same word for "leopard" increases. Payoffs also increase as more players adopt the same vocabulary and associations. With this kind of payoff structure in place, there is an incentive for players to understand each other, which can lead, over time, to the development of a common language.

PASS IT ON

- Languages are passed on to younger generations in a variety of ways.

Long-term development of language requires agreed-upon, object-signal associations to be passed down to new generations of speakers and listeners. How is this language transmitted to new generations?

Nowak, Plotkin, and Krakauer identified three main methods of language transmission. The first, and perhaps most intuitive, is parental transmission. Children tend to acquire the language of their parents and in this mode of transmission, greater language "fitness" (average payoff of one's list of associated signals and objects) would correlate with greater biological "fitness." In other words, the successful use of language can affect one's chances of passing on genes successfully to the next generation.

The second mode of transmission was identified to be through a role model outside of the family. In this mode, a high-status member of the group gains many young imitators. High-ranking role models illustrate the connection between language and status. So, if a child imitates the language profile of a high-status individual, that child will, on average, out-compete the children who do not imitate high-status individuals.

SECTION 9.7

LANGUAGE CONTINUED

The final mode of transmission is simply random learning. In this scenario, there are no clear incentives for learning language from any particular individual; instead, children imitate a random mixture of adults without regard to status or payoff. This tends to maximize confusion and, thus, minimize the payoffs that can accrue from mutual understanding. Groups who transmit language via this method tend to take a significantly greater amount of time to develop a common language, as opposed to groups that use the other two transmission methods.

SECTION 9.2

ORIGINS OF GAME THEORY

- Mathematicians attempt to analyze types of interactions by treating them as games in which players use strategies to obtain payoffs.
- Games like checkers and chess are games of perfect information; the board shows both players all the information needed to make the right decision.
- Games like poker are games of imperfect information; no player has enough information (the other players' cards are hidden) to make a guaranteed right decision.
- Zero-sum games are games in which a winner's payoff is equal to a loser's loss.
- Utility is an imprecise concept, but biological payoffs are measurable.
- The minimax theorem, attributable to Von Neumann and Morgenstern, states that players seek the strategy that minimizes their maximum loss.

SECTION 9.3

SIMPLE GAMES

- Cake division is a very simple zero-sum game, modeled with a 2×2 payoff matrix.
- If each player is greedy, we assume that each will choose the strategy with the best worst-case scenario.
- Equilibrium is reached when both players have no incentive to change their strategies; these can be considered the "best" strategies.
- Matching pennies, in its one-off form, is an example of a game in which there is no clear "best" strategy.
- Playing multiple rounds of matching pennies does have a clear equilibrium.

SECTION 9.4

PRISONER'S DILEMMA

- The Prisoner's Dilemma is a classic example of a non-zero-sum game.
- The equilibrium in Prisoner's Dilemma is not the optimum solution.
- The Iterated Prisoner's Dilemma admits a wide variety of equilibrium outcomes, depending on the mix of strategies adopted by the players.
- In computer tournaments, strategies that are neither always generous nor always punitive tend to fare the best.

SECTION 9.5

HAWKS AND DOVES

- Hawks and Doves is a rudimentary model of game theory in a biological context in which creatures compete for resources by being either aggressive or passive.
- Aggressive players beat passive players, but they incur costs when they compete against other aggressive players.
- Pure strategies can be a bad idea.
- The optimum mix of passive and aggressive behavior depends on the exact values of the costs and benefits.

SECTION 9.6

FAIRNESS IN DIFFERENT CULTURES

- Games can be used as a sociologist's measuring stick to quantify notions of fairness in human cultures.
- The Ultimatum Game gives one player a sum of resources to be shared with another player, who can accept or reject the offer.
- The notion of what is fair depends on cultural norms.

SECTION 9.7

LANGUAGE

- Language development can be modeled as a game in which players receive a payoff when they are in agreement about a particular word referring to a particular object.
- The payoff depends on how many players are in agreement.
- Languages are passed on to younger generations in a variety of ways.

BIBLIOGRAPHY

PRINT

Axelrod, Robert. *The Evolution of Cooperation*. USA: Basic Books (Perseus Books Group), 1984.

Benjamin, Arthur T. and A.J. Goldman. "Analysis of N-Card le Her," *Journal of Optimization Theory and Applications*, vol. 114, no. 3 (September 2002).

Barash, David P. *The Survival Game: How Game Theory Explains the Biology of Cooperation and Competition*. New York: Times Books: Henry Holt and Co., 2003.

Bernstein, Peter L. *Against the Gods: The Remarkable Story of Risk*. New York: John Wiley and Sons, 1996.

Dutta, Prajit K. *Strategies and Games: Theory and Practice*. Cambridge, MA: MIT Press, 1999.

Eshel, Ilan and L.L. Cavalli-Sforza. "Assortment of Encounters and Evolution of Cooperativeness," *Proceedings of the National Academy of Sciences, USA*, vol. 79, no. 4 (1982).

Ficici, S.G., O. Melnik, and J.B. Pollack. "A Game-Theoretic Investigation of Selection Methods Used in Evolutionary Algorithms," *Proceedings of the 2000 Congress on Evolutionary Computation CEC00*, (2000).

Fisher, Sir Ronald Aylmer. "Randomisation and an Old Enigma of Card Play," *Mathematical Gazette*, vol. 18 (1934).

Grassly, N.C., A. von Haeseler, and D. Krakauer. "Error, Population Structure, and the Origin of Diverse Sign Systems," *Journal of Theoretical Biology*, vol. 206, no. 3 (2000).

Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath. "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies," *The American Economic Review*, vol. 91, no. 2, Papers and Proceedings of the 113th Annual Meeting of the American Economic Association (May 2001).

Henrich, Joseph. "Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga of the Peruvian Amazon," *The American Economic Review*, vol. 90, no. 4 (September 2000).

Jager, G. "Evolutionary Game Theory and Linguistic Typology: A Case Study," In P. Dekker, editor, *Proceedings of the 14th Amsterdam Colloquium*. ILLC, University of Amsterdam, (2003).

Mero, Laszlo. [Translated by Anna C. Gosi-Greguss. English version edited by David Kramer.] *Moral Calculations: Game Theory, Logic, and Human Frailty*. New York: Copernicus Springer-Verlag New York, Inc., 1998.

Nowak M.A. and R.M. May. "Evolutionary Games and Spatial Chaos," *Nature* 359 (1992).

Nowak M.A., S. Bonhoeffer, and R.M. May. "More Spatial Games," *International Journal of Bifurcation and Chaos*, vol. 4, issue 1 (February 1994).

Nowak, M.A., J.B. Plotkin, and D. Krakauer. "The Evolutionary Language Game," *Journal of Theoretical Biology*, vol. 200, Issue 2 (21 September 1999).

Nowak, M.A., Karen M. Page, and Karl Sigmund. "Fairness Versus Reason in the Ultimatum Game," *Science*, vol. 289, issue 5485 (2000).

Poundstone, William. *Prisoner's Dilemma: John Von Neumann, Game Theory, and the Puzzle of the Bomb*. New York: Doubleday, 1992.

Siegfried, Tom. *A Beautiful Math: John Nash, Game Theory, and the Modern Quest for a Code of Nature*. Washington, D.C.: John Henry Press, 2006.

Von Neumann, John and Oskar Morgenstern. *Theory of Games and Economic Behavior* (6th paperback edition). Princeton, NJ: Princeton University Press, 1990.

¹ Calvino, Italo. (translated by William Weaver) *If on a winter's night a traveler*. San Diego-New York – London: Harvest Books/Harcourt Inc., 1981; p 4.

UNIT 9

GAME THEORY
TEXTBOOK

NOTES