13

Introduction to Game Theory

In Chapter 10 we examined the importance of interdependence among firms in pricing and output decisions in connection with duopolistic and oligopolistic market structure. In particular, we saw how the output and pricing decisions of one firm affect, and are affected by, the pricing and output decisions of other firms in the same industry. Moreover, we saw that this interdependency in the managerial decision-making process tends to become more pronounced the smaller the number of firms in the industry or, which is nearly the same thing, as two or more firms grow large enough to dominate industry supply.

In this chapter we take a closer look at a very important analytical tool that was only briefly examined in our discussion of oligopolistic behavior. This chapter is devoted to a more detailed examination of *game theory*. As we mentioned in Chapter 10, game theory is perhaps the most important tool in the economist's analytical kit for analyzing *strategic behavior*. Strategic behavior is concerned with how individuals make decisions when they recognize that their actions affect, and are affected by, the actions of other individuals or groups. In other words, strategic behavior recognizes that the decision-making process is frequently mutually interdependent.

Definition: Strategic behavior reflects recognition that decisions of competing individuals and groups are mutually interdependent.

As we noted in our discussion of oligopolistic markets in Chapter 10, game theory has numerous and widespread applications for analyzing the managerial decision-making process. It is a topic without which no textbook in managerial economics would be complete.

GAMES AND STRATEGIC BEHAVIOR

Most of our treatment of the behavior of profit-maximizing firms has been rather mechanistic in the sense that managers make pricing and output decisions without regard to the actions of their competitors. While this argument may be more or less correct for firms operating at the extreme end of the competitive spectrum (perfect competition and monopoly), it is far less likely apply to intermediate market structures, such as monopolistic competition, duopoly, and oligopoly. As we noted in Chapter 10, decision making by firms in oligopolistic industries is characterized by strategic behavior: the actions that result because individuals and groups that make decisions recognize that their actions affect, and are affected by, the actions of other individuals or groups.

In many respects, running a business is like playing a game of football or chess. The object of the game is to achieve an optimal outcome. But unlike these games, the "best" outcome does not always mean that your opponent loses. As we will see, the best outcome often results when the players cooperate. When cooperation is not possible, or illegal, then the objective is to win the game. But, victory does not always go to the strongest, or the fastest, or the most talented. Very often, victory belongs to the player who best understands the rules and has the superior game plan. The purpose of this chapter is to learn how to play a good game, regardless of whether cooperation and mutually beneficial outcomes are possible.

What is a game? Most of us think of a game as an activity involving two or more individuals, or teams, hereinafter referred to as players, in competition with each other. In general, the objective is to win the game because "to the winner go the spoils." Sometimes the spoils are little more than "bragging rights," often symbolized by a metal or glass artifact (trophy) of undistinguished design. Sometimes the spoils are monetary. Sometimes the winner wins both cash and trinkets—for example, to the owner of the winning team in the Super Bowl is presented the sterling silver Vince Lombardi Trophy and each player receives a gold and diamond ring and a cash award. After winning Super Bowl XXXV, each member of the Baltimore Ravens, received \$58,000, while the losing New York Giants, received \$34,500 per player. In business, we tend to think of the winning "team" as the firm that earns the greatest profits, or captures the largest market share, or achieves some objective more successfully than its rivals. Unlike football games, however, sometimes it is in the best interest of the teams to cooperate to achieve a mutually advantageous outcome.

In general, all games involve social and economic interactions in which the decisions made by one player affect, and are affected by, the decisions made by other players. It would be foolish for a chess player to make a move without first considering the prior play of his or her opponent. It would not make much business sense for the owner of a gasoline station to set prices for the various grades of gasoline without first considering the prices charged by other gas station owners in the neighborhood. It is this interdependency in the decision-making process that is at the heart of all games.

In game theory, decision makers are called players. Players make decisions based on strategies. These decisions dictate the players' moves. Players with the best strategies very often win the game, although this does not always happen. The rules of the game dictate the manner in which the players move. In a *simultaneous-move game*, it is useful to think of players moving at the same time. Simultaneous-move games are sometimes referred to as *static games*. Examples of simultaneous-move games are the children's games of war and rock–scissors–paper. The distinguishing characteristic of a simultaneous-move game is that no player is aware of the decisions of any other player until after the moves have been made. In a two-player game, player A is unaware of the decisions of player B, and vice versa, until both have moved.

Definition: In a simultaneous-move game the players effectively move at the same time.

In a simultaneous-move game the players are not required to actually move at the same time. In the card game called war, a standard deck of cards is shuffled and dealt out equally between two players. The players then recite the phrase "w-a-r spells war." When they say the word "war," in unison they place a card, face up, on the table. The cards are valued from lowest to highest 2–10, jack, queen, king, ace. Suits (clubs, diamonds, hearts, and spades) in this game are irrelevant. The player who shows the highest valued card wins the other player's card. If both players show a card with the same value the move is repeated until a player wins. The game ends when the deck is exhausted. The player with the greatest number of cards at the end of the game wins.

It is reemphasized that the players of simultaneous-move games are not actually required to move at the same time. "War" could also be played, for example, by isolating the players in separate rooms. Communication between the players is prohibited. A third individual, the referee, asks both players to reveal the top card on their respective portions of the deck and declares a winner accordingly. It is assumed that both players are honest and do not attempt to rearrange the order of the cards in the deck when no one is looking. The essential element of this game is that each player must move without prior knowledge of the move of the other player.

In a *sequential-move game* the players take turns. Sequential-move games are sometimes referred to as *multistage* or *dynamic games*. In a two-player game, player A moves first, followed by player B, followed again by player A, and so on. Unlike a simultaneous-move game, player B's move is based on the knowledge of how player A has already moved. Moreover, player A's next move will be based on the knowledge of how player B

moved in response to player *A*'s last move, and so on. Examples of sequential-move games include most board games, such as chess, checkers, and Monopoly. The model of duopoly developed by Augustin Cournot (1897), which was discussed in Chapter 10, is an example of a sequential-move game. Although Cournot's work was later criticized by Joseph Bertrand in the *Journal des Savants* (September 1883), both models attempt to explain the dynamic interaction of firms in a market setting.

Definition: In a sequential-move game, the players move in turn.

In addition to the manner in which the players move, games are defined by the number of games played. *One-shot games* are played only once. *Repeated games* are played more than once. If, for example, you agree with a friend to play just one game of backgammon, then you are playing a one-shot game. If you agree to play more than one game, then you are playing a repeated game.

Definition: A one-shot game is a game that is played only once.

Definition: A repeated game is a game that is played more than once.

In many ways running a business is like playing a game. In a competitive environment, the objective is to win the game. In the paragraphs that follow we will develop the basic principles of *game theory*. Game theory is the study of how rivals make decisions in situations involving strategic interaction (move and countermove). In other words, game theory refers to process by which the strategic behavior of the players is modeled. The modern version of game theory can be traced to the groundbreaking work of mathematician John von Neumann and economist Oskar Morgenstern in their 1944 classic, *Theory of Games and Economic Behavior*. As we will see, game theory is a very powerful tool for analyzing a wide variety of competitive business situations.

Definition: Game theory is the study of how rivals make decisions in situations involving strategic interaction (i.e., move and countermove) to achieve an optimal outcome.

NONCOOPERATIVE, SIMULTANEOUS-MOVE, ONE-SHOT GAMES

In this section we will examine two-person, noncooperative, non-zero-sum, simultaneous-move, one-shot games. Although the description of games of these types sounds rather daunting, it is the most basic of all game theoretic scenarios. We will begin by assuming that only two players will be playing. We will also assume that the games are noncooperative. In a *non-cooperative game* the two players do not engage in collusive behavior. In other words, the two players do not conspire to "rig" the final outcome.

We will also consider non-zero-sum games. A *zero-sum game* is one in which one player's gain is exactly the other player's loss. Poker and lotteries are zero-sum games. We will consider games in which the final solution is mutually advantageous. Each player has one, and only one, move, and both players move simultaneously. The significance of this assumption is that neither player enjoys the benefit of knowing the intentions of the other player, although each player knows the resulting payoffs from any combination of moves by both players. The Prisoners' Dilemma, discussed in Chapter 10, is an example of a two-person, noncooperative, non-zero-sum, simultaneous-move, one-shot game.

Definition: A noncooperative game is one in which the players do not engage in collusive behavior. In other words, the players do not conspire to "rig" the final outcome.

Definition: A zero-sum game is one in which one player's gain is exactly the other player's loss.

Moves are based on strategies. A strategy is a game plan. It is a kind of decision rule that a player will apply to situations in which choices need to be made. Knowledge of a player's strategy should allow us to predict what course of action a player will take when confronted with options.

Definition: A strategy is a game plan. It is decision rule that indicates what action a player will take when confronted with the need to make a decision.

Before presenting an example of a simultaneous-move, one-shot game, it is important to distinguish between *risk takers* and *risk avoiders*. The strategy selected reflects the personality of the player. Gamblers, for example, are risk takers. Risk takers have an "all or nothing" mentality; they prefer situations in which the prospect of winning results in a big payoff, even though the probability of losing is greater, and sometimes considerably so, than the probability of winning. In the parlance of probability theory, individuals are said to be risk takers (sometimes called risk lovers), when they prefer the expected value of a payoff to its certainty equivalent. Risk takers are commonly found Las Vegas, Atlantic City, and the New York Stock Exchange.[1]

Definition: Risk takers are individuals who prefer risky situations in which the expected value of a payoff is preferred to its certainty equivalent.

Risk avoiders, on the other hand, prefer a certain payoff to a risky prospect with the same expected value. Risk-averse individuals seek to minimize uncertainty. Risk avoiders prefer predictable behavior to probabilistic outcomes. When probabilistic outcomes are unavoidable, risk avoiders

[1] An expected value is defined as the weighted average of all possible outcomes, with the weights being the probability of each outcome, this is,

$$E(x) = \sum_{i=1 \to n} x_i p_i$$

where x_i is the value of the outcome and p_i is the probability of its occurrence.

will choose the "safer" outcome. Risk avoiders are loss minimizers. Depending of the level of risk aversion, for example, risk avoiders would prefer to invest in a mutual fund rather than in the individual stocks that make up the mutual fund. The reason for this is that although the expected rate of return may be lower, so too is the probability of loss. Of course, risk aversion is a relative concept. For extremely risk-averse individuals, investing in mutual funds may seem like a risky proposition. For these individuals, investing in high-grade corporate bonds or commercial bank certificates of deposits may be the way to go.

Definition: Risk avoiders prefer a certain payoff to a risky prospect with the same expected value. Risk avoiders prefer predictable outcomes to probabilistic expectations.

A player's strategy will reflect the individual's attitude toward risk. Since risk avoidance would appear to be the dominant manifestation of human behavior, we will assume in our game theoretic scenarios that the players are risk avoiders. Consider, for example, the two-player, noncooperative, non-zero-sum, simultaneous-move, one-shot game presented in Figure 13.1.

Figure 13.1 summarizes the players in the game (player *A* and player *B*), the possible strategies of each player (*A1*, *A2*, *B1*, and *B2*), and the payoffs to each player from each strategic combination. The list of strategies of each player in a game is referred to as a *strategy profile*. Strategy profiles are often depicted within curly braces. In the game depicted in Figure 13.1 there were four strategy profiles: {*A1*, *B1*}, {*A1*, *B2*}, {*A2*, *B1*}, and {*A2*, *B2*}.

The entries in the cells of the matrix refer to the payoffs to each player from each combination of strategies. The first entry in each cell of the *payoff matrix* refers to the payoff to player *A* and the second entry refers to the payoff to player *B*. Payoffs are often depicted in parentheses. We will adopt the convention that the first entry in each cell refers to the payoff to the player indicated on the left of the payoff matrix and the second entry refers in each cell refers to the payoff to the player indicated at the top. There are four payoffs depicted in Figure 13.1: (100, 200), (150, 75), (50, 50), and (100, 100). For example, if player *A* follows strategy *A1* while player *B* follows

		Player *B*	
		B1	B2
Player *A*	A1	(100, 200)	(150, 75)
	A2	(50, 50)	(100, 100)

Payoffs: (Player *A*, Player *B*)

FIGURE 13.1 Payoff matrix for a two-player, simultaneous-move game.

strategy *B2*, {*A1*, *B2*}, then the payoff to player *A* is *150* and the payoff to player *B* is *75*, (150, 75). The representation, which is depicted in Figure 13.1, is referred to as a *normal-form game*.

Definition: A normal-form game summarizes the players, possible strategies, and payoffs from alternative strategies in a simultaneous-move game.

Games such as the one presented in Figure 13.1 can apply to almost any situation involving decisions between two or more "players." In Chapter 10 the Cournot duopoly model was introduced as an example of two firms interacting in a market setting. The Cournot duopoly model is an attempt to explain the process by which two firms decide whether to charge a high price or a low price for its product, and the likely effect on each firm's profits from alternative combinations of strategies. The game might involve two competing firms trying to decide whether to advertise on television or in magazines, to introduce an entirely new product into the market, or to improve on an existing product.

STRICTLY DOMINANT STRATEGY

What is the optimal strategy for each player in the game depicted in Figure 13.1? Consider the strategies open to player *A*. If player *A* chooses strategy *A1*, then the payoffs will be 100 if player *B* chooses strategy *B1* and 150 if player *B* chooses strategy *B2*. On the other hand, if player *A* chooses strategy *A2*, then the payoffs will be 50 if player *B* chooses strategy *B1* and 100 if player *B* chooses *B2*. In this case, there is no question about what player *A* will do. Since the highest payoff that player *A* can expect by following strategy *A2* is the same as the lowest payoff from following strategy *A1*, player *A* will obviously choose strategy *A1*. Here, strategy *A1* is referred to as a *strictly dominant strategy* because that strategy will result in the largest payoff for each action that can be taken by player *B*.

Definition: A strictly dominant strategy is a strategy that results in the largest payoff regardless of the strategy adopted by other players.

What is the optimal strategy for player *B*? The reader should verify that player *B* does not have a strictly dominant strategy. Nevertheless, the fact that player *A* does have a strictly dominant strategy determines player *B*'s next move. Since player *A*'s strictly-dominant strategy is *A*, the best player *B* can do is choose strategy *B1*, which yields the largest payoff.

NASH EQUILIBRIUM

The final solution to the game in Figure 13.1 is the strategy profile {*A1*, *B1*}. An interesting aspect of this solution is that even though player *A* has the strictly dominant strategy, player *A* does not receive the maximum payoff of 150. Thus, having a dominant strategy does not guarantee that a player will receive the largest payoff. What is more, unless there is a fundamental change in the condition of the game, this solution constitutes equilibrium, at which time the game ends. Why? The answer is that no player can unilaterally improve his or her payoff by changing strategies. In Figure 13.1, if either player *A* or player *B* changes strategy, the payoff to both players is reduced. This resolution is called a *Nash equilibrium*, named for John Forbes Nash Jr., who, along with John Harsanyi and Reinhard Selten, received the 1994 Nobel Prize in economics for pioneering work in game theory.

Nash created quite a stir in the economics profession when he first proposed his now famous solution, which he called a "fixed-point equilibrium," in 1950. The reason was that his result often contradicts Adam Smith's famous metaphor of the invisible hand, according to which the welfare of society as a whole is maximized when each individual pursues his or her own private interests. This was illustrated, for example, in the final solution to the pricing game depicted in Figure 10.9. In that case, the strictly-dominant strategy of each firm was to charge a "low" price. The resulting payoff to each firm was $250,000. Yet, the strategy profile {low price, low price} did not result in the largest payoff to both firms. The best outcome would have required both firms to engage in cooperative, or collusive, behavior by charging a "high" price. In that case, the strategy profile {high price, high price} would have resulted in payoffs to both firms of $1,000,000.

Definition: A Nash equilibrium occurs when each player adopts the strategy believed to be the best response to the other player's strategy. When a game is in Nash equilibrium, the players' payoffs cannot be improved by changing strategies.

Nash equilibria are appealing precisely because they are self-fulfilling solutions to a game theoretic problem. In particular, if each player expects the other to adopt a Nash equilibrium strategy, both parties will, in fact, choose a Nash equilibrium strategy. For Nash equilibria, actual and anticipated behavior are one and the same.

EXAMPLE: OIL DRILLING GAME

Bierman and Fernandez (1998, Chapter 1) illustrate the concepts of dominant strategy and Nash equilibrium in the Oil Drilling game. The game begins by assuming that the Clampett Oil Company owns a 2-year lease on land that lies above a 4-million-barrel crude oil deposit with an estimated market value of $80 million, or $20 per barrel. The price per barrel of crude oil is not expected to change in the foreseeable future. To extract the oil, Clampett has the option of drilling a "wide" well or a "narrow" well. If Clampett drills a "wide" well, the entire deposit can be extracted in a year at a profit of $31 million. On the other hand, if Clampett drills a "narrow" well it will take 2 years to extract the oil but the profit will be $44 million.

STRICTLY DOMINANT STRATEGY EQUILIBRIUM

Enter the Texas Exploration Company (TEXplor). TEXplor has purchased a 2-year lease on land adjacent to the land leased by Clampett. The land leased by TEXplor lies above the same crude oil deposit. If both companies sink wells of the same size at the same time, each company will receive half the total crude oil reserve. For example, if both companies sink "wide" wells, each will extract 2 million barrels in 6 months, but each will earn profits of only $1 million. On the other hand, if each company sinks a "narrow" well, it will take a year for Clampett and TEXplor to extract their respective shares, but their profits will be $14 million apiece. Finally, if one company drills a "wide" well while the other company drills a "narrow" well, the first company will extract 3 million barrels and the second company will extract only 1 million barrels. In this case the first company will earn profits of $16 million and the second company will actually lose $1 million. The payoff matrix ($ millions) for this game is illustrated in Figure 13.2.

In the Oil Drilling game, the strategy profiles are {*Narrow*, *Narrow*}, {*Narrow*, *Wide*}, {*Wide*, *Narrow*}, and {*Wide*, *Wide*}. The respective payoffs from each strategy profile are (14, 14), (–1, 16), (16, –1), and (1, 1). Unlike the game theoretic scenario depicted in Figure 13.1, the payoffs depicted in Figure 13.2 are symmetrical. Which strategy should each player adopt? First consider the decision choices faced by Clampett. Whether Clampett will drill a "narrow" well or a "wide" well depends on the kind of well the firm thinks TEXplor will sink. If Clampett believes that TEXplor will drill a "narrow" well, then Clampett's best strategy is to sink a "wide" well because of its higher payoff. If, on the other hand, Clampett believes that TEXplor will sink a "wide" well, then once again Clampett's best strategy is to sink a "wide" well. In other words, regardless of the decision make by TEXplor, Clampett best strategy is to sink a "wide" well. The Oil Drilling game depicted in Figure 13.2 may look familiar. It is a variation of the Prisoners' Dilemma, which was discussed in Chapter 10. The distinguishing characteristic of both games is that each player has a strictly dominant strategy.

		Clampett	
		Narrow	Wide
TEXplor	Narrow	(14, 14)	(–1, 16)
	Wide	(16,–1)	(1, 1)

Payoffs: (TEXplor, Clampett)

FIGURE 13.2 Payoff matrix with a strictly dominant strategy equilibrium.

Since a "wide" strategy will be chosen by Clampett regardless of the strategy adopted by TEXplor, it may be said that a "wide" strategy *strictly dominates* a "narrow" strategy. Stated differently, a "narrow" strategy is *strictly dominated* by a "wide" strategy. Because of the symmetrical nature of the problem, the same must hold true for TEXplor. In this case, the strategy profile is {*Wide*, *Wide*} and with payoffs of (1, 1). Since both companies have the same strictly dominant strategy, the Nash equilibrium for this problem is called a *strictly dominant strategy equilibrium*.

Definition: A strictly dominant strategy equilibrium is a Nash equilibrium that results when all players have a strictly dominant strategy.

WEAKLY DOMINANT STRATEGY

Consider the variation of the Oil Drilling game summarized in Figure 13.3 (Bierman and Fernandez, 1998, Chapter 1). In this variation, the reader will quickly verify that if TEXplor chooses a *Wide* strategy, Clampett will be indifferent between a *Narrow* and a *Wide* strategy. In this case, *Wide* is no longer a strictly dominant strategy because both strategies yield a zero profit for Clampett if TEXplor chooses to drill a "wide" well. In this case, *Wide* is referred to as a *weakly dominant strategy*.

Definition: A weakly dominant strategy is a strategy that results in a payoff that is no lower than any other payoff regardless of the strategy adopted by the other player.

A rational player will always play a weakly dominant strategy. In the symmetrical game depicted in Figure 13.3, this means that the *weakly dominant strategy equilibrium* is for both players to drill a wide well. The reason for this is simple. Playing a weakly dominant strategy will never result in a lower payoff, while playing a weakly dominated strategy might. Suppose, for example, that TEXplor chooses to drill a narrow well. By drilling a narrow well, the best that Clampett can expect is a payoff of $14 million. The lowest payoff is $0. On the other hand, by drilling a wide well, Clampett's highest possible payoff is $16 million. The lowest possible payoff is still $0. If Clampett is rational, there is no reason ever to adopt a

		Clampett	
		Narrow	Wide
TEXplor	Narrow	(14, 14)	(0, 16)
	Wide	(16, 0)	(0, 0)

Payoffs: (TEXpbr, Clampett)

FIGURE 13.3 Payoff matrix with a weakly dominant strategy ($ millions).

Narrow strategy. Since the game is symmetrical, the same is true for TEXplor. Finally, the reader should note that the strategy profile {*Wide*, *Wide*} is a Nash equilibrium because neither player can improve the payoff by switching strategies.

Definition: A weakly dominant strategy equilibrium is a Nash equilibrium that results when all players have a strictly dominant strategy.

ITERATED STRICTLY DOMINANT STRATEGY

When both players have a strictly dominant strategy, the solution to a noncooperative, simultaneous-move, one-shot game is fairly straightforward. In the following version of the Oil Drilling game, which is also taken from Bierman and Fernandez (1998, Chapter 1), however, neither TEXplor nor Clampett has a strictly dominant strategy.

The payoff matrix in Figure 13.4 introduces a third strategy—*Don't drill*. An examination of the payoff matrix reveals that *Wide* strictly dominates *Don't drill*, but no longer dominates *Narrow*. For example, if TEXplor chooses a *Don't drill* strategy, Clampett should choose *Narrow*. On the other hand, if TEXplor chooses *Narrow* or *Wide*, Clampett should choose *Wide*. Moreover, *Narrow* does not dominate either *Don't drill* or *Wide*. The only thing that is absolutely certain is that Clampett will not adopt a *Don't drill* strategy. *Don't drill* is called a strictly dominated strategy. A strictly dominated strategy is a strategy that is dominated by every other strategy.

Definition: A strictly dominated strategy is a strategy that is dominated by every other strategy and will always result in a lower payoff (i.e., regardless of the strategy adopted by other players).

Since the payoff matrix in Figure 13.4 is symmetrical, the same is true for TEXplor. Since neither TEXplor nor Clampett will ever choose *Don't drill*, this strategy may be eliminated from consideration. The resulting payoff matrix reduces to the two-strategy game in Figure 13.2, which had the strictly dominant strategy equilibrium {*Wide*, *Wide*}. Thus, *Wide* is the solution to the three-strategy game summarized in Figure 13.4. Wide is

		Clampett		
		Don't drill	Narrow	Wide
	Don't drill	(0, 0)	(0, 44)	(0, 31)
TEXplor	Narrow	(44, 0)	(14, 14)	(!1, 16)
	Wide	(31, 0)	(16, –1)	(1, 1)

Payoffs: (TEXplor, Clampett)

FIGURE 13.4 Payoff matrix with a iterated strictly dominant strategy ($ millions).

called an *iterated strictly dominant strategy* because it was obtained after systematically eliminating *Don't drill* from both players' strategy profiles.

Games with a large number of strategies and players may require several iterations before an iterated strictly dominant strategy equilibrium is achieved. As long as each player has a strictly dominant strategy, the order in which strictly dominated strategies are eliminated is irrelevant. We will still end up with a strictly dominant strategy equilibrium. On the other hand, if strictly dominant strategies are replaced by weakly dominant strategies, it may be demonstrated that the order in which weakly dominated strategies are removed could change the outcome of the game.

NON–STRICTLY DOMINANT STRATEGY

Finally, consider yet another variation on the Oil Drilling game, also taken from Bierman and Fernandez (1998, Chapter 1). Suppose that instead of losing $1 million from a {*Narrow*, *Wide*} strategy Clampett and TEXplor earn a positive profit of $2 million. The revised payoff matrix is illustrated in Figure 13.5.

An examination of the payoff matrix in Figure 13.5 reveals that no strictly dominant strategy exists for either player. The optimal strategy for both players depends on what each player *believes* the other player will do. To see this, suppose that Clampett believes that TEXplor will drill a "narrow" well. Clearly, it will be in Clampett's best interest to drill a "wide" well, since that strategy will generate $16 million in profits, which is greater than $14 million if it drills a narrow well. From Clampett's perspective a {*Narrow*, *Wide*} strategy is rational. Similarly, if Clampett *believes* that TEXplor intends drill a wide well, Clampett will drill a narrow well and earn only $2 million. In this case, a {*Wide*, *Narrow*} strategy is rational. Since both Clampett and TEXplor believe that this strategy profile is in their best interest, it will be adopted. Thus, the strategy profile {*Wide*, *Narrow*} leads to a Nash equilibrium. The important thing is that if either firm believes that the other will adopt a particular strategy, it will be in the best interest of that firm to adopt the same strategy. When this occurs, the strategy profile is said to be *self-confirming*.

		Clampett	
		Narrow	Wide
TEXplor	Narrow	(14, 14)	(2, 16)
	Wide	(16, 2)	(1, 1)

Payoffs: (TEXpbr, Clampett)

FIGURE 13.5 Payoff matrix with a non–strictly dominant strategy ($ millions).

Definition: In a two-person game, a non–strictly dominant strategy exists when a strictly dominant strategy does not exist for either player. In this case, the optimal strategy for either player depends on what each player believes to be the strategy of the other player.

Although the concept of a Nash equilibrium is virtually unchallenged as a solution to noncooperative games, it is not without controversy. The reason for this is that Nash equilibria are often not unique. As the reader will readily verify, the strategy profiles {*Narrow*, *Wide*} and {*Wide*, *Narrow*} constitute a Nash equilibrium to the game depicted in Figure 13.5. When there is one than more than Nash equilibrium, it will be difficult to predict the strategies of the other players without more information. In other words, the existence of a Nash equilibrium does not always guarantee a solution to a game. One proposed solution to games involving multiple Nash equilibria is a *focal-point equilibrium*, which will be discussed later in this chapter.

MAXIMIN STRATEGY

In the game depicted in Figure 13.1 we saw that the strictly dominant strategy of player *A* determined the optimal strategy of player *B*. But what if neither player has a strictly dominant strategy? Will it still be possible to determine the optimal strategy profile for the game? Will the game still have a Nash equilibrium? If we assume that both players are risk averse, an optimal strategy profile may be determined when both players choose a *maximin strategy*. Sometimes referred to as a *secure strategy*, a maximin strategy selects the highest payoff from the worst possible scenarios.

Definition: A maximin strategy selects the largest payoff from among the worst possible payoffs.

Consider, again, the game depicted in Figure 13.1. Player *A* has a strictly dominant strategy (*A1*), which determined player *B*'s strategy (*B1*). But suppose that player *B* had no knowledge of the payoffs to player *A*. What would have been player *B*'s secure (maximin) strategy? In that game, had player *B* opted for strategy *B1*, the worst possible payoff would have been 50. Had player *B* selected strategy *B2*, then the worst possible payoff would have been 75. Thus, player *B*'s maximin strategy would have been strategy *B2*, since it would have resulted in the largest payoff from among the worst possible payoffs. Of course, player *B* did not play the secure strategy because player *A*'s strictly dominant strategy determined player *B*'s next move.

To underscore the logic underlying a maximin strategy, consider the following variation of the game depicted in Figure 13.1. The reader will immediately verify that neither player *A* or player *B* has a dominant strategy. If player *A*, selects strategy *A1*, then player *B* will select strategy *B1*. If player *A* selects strategy *A2*, then player *B* will select strategy *B2*. On the other

		Player *B*	
		B1	B2
Player *A*	A1	(100, 100)	(75, 75)
	A2	(–100, 100)	(200, 200)

Payoffs: (Player *A*, Player *B*)

FIGURE 13.6 Payoff matrix and a maximin strategy.

hand, if player *B* selects strategy *B1*, then player *A* will select strategy *A1*. If player *B* selects strategy *B2*, then player *A* will select strategy *A2*. What strategy will players *A* and *B* choose?

Suppose that in the game depicted in Figure 13.6 both players follow a maximin strategy. If player *B* plays strategy *B1*, then the minimum payoff for player *A* is –100 by playing strategy *A2*. If player *B* plays strategy *B2*, then the minimum payoff for player *A* is 75 by playing strategy *A1*. Following a maximin strategy, player *A* will choose the strategy with the largest of the two worst payoffs. In this case, player *A*'s secure strategy is to play strategy *A1*.

What about player *B*? If player *A* plays strategy *A1*, the minimum payoff for player *B* is 75 by choosing strategy *B2*. If player *A* plays strategy *A2*, then the minimum payoff for player *B* is 100 by playing strategy *B1*. The maximin (secure) strategy for player *B* is to play strategy *B1*. Thus, the strategy profile for this game is {*A1*, *B1*}. The reader should verify that this strategy profile constitutes a Nash equilibrium, but not the only Nash equilibrium.

Regrettably, solutions to games using a maximin strategy may not be as simple as they appear. Consider another variation on the game depicted in Figure 13.1. In Figure 13.7 the reader should verify that player *B* has the dominant strategy *B2*. Player *A* knows that player *B* has a dominant strategy and that player *B* is likely to play that strategy. In this case, player *A*'s best move is to play strategy *A2*, which results in a payoff of 200. Thus, the strategy profile in this game is {*A2*, *B2*} for payoffs of (200, 100).

Suppose, however, that in the game depicted in Figure 13.7 player *A* believes that player *B* might not, in fact, play his or her dominant strategy. This possibility might arise if player *B* has a history of making mistakes. When risk and uncertainty are introduced, the game changes. Depending on the level of risk aversion, it might be in player *A*'s best interest to follow a maximin strategy, especially if the potential loss by choosing the wrong strategy is great. In this case, player *A* believes that player *B* might adopt either strategy *B1* or *B2*. If player *B* follows strategy *B1*, the lowest payoff for player *A* will be –1,000 by following strategy *A2*. If player *B* follows strategy *B2*, the lowest payoff to player *A* is 100 by following strategy *A1*.

		Player *B*	
		B1	B2
Player *A*	A1	(100, 0)	(100, 100)
	A2	(–1,000, 0)	(200, 100)

Payoffs: (Player *A*, Player *B*)

FIGURE 13.7 Risk aversion and a maximin strategy.

		Baltimore Ravens	
		Pass	Run
New York Giants	Pass defense	(50, 50)	(40, 60)
	Run defense	(20, 80)	(80, 20)

Payoffs: (New York Giants, Baltimore Ravens)

FIGURE 13.8 Payoff matrix and the Touchdown game.

Being risk averse, player *A* might decide that a guaranteed payoff of 100 by following a secure (maximin) strategy is preferable to an uncertain payoff of 200, or possible loss of –1,000.

Example: Touchdown Game

Suppose that the New York Giants and the Baltimore Ravens are in the fourth quarter of the Super Bowl with seconds remaining on the clock. It is the last play of the game. The score is Ravens 13 and the Giants 17. The Ravens have the ball on the Giants' 8 yard line. There are no time-outs for either side. A field goal for 3 points will not help Baltimore. To win the game, Baltimore must score a touchdown for 6 points. Both sides must decide on a strategy for the final play of the game. The objective of both teams is to maximize the probability of winning the game. Both head coaches are aware of the strengths and weaknesses of the other team. The probabilities of either team winning the game from alternative offensive and defensive strategies are summarized in Figure 13.8. The student should note that the sum of the probabilities in each cell is 100%.

An examination of the payoff matrix in Figure 13.8 will verify that neither team has a strictly dominant strategy. If the Giants, for example, adopt a pass defense, the best offensive play for the Ravens is to run the ball. If the Giants adopt a run defense, the best offensive play for the Ravens is to pass. On the other hand, if the Ravens decide to pass the ball, the best strategy for the Giants is a pass defense. If the Ravens decide to run the ball, then the best strategy for the Giants is a run defense. The student should also verify that a Nash equilibrium does not exist here because there is no strategy profile for which a change in strategy will result in a lower payoff for either team.

Since neither team has a strictly dominant strategy, what strategy should be adopted by each coach? If we assume that both head coaches are risk averse, both teams should adopt a maximin strategy. If the coach of the New York Giants decides to play a pass defense, the worst the team can do is a 40% probability of winning the game. If the Giants decide to play a run defense, the worst they can do is a 20% chance of winning. Since a 40% probability of winning the game is the largest payoff from among the worst possible scenarios, playing a pass defense is the Giants secure strategy.

Now consider the secure strategy of the Baltimore Ravens. If the Ravens play a pass offense, the worst the team can do is a 50% probability of winning the game. If the Ravens decide to play a run offense, the worst they can do is a 20% probability of winning the game. Since a 50% probability of winning the game is the largest payoff from among the worst possible scenarios, playing a pass offense is the Ravens's secure strategy. Thus, the solution profile for this version of the Touchdown game is {*Pass defense*, *Pass*}.

Problem 13.1. Suppose that in the Touchdown game the probabilities of either team winning from alternative offensive and defensive strategies are as shown in Figure 13.9.

a. Does either firm have a strictly dominant strategy?
b. Assuming that both coaches are risk averse, what strategy will each coach likely adopt for the last play of the game?
c. Does this game have a Nash equilibrium?

Solution

a. Neither team has a strictly dominant strategy. If the New York Giants adopt a pass defense, the best offensive play for the Baltimore Ravens is to run the ball. If the Giants adopt a run defense, the best offensive play for the Ravens is to pass. On the other hand, if the Ravens decide to pass the ball, the best strategy for the Giants is a pass defense. If the

		Baltimore Ravens	
		Pass	Run
New York Giants	Pass defense	(70, 30)	(20, 80)
	Run defense	(10, 90)	(50, 50)

Payoffs: (New York Giants, Baltimore Ravens)

FIGURE 13.9 Payoff matrix for problem 13.1.

Ravens decide to run the ball, the best strategy for the Giants is a run defense.

b. Assuming again that both head coaches are risk averse, both teams should adopt a secure strategy. If the coach of the New York Giants decides to play a pass defense, the worst the team can do is a 20% chance of winning the game. If the Giants decide to play a run defense, the worst they can do is a 10% chance of winning the game. Since a 20% probability of winning the game is the larger of the two worst payoffs, the Giants' secure strategy is to play a pass defense.

Now consider the secure strategy of the Baltimore Ravens. If the Ravens play a pass offense, the worst the team can do is a 30% chance of winning. If the Ravens decide to play a run offense, the worst they can do is a 50% chance of winning. Since a 50% chance of winning is the larger of the two worst payoffs, the Ravens's secure strategy is to run the ball. Thus, the solution profile for this version of the Touchdown game is {*Pass defense*, *Run*}. The reader should verify once again that this strategy profile does not constitute a Nash equilibrium.

c. A Nash equilibrium does not exist for this game. Either team can improve its payoff by switching strategies.

Problem 13.2. The two leading firms in the highly competitive running shoe industry, Treebark and Adios, are considering an increase in advertising expenditures. Both companies are considering buying advertising space in *Joggers World*, the leading national magazine about recreational, long-distance running, or buying air time with KNUT, an all-talk, all-sports, all-the-time radio station. Figure 13.10 summarizes the payoffs associated with the advertizing strategy of each firm.

a. Do either Treebark or Adios have a dominant strategy?
b. Based on your answer to part a, what is the strategy of the other firm?
c. What is the Nash equilibrium for this problem?

Solution

a. While Treebark has a dominant strategy, which is to advertise in *Joggers World*, Adios does not. To see this, suppose that Adios advertises in

		Treebark: Joggers World	Treebark: KNUT
Adios	Joggers World	(\$1,000,000, \$2,000,000)	(\$300,000, \$350,000)
	KNUT	(\$750,000, \$750,000)	(\$2,500,000, \$500,000)

Payoffs: (Adios,Treebark)

FIGURE 13.10 Payoff matrix for problem 13.2.

Joggers World, Treebark's best response is to advertise in *Joggers World* as well. If Adios advertises on KNUT, again Treebark's best response is to advertise in *Joggers World*. In both instances, Treebark should advertise in *Joggers World*.

b. Adios does not have a dominant strategy. As can be seen from Figure 13.10, if Treebark advertises in *Joggers World*, then Adios's best response is to advertise in *Joggers World*. On the other hand, if Treebark advertises on KNUT, Adios's best strategy is to advertise on KNUT. In other words, Adios's strategy will depend on what it thinks Treebark will do. This is not the case for Treebark, which will advertise in *Joggers World* regardless of the strategy adopted by Adios. Since Treebark will advertise in Joggers World regardless of the strategy adopted by Adios, then it will be in the best interest of Adios to advertise in Joggers World as well. Thus, the solution profile for this game is {Joggers World, Joggers World} with payoffs to Adios and Treebark of $1,000,000 and $2,000,000, respectively.

c. A Nash equilibrium occurs when both Treebark and Adios advertise in *Joggers World*. The reason for this is that if Treebark changes its strategy to buying air time on KNUT, its payoff falls to $350,000. If Adios changes its strategy to buying air time on KNUT, its payoff falls to $750,000. This is a Nash equilibrium, since neither player can unilaterally improve its payoff by changing strategies.

COOPERATIVE, SIMULTANEOUS-MOVE, INFINITELY REPEATED GAMES

Consider, again, the duopoly problem discussed in Chapter 10 in which firms *A* and *B* are confronted with the decision to charge a "high" price or a "low" price for their product. The payoff matrix for that problem, Figure 10.9, is reproduced here.

We saw in this game that the strictly dominant strategy of both firm *A* and firm *B* was to charge a "low" price. The reason is that if firm *B* chooses

		Firm *B*	
		High price	*Low price*
Firm *A*	*High price*	($1,000,000, $1,000,000)	($1,000,000, $5,000,000)
	Low price	($6,000,000 $100,000)	($250,000, $250,000)

Payoffs: (Player*A*, Player*B*)

FIGURE 10.9 Game theory and interdependent pricing behavior.

a "low" price with a payoff of $5,000,000, then, from firm *B*'s perspective, it will be rational for firm *A* to choose a "high" price, where the payoff is $1,000,000. From firm *B*'s perspective the rational combination of strategies is (*Low*, *High*). Since the payoff matrix is symmetrical, the same reasoning pertains to firm *A*. From firm *B*'s perspective the rational combination of strategies is also (*Low*, *High*). The result, however, is an "irrational" (*Low*, *Low*) combination of strategies in which firms *A* and *B* earn profits of $250,000. This is a Nash equilibrium because neither player can unilaterally improve its payoff by changing strategies. For example, if firm *A* were to switch to a "high" price strategy while firm *B* continues to charge a "low" price, firm *B*'s profits will drop to $100,000, while firm *A*'s profits will increase to $5,000,000. Precisely the same thing would occur should firm *B* attempt to switch to a "high" price strategy while firm *A* continued to charge a "low" price.

The problem summarized in Figure 10.9 illustrates the Bertrand duopoly pricing model discussed in Chapter 10. The reader will recall that in the Bertrand model each firm will set the price of its product duopoly to maximize profits while ignoring its rival's output level. In the problem, both firms can clearly maximize profits by agreeing to charge a "high" price for their product. For this to occur, however, both firms must agree to collude on their pricing decisions (see Chapter 10). The problem with collusive behavior, however, is that for a two-person, non-zero-sum, simultaneous-move, one-shot game, there is an incentive for either firm to "cheat." In other words, it is in the best interest to either firm to violate any formal pricing agreement by charging a low price at the expense of its rival.

The game theoretic scenario summarized in Figure 10.9 illustrates the fragile nature of collusion in a two-person, non-zero-sum, simultaneous-move, one-shot game. It was demonstrated that in the absence of a cooperative pricing strategy, a Nash equilibrium occurs when both firms charged a "low" price, with each firm earning $250,000 in profits. On the other hand, if the firms collude, both will charge a "high" price and each will earn $1,000,000 in profits. In a one-time game, however, if firm *A* were to violate the agreement and charge a "low" price, many consumers would switch their purchases away from firm *B*. Firm *A*'s profit would soar from $250,000 to $5,000,000, while firm *B*'s profits would fall from $250,000 to $100,000. The result would be precisely the same if firm *B* violated the agreement.

The duopoly problem illustrates the situation in which, in the absence of collusion, a Nash equilibrium results in an inferior solution. Why, then, will the two firms not engage in collusive behavior? For one thing, collusion is illegal in the United States. For another, in a two-person, non-zero-sum, simultaneous-move, one-shot game, the incentive to cheat will ultimately undermine any such agreement. In fact, since both firms are aware of the inherent weakness of collusive behavior, it is unlikely that the cooperative pricing arrangement would have been entered into in the first place.

A cartel is an example of a collusive arrangement. A cartel is a formal agreement among firms in oligopolistic industries to allocate market share and/or industry profits. To take perhaps the most famous example, the members of the Organization of Petroleum Exporting Countries are able to influence world oil prices by jointly agreeing on production levels. For the reasons cited, however, cartel arrangements have historically proven to be very short-lived precisely because of the sometimes irresistible temptation to cheat. In the case of OPEC, Venezuela has repeatedly cheated on almost every production agreement it has entered into. The government in Caracas would encourage OPEC members, especially the cartel's "swing" producer, Saudi Arabia, to restrict output to raise oil prices, after which it would increase its own production levels to bolster profits.

Is "cheating" inevitable? Under the appropriate conditions, for a two-person, non-zero-sum, simultaneous-move, one-shot game, the answer is more than likely to be yes. But, what if the game were played more than once? What if the game were infinitely repeated, as is seemingly the case with OPEC production agreements? As we will see, naughty behavior may be punishable, which may affect the manner in which the players play future games. It is this possibility that we will consider in the next section.

In our discussion of two-person, non-zero-sum, simultaneous-move, one-shot games, we observed that collusive behavior among the players was inherently unstable because of the incentive to cheat. Does this conclusion hold, however, if the game infinitely repeated? In this section we will consider the game theoretic scenario of two cooperating players engaged in a simultaneous, non-zero-sum game, which is played over and over again. As the reader may have guessed, with infinitely repeated games, cheating may have consequences for how future games are played.

Definition: Infinitely repeated games are games that are played over and over again with no end.

Suppose that instead of the game being played just once, it is infinitely repeated. Do the conclusions reached with respect to the infeasibility of collusive behavior for two-person, non-zero-sum, simultaneous-move, one-shot games continue to hold? Maybe not. Past naughty behavior by one player may cause the other player to adopt a different strategy for future play. Such contingent game plans are referred to as *trigger strategies*. A trigger strategy is a game plan that is adopted by one player in response to unanticipated moves by the other player. Once adopted, a trigger strategy will continue to be used until the other player initiates yet another unanticipated move.

Definition: A trigger strategy is a game plan that is adopted by one player in response to unanticipated moves by the other player. A trigger strategy will continue to be used until the other player initiates yet another unanticipated move.

For two-person, cooperative, non-zero-sum, simultaneous-move, infinitely repeated games, trigger strategies may, in fact, introduce stability into collusive arrangements. The reason for this is the so-called *credible threat*. To see what is involved, consider again the game theoretic scenario summarized in Figure 10.9. Suppose that firms *A* and *B* agree, perhaps illegally, to charge a "high" price for their products. Unlike the one-shot game, if either player "cheats," future game plays will be changed. If firm *A* were to charge a "low" price in violation of its agreement, firm *B* might seek to punish firm *A* by ruling out any future cooperation. In other words, a violation of firm *A*'s agreement with firm *B* could "trigger" a strategy change by firm *B*. Firm *B*'s promise to retaliate may prevent firm *A* from violating the agreement, but only if this threat is considered to be credible. A threat is credible only if it is in the best interest of the player making the threat to follow through when the trigger situation presents itself.

Definition: A threat is credible only if it is in a player's best interest to follow through with the threat.

Does the knowledge that naughty behavior will result in punishment eliminate the possibility of cheating? Not necessarily. To begin with, if the threat of retaliation is not credible, it will be ignored. Moreover, even if threats are credible, cheating may still occur in infinitely repeated games if naughty behavior (cheating) is more profitable than honest behavior. To see this, it is necessary to compare the present value of the stream of profits resulting from cheating to the present value of profits earned by adhering to the agreement.

Recall from Chapter 12 the present value of an annuity due (*PVAD*), which is summarized in Equation (12.21). If we assume that the profits earned by a firm are the same in each period, then Equation (12.21) may be rewritten as

$$\begin{aligned} PVAD &= \frac{\pi}{(1+i)^{n-1}} + \frac{\pi}{(1+i)^{n-2}} + \ldots + \frac{\pi}{(1+i)^{0}} \\ &= \sum_{t=0\to n} \left(\frac{1}{1+i}\right)^{n-1} \end{aligned} \tag{13.1}$$

where *PVAD* is the present value of an annuity due and i is the nominal (market) interest rate. For infinitely repeated games ($n = \infty$), it can be easily demonstrated that

$$PVAD = \frac{\pi(1+i)}{i} \tag{13.2}$$

since Equation (13.2) is the sum of a geometric progression (see Chapter 2).[2]

COLLUSION

Consider, once again, the simultaneous-move game summarized in Figure 10.9. We saw that in a one-shot game the Nash equilibrium was reached when both firms charged a "low" price. Now suppose that firms *A* and *B* collude, perhaps illegally, to charge a "high" price. Suppose, further, that cheating by one firm "triggers" a change in strategy by the other firm. In particular, suppose that firm *B* violates the agreement by charging a "high" price. This action would cause firm *A* to punish firm *B* by charging a "low" price in all future periods. If both firms adopt the same trigger strategy, will the cartel hold together? The answer to this question depends on a comparison of the economic incentives to cheat and to maintain the agreement.

The economic benefit of maintaining the agreement is the present value of all future profits earned by remaining "honest" ($PVAD_{\mathrm{H}}$) to the terms of the agreement, which is given by the Equation (13.3).

[2] Equation (13.1) may be rewritten as

$$\begin{aligned} PVAD &= \pi + \pi\left[\frac{\pi}{1+i} + \frac{1}{(1+i)^2} + \frac{1}{(1+i)^3} + \ldots\right] \\ &= \pi\left[\frac{1+1}{1+i} + \frac{1}{(1+i)^2} + \frac{1}{(1+i)^3} + \ldots\right] \\ &= \pi S \end{aligned} \tag{F13.1}$$

where

$$S = \frac{1+1}{1+i} + \frac{1}{(1+i)^2} + \frac{1}{(1+i)^3} + \ldots \tag{F13.2}$$

After multiplying both sides of Equation (F13.2) by $1/(1+i)$, we get

$$S\left(\frac{1}{1+i}\right) = \frac{1}{1+i} + \frac{1}{(1+i)^2} + \frac{1}{(1+i)^3} + \frac{1}{(1+i)^4} + \ldots \tag{F13.3}$$

Subtracting Equation (F13.3) from Equation (F13.2) yields

$$S - S\frac{1}{1+i} = 1$$

which simplifies to

$$S = \frac{1+i}{i} \tag{F13.4}$$

Subtracting Equation (F13.3) from Equation (F13.2) yields

$$PVAD = \frac{\pi(1+i)}{i} \qquad \text{Q.E.D.}$$

$$PVAD_{\mathrm{H}} = \frac{\pi_{\mathrm{H}}(1+i)}{i} \tag{13.3}$$

If, on the other hand, a firm were to violate the agreement, the economic benefit would be the immediate (one-shot) gain from cheating (π_{C}) plus the present value of the per-period profits (π_{N}) earned in the absence of a collusive agreement ($PVAD_{\mathrm{N}}$). The economic benefit of violating the agreement is summarized in Equation (13.4).

$$\pi_{\mathrm{C}} + PVAD_{\mathrm{N}} = \pi_{\mathrm{C}} + \frac{\pi_{\mathrm{N}}(1+i)}{i} \tag{13.4}$$

When will it pay to cheat? Cheating will occur when the present value of violating the agreement is greater than the present value of remaining "honest." This condition is summarized in the inequality (13.5):

$$\pi_{\mathrm{C}} + PVAD_{\mathrm{N}} > PVAD_{\mathrm{H}}$$

or

$$\pi_{\mathrm{C}} + \pi_{\mathrm{N}}\frac{(1+i)}{i} > \frac{\pi_{\mathrm{H}}(1+i)}{i} \tag{13.5}$$

Problem 13.3. Consider, again, the payoff matrix summarized in Figure 10.9. Suppose, further, that this is an infinitely repeated game and that the interest rate at which profits may be reinvested is 5%.

a. What is the economic benefit to firm *A* and firm *B* from a Nash equilibrium (no collusion) in an infinitely repeated game?
b. What is the economic benefit to firm *A* and to firm *B* from a collusive agreement?
c. What is the economic benefit to firm *A* or firm *B* from violating (cheating) the agreement?
d. Based on your answers to parts a and b, is the collusive agreement stable? That is, is the cartel likely to last?

Solution

a. A Nash equilibrium occurs when both firms *A* and *B* charge a "low" price. From Equation (13.5), the economic benefit of a Nash equilibrium for an infinitely repeated game is

$$\frac{\pi_{\mathrm{N}}(1+i)}{i} = \frac{\$250{,}000(1.05)}{0.05} = \$250{,}000(21) = \$5{,}250{,}000$$

b. In a collusive agreement, both firms will charge a "high" price. From Equation (13.3), the economic benefit of charging a "high" price in an infinitely repeated game, and remaining "honest" to the agreement is

$$\frac{\pi_H(1+i)}{i} = \frac{\$1,000,000(1.05)}{0.05} = \$1,000,000(21) = \$21,000,000$$

c. From Equation (13.5), the economic benefit from violating the agreement is the immediate (one-shot) gain from cheating (π_C) plus the present value of all profits that will be earned from the Nash equilibrium ($PVAD_N$) thereafter:

$$\pi_C + \frac{\pi_N(1+i)}{i} = \$5,000,000 + \$5,250,000 = \$10,250,000$$

d. Since $\pi_C + \pi_N(1 + i)/i < \pi_H(1 + i)/i$, there is no incentive to cheat. In other words, since the economic benefit to both firms to remain "honest" is greater than the economic benefit to either firm of cheating (\$21,000,000 > \$10,250,000), there is no incentive for either firm to cheat.

Problem 13.4. Suppose that in Problem the interest rate is 20%.

a. What is the economic benefit to both firms from a Nash equilibrium in an infinitely repeated game?
b. What is the economic benefit to both firms from a collusive agreement?
c. What is the economic benefit to either firm of cheating?
d. Based on your answers to parts a and b, is the collusive agreement stable?

Solution

a. From Equation (13.5), the economic benefit of a Nash equilibrium for an infinitely repeated game is

$$\frac{\pi_N(1+i)}{i} = \frac{\$250,000(1.20)}{0.20} = \$250,000(6) = \$1,500,000$$

b. In a collusive agreement, both firms will charge a "high" price. From Equation (13.3), the economic benefit of charging a "high" price in an infinitely repeated game, and remaining "honest" with respect to the agreement is

$$\frac{\pi_H(1+i)}{i} = \frac{\$1,000,000(1.20)}{0.20} = \$1,000,000(6) = \$6,000,000$$

c. From Equation (13.4), the economic benefit from violating the agreement is the immediate (one-shot) gain from cheating (π_C) plus the present value of all profits that will be earned from the Nash equilibrium ($PVAD_N$) thereafter:

$$\pi_C + \frac{\pi_N(1+i)}{i} = \$5,000,000 + \$1,500,000 = \$6,500,000$$

d. Since $\pi_C + \pi_N(1 + i)/i > \pi_H(1 + i)/i$ (i.e., \$6,500,000 > \$6,000,000), there is an incentive to cheat. In other words, the cartel is unstable and the collusive agreement is likely to break down.

CHEATING RULE FOR INFINITELY-REPEATED GAMES

Admittedly, the foregoing assumptions of unchanged profits and interest rates for a two-person, cooperative, non-zero-sum, simultaneous-move, infinitely repeated game are simplistic. Nevertheless, with these assumptions it is possible to summarize the conditions under which a cartel is likely to be unstable. Inequality (13.5) may be rearranged to yield

$$\frac{\pi_H - \pi_N}{\pi_C + \pi_N - \pi_H} < i \tag{13.6}$$

Inequality (13.6) admits to a straightforward interpretation. If the net rate of return from adhering to the collusive agreement relative to the net rate of return from cheating is less than the prevailing rate of interest, there is an incentive to violate the agreement. If the inequality (13.6) is satisfied, the cartel will be unstable, since the incentive to "cheat" is greater than the incentive to be "honest." If the inequality (13.6) is not satisfied—that is, if $(\pi_H - \pi_N)/(\pi_C + \pi_N - \pi_H) > i$, then a trigger strategy by one firm that punishes the cheater by refusing to enter into future collusive agreements will be sufficient to hold the cartel together. Finally, if $(\pi_H - \pi_N)/(\pi_C + \pi_N - \pi_H)/i$, then, *ceteris paribus*, each firm will be indifferent between cheating and remaining honest.

Problem 13.5. Consider, again, the payoff matrix summarized in Figure 10.9. Suppose that each firm adopts the trigger strategy such that it will respond to cheating by the other firm by choosing a one-shot Nash equilibrium for all future plays. Assuming that the payoffs in Figure 10.9 are expected to be infinitely repeated, below what interest rate can we expect the cartel to break down?

Solution. Substituting the data from Figure 10.9 into the left-hand side of inequality (13.6) yields

$$\frac{\pi_H - \pi_N}{\pi_C + \pi_N - \pi_H} = \frac{\$1{,}000{,}000 - \$250{,}000}{\$5{,}000{,}000 - \$250{,}000 - \$1{,}000{,}000}$$
$$= \frac{\$750{,}000}{\$4{,}750{,}000} = 0.1765$$

Thus, from inequality (13.6), if the prevailing rate of interest is greater than 17.65%, each firm will have an incentive to violate the collusive agreement, rendering the cartel unstable. If the rate of interest is less than

17.65%, then it will be in the best interest for both firms to honor the agreement and the cartel will be stable. Finally, if the interest rate is exactly 17.65%, then, *ceteris paribus*, each firm will be indifferent between cheating and remaining honest.

Problem 13.6. Consider the payoff matrix in Figure 13.11: two firms that must decide whether to charge \$30 or \$50 for their product. The first entry in each cell of the matrix represents the profit earned by firm *a* and the second entry represents the profit earned by firm *B*. Thus, if firm *A* charges \$30 while firm *B* charges \$50, the first firm's profit is \$100,000 and the second firm will get \$30,000.

a. For a noncooperative, simultaneous-move, one-shot game, does either firm have a dominant strategy? If not, what is each firm's secure strategy? What is the Nash equilibrium for this problem? Why?
b. If this were a cooperative, simultaneous-move, one-shot game, what price should each firm charge? Why?
c. Suppose that the interest on reinvested profits is 20%. What is the economic benefit to firm *A* and to firm *B* from a Nash equilibrium (no collusion) in an infinitely repeated game?
d. Find the economic benefit to firm *A* and to firm *B* from a collusive agreement.
e. Find the economic benefit to firm *A* or firm *B* from violating (cheating) the agreement.
f. Based on your answers to parts a and b, is the collusive agreement stable? That is, is the cartel likely to last?
g. Suppose that the interest rate was 30%. Is the collusive agreement stable?
h. Suppose that each firm adopts the trigger strategy such that it will respond to cheating by the other firm by choosing a one-shot Nash equilibrium for all future plays. Above what interest rate can we expect the cartel to break down?

Solution

a. The dominant strategy of both firms is to charge \$30 for their product. The strategy profile {*\$30*, *\$30*} is a strictly dominant strategy equilibrium,

		Firm *B* \$30	\$50
Firm *A*	\$30	(\$60,000, \$60,000)	(\$100,000, \$30,000)
	\$50	(\$30,000, \$100,000)	(\$80,000, \$80,000)

Payoffs: (Player A, Player B)

FIGURE 13.11 Payoff matrix for problem 13.6.

which is a Nash equilibrium because neither player can improve its payoff by switching strategies.

b. If this was a cooperative, simultaneous-move, one-shot game, it would pay for firm *A* and firm *B* to enter into a collusive agreement and charge \$50, since each firm would earn profits of \$80,000.

c. From Equation (13.5), the economic benefit of a Nash equilibrium for an infinitely repeated game is

$$\frac{\pi_N(1+i)}{i} = \frac{\$60{,}000(1.20)}{0.20} = \$60{,}000(6) = \$360{,}000$$

d. In a collusive agreement, both firms will charge \$50. From Equation (13.3), the economic benefit is

$$\frac{\pi_H(1+i)}{i} = \frac{\$80{,}000(1.20)}{0.20} = \$80{,}000(6) = \$480{,}000$$

e. From Equation (13.4), the economic benefit from violating the agreement is

$$\frac{\pi_C + \pi_N(1+i)}{i} = \$100{,}000 + \$360{,}000 = \$460{,}000$$

f. Since $\pi_C + \pi_N(1 + i)/i > \pi_H(1 + i)/i$, (i.e., \$460,000 > \$480,000), there is no incentive to cheat. In other words, the cartel is stable and the collusive agreement is not likely to break down.

g.
$$\frac{\pi_N(1+i)}{i} = \frac{\$60{,}000(1.30)}{0.30} = \$60{,}000(4.33) = \$260{,}000$$

$$\frac{\pi_H(1+i)}{i} = \frac{\$80{,}000(1.30)}{0.30} = \$80{,}000(4.33) = \$346{,}666.67$$

$$\pi_C = \$100{,}000$$

Since $\pi_C + \pi_N(1 + i)/i > \pi_H(1 + i)/i$ (\$360,000 > \$346,666.67), then there is an incentive to cheat and the cartel is unstable. The collusive agreement is likely to break down.

h. Substituting the information in the payoff matrix into the left-hand side of inequality (13.6) yields

$$\frac{\pi_H - \pi_N}{\pi_C + \pi_N - \pi_H} = \frac{\$80{,}000 - \$60{,}000}{\$100{,}000 + \$60{,}000 - \$80{,}000}$$
$$= \frac{\$20{,}000}{\$80{,}000} = 0.25$$

Thus, from inequality (13.6), if the rate of return of 25% from remaining “honest” is less than the prevailing rate of interest, there will be an incentive to violate the collusive agreement and the cartel will be unstable. If the rate of return of 25% is greater than the prevailing rate of interest, it will be in the best interest for both firms to honor the agreement, in which case the cartel will be stable. These conclusions were verified by the answers to parts f and g.

Finally, if the interest rate is exactly 25%, then, *ceteris paribus*, each firm will be indifferent between cheating and remaining honest. This result may be verified by substituting 25% into the following expressions:

$$\frac{\pi_N(1+i)}{i} = \frac{\$60,000(1.25)}{0.25} = \$60,000(5) = \$300,000$$

$$\frac{\pi_H(1+i)}{i} = \frac{\$80,000(1.25)}{0.25} = \$80,000(5) = \$400,000$$

$$\pi_C = \$100,000$$

Since $\pi_C + \pi_N(1 + i)/i = \pi_H(1 + i)/i$ (i.e., \$400,000 = \$400,000), then, other things equal, each firm will be indifferent between cheating and remaining honest.

DETERMINANTS OF COLLUSIVE AGREEMENTS

The collusive agreements discussed thus for involved only two firms. The success of the collusion depended on the economic benefit of violating the agreement. If the economic benefit of violating the agreement is greater than the economic benefit of remaining faithful to the agreement, the cartel is likely to collapse. If the economic benefit of violating the agreement is less than the economic benefit of adhering to the collusive agreement, the viability of the cartel will depend on the existence of an effective trigger strategy to punish the cheater. Thus it would be useful to be ask to determine, in general, when collusive agreements are likely to be entered into and under what circumstances they are likely to succeed.

Number of Firms

Collusive agreements are more likely when the number of firms with similar interests and objectives is small. Collusive agreements are difficult to achieve among a large number of firms with widely divergent interests. Nevertheless, similarity of interests is no guarantee of success. In fact, as the number of firms that are party to the agreement increases, the probability of its success declines.

As the membership of a collusive agreement increases, it becomes increasingly difficult to monitor the behavior of each member. To see this, suppose that there are n parties to the agreement. Each member of the cartel must monitor the behavior of the other $(n - 1)$ members. Thus, the total number of monitoring arrangements necessary to police the cartel is $n(n - 1)$. In the two-firm case, only $2(2 - 1) = 2$ monitoring arrangements were needed to police the cartel. In the case of OPEC, which has 11 members (Algeria, Indonesia, Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi Arabia, the United Arab Emirates, and Venezuela), 110 monitoring arrangements are necessary to police compliance. Policing is made more difficult in the case of OPEC because of the widely divergent cultural, economic, and political characteristics of the members. Is it any wonder the membership of OPEC meets as frequently as it does to hammer out new production agreements? The incentive to cheat, especially by members with low production quotas, is very strong.

In addition to the difficulty of forming a collusive agreement, as the number of parties to the agreement increases, rising monitoring costs may make continuation of the cartel impractical. Under these circumstances the threat of sanctions being levied against the offending member is an empty one, and the cartel is likely to break down.

Firm Size

Economies of scale exist in the monitoring and policing of cartel arrangements. It is relatively less expensive for large firms to monitor the behavior of a relatively small number of large rivals, or a large number of relatively small rivals, than it is for small firms to monitor the behavior of a relatively large number of small rivals, or a small number of large rivals.

Explicit Versus Tacit Collusion

An important factor determining the existence and durability of collusive agreements is the manner in which such arrangements are entered into. Collusions may be either explicit or tacit. If a collusive agreement is explicit, the firms actually meet to hammer out details. An explicit collusive agreement will specify the responsibilities of each member. For example, explicit collusive agreements will specify production quotas for each member, collective pricing policies, and market shares. Moreover, to be effective, the collusive agreement will also specify the penalties for violating the agreement.

When an explicit agreement is not possible, perhaps because such an arrangement is illegal, firms may engage in tacit collusion. Tacit collusion occurs when firms do not explicitly conspire but, instead, come to an agreement indirectly. Such implicit agreements are possible only when firms in an industry develop an understanding of how the game is played. Firms develop this understanding by observing the behavior of rivals over time.

In the case of tacit collusions, firms also learn the most likely penalties levied for violating such "gentlemen's agreement". The reader might recall the "kinked" demand curve model of oligopolistic behavior discussed in Chapter 10. A central feature of that model was the anticipated reaction of firms to a price change by a rival. In the model, if a maverick firm lowered the price of its product to increase its market share, the price reduction would be matched by other firms in the industry, thereby thwarting the intentions of the initiator of the "price war." On the other hand, if a firm raised the price of its product, that increase would not be matched by its rivals, and the maverick firm would lose market share. This, of course, does not mean that prices are never raised or lowered. In the "kinked" demand curve model, for example, changes in collective price and output policy occurred only after changes in market and cost conditions common to all firms indicated that such changes were appropriate.

Finally, the threat of punishment for violating a collusive agreement will be meaningful only if threats are actually carried out. If member firms are unwilling or unable to punish violators, explicit and implicit collusions will be unstable and will ultimately break down. On the other hand, if the threat of sure, swift punishment is credible, collusive agreements will be stable.

Discriminatory Pricing

In the case of the "kinked" demand curve model, it was assumed that all firms in the industry charged customers the same price. Thus effective punishment of attempts by one firm to capture a larger market share by lowering price requires all firms in the industry to lower their prices as well. Clearly, in this case the cost of policing a collusive agreement will be quite high. On the other hand, if the industry is characterized by discriminatory pricing (i.e., charging a different price to different customers), member firms can punish violators by charging the lower price to the rival's customers while continuing to charge the higher price to its own customers. In this case, the cost of policing a collusive agreement is considerably reduced.

COOPERATIVE, SIMULTANEOUS-MOVE, FINITELY REPEATED GAMES

We have thus far considered games played only once and games played an infinite number of times. In this section we will examine games that are repeated a finite number of times.

Definition: A finitely repeated game is a game that is repeated a limited number of times.

There are two classes of finitely repeated games: those in which the players are uncertain about when the game will end and those in which the last play of the game is known to each player.

FINITELY REPEATED GAMES WITH AN UNCERTAIN END

Analytically, the only difference between infinitely repeated games and finitely repeated games with an uncertain end is the probability that the game will end after each play is $0 < \theta < 1$. Thus, the undiscounted, expected profit stream may be written as

$$E(\pi) = \pi + (1-\theta)\pi + (1-\theta)^2\pi + (1-\theta)^3\pi + \cdots$$

Thus, the discounted value of the expected profit stream may be written as

$$\begin{aligned} PVAD &= \pi\frac{1-\theta}{(1+i)^{n-1}} + \pi\frac{1-\theta}{(1+i)^{n-2}} + \cdots + \pi\frac{(1-\theta)^0}{(1+i)^0} \\ &= \sum_{t=1\to n} \pi\left(\frac{1-\theta}{1+i}\right)^{n-1} \end{aligned} \tag{13.7}$$

where $PVAD$ is the present value of an annuity due, i is the nominal (market) interest rate, and θ is the probability that the game will end. For infinitely repeated games ($n = \infty$), it can be demonstrated that

$$PVAD = \frac{\pi(1+i)}{1+\theta} \tag{13.8}$$

The economic benefit of maintaining the agreement is the present value of all expected future profits earned by remaining "honest" ($PVAD_H$) with respect to the terms of the agreement, which is given by

$$PVAD_H = \frac{\pi_H(1+i)}{1+\theta} \tag{13.9}$$

If, on the other hand, a firm were to violate the agreement, the economic benefit is the immediate (one-shot) gain from cheating (π_C) plus the present value of all expected profits that will be earned from the Nash equilibrium ($PVAD_N$), that is, the profits earned in the absence of a collusive agreement (π_N). The economic benefit of violating the agreement is summarized as follows:

$$\pi_C + PVAD_N = \pi_C + \frac{\pi_N(1+i)}{1+\theta} \tag{13.10}$$

When will it pay to cheat for finitely repeated games? Cheating will occur when the expected present value of violating the agreement is greater than the expected present value of remaining "honest." This condition is summarized in the following inequality:

$$\pi_C + PVAD_N > PVAD_H$$

or

$$\pi_C + \frac{\pi_N(1+i)}{i+\theta} > \frac{\pi_H(1+i)}{1+\theta} \tag{13.11}$$

CHEATING RULE FOR FINITELY REPEATED GAMES WITH AN UNCERTAIN END

Inequality (13.11) may be rearranged to yield the cheating rule for finitely repeated games with an uncertain end. The interpretation of inequality (13.12) is similar to the interpretation of inequality (13.6). If the expected rate of return from adhering to the collusive agreement is less than the prevailing rate of interest, there will be an incentive to cheat and the cartel will break down.

$$\frac{\pi_H - \pi_N - \theta\pi_C}{\pi_C + \pi_N - \pi_H} < i \tag{13.12}$$

It should also be noted that inequality (13.12) differs from inequality (13.6) by the addition of $-\theta\pi_c$ in the numerator. This indicates that in the presence of uncertainty about the duration of the game, the threshold for violating the agreement is lowered. In other words, in the presence of uncertainty about the duration of the game, the likelihood that a player will violate the agreement will be greater than for an infinitely repeated game.

The reader should note that when the probability that a game will end is zero, the solution to inequality (13.12) is identical to that of inequality (13.6). Finally, it is interesting to ask at what probability any positive interest rate will result in a breakdown of the collusive agreement. This probability may be determined by setting the left-hand side of inequality (13.12) equal to zero and solving for θ:

$$\frac{\pi_H - \pi_N - \theta\pi_C}{\pi_C + \pi_N - \pi_H} = 0 \tag{13.13}$$

Assuming that the denominator of Equation (13.13) is nontrivial, then this reduces to

$$\pi_H - \pi_N - \theta\pi_C = 0$$

or

$$\theta = \frac{\pi_H - \pi_N}{\pi_C} \tag{13.14}$$

Problem 13.7. Consider once again, the payoff matrix summarized in Figure 10.9.

a. Suppose that the probability that will the game end after each play is *0.1*. Above what interest rate can we expect the cartel to break down?
b. Above what probability will the game end on the next play for any positive interest rate result if at least one player violates the collusive agreement?

Solution

a. Substituting the data from Figure 10.9 into the left-hand side of inequality (13.12) yields

$$\frac{\pi_H - \pi_N - \theta\pi_C}{\pi_C + \pi_N - \pi_H} = \frac{\$1{,}000{,}000 - \$250{,}000 - 0.1(\$5{,}000{,}000)}{\$5{,}000{,}000 - \$250{,}000 - \$1{,}000{,}000}$$
$$= \frac{\$250{,}000}{\$4{,}750{,}000} = 0.0526$$

Thus, from inequality (13.6), if the prevailing rate of interest is greater than 5.26%, each firm will have an incentive to violate the collusive agreement, in which case the cartel will be unstable. If the rate of interest is less than 5.26%, it will be in the best interest for both firms to honor the agreement and the cartel will be stable. Finally, if the interest rate is exactly 5.26%, then *ceteris paribus*, each firm will be indifferent between cheating and remaining honest.

b. Substituting the data from Figure 10.9 into the right-hand side of Equation (13.14) yields

$$\theta = \frac{\pi_H - \pi_N}{\pi_C} = \frac{\$1{,}000{,}000 - \$250{,}000}{\$5{,}000{,}000}$$
$$= \frac{\$750{,}000}{\$5{,}000{,}000} = 0.15$$

That is, if the probability that the game will end on the next play is 15% or greater, any positive interest rate will result in cheating. To verify this result, substitute θ = 0.15 into the left-hand side of inequality (13.12). This yields

$$\frac{\pi_H - \pi_N - \theta\pi_C}{\pi_C + \pi_N - \pi_H} = \frac{\$1{,}000{,}000 - \$250{,}000 - 0.15(\$5{,}000{,}000)}{\$5{,}000{,}000 - \$250{,}000 - \$1{,}000{,}000}$$
$$= \frac{\$0}{\$4{,}750{,}000} = 0.0$$

FINITELY REPEATED GAMES WITH A CERTAIN END

Surprisingly, when a finitely repeated game has a certain end, the solution collapses into a series of noncooperative one-shot games. To see this, consider once again the game in Figure 10.9. Let us assume initially that this game is played just twice. Suppose, further, that firms *A* and *B* agree to charge a "high" price for their product. In an infinitely repeated game, violation of the agreement by either player will alter future game plays. For example, if firm *A* were to charge a "low" price in violation of its agreement, firm *B*, might seek to punish firm *A* by ruling out any future cooperation. In other words, a violation of firm *A*'s agreement with firm *B* could "trigger" a strategy change by firm *B*.

With a finitely repeated game with a certain end, however, the use of trigger strategies to enforce a collusive agreement will fail. The reason for this is relatively straightforward. Since each player realizes that there can be no punishment for "dishonest" actions in subsequent periods, each has the incentive to adopt a strategy that is consistent with a one-shot noncooperative game. This is known as the end-of-period problem. In the game depicted in Figure 10.9, the end-of-period problem means that each player will adopt a "low" price strategy in the second period. Even if firm *A* believed that firm *B* would continue to charge a "high" price, it would be in firm *A*'s best interest to charge a "low" price because there is nothing firm *B* can do to punish firm *A* for deviant behavior. The same line of reasoning, of course, holds true for firm *B*. This result is not surprising, nor is it particularly interesting. What is interesting, however, is how the play in the second period affects the play in the first period.

In the game depicted in Figure 10.9, since each firms realizes that its rival will charge a "low" price in the second period, the first period, in effect, becomes the last period, in which case both firms again have an incentive to adopt the same strategy as in a one-shot game! In terms of the game in Figure 10.9, each firm will charge a "low" price in the first period as well. In other words, the Nash equilibrium for both periods is for both firms to charge a "low" price, with each player earning profits of $250,000.

What is remarkable about finitely repeated games with a certain end is that regardless of the number of periods, the process reduces to a series of noncooperative one-shot games. In other words, once the last period has been identified, each player has an incentive to view the next-to-the-last period as a one-shot game, which transforms that period into the last period, and so on. This process, which has been described as "backward unraveling," effectively renders collusive agreements unworkable.

Definition: The end-of-period problem arises in a finitely repeated game with a certain end because each period effectively becomes the final period, in which case the game reduces to a series of noncooperative one-shot games.

Two interesting examples of the end-of-game problem may be found in Baye and Beil (1994, Chapter 10). In the first, a worker announces the intention to quit on a specified date. In general, it is reasonable to assume that at least one reason people work hard is the fear of being fired if they are caught "goofing off." In fact, if the net benefit of being diligent is greater than the net benefit of being a "gold brick," workers will find it in their best interest to do a good job.

Now, suppose that a worker announces on Monday the intention quit on Tuesday. How seriously will the worker take his or her job on Tuesday?

Probably not very seriously. Why? Since the worker has no intention of showing up on Wednesday, any threat by management to "fire" the employee on Tuesday is meaningless. In terms of the preceding discussion, the worker's choice of working hard or goofing off reduces to a noncooperative, one-shot game. Since the threat of being fired no longer has any meaning, the net benefit of working hard is considerably reduced, suggesting that it may be in the worker's best interest to goof off on Tuesday.

Will the employee's attitude toward work on Monday be affected by the decision to quit on Tuesday? Probably, yes. As noted earlier, since the worker has identified Tuesday as the final period, there is an incentive to view Monday as a one-shot game as well, which suggests that the worker will goof off on Monday. Even if the employee had given 2 weeks notice, the remaining days on the job will collapse into a series of noncooperative, one-shot games. Does this scenario sound unrealistic? If it does, consider how your attitude toward work might change if you handed in your 2-week notice. Would you be disposed to work just as hard as before, or would you instead "take it easy" while counting down your final days?

What, if anything, should management do under these circumstances? Management could, of course, fire the worker immediately upon learning of the worker's intention to quit, but such a move would be counterproductive. The reason for this is that workers would change their strategy of giving "2 weeks notice" to a strategy of advising management at the close of business on the day of the planned resignation. This would present management with the extremely difficult task of finding replacement workers at short notice without disrupting the production process.

On discussion of a finitely repeated game with a certain end, however, suggests a possible solution to management's dilemma. The answer lies in extending the game beyond the resignation date. For example, management could offer the employee assistance in identifying new employment opportunities, or perhaps provide letters of recommendation to potential future employers. By extending the game into the future, it will be in the worker's best interest to avoid "burning bridges" by goofing off during the final days of employment.

Baye and Beil's second example of the end-of-game problem deals with the so-called snake oil salesman. During the American westward expansion of the late nineteenth century, "snake oil" salesmen traveled from frontier town to frontier town selling bottles of elixirs promising everything from a cure for toothaches to a remedy for baldness. Of course, these claims were bogus, but by the time customers realized that they had been "had," the salesman would be long gone. By contrast, had a local merchant attempted to pull a similar scam, there was a very good chance that the merchant would soon be decorating the nearest tree—from the neck! In fact, it is precisely because the local merchant is playing a finitely repeated game (assuming that he or she does not live forever) with an uncertain future that the threat of punishment for unethical behavior ensures that the person sells products of reliable quality. For the snake oil salesman there is no tomorrow, and so the transaction is played as a noncooperative, one-shot game.

FOCAL-POINT EQUILIBRIUM

As the preceding paragraphs testify, the idea of a Nash equilibrium is a very powerful concept. Unfortunately, the existence of a Nash equilibrium does not guarantee that solutions to game-theoretic problems will be unique. The game summarized earlier in Figure 13.5 is one such example. Although it was argued that TEXplor and Clampett acted rationally by adopting the strategy profile {*Wide*, *Narrow*}, the result was an "irrational" {*Wide*, *Wide*} strategy profile, which was nonetheless a Nash equilibrium. The strategy profile {*Wide*, *Wide*} is a Nash equilibrium because neither firm can improve its payoff by switching strategies. It can be readily verified, however, that if the firms had adopted a {*Narrow*, *Wide*} strategy profile, the result would have been a {*Narrow*, *Narrow*} Nash equilibrium.

In general, it may be demonstrated that if both players have a strictly dominant strategy, the result will be a strictly dominant strategy equilibrium, in which case there is a unique Nash equilibrium. If neither player has a strictly dominant strategy, or if the strategy profile results in a weakly dominant strategy equilibrium, then the Nash equilibrium may not be unique. When there are multiple Nash equilibria, then without additional information regarding the terms of the game, it will be difficult to predict the strategy profiles that will be adopted by players, the context within which the game is being played, or the interactions between players. One possible solution to a game theoretic problem in the presence of multiple Nash equilibria is the *focal-point equilibrium*, suggested by Thomas Schelling (1960). Schelling has suggested that in the presence of multiple Nash equilibria, a single solution may "stand out" because the players share a common "understanding" of the problem.

Definition: A focal-point equilibrium may exist in the presence of multiple Nash equilibria when a single solution "stands out" because the players share a common "understanding" of the problem.

To illustrate the concept of a focal-point equilibrium, suppose that a father and his son become separated in an amusement park, and no prior arrangement had been made to set a place to meet in the event that this happened. Is it not likely, however, that in the event of separation both would think of the same place to try and find the other, such as the park's main gate or the office of park security? Bierman and Fernandez (1998, Chapter 1) referred to this common understanding of the problem as "conventional wisdom." Shelling illustrated the concept of focal-point equilibria with the following "abstract puzzles."

1. A coin is flipped and two players are instructed to call "heads" or "tails." If both players call "heads," or both call "tails," then both win a prize. If one player calls "heads" and the other calls "tails," then neither wins a prize.
2. A player is asked to circle one of the following six numbers: 7, 100, 13, 261, 99, and 555. If all of the players circle the same number, then each wins a prize; otherwise no one wins anything.

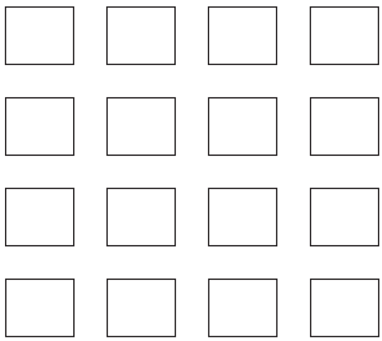

3. A player is asked to put a check mark in one sixteen squares, arranged as shown. If all the players check the same square, each wins a prize; otherwise no one wins anything.
4. Two players are told to meet somewhere in New York City, but neither player has been told where the meeting is to occur. Neither player has ever been placed in this situation before, and the two are not permitted to communicate with each other. Each player must guess the other's probable location.
5. In the preceding scenario each player is told the date, but not the time, of the meeting. Each player must guess the exact time that the meeting is to take place.
6. A player is told to write down a positive number. If all players write the same number, each player wins a prize; otherwise no one wins anything.
7. A player is told to name an amount of money. If all players name the same amount, each wins that amount.
8. A player is asked to divide \$100 into two piles labeled pile *A* and pile *B*. Another player is asked to do the same. If the amounts in all four piles coincide, player each receives \$100, otherwise, neither player wins anything.

9. The results of a first ballot in an election were tabulated as follows:

Smith	19 votes
Jones	28 votes
Brown	15 votes
Robinson	29 votes
White	9 votes

A second ballot is to be taken. A player is asked to predict which candidate will receive a majority of votes on the second ballot. The player has no interest in the outcome of the second ballot. The player who correctly predicts the candidate receiving the majority of votes will win a prize, and everyone knows that a correct prediction is in everyone's best interest. If the player incorrectly predicts the "winner" of the second ballot, he or she will win nothing.

In each of these nine scenarios there are multiple Nash equilibria. Schelling found, however, that in an "unscientific sample of respondents," people tended to focus (i.e., to use focal points) on just a few such equilibria. Schelling found, for example, that 86% of the respondents chose "heads" in problem 1. In problem 2 the first three numbers received 90% of the votes, with the number 7 leading the number 100 by a slight margin and the number 13 in third place. In problem 4, an absolute majority of the respondents, who were sampled in New Haven, Connecticut, proposed meeting at the information booth in Grand Central Station, and virtually all of them agreed to meet at 12 noon. In problem 6, two-fifths of all respondents chose the number 1. In problem 7, 29% of the respondents chose \$1 million, and only 7 percent chose cash amounts that were not multiples of 10. In problem 8, 88% of the respondents put \$50 into each pile. Finally, in problem 9, 91% of the respondents chose Robinson.

Schelling also found that the respondents chose focal points even when these choices where not in their best interest. For example, consider the following variation of problem 1. Players *A* and *B* are asked to call "heads" or "tails." The players are not permitted to communicate with each other. If both players call "heads," player *A* gets \$3 and player *B* gets \$2. If both players call "tails," then player *A* gets \$2 and player *B* gets \$3. Again, if one player calls "heads" and the other calls "tails," neither player wins a prize. In this scenario Schelling found that 73% of respondents chose "heads" when given the role of player *A*. More surprising is that 68% of respondents in the role of player *B* still chose "heads" in spite of the bias against player *B*. The reader should verify that if both players attempt to win \$3, neither one will win anything.

The economic significance of focal-point equilibria becomes readily apparent when we consider cooperative, non-zero-sum, simultaneous-move, infinitely repeated games. Where explicit collusive agreements are prohibited, the existence of focal-point equilibria suggests that tacit collusion, coupled with the policing mechanism of trigger strategies, may be possible. A fuller discussion of these, and other related matters, is deferred to the next section.

MULTISTAGE GAMES

The final scenario we will consider in this brief introduction to game theory is that of the multistage game. Multistage games differ from the games considered earlier in that play is sequential, rather than simultaneous. Figure 13.12, which is an example of an *extensive-form game*, summarizes the players, the information available to each player at each stage, the order of the moves, and the payoffs from alternative strategies of a multistage game.

Definition: An extensive-form game is a representation of a multistage game that summarizes the players, the stages of the game, the information available to each player at each stage, player strategies, the order of the moves, and the payoffs from alternative strategies.

The extensive-form game depicted in Figure 13.12 has 2 players: player *A* and player *B*. The boxes in the figure are called *decision nodes*. Inside each box is the name of the player who is to move at that decision node. At each decision node the designated player must decide on a strategy, which is represented by a *branch*, which represents a possible move by a player. The arrow indicates the direction of the move. The collection of decision nodes and branches is called a *game tree*. The first decision node is called the *root* of the game tree. In the game depicted in Figure 13.12, player *A* moves first. Player *A*'s move represents the first stage of the game. Player *A*, who is at the root of the game tree, must decide whether to adopt a *Yes* or a *No* strategy. After player *A* has decided on a strategy, player *B* must

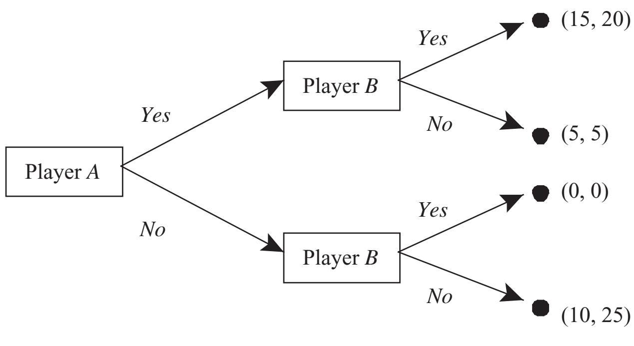

Payoffs: (Player *A*, Player *B*)

FIGURE 13.12 Extensive-form game.

decide how to respond in the second stage of the game. For example, if player A's strategy is *Yes*, then player B must decide whether to respond with a *Yes* or a *No*.

At the end of each arrow are small circles called *terminal nodes*. The game ends at the terminal nodes. To the right of the terminal notes are the payoffs. In Figure 13.12, the first entry in parenthses is the payoff to player A and the second entry is the payoff to player B. If player B adopts a *Yes* strategy, the payoff for player A is 15 and the payoff for player B is 20. In summary, an extensive-form game is made up of a game tree, terminal nodes, and payoffs.

As with simultaneous-move games, the eventual payoffs depend on the strategies adopted by each player. Unlike simultaneous-move games, in multistage games the players move sequentially. In the game depicted in Figure 13.12, player A moves without prior knowledge of player B's intended response. player B's move, on the other hand, is conditional on the move of player A. In other words, while player B moves with the knowledge of player A's move, player A can only anticipate how player B will react. The ideal strategy profile for player A is {*Yes*, *Yes*}, which yields payoffs of (15, 20). For player B, the ideal strategy profile is {*No*, *No*}, which yields payoffs of (10, 25). The challenge confronting player B is to get player A to say *No* on the first move. As we will see, the solution is for player B to convince player A that regardless of what player A says, player B will say *No*. To see this, consider the following scenario.

Suppose that player B announces that he or she has adopted the following strategy: if player A says *Yes*, then player B will say *No*; if player A says *No*, player B will also say *No*. With the first strategy profile {*Yes*, *No*} the payoffs are (5, 5). With the second strategy profile the payoffs are (10, 25). In this case, it would be in player A's best interest to say *No*. Of course, the choice of strategies is a "no brainer" if player A believes that player B will follow through on his or her "threat." player A's first move will be *No* because the payoff to player A from a {*No*, *No*} strategy is greater than from a {*Yes*, *No*} strategy. In fact, the strategy profile {*No*, *No*} is a Nash equilibrium. Why? If player B's threat to always say *No* is credible, then player A cannot improve his or her payoff by changing strategies.

As the reader may have already surmised, the final outcome of this game depends crucially on whether player A believes that player B's threat to always say *No* is credible. Is there a reason to believe that this is so? Probably not. To see this, assume again that the optimal strategy profile for player A is {*Yes*, *Yes*}, which yields the payoff (15, 20). If player A says *Yes*, the payoff to player B from saying *No* is 5, but the payoff for saying *Yes* is 20. Thus, if player B is rational, the threat to say *No* lacks credibility and the resulting strategy profile is {*Yes*, *Yes*}.

Note that strategy profile {*Yes*, *Yes*} is also a Nash equilibrium. Neither player can improve his or her payoff by switching strategies. In particular, if player B's strategy was to say *Yes* if player A says *Yes* and say *No* if player A says *No*, then player A's payoff is 15 by saying *Yes* and 10 by saying *No*. Clearly, player A's best strategy, given player B's move, is to say *Yes*.

We now have two Nash equilibria. Which one is the more reasonable? It is the Nash equilibrium corresponding to the strategy profile {*Yes*, *Yes*} because player B has no incentive to carry through with the threat to say *No*. The Nash equilibrium corresponding to the strategy profile {*Yes*, *Yes*} is referred to as a *subgame perfect equilibrium* because no player is able to improve on his or her payoff at any stage (decision node) of the game by switching strategies. In a subgame perfect equilibrium, each player chooses at each stage of the game an optimal move that will ultimately result in optimal solution for the entire game. Moreover, each player believes that all the other players will behave in the same way.

Definition: A strategy profile is a subgame perfect equilibrium if it is a Nash equilibrium and allows no player to improve on his or her payoff by switching strategies at any stage of a dynamic game.

The idea of a subgame perfect equilibrium may be attributed to Reinhard Selten (1975). Selten formalized the idea that a Nash equilibrium with incredible threats is a poor predictor of human behavior by introducing the concept of the *subgame*. In a game with perfect information, a subgame is any subset of branches and decision nodes of the original multistage game that constitutes a game in itself. The unique initial node of a subgame is called a *subroot* of the larger multistage game. Selten's essential contribution is that once a player begins to play a subgame, that player will continue to play the subgame until the end of the game. That is, once a player begins a subgame, the player will not exit the subgame in search of an alternative solution. To see this, consider Figure 13.13, which recreates Figure 13.12.

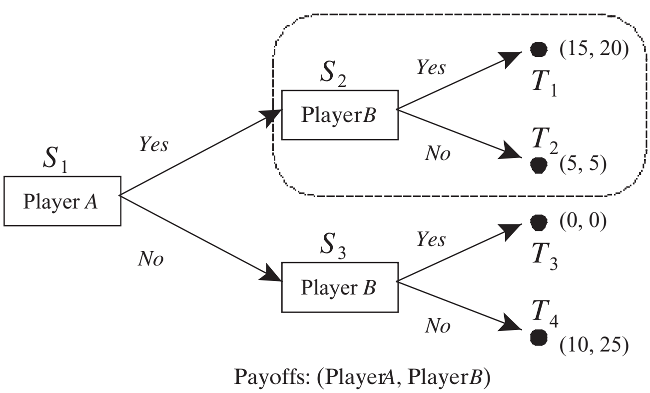

FIGURE 13.13 A subgame.

Figure 13.13 is a multistage game consisting of two subgames. The multistage game itself begins at the initial node, S_1. The two subgames begin at subroots S_2 and S_3. The subgame that begins at subroot S_2, which is highlighted by the dashed, rounded rectangle, has two terminal nodes, T_1 and T_2, with payoffs of (15, 20) and (5, 5), respectively. In games with perfect information, every decision node is the subroot of a larger game. A player who begins a subgame is common knowledge to all the other players. The student should verify that this subgame has a unique Nash equilibrium. At this Nash equilibrium player *B* says *Yes*. The reader should also verify that the subgame with subroot S_3 also has a unique Nash equilibrium.

As we have seen, the final outcome of the multistage game depicted in Figure 13.12 depends on whether player *A* believes that player *B*'s threat to say *No* is credible. If player *B* is rational, the threat to say *No* lacks credibility and the resulting strategy profile is {*Yes*, *Yes*}. Thus, the nonoptimality of the strategy profile {*No*, *No*} makes player *B*'s threat incredible. Thus, this strategy profile is eliminated by the requirement that Nash equilibrium strategies remain when applied to any subgame. A Nash equilibrium with this property is called a *subgame perfect equilibrium*. The Nash equilibrium corresponding to the strategy profile {*Yes*, *Yes*} is referred to as a subgame perfect equilibrium because no player is able to improve on his or her payoff at any stage (decision node) of the game by switching strategies. As we will soon see, the concept of a subgame perfect equilibrium is essential element of the backward induction solution algorithm.

EXAMPLE: SOFTWARE GAME

As we have already seen, one of the problems with multistage games is the selection of an optimal strategy profile in the presence of multiple Nash equilibria. This issue will be addressed in later sections. For now, consider the following example of a subgame perfect equilibrium, which comes directly from Bierman and Fernandez (1998, Chapter 6).

Macrosoft Corporation is a computer software company that is planning to introduce a new computer game into the market. Macrosoft's management is considering two marketing approaches. The first approach involves a "Madison Avenue" type of advertising campaign, while the second approach emphasizes word of mouth. Bierman and Fernandez described the first approach as "slick" and the second approach as "simple."

The timing involved in both approaches is all-important in this example. Although expensive, the "slick" approach will result in a high volume of sales in the first year, while sales in the second year are expected to decline dramatically as the market becomes saturated. The inexpensive "simple" approach, on the other hand, is expected to result in relatively low sales volume in the first year, but much higher sales volume in the second year as "word gets around." Regardless of the promotional campaign adopted,

TABLE 13.1 Macrosoft's Profits if Microcorp Does Not Enter the Market

	Slick	*Simple*
Gross profit in year 1	\$900,000	\$200,000
Gross profit in year 2	\$100,000	\$800,000
Total gross profit	\$1,000,000	\$1,000,000
Advertising cost	–\$570,000	–\$200,000
Total net profit	\$430,000	\$800,000

TABLE 13.2 Macrosoft's Profits if Microcorp Enters the Market

	Slick	*Simple*
Gross profit in year 1	\$900,000	\$200,000
Gross profit in year 2	\$50,000	\$400,000
Total gross profit	\$950,000	\$600,000
Advertising cost	–\$570,000	–\$200,000
Total net profit	\$380,000	\$400,000

no significant sales are anticipated after the second year. Macrosoft's net profits from both campaigns are summarized in Table 13.1.

The data presented in Table 13.1 suggest that Macrosoft should adopt the inexpensive "simple" approach because of the resulting larger total net profits. The problem for Macrosoft, however, is the threat of a "legal clone," that is, a competing computer game manufactured by another firm, Microcorp, that is, to all outward appearances, a close substitute for original. The difference between the two computer games is in the underlying programming code, which is sufficiently different to keep the "copycat" firm from being successfully sued for copyright infringement. In this example, Microcorp is able to clone Macrosoft's computer game within a year at a cost of \$300,000. If Microcorp decides to produce the clone and enter the market, the two firms will split the market for the computer game in the second year. The payoffs to both companies in years 1 and 2 are summarized in Tables 13.2 and 13.3.

Given the information provided in Tables 13.2 and 13.3 what is the optimal marketing strategy for each player, Macrosoft and Microcorp? Since the decisions of both companies are interdependent and sequential the problem may be represented as the extensive-form game in Figure 13.14.

It should be obvious from Figure 13.14 that Macrosoft moves first and has just one decision node. The choices facing Macrosoft consist of "slick"

TABLE 13.3 Microcorp's Profits after Entering the Market

	Slick	*Simple*
Gross profit in year 1	$0	$0
Gross profit in year 2	$50,000	$400,000
Total gross profit	$50,000	$400,000
Cloning cost	–$300,000	–$300,000
Total net profit	$250,000	$100,000

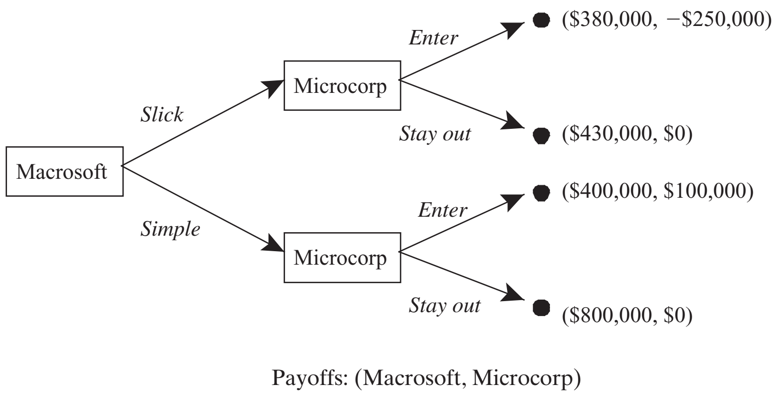

FIGURE 13.14 The software game.

and "simple." Microcorp, on the other hand, has two decision nodes. Microcorp's strategy is conditional on Macrosoft's decision of a promotional campaign. For example, if Macrosoft decides upon a "slick" campaign, Microcorp might decide to "stay out" of the market. On the other hand, if Macrosoft decides on a "simple" campaign, Microcorp might decide that its best move is to "enter" the market. This strategy profile for Microcorp might be written {*Stay out*, *Enter*}. As the reader will readily verify, there are four possible strategy profiles available to Microcorp. These strategy profiles represent Microcorp's contingency plans. Which strategy is adopted will depend on Macrosoft's actions. Since different strategies will often result in the same sequence of moves, it is important not to confuse strategies with actual moves.

NASH EQUILIBRIUM AND BACKWARD INDUCTION

At this point we naturally are interested in the strategic choices of each player. As we will soon see, finding an optimal solution for multistage games

		Macrosoft Slick	Macrosoft Simple
Microcorp	Enter	(–$250,000, $380,000)	($100,000, $400,000)
	Stay out	($0, $430,000)	($0, $800,000)

Payoffs: (Microcorp, Macrosoft)

FIGURE 13.15 Payoff matrix for a two-player, simultaneous-move game.

is not nearly as simple as it might seem at first glance. This is because multistage noncooperative games are often plagued with multiple Nash equilibria. A *solution concept* is a methodology for finding solutions to multistage games. There is no universally accepted solution concept that can be applied to every game. Bierman and Fernandez (1998, Chapter 6) have proposed the *backward induction* concept for finding optimal solutions to multistage games involving multiple Nash equilibria. The backward induction method is sometimes referred to as the *fold-back method*.

Definition: Backward induction is a methodology for finding optimal solutions to multistage games involving multiple Nash equilibria.

The solution concept of backward induction will be applied to the multistage game depicted in Figure 13.14, which assumes that Macrosoft and Microcorp have *perfect information*. Perfect information consists of player awareness of his or her position on the game tree whenever it is time to move. Before discussing the backward induction methodology, consider again the payoffs (in $000's) in Figure 13.14, which is summarized as the normal-form game in Figure 13.15.

Now consider the noncooperative solution to the game depicted in Figure 13.15. The reader should verify that a Nash equilibrium to this game is the strategy profile {*Enter*, *Simple*}. It will be recalled that in a Nash equilibrium, each player adopts a strategy it believes is the best response to the other player's strategy and neither player's payoff can be improved by changing strategies.

The limitation of a Nash equilibrium as a solution concept is that changing the strategy of any single player may result in a new Nash equilibrium, which may be not be an optimal solution. To see this, consider Figure 13.16, which is the *strategic form* of the multistage game in Figure 13.14. Strategic-form games illustrate the payoffs to each player from every possible strategy profile. Macrosoft, for example, may adopt one of two promotional campaigns—*Slick* or *Simple*. Microcorp, on the other hand, may adopt one of four strategic responses: (*Enter*, *Enter*), (*Enter*, *Stay out*), (*Stay out*, *Enter*), or (*Stay out*, *Stay out*).

Definition: The strategic form of a game summarizes the payoffs to each player arising from every possible strategy profile.

		Macrosoft	
		Slick	Simple
Microcorp	(Enter, Enter)	(−$250,000, $380,000)	**($100,000, $400,000)**
	(Enter, Stay out)	(−$250,000, $380,000)	($0, 800,000)
	(Stay out, Enter)	**($0, $430,000)**	($100,000 $400,000
	(Stay out, Stay out)	($0, $430,000)	($0, $800,000

Payoffs: (Microcorp, Macrosoft)

FIGURE 13.16 Payoff matrix for a strategic-form game.

The cells in Figure 13.16 summarize the payoffs from all possible strategic combinations. For example, suppose that Microcorp decides to "enter" regardless of the promotional campaign adopted by Macrosoft. In this case, Macrosoft will select a "simple" campaign, which is the Nash equilibrium of the normal-form game illustrated in Figure 13.15. The strategy profile for this game may be written {*Simple*, (*Enter*, *Enter*)}. On the other hand, if Macrosoft adopts a "slick" strategy, Microcorp can do no better than to adopt the strategy (*Stay out*, *Enter*). The strategy profile for this game may be written {*Slick*, (*Stay out*, *Enter*)}. This is a Nash equilibrium for the strategic-form game in Figure 13.16 but is not a Nash equilibrium for the normal-form game in Figure 13.15!

Finding an optimal solution to a multistage game using the backward induction methodology involves five steps:

1. Start at the terminal nodes. Trace each node to its immediate predecessor node. The decisions at each node may be described as "basic," "trivial," or "complex." Basic decision nodes have branches that lead to exactly one terminal node. Basic decision nodes are trivial if they have only one branch. A decision node is complex if it is not basic, that is, if at least one branch leads to more than one terminal node. If a trivial decision node is reached, continue to move up the decision tree until a complex or a nontrivial decision node is reached.
2. Determine the optimal move at each basic decision node reached in step 1. A move is optimal if it leads to the highest payoff.
3. Disregard all nonoptimal branches from decision nodes reached in step 2. With the nonoptimal branches disregarded, these decision nodes become trivial (i.e., they now have only one branch). The resulting game tree is simpler than the original game tree.
4. If the root of the game tree has been reached, then stop. If not, repeat steps 1–3. Continue in this manner until the root of the tree has been reached.

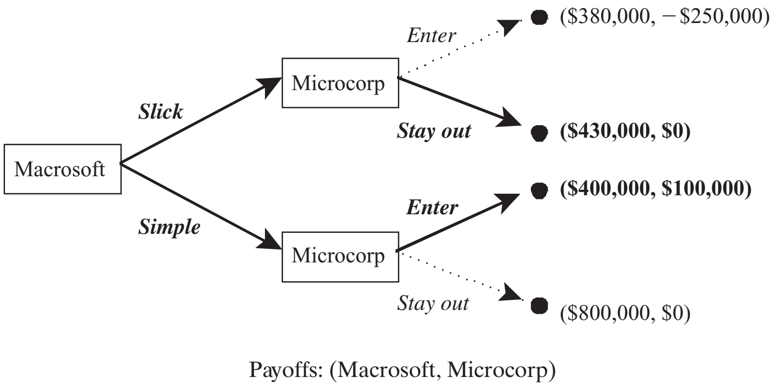

FIGURE 13.17 Using backward induction to find a Nash equilibrium.

5. After the root of the game tree has been reached, collect the optimal decisions at each player's decision nodes. This collection of decisions comprises the players' optimal strategies.

The backward induction solution concept will now be applied to the multistage game depicted in Figure 13.14. From each terminal node, move to the two Microcorp decision nodes. Each of these decision nodes is basic, since the branches lead to exactly one terminal node. If Macrosoft chooses a "slick" campaign, the optimal move for Microcorp is to stay out, since the payoff is \$0 compared with a payoff of –\$250,000 by entering. The "enter" branch should be disregarded in future moves. If Macrosoft chooses a "simple" campaign, the optimal move for Microcorp is to enter, since the payoff is \$100,000 compared with a payoff of \$0 by staying out. This "stay out" branch should be disregarded in future moves. The resulting extensive-form game is illustrated in Figure 13.17.

An examination of Figure 13.17 will reveal that the optimal strategy for Microcorp is (*Stay out*, *Enter*). The final optimal strategy profile is {*Slick*, (*Stay out*, *Enter*)}, which yields payoffs of \$430,000 for Macrosoft and \$0 for Microcorp. The reader should note that the choice of this Nash equilibrium (\$0, \$430,000) from Figure 13.16 differs from the Nash equilibrium (\$100,000, \$400,000) in Figure 13.12. The implication of the backward induction method is straightforward. By taking Microcorp's entry decision into account, Macrosoft avoided making a strategy decision that would have cost it \$30,000.

Problem 13.8. Consider, again, the strategy for the software game summarized in Figure 13.17. Suppose that the cost of cloning Macrosoft's computer game is \$10,000 instead of \$300,000.

TABLE 13.4 Microcorp's Profits after Entering the Market

	Slick	*Simple*
Gross profit in year 1	$0	$0
Gross profit in year 2	$50,000	$400,000
Total gross profit	$50,000	$400,000
Cloning cost	–$10,000	–$10,000
Total net profit	$40,000	$390,000

a. Diagram the new extensive-form for this multistage game.
b. Use the backward induction solution concept to determine the new optimal strategy profile for this game. Illustrate your answer.

Solution

a. Microcorp's profits at the lower cost of cloning Macrosoft's computer game and entering the market are presented in Table 13.4.
Assuming that Macrosoft's net profits remain unchanged, the extensive form of this game is as shown in Figure 13.18.
b. Using the backward induction solution methodology, from each terminal node move to Microcorp's two decision nodes. Each of these decision nodes is basic. If Macrosoft chooses a *Slick* campaign, the optimal move for Microcorp is to *Enter*, since the payoff is $40,000 compared with a payoff of $0 by staying out. The *Stay out* branch should be disregarded in future moves. If Macrosoft chooses a *Simple* campaign, the optimal move for Microcorp is to Enter, since the payoff is $390,000 compared with a payoff of $0 if it adopts a *Stay out* strategy. The *Stay out* branch should be disregarded in future moves. In the resulting extensive-form game, diagrammed in Figure 13.19, we see that the optimal strategy for Microcorp is (*Enter*, *Enter*). The final optimal strategy profile is {*Simple*, (*Enter*, *Enter*)}, which yields payoffs of $400,000 for Macrosoft and $390,000 for Microcorp.

Problem 13.9. Suppose that in Problem 13.8 the cost of cloning Macrosoft's computer game is $500,000 instead of $300,000.
a. Diagram the new extensive-form for this multistage game.
b. Use the backward induction solution concept to determine the new optimal strategy profile for this game. Illustrate your answer.

Solution

a. Microcorp's profits at the higher cost of cloning Macrosoft's computer game and entering the market are presented in Table 13.5.
The extensive form of this game, assuming that Macrosoft's net profits remain unchanged, is diagrammed in Figure 13.20.

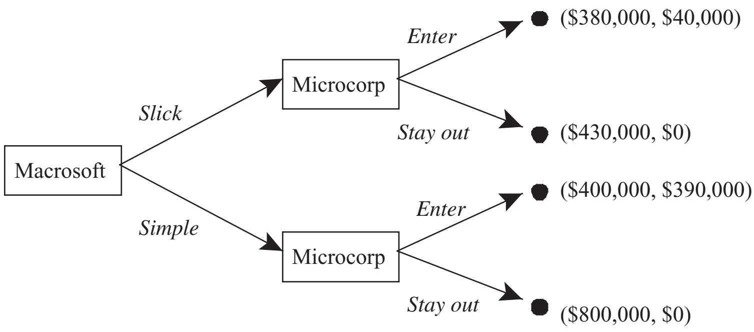

FIGURE 13.18 Game tree for problem 13.8.

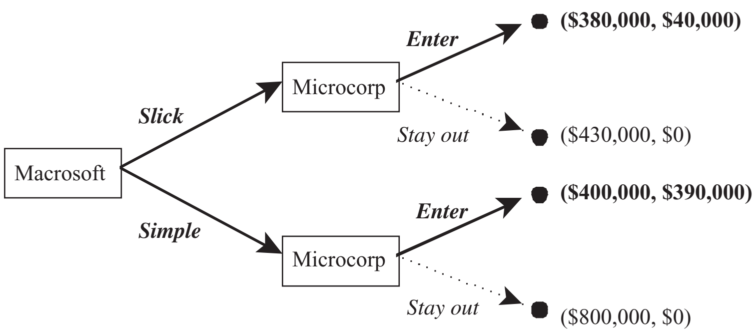

FIGURE 13.19 Solution to problem 13.8 using backward induction.

TABLE 13.5 Microcorp's Profits after Entering the Market

	Slick	*Simple*
Gross profit in year 1	$0	$0
Gross profit in year 2	$50,000	$400,000
Total gross profit	$50,000	$400,000
Cloning cost	$500,000	$500,000
Total net profit	–$450,000	–$100,000

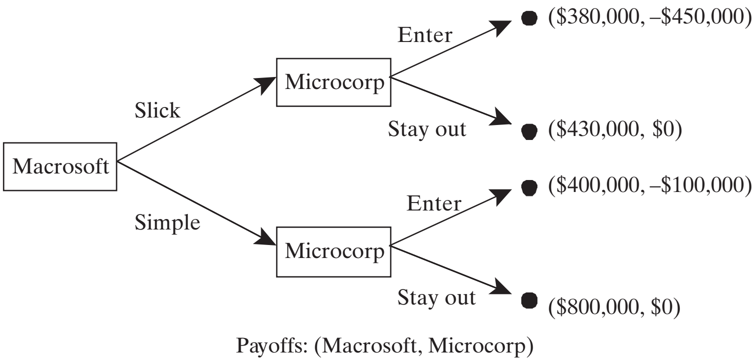

FIGURE 13.20 Game tree for problem 13.9.

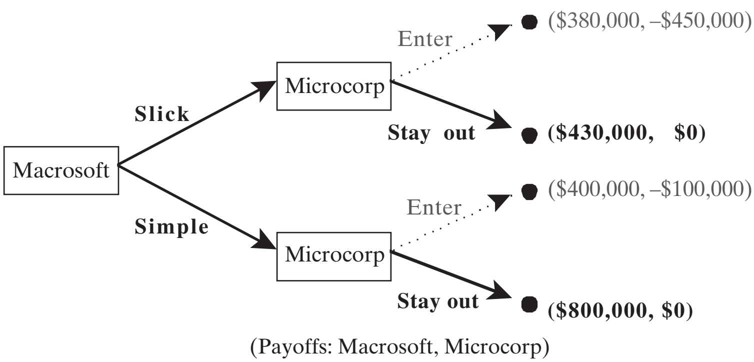

FIGURE 13.21 Solution to problem 13.9 using backward induction.

b. Using the backward induction solution concept, from each terminal node move to Microcorp's two decision nodes. Each of these decision nodes is basic. If Macrosoft chooses a "slick" campaign, the optimal move for Microcorp is to stay out, since the payoff is $0 compared with a payoff of –$450,000 by entering. The "enter" branch should be disregarded in future moves. If Macrosoft chooses a "simple" campaign, again the optimal move for Microcorp is to stay out, since the payoff is $0 compared with a payoff of –$100,000. The "enter" branch should be disregarded in future moves. In the resulting extensive-form game, diagrammed in Figure 13.21, we see that the optimal strategy for Microcorp is (*Stay out*, *Stay out*). The final optimal strategy profile is {*Simple*, (*Stay out*, *Stay out*)}, which yields payoffs of $800,000 for Macrosoft and $0 for Microcorp.

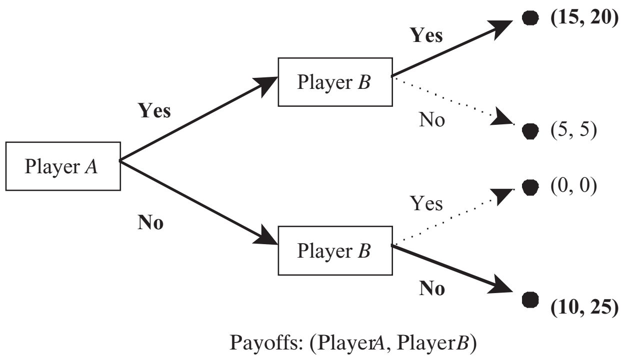

FIGURE 13.22 Solution to problem 13.10 using backward induction.

Problem 13.10. Consider again the multistage game in Figure 13.12. Use the backward induction solution concept to determine the optimal strategy profile for this game. Illustrate your answer.

Solution. Using the backward induction solution concept, from each terminal node move to the two Microcorp decision nodes. Each of these decision nodes is basic. If player *A* says "yes," the optimal move for player *B* is to say "yes," since the payoff is $20 compared with $5 by saying "no." Thus, the "no" branch should be disregarded in future moves. If player *A* says "no," the optimal move for player *B* is to say "no" since the payoff is $25 compared with $0 by saying "yes." The "yes" branch should be disregarded in future moves. In the resulting extensive-form game, diagrammed in Figure 13.22, we see that the optimal strategy for player *B* is (*Yes*, *No*). The final optimal strategy profile is {*Yes*, (*Yes*, *No*)}, which yields payoffs of 15 for player *A* and 20 for player *B*. The student is encouraged to compare this result with the earlier discussion of the selection of Nash equilibria with credible threats.

BARGAINING

In Chapter 8, perfectly competitive markets were characterized by large numbers of buyers and sellers. Firms in perfectly competitive industries were described as "price takers" because of their inability in influence the market price through individual production decisions. Consumers in such markets may similarly be described as price takers because they are individually incapable of extracting discounts or better terms from sellers. Since neither the buyer nor seller has "market power," the theoretical ability to "haggle" over the terms of the sale, or product content, is nonexistent. In the case of a monopolist selling to many small buyers, which was also discussed in Chapter 8, it was assumed that firms set the selling price of the product, and buyers, having no place else to go, accept that price without question. Even when of a neither the buyer or the seller may be thought of as a "price taker," such as the case monopsonist selling to an oligopolist, economists have had little to say about the possibility of negotiating, or "bargaining" over the contract terms.

Yet, bargaining is a fact of life. Whether bargaining with the boss for an increase in wages and benefits or haggling over the price of a new car, such interactions between buyer and seller are commonplace. In many instances, contract negotiations between producer and supplier, contractor and subcontractor, wholesaler and distributor, retailer and wholesaler, and so on, are the norm, rather than the exception. As an exercise, the reader is asked to consider why market power and the ability to bargain with product suppliers allow large retail outlets, such as Home Depot, Sports Authority, or Costco, to offer prices that are generally lower than those featured at the local hardware store, sporting goods store, or other retailer. Even in markets characterized by many buyers and sellers, it is often possible to find "pockets" of local monopoly or monopsony power that permits limited bargaining over contract terms to take place. Game theory is a useful tool for analyzing and understanding the dynamics of the bargaining process.

BARGAINING WITHOUT IMPATIENCE

We will begin our discussion of the bargaining process by considering the following scenario. Suppose that Andrew wishes to purchase an annual service contract from Adam. It is known by both parties that Andrew is willing to pay up to $100 for the service contract and that Adam will not accept any offer below $50. The maximum price that Andrew is willing to pay is called the *buyer's reservation price* and the minimum price that Adam is willing to accept is called the *seller's reservation price*. If Andrew and Adam can come to an agreement, the gain to both will add up to the difference between the buyer's and the seller's reservation prices, which in this case is $50.

Negotiations between Andrew and Adam may be modeled as the extensive-form game illustrated in Figure 13.23. We will assume for simplicity that negotiations involve only two offers and that Andrew makes the first offer, which is denoted as P_1. This is indicated as the first branch of the decision tree. After Andrew has made the offer, Adam can either accept or reject it. If Adam accepts the offer, the bargaining process is completed and the payoffs for Andrew and Adam are $(100 - P_1, P_1 - 50)$, respectively. For example, if Adam accepts Andrew's offer of, say, $80, then Andrew's gain from trade is $20 and Adam's gain from trade is $30, which sum to the difference between the respective parties' reservation prices. If Adam

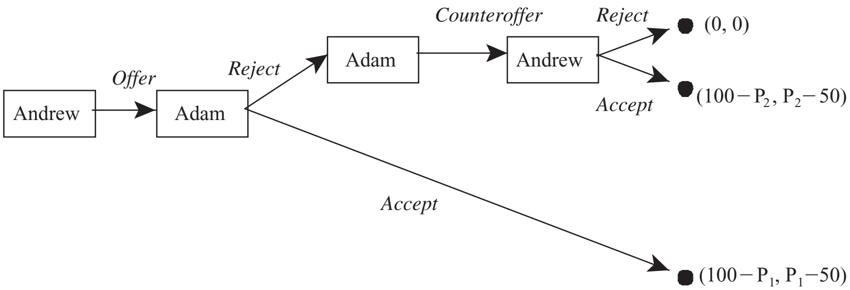

Payoffs: (Andrew, Adam)

FIGURE 13.23 Bargaining without impatience.

rejects Andrew's offer, Adam can come back with a counteroffer, which is denoted as P_2. If Andrew accepts Adam's counteroffer, the payoffs to Andrew and Adam are $(100 - P_2, P_2 - 50)$, respectively. If, on the other hand, Andrew rejects Adam's counteroffer, this game comes to an end and no agreement is reached, in which case the payoffs are (0, 0).

Earlier we discussed the procedure of backward induction for finding solution values to multistage games with multiple equilibria. Applying this approach to the present bargaining game, it is easy to see that as long as Adam's counteroffer is not greater than \$100, Andrew will accept. The reason for this is that Andrew cannot do any better than to accept an offer that does not exceed \$100. Moving up the game tree to another node, it is equally apparent that Adam will reject any offer by Andrew that is less than \$100. Moreover, accepting the offer ignores the fact that Adam has the ability to make a more advantageous (to him) counteroffer in the next round of negotiations. What all this means is that no matter what Andrew's initial offer was, he will end up paying Adam \$100. In other words, as long as Adam has the ability to make a counteroffer, Adam will never accept Andrew's offer as final! Thus, in the two rounds of negotiation in this game, since Adam has the last move, then Adam "holds all the cards." The ability of Adam to dictate the final terms of the negotiations is referred to as the *last-mover's advantage*. Andrew might just as well save his breath and offer Adam \$100 at the outset of the bargaining process.

As the scenario illustrates, the final outcome of this class of bargaining processes depends crucially on who makes the first offer, and on the number of rounds of offers. The reader can verify, for instance, that if Andrew makes the first offer, and there are an odd number of rounds of negotiations, Andrew has the last-mover's advantage, in which case Andrew will be able to extract the entire surplus of \$50. If such is the case, it will be in both parties' best interest for Adam to accept Andrew's initial offer of \$50, thereby saving both individuals the time, effort, and aggravation of an extended bargaining process. Similarly, if Adam has the first move and there are an even number of rounds of negotiations, it will in both parties' interest for Andrew to accept Adam's initial offer $100. In this case, Adam will extract the entire surplus of $50.

BARGAINING WITH SYMMETRIC IMPATIENCE

If negotiations of the type just described were that simple, bargaining would never take place. Of course, bargaining is a fact of life, so something must be missing. In this section we will make the underlying conditions of the bargaining process somewhat more realistic by assuming that there are multiple rounds of offers and counteroffers and costs associated with not immediately reaching an agreement. In the terminology of capital budgeting, this section will introduce the time value of money by discounting to the present future payoffs from negotiations.

In the example of bargaining without impatience, it was assumed that there were only two rounds of bargaining. In fact, the bargaining process is likely to involve multiple rounds of offer and counteroffer lasting days, weeks, or months. Failure to reach an agreement immediately may impose considerable costs on the bargainers. Consider, for example, the rather large opportunity costs incurred by a person who discovers that his or her car has been stolen. It is Saturday and the person needs to be able to drive to work on Monday. Although the stolen car was old, and the person was planning to buy another car anyway, the theft has the introduced a higher than usual level of anxiety into the situation. Failure to quickly come to terms on the purchase price of a replacement car may result not only in high psychological opportunity costs but in lost income, as well.

In this scenario, the buyer can take one of two possible approaches in negotiations with the used-car salesman. On the one hand, the buyer can withhold from the seller the details of his or her ill fortune and negotiate with a "cool head." Alternatively, the buyer may be unable, or unwilling, to withhold knowledge of the theft, preferring to attempt to garner understanding and sympathy. As we will soon see, sympathy in the bargaining process is not without cost: when one person's gain is another's loss, a buyer seeking sympathy will be better off visiting a psychiatrist, not a used-car salesman. To see this, let us consider the situation in which the buyer and the seller enter into negotiations without any knowledge of the opportunity costs that may be imposed on the other because of a failure to immediately reach an agreement. This situation is equivalent to the situation of the buyer who negotiates with the used-car salesperson with a "cool head."

Suppose, once again, that Andrew wishes to purchase an annual service contract from Adam, that Andrew is willing to pay up to $100 for the service contract and that Adam will not accept any offer below $50. Instead of only two negotiating rounds, however, suppose that there are 50 offer–

TABLE 13.6 Nash Equilibrium with Symmetric Impatience

Round	Offer maker	Offer price	Adam's surplus	Andrew's surplus
50	Seller	$100.00	$50.00	$0.00
49	Buyer	$97.50	$47.50	$2.50
48	Seller	$97.62	$47.62	$2.38
47	Buyer	$95.24	$45.24	$4.76
46	Seller	$95.48	$45.48	$4.52
⋮	⋮	⋮	⋮	⋮
5	Buyer	$76.78	$26.78	$23.22
4	Seller	$77.94	$27.94	$22.06
3	Buyer	$76.55	$26.55	$23.46
2	Seller	$77.72	$27.72	$22.28
1	Buyer	$76.33	$26.33	$23.67

counteroffer rounds. Since neither Andrew nor Adam knows anything about the other's personal circumstances, let us further assume that any delay in reaching an agreement reduces the gains from trade to both by 5% per round. This assumption is equivalent to assuming that both players have *symmetric patience*. We will assume that both players are aware of the cost imposed on the other by failing to come to an agreement immediately. With 50 rounds of negotiations, it is impractical to illustrate the bargaining process as an extensive-form game. Nevertheless, it is still possible to use backward induction to determine Andrew's and Adam's negotiating strategies. Consider the information summarized in Table 13.6.

We know that since Andrew makes the first offer and there are an even number of negotiating rounds, Adam has the last-mover's advantage. Thus, if negotiations drag on to the 50th round, Adam will sell the service contract for $100 and extract the entire surplus of $50. Andrew, of course, knows this. Andrew also knows that Adam will be indifferent between receiving $100 in the 50th round and receiving the entire surplus of $50, or receiving $97.50 in the 49th round because delays in reaching an agreement reduce Adam's gain by 5% per round. Thus, Adam will accept any offer from Andrew of $97.50 or more in the 49th round, which results in a surplus of $47.50, and reject any offer that is less than that. In capital budgeting terminology, the time value of $97.50 in the 49th round for Adam is the same as the time value of $100 in the 50th round. But this is not the end of the game.

Adam also knows that delays in reaching an agreement will reduce Andrew's gain from trade by 5% per round. Thus, Andrew is indifferent between receiving a surplus of $2.50 in the 49th round or receiving 5% less ($2.38) in the 48th round. Thus, Adam should offer to sell the service contract for $97.62 in the 48th round, thereby receiving a surplus of $47.62.

Once again, Andrew knows that Adam is indifferent between a price of \$97.62 in the 48th round and \$95.24 in the 47th round, which reduces Adam's surplus by 5%, to \$45.24. Andrew's surplus, on the other hand, will increase to \$4.76. Continuing in the same manner, the reader can verify through the use of backward induction that Andrew's best offer in the first round is \$76.33, which Adam should accept. Adam's and Andrew's gains from trade are \$26.33 and \$23.67, respectively. The reader might suspect that if this process is continued, eventually Andrew and Adam will evenly divide the surplus; but as long as Adam moves last, he will enjoy an advantage, however slight, over Andrew.

BARGAINING WITH ASYMMETRIC IMPATIENCE

Suppose that instead of maintaining an "even keel" the buyer reveals to the used-car salesman the importance of quickly replacing the stolen car. The used-car salesman will immediately recognize the higher opportunity cost to the buyer from delaying a final agreement. To demonstrate the impact that his knowledge has on the bargaining process, consider again the negotiations between Andrews and Adam. We will continue to assume that there are 50 rounds of negotiations, but that the opportunity cost to Andrew from delaying an agreement reduces the gain from trade by 10% per round, while the opportunity cost to Adam continues to be 5% per round. Proceeding as before, the information in Table 13.7 summarizes the gains from trade to both Andrew and Adam that result from bargaining in the presence of asymmetric impatience (i.e., different opportunity costs for each player).

Utilizing backward induction, the reader will readily verify from Table 13.7 that Andrew's best first round offer is \$83.10. This will result in a

TABLE 13.7 Nash Equilibrium with Asymmetric Impatience

Round	Offer maker	Offer price	Adam's surplus	Andrew's surplus
50	Seller	\$100.00	\$50.00	\$0.00
49	Buyer	\$97.50	\$47.50	\$2.50
48	Seller	\$97.75	\$47.75	\$2.25
47	Buyer	\$95.36	\$45.36	\$4.64
46	Seller	\$95.83	\$45.83	\$4.17
⋮	⋮	⋮	⋮	⋮
5	Buyer			
4	Seller	\$84.91	\$34.91	\$15.09
3	Buyer	\$83.16	\$33.16	\$16.84
2	Seller	\$84.84	\$34.84	\$15.16
1	Buyer	\$83.10	\$33.10	\$16.90

surplus to Adam of $33.10, which is nearly twice the gain from trade enjoyed by Andrew. The results presented in Table 13.7 demonstrate that the negotiating party with the lowest opportunity cost has the clearest advantage in the negotiating process. Within the context of the stolen car example, clearly patience and secrecy are virtues. By "crying the blues" to the used-car salesman, the buyer placed himself or herself at a bargaining disadvantage. Unless the buyer is dealing with a paragon of rectitude and virtue, looking for sympathy from a rival during negotiations will clearly result in a disadvantageous division of the gains from trade.

If effect, impatience has been used as the discount rate for finding the present value of gains from trade in bargaining. The greater the players' impatience (the higher the discount rate), the less advantageous will be the gains from bargaining. Ariel Rubinstein (1982) has demonstrated that in this type of two-player bargaining game there exists a unique subgame perfect equilibrium. Assume that, player A and player B are bargaining over the division of a surplus and player B makes the first offer. Assume further that there is no limit to the number of rounds of offer and counteroffer and that both players accept offers when indifferent between accepting and rejecting the offer. Denote player A's discount rate as δ_A and Player B's discount rate as δ_B. A bargaining game has a unique subgame perfect equilibrium if in the first round player B offers player A

$$\omega_A = \frac{\theta_A(1-\theta_B)}{1-\theta_A\theta_B} \tag{13.15}$$

as a share of the surplus, where $\theta_A = 1 - \delta_A$ and $\theta_B = 1 - \delta_B$. Player B's share of the surplus is

$$\omega_B = \frac{1-\theta_A}{1-\theta_A\theta_B} \tag{13.16}$$

Problem 13.11. Andrew and Adam are bargaining over a surplus of \$50. Assume that there is no limit to the number of rounds of offer and counteroffer, and that the discount rates for both players are $\delta_A = 0.05$ and $\delta_B = 0.05$.

a. For a subgame perfect equilibrium to exist, what portion of the surplus should Adam offer Andrew in the first round? What portion of the surplus should Adam keep for himself?
b. Suppose that Adam's discount rate is $\delta_A = 0.05$ and Andrew's discount rate is $\delta_B = 0.10$. What portion of the surplus should Adam offer Andrew in the first round and what portion should he keep for himself?

Solution

a. $\theta_A = 1 - \delta_A = 0.95$; $\theta_B = 1 - \delta_B = 0.95$. Substituting these values into expression (13.15) we obtain

$$\omega_A = \frac{\omega_A(1-\theta_B)}{1-\theta_A\theta_B} = \frac{(0.95)(1-0.95)}{1-0.9025} = \frac{0.0475}{0.0975} = 0.4872$$

The amount of the surplus that Adam should offer Andrew is

$$\omega_A(\$50) = 0.4872(\$50) = \$24.36$$

From equation (3.16) we obtain,

$$\omega_B = \frac{1-\theta_A}{1-\theta_A\theta_B} = \frac{(1-0.95)}{1-(0.95)(0.95)} = \frac{0.05}{0.0975} = 0.5128$$

The share of the surplus that Adam should keep is, therefore,

$$\omega_B(\$50) = 0.5128(\$50) = \$25.64$$

Of course, the sum of the shared surpluses is \$50. The student should note that as the last mover, Adam earns slightly more of the surplus than Andrew. The student is urged to compare these results with those found in Table 13.6. For the same discount rates and 50 negotiating rounds Adam received \$26.33 and Andrew received \$23.67.

b. $\theta_A = 1 - \delta_A = 0.90$; $\theta_B = 1 - \delta_B = 0.95$. Substituting these values into expression (13.15) we obtain

$$\omega_A = \frac{\omega_A(1-\theta_B)}{1-\theta_A\theta_B} = \frac{(0.90)(1-0.95)}{1-(0.90)(0.95)} = \frac{0.045}{0.145} = 0.3103$$

The amount of the surplus that Adam should offer Andrew is

$$0.3103(\$50) = \$15.52$$

The share of the surplus that Adam should keep for himself can be found by first substituting the information provided into expression (13.16), or

$$\omega_B = \frac{1-\theta_A}{1-\theta_A\theta_B} = \frac{1-0.90}{1-(0.90)(0.95)} = \frac{0.10}{0.145} = 0.6897$$

The share of the surplus that Adam should keep is, therefore,

$$\omega_B = 0.6897(\$50) = \$34.48$$

Once again, the sum of the shared surpluses is \$50. The student should note that, as the last mover, Adam retains more of the surplus then Andrew.

CHAPTER REVIEW

Game theory is the study of the strategic behavior involving the interaction of two or more individuals, teams, or firms, usually referred to as players. Two game theoretic scenarios were examined in this chapter: *simultaneous-move* and *multistage* games. In simultaneous-move games the players effectively move at the same time. A *normal-form game* summarizes the players' possible strategies and payoffs from alternative strategies in a simultaneous-move game.

Simultaneous-move games may be either *noncooperative* or *cooperative*. In contrast to noncooperative games, players of cooperative games engage in collusive behavior (i.e., they conspire to "rig" the final outcome). A *Nash equilibrium*, which is a solution to a problem in game theory, occurs when the players' payoffs cannot be improved by changing strategies.

Simultaneous-move games may be either *one-shot* or *repeated games*. One-shot games are played only once. Repeated games are played more than once. *Infinitely repeated games* are played over and over again without end. *Finitely repeated games* are played a limited number of times. Finitely repeated games can have certain or uncertain ends.

Analytically, there is little difference between infinitely repeated games and finitely repeated games with an uncertain end. With infinitely repeated games and finitely repeated games with an uncertain end, collusive agreements between and among the players are possible, although not necessarily stable. The solution to a finitely repeated game with a certain end collapses into a series of noncooperative, one-shot games. Collusive agreements between and among players of finitely repeated games are inherently unstable.

Multistage games differ from simultaneous-move games in that the play is sequential. An *extensive-form game* summarizes the players, the information available to each player at each stage, the order of the moves, and the payoffs from alternative strategies of a multistage game. A Nash equilibrium in a multistage game is a *subgame perfect equilibrium*. In this case, no player is able to improve on his or her payoff at any stage of the game by switching strategies. *Backward induction* is a solution concept proposed by Bierman and Fernandez for finding optimal solutions to multistage games involving multiple Nash equilibria.

Bargaining is a version of a multistage game. In *bargaining without impatience*, players assume that negotiators incur no costs by not immediately reaching an agreement. To use capital budgeting terminology, the discount rate for finding the present value of future payoffs is zero. The final outcome of this class of bargaining processes depends crucially on who makes the first offer, and on the number of rounds of offers. Players who make the final offer in negotiations have *last-mover's advantage* and are able to extract the entire gains from trade.

In *bargaining with impatience*, players assume that negotiators do incur costs when agreements are not immediately reached. Impatience may be symmetric or asymmetric. In *symmetric impatience*, players assume that the costs to the negotiators from not immediately reaching an agreement are identical. In this case, the discount rate for finding the present value of a future settlement is the same for both players. In *asymmetric impatience*, players assume that this discount rate is different for each player. Players with greater patience (lower discount rate) have the advantage in the negotiating process. In both cases, the player with the final move will receive most of the gains from trade. The extent of this gain will depend on the relative degrees of impatience of the negotiators. The greater a negotiator's patience, the larger will be that player's gain from trade.

KEY TERMS AND CONCEPTS

Backward induction A methodology for finding optimal solutions to multistage games involving multiple Nash equilibria.

Cheating rule for infinitely repeated games For a two-person, cooperative, non-zero-sum, simultaneous-move, infinitely repeated game, where future payoffs and interest rates are assumed to be unchanged, a collusive agreement will be unstable if $(\pi_H - \pi_N)/(\pi_C + \pi_N - \pi_H) < i$, where π_H is the one-period payoff from adhering to the agreement, π_C is the first-period payoff from violating a collusive agreement, π_N is the per-period payoff in the absence of a collusive agreement, and i is the market interest rate. For a two-person, cooperative, non-zero-sum, simultaneous-move, finitely repeated game with an uncertain end, a collusive agreement will be unstable if $(\pi_H - \pi_N - \theta\pi_C)/(\pi_C + \pi_N - \pi_H) < i$, where $0 < \theta < 1$ is the probability that the game will end after each play.

Cooperative game A game in which the players engage in collusive behavior to "rig" the final outcome.

Credible threat A threat is credible only if it is in a player's best interest to follow through with the threat when the situation presents itself.

Decision node A point in a multistage game at which a player must decide upon a strategy.

End-of-period problem For finitely repeated games with a certain end, each period effectively becomes the final period, in which case the game reduces to a series of noncooperative one-shot games.

Finitely repeated game A game that is repeated a limited number of times.

Focal-point equilibrium When a single solution to a problem involving multiple Nash equilibria "stands out" because the players share a common "understanding" of the problem, focal-point equilibrium has been achieved.

Game theory The study of how rivals make decisions in situations involving strategic interaction (i.e., move and countermove) to achieve an optimal outcome.

Infinitely repeated game A game that is played over and over again without end.

Maximin strategy A strategy that selects the largest payoff from among the worst possible payoffs.

Nash equilibrium Is reached when each player adopts a strategy it believes to be the best response to the other players' strategy. When a game is in a Nash equilibrium, the players' payoffs cannot be improved by changing strategies.

Noncooperative game A game in which the players do not engage in collusive behavior. In other words, the players do not conspire to "rig" the final outcome.

Non–strictly dominant strategy When a strictly dominant strategy does not exist for either player and the optimal strategy for either player depends on what each player believes to be the strategy of the other players, the result is a non–strictly dominant strategy.

Normal-form game A game in which each player is aware of the strategy of every other player as well as the possible payoffs resulting from alternative combinations of strategies.

One-shot game A game that is played only once.

Repeated game A game that is played more than once.

Risk avoider An individual who prefers a certain payoff to a risky prospect with the same expected value. A risk avoider prefers predictable outcomes to probabilistic expectations.

Risk taker An individual who prefers a risky situations in which the expected value of a payoff is preferred to its certainty equivalent.

Sequential-move game A game in which the players move in turn.

Simultaneous-move game A game in which the players move at the same time.

Strategic behavior The actions of those who recognize that the behavior of an individual or group affect, and are affected by, the actions of other individuals or groups.

Strategic form of a game A summary of the payoffs to each player arising from every possible strategy profile.

Strategy A game plan or a decision rule that indicates what action a player will take when confronted with the need to make a decision.

Strictly dominant strategy A strategy that results in the largest payoff regardless of the strategy adopted by another player.

Strictly dominant strategy equilibrium A Nash equilibrium that results when all players have a strictly dominant strategy.

Subgame perfect equilibrium A strategy profile in a multistage game that is a Nash equilibrium and allows no player to improve on his or her payoff by switching strategies at any stage of the game.

Trigger strategy A game plan that is adopted by one player in response to unanticipated moves by the other player. A trigger strategy will continue to be used until the other player initiates yet another unanticipated move.

Weakly dominant strategy A strategy that results in a payoff that is no lower than any other payoff regardless of the strategy adopted by the other players.

Zero-sum game A game in which one player's gain is exactly the other player's loss.

CHAPTER QUESTIONS

13.1 In the game "rock–scissors–paper" two players in unison show a fist (rock), two fingers (scissors), or an open hand (paper). The winner of each round is determined by what hand signals the players shows. If one player shows a fist, while another shows two fingers, the first player wins because "rocks break scissors." If, on the other hand, the second player shows an open hand, then that player wins because "paper covers rock." Finally, if one player shows two fingers and the other player shows an open hand, then the second player wins because "scissors cut paper," and so on. An alternative way to play this game is to isolate the players in separate rooms, prohibiting communication between them. A third individual, the referee, goes to each room and asks the player to reveal his or her hand. After inspecting the hand of each player, the referee declares a winner. Both versions of this game may be called simultaneous-move games. Do you agree? If not, then why not?

13.2 A subgame perfect equilibrium is impossible in a game with multiple Nash equilibria. Do you agree or disagree? Explain.

13.3 Explain the difference between moves and strategies.

13.4 Suppose you and a group of your coworkers have decided to have lunch at a Japanese restaurant. It has been decided in advance that the lunch bill will be divided equally. Each person in the group is concerned about his or her share of the bill. Without explicitly agreeing to do so, each person will order from among the least expensive items on the menu. Comment.

13.5 Explain the difference between a strictly dominant strategy and a non–strictly dominant strategy equilibrium. Under what circumstances will a strictly dominant strategy lead to a non–strictly dominant strategy equilibrium?

13.6 The existence of a Nash equilibrium confirms Adam Smith's famous metaphor of the invisible hand. Do you agree with this statement? If not, then why not?

13.7 Explain the difference between a strictly dominant strategy and an iterated strictly dominant strategy.

13.8 In a two-player, simultaneous-move game with a strictly dominant strategy equilibrium, at least one of the players will adopt a secure strategy. Do you agree? If not, why not?

13.9 Explain the difference between a strictly dominant strategy and a weakly dominated strategy.

13.10 It is not possible to have multiple Nash equilibria in the presence of a subgame perfect equilibrium. Do you agree with this statement? If not, why not?

13.11 In a two-player, one-shot game, if one player has a dominant strategy, the second player will never adopt a maximin strategy. Do you agree? Explain.

13.12 Explain the difference between a strictly dominant strategy and a weakly dominant strategy.

13.13 If neither player in a noncooperative, one-shot game has a strictly dominant strategy, or if the strategy results in a weakly dominant strategy equilibrium, explain how the concept of a focal-point equilibrium might lead to a solution in game theory.

13.14 Under what conditions will trigger strategies be successful in maintaining the integrity of a collusive agreement?

13.15 The existence of a trigger strategy that punishes a violator of a cooperative agreement will eliminate the problem of cheating in a simultaneous-move, infinitely repeated game. Do you agree? Explain.

CHAPTER EXERCISES

13.1 Argon Airlines and Boron Airways are two equal-sized commercial air carriers that compete for passengers along the lucrative Boston–Albany–Buffalo route. Both firms are considering offering discount air fares during the traditionally slow month of February. The payoff matrix ($ millions) for this game is illustrated in Figure E13.1.

a. Does either firm have a strictly dominant strategy?

b. What is the Nash equilibrium for this game?

		Boron	
		Discount	No discount
Argon	Discount	(2, 3)	(7.5, 1)
	No discount	(1.5, 6)	(3, 2)

Payoffs: (Argon, Boron)

FIGURE E13.1 Payoff matrix for chapter excercise 13.1.

		Firm B: Don't cheat	Firm B: Cheat
Firm A	Don't cheat	(10, 10)	(–5, 20)
	Cheat	(20, –5)	(5, 5)

Payoffs: (Firm A, Firm B)

FIGURE E13.2 Payoff matrix for chapter excercise 13.2.

		Firm *B*: High price	Firm *B*: Low price
Firm *A*	High price	(10, 10)	(−5, 20)
	Low price	(20, −5)	(5, 5)

Payoffs: (Firm *A*, Firm *B*)

FIGURE E13.3 Payoff matrix for chapter exercise 13.3.

13.2 Consider the normal-form, one-shot game shown in Figure E13.2, involving two firms that have entered into a collusive agreement. The payoffs in the parentheses are in millions of dollars. Having entered into the agreement, both firms must decide whether to remain faithful to the agreement (*Don't cheat*) or to violate the agreement (*Cheat*).

a. Does either firm have a dominant strategy?
b. If both firms follow a maximin strategy, what is the strategy profile for this game? Is this strategy profile a Nash equilibrium?
c. Suppose that firm *B* were to cheat on the agreement. What would firm *A* do?
d. How might your answer be different if this were an infinitely repeated game? What factors not presented here must be considered?

13.3 Consider the two-person, noncooperative, non-zero-sum, simultaneous-move, one-shot pricing game shown in Figure E13.3. The numbers in the parentheses are in millions of dollars.

a. Does either player have a strictly dominant strategy? If so, what is the dominant strategy equilibrium? Is this a Nash equilibrium?
b. If this game were repeated an infinite number of times, would either player change strategies?

13.4 Consider Figure E13.4, a normal-form game describing the interaction between labor and management. The payoff matrix reflects management's desire for labor to work hard and labor's desire to take it easy. Management has two options. Managers can either secretly monitor

		Labor	
		Work hard	Goof off
Management	Observe	(–1, 1)	(1, –1)
	Don't observe	(1, –1)	(–1, 1)

Payoffs: (Management, Labor)

FIGURE E13.4 Payoff matrix for chapter exercise 13.4.

worker performance or they can trust employees to work hard on their own. Labor also has two options: to work or to goof off. The payoff matrix may be read as follows. If management secretly observes labor, management "loses" because of the time spent monitoring workers already working. Presumably labor "wins" because hard work will be rewarded with extra pay, benefits, and so on. In this case, the strategy profile {*Observe*, *Work hard*} has a payoff of (–1, 1). Note that the payoff is the same for the strategy profile {*Don't observe*, *Goof off*} because management continues to employ a "goldbrick" while the workers gain leisure time. When the strategy profile is {*Don't observe*, *Work hard*}, management wins because it did not incur the expense of monitoring the performance of a hard-working employee, while the worker loses because he or she could have goofed off without penalty. Finally, the strategy profile {*Observe*, *Goof off*} has a payoff of (1, –1) because management discovers, and presumable fires, the shirker.

a. Does either player in this game have a dominant strategy? Explain.
b. Does this game have a Nash equilibrium? If not, then why not?
c. What would the absence of a Nash equilibrium suggest for optimal management–employee relations in the present context?

13.5 Consider the normal-form, simultaneous-move, one-shot game shown in Figure E13.5. Suppose that an industry consists of two firms, Magna Company and Summa Corporation. The firms produce identical products. Magna and Summa are trying to decide whether to expand (*Expand*) or not to expand (*None*) production capacity for the next operating period. Assume that each firm produces at full capacity. The trade-off facing each firm is that expansion will result in a larger market share, but increased output will put downward pressure on price. Expected profits are summarized in Figure E13.5, when the numbers in the parentheses are in millions of dollars. The first payoff is Magna's.

a. Does either firm have a dominant strategy?
b. What is the Nash equilibrium for this game?

13.6 Suppose that in Exercise 13.5 Magna and Summa have three options: no expansion (*None*), moderate expansion (*Moderate*), and extensive expansion (*Extensive*). Expected profits are summarized in the normal-

		Summa: None	Summa: Expand
Magna	None	(25, 25)	(15, 30)
	Expand	(30, 15)	(20, 20)

Payoffs: (Magna, Summa)

FIGURE E13.5 Payoff matrix for chapter exercise 13.5.

		Summa: None	Summa: Moderate	Summa: Extensive
Magna	None	(25, 25)	(15, 30)	(10, 25)
	Moderate	(30, 15)	(20, 20)	(8, 13)
	Extensive	(25, 10)	(12, 8)	(0, 0)

Payoffs: (Magna, Summa)

FIGURE E13.6 Payoff matrix for chapter exercise 13.6.

form game shown in Figure E13.6. What is the Nash equilibrium for this game?

13.7 Suppose that the simultaneous move game in Exercise 13.6 was modeled as a sequential-move game, with Magna moving first.

a. Illustrate the extensive form of this game.
b. What are the subgames for this game?
c. What is the Nash equilibrium for each subgame?
d. Use backward induction to find the subgame perfect equilibrium.

13.8 Consider the simultaneous-move, one-shot game shown in Figure E13.8.

a. If player B believes that player A will play strategy A, what strategy should player B adopt?
b. If player B believes that player A will play strategy B, what strategy should player B adopt?
c. Does this game have a Nash equilibrium?
d. Does this game have a unique solution?

13.9 Tom Teetotaler and Brandy Merrybuck are tobacconists specializing in three brands of pipe-weed: Barnacle Bottom, Old Toby, and Southern Star. Both Teetotaler and Merrybuck are trying to decide what brands to carry in their shops, Red Pony and Blue Dragon, respectively. Expected earnings in this simultaneous, one-shot game are summarized in the normal-form game shown in Figure E13.9.

		Player *B*	
		Strategy A	Strategy B
Player *A*	Strategy A	(20, 20)	(5, 25)
	Strategy B	(25, 5)	(2, 2)

Payoffs: (Player *A*, Player *B*)

FIGURE E13.8 Payoff matrix for chapter exercise 13.8.

		Blue Dragon		
		Narrow	Medium	Wide
Red Pony	Narrow	(150, 150)	(100, 200)	(50, 250)
	Medium	(200, 100)	(200, 200)	(150, 300)
	Wide	(250, 50)	(300, 150)	(300, 300)

Payoffs: (Red Pony, Blue Dragon)

FIGURE E13.9 Payoff matrix for chapter exercise 13.9.

		Blue Dragon		
		Narrow	Medium	Wide
Red Pony	Narrow	(150, 150)	(200, 250)	(250, 350)
	Medium	(250, 125)	(175, 200)	(270, 245)
	Wide	(350, 250)	(150, 275)	(200, 300)

Payoffs: (Red Pony, Blue Dragon)

FIGURE E13.10 Payoff matrix for chapter exercise 13.10.

a. What is the solution to this game?
b. Is this solution a Nash equilibrium?

13.10 Suppose that the payoffs for the game in Exercise 13.9 were as shown in Figure E13.10.

a. Does either firm have a strictly dominant strategy?
b. Is the solution for this game a Nash equilibrium?

13.11 Suppose that the simultaneous-move game in Exercise 13.10 was modeled as a sequential-move game with Red Pony moving first.

a. Illustrate the extensive form of this game.
b. What are the subgames for this game?

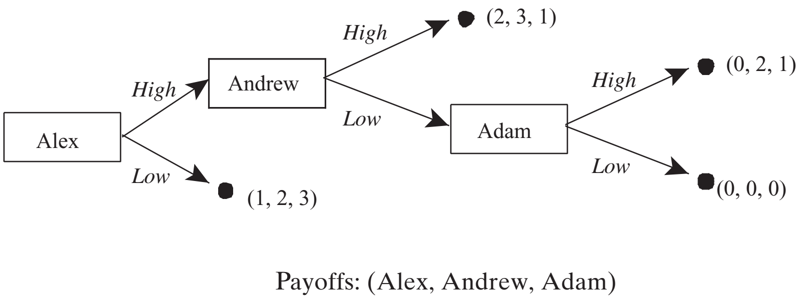

FIGURE E13.12 Game tree for chapter exercise 13.12.

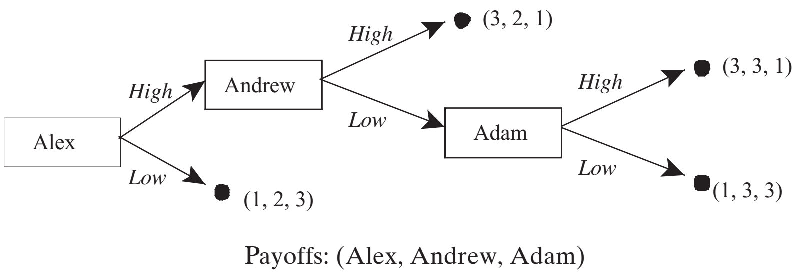

FIGURE E13.13 Game tree for chapter exercise 13.13.

c. What is the Nash equilibrium for each subgame?
d. Use backward induction to find the subgame perfect equilibrium.

13.12 Alex, Andrew, and Adam are playing the multistage game shown in Figure E13.12.

a. What are the subgames for this game?
b. What is the Nash equilibrium for each subgame?
c. Use backward induction to find the subgame perfect equilibrium.

13.13 Suppose that the multistage game for Alex, Andrew, and Adam is as shown in Figure E13.13.

a. What are the subgames for this game?
b. What is the Nash equilibrium for each subgame?
c. Use backward induction to find the subgame perfect equilibrium.

13.14 At the Hemlock Bush Tavern, Jethro (Jellyroll) Bottom announces that he will auction off an envelope containing \$35. Clem and Heathcliff are the only two bidders, and each has \$40. The rules of the auction are as follows:

(1) The bidders take turns. After a bid is made, the next bidder can make either another bid or pass. The opening bid must be \$10.
(2) Succeeding bids must be in \$10 increments.
(3) Bidders cannot bid against themselves.

(4) The bidding comes to an end when either bidder passes, except on the first bid. If the first bidder passes, the second bidder is given the option of accepting the bid.
(5) The highest bidder wins.
(6) All bidders must pay Jethro the amount of their last bid.
(7) Assume that Clem bids first.
  a. Diagram the game tree for this game.
  b. Determine the subgame perfect equilibrium strategies for Clem and Heathcliff using the method of backward induction.
  c. What is the outcome of the auction?

SELECTED READINGS

Axelrod, R. *The Evolution of Cooperation.* New York: Basic Books, 1984.

Baye, M. R., and R. O. Beil. *Managerial Economics and Business Strategy.* Boston: Richard D. Irwin, 1994.

Besanko, D., D. Dranove, and M. Shanley. *Economics of Strategy*, 2nd ed. New York: John Wiley & Sons, 2000.

Bierman, H. S., and L. Fernandez. *Game Theory with Economic Applications*, 2nd ed. New York: Addison-Wesley, 1998.

Brams, S., and M. Kilgour. *Game Theory and National Security.* Oxford: Basil Blackwell, 1988.

Cournot, A. *Researches into the Mathematical Principles of the Theory of Wealth*, translated Nathaniel T. Bacon. New York: Macmillan, 1897.

Davis, O., and A. Whinston. "Externalities, Welfare, and the Theory of Games." *Journal of Political Economy*, 70 (June 1962), pp. 241–262.

de Fraja, G., and F. Delbono. "Game Theoretic Models of Mixed Oligopoly." *Economic Surveys*, 4 (1990), pp. 1–17.

Dixit, A., and B. Nalebuff. *Thinking Strategically.* New York: W. W. Norton, 1991.

Horowitz, I. "On the Effects of Cournot Rivalry between Entrepreneurial and Cooperative Firms." *Journal of Comparative Economics*, 15 (March 1991), pp. 115–211.

Luce, D. R., and H. Raiffa. *Games and Decisions: Introduction and Critical Survey.* New York: John Wiley & Sons, 1957.

Nash, J. "The Bargaining Problem." *Econometrica*, 18 (1950a), pp. 155–162.

———. "Equilibrium Points in *n*-Person Games." *Proceedings of the National Academy of Sciences, USA*, 36 (1950b), pp. 48–49.

———. "A Simple Three-Person Poker Game" (with Lloyd S. Shapley). *Annals of Mathematics Study*, 24 (1950c).

———. "Noncooperative Games." *Annals of Mathematics*, 51 (1951), pp. 286–295.

———. "A Comparison of Treatments of a Duopoly Situation" (with J. P. Mayberry and M. Shubik). *Econometrica*, 21 (1953a), pp. 141–154.

———. "Two-Person Cooperative Games." Econometrica, 21 (1953b), pp. 405–421.

Nasar, S. *A Beautiful Mind.* New York: Simon & Schuster, 1998.

Poundstone, W. *Prisoners' Dilemma.* New York: Doubleday, 1992.

Rasmusen, E. *Games and Information: An Introduction to Game Theory.* New York: Basil Blackwell, 1989.

Rubinstein, A. "Perfect Equilibrium in a Bargaining Model." *Econometrica*, 61 (1) (1982), pp. 97–109.

Schelling, T. *The Strategy of Conflict.* London: Oxford University Press, 1960.

Schotter, A. *Free Market Economics: A Critical Appraisal.* New York: St. Martin's Press, 1985.

———. Microeconomics: A Modern Approach. New York: Addison-Wesley, 1998.

Selten, R. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory*, 4 (1975), pp. 25–55.

Von Neumann, J., and O. Morgenstern. *Theory of Games and Economic Behavior.* New York: John Wiley & Sons, 1944.

14

Risk and Uncertainty

We have assumed throughout most of this book that the economic decisions were made under conditions of complete certainty. It was assumed that the decisions of both consumers and producers were based on complete and accurate knowledge of consumer, firm, and market conditions. In fact, however, most economic decisions are made with something less than perfect information, and the consequences of these decisions cannot, therefore, be known beforehand with any degree of precision. A manager cannot know, for example, whether the introduction of a new product will be profitable because of the uncertainty of macroeconomic conditions, consumer tastes, reactions by competitors, resource availability, input prices, labor unrest, political instability, and so forth.

In addition to the uncertainty associated with decisions made at any point in time, the uncertainty of outcomes associated with those decisions tends to increase the further we project into the future. An automobile company that plans to introduce a new model within 2 years is more likely to successfully satisfy prevailing consumer tastes in terms of styling and options, and therefore to be better able to capture a significant market share, than a company that takes 5 years to bring a new product to market. After 5 years, consumer tastes could significantly change, reducing the probability of the product's success.

A formal treatment of the decision-making process under conditions of uncertainty is well beyond the scope of this book. Nevertheless, this chapter will introduce some of the more essential elements of decision making in the absence of complete information. We begin with a formal distinction between *risk* and *uncertainty* and move on to a discussion of decision making with uncertain and risky outcomes.

RISK AND UNCERTAINTY

When one is examining the decision-making process under conditions of imperfect information, it is important to distinguish between the closely related concepts of risk and uncertainty. Risky situations involve multiple outcomes (or payoffs), where the probability of each outcome is known or can be estimated. An example of a risky situation is the flipping of a fair coin. The probability that either a head or a tail will result from flipping a fair coin is 50%. Investing in the stock market is another risky situation. While the investor cannot know with certainty the rate of return on the investment, it is possible to estimate an expected rate of return based on a company's past performance.

Definition: Risk involves choices involving multiple possible outcomes in which the probability of each outcome is known or may be estimated.

Uncertainty also involves multiple-outcomes situations. What distinguishes risk from uncertainty, however, is that with uncertainty the probability of each outcome is unknown and cannot be estimated. In many cases, these probabilities cannot be estimated because of the absence of historical evidence about the event. Nevertheless, there is a fine line between decision making under conditions of risk and of uncertainty.

Definition: Uncertainty involves choices involving multiple possible outcomes in which the probability of each outcome is unknown and cannot be estimated.

When one is considering the different ways in which managers deal with uncertain outcomes it is important to distinguish between two types of uncertainty. In situations of *complete ignorance*, the decision maker is unable to make any assumptions about the probabilities of alternative outcomes under different states of nature. In these situations, the decision maker may adopt any of a number of rational criteria to facilitate the decision-making process.

Situations involving partial ignorance, on the other hand, assume that the decision maker is able to assign subjective probabilities to multiple outcomes. Whenever the decision maker is able to use personal knowledge, intuition, and experience to assign subjective probabilities to outcomes, then decision making under uncertainty is effectively transformed into decision making under risk. In the next section, we will examine the most commonly used statistical measures of risk.

Much of the discussion that follows will deal with decision making under risk, uncertainty involving partial ignorance, or uncertainty involving complete ignorance. While the procedures for evaluating outcomes of decisions made under conditions of risk, or uncertainty involving partial ignorance, are identical, the process of evaluating outcomes under conditions of complete ignorance requires alternative approaches to the decision-making process. In spite of these distinctions, we will refer to all situations in which the probability of each outcome is not known and cannot be estimated as conditions of uncertainty. It will be clear from the context of each situation whether this involves risk or uncertainty from partial or complete ignorance.

MEASURING RISK: MEAN AND VARIANCE

MEAN (EXPECTED VALUE)

The most commonly used summary measures of risky, random payoffs are the *mean* and the *variance*. These random payoffs may refer to profits, capital gains, prices, unit sales, and so on. In risky situations, the expected value of these random payoffs is called the mean. The mean is the weighted average of all possible random outcomes, with the weights being the probability of each outcome. For discrete random variables, the expected value may be calculated using Equation (14.1)

$$E(x) = \mu = \sum_{i=1 \to n} x_i p_i \tag{14.1}$$

where x_i is the value of the outcome, p_i is the probability of its occurrence, and $\Sigma_{i=1\to n}\, p_i = 1$.

When the probability of each outcome is the same as the probability of every other outcome, then the expected value is the sum of the outcomes divided by the number of observations. In this case, the expected value of a set of uncertain outcomes may be calculated using Equation (14.2)

$$E(x) = \mu = \left(\frac{1}{n}\right) \sum_{i=1 \to n} x_i \tag{14.2}$$

Definition: The mean is the expected value of a set of random outcomes. The mean is the sum of the products of each outcome and the probability of its occurrence. When the probability of the occurrence of each outcome is the same as the probability of every other outcome, the mean is the sum of the outcomes divided by the number of observations.

Problem 14.1. Suppose that the chief economist of Silver Zephyr Ltd. believes that there is a 40% ($p_1 = 0.4$) probability of a recession in the next operating period and a 60% ($p_2 = 0.6$) probability that a recession will not occur. The COO of Silver Zephyr believes that the firm will earn profits of π_1 = \$100 in the event of a recession and π_2 = \$1,000 otherwise. What are Silver Zephyr's expected profits?

Solution. Silver Zephyr's expected profits are

$$\begin{aligned} E(\pi) &= \sum_{i=1\to 2} \pi_i p_i = \pi_1 p_1 + \pi_2 p_2 \\ &= 0.4(100) + 0.6(1{,}000) \\ &= 40 + 600 = \$640 \end{aligned}$$

Thus, Silver Zephyr's expected profits for the next operating period are \$640.

Problem 14.2. Suppose that Bob mates his brother Nob the following offer. For a payment of \$3.50, Bob will pay Nob the dollar value of any roll, v, of a fair die. For example, for a roll of 1, Bob will pay Nob \$1. For a roll of 6, Bob will pay Nob \$6. How much can Nob expect to earn if he accepts Bob's offer?

Solution. Since the probability of any number between 1 and 6 is 1/6, then Bob's expected payout is

$$\begin{aligned} E(v) &= \left(\frac{1}{n}\right) \sum_{i=1\to n} v_i \\ &= \left(\frac{1}{6}\right)(1+2+3+4+5+6) = \left(\frac{1}{6}\right)(21) = \$3.50 \end{aligned}$$

Since it will cost \$3.50 to play this game, Nob's can expect to earn $E(v)$ – 3.50 = \$0. Whether Nob should accept Bob's offer will depend on Nob's attitude toward risk. An individual's attitude toward risk will be discussed in the paragraphs to follow.

VARIANCE

The strength of the mean is its simplicity. In a single number, the mean (expected value) summarizes important information about the most likely outcome of a set of random payoffs. Unfortunately, this strength hides other important information that is valuable to the decision maker. For example, suppose that an individual is offered the following fair wager. If the individual flips a coin and it comes up heads, then the individual wins \$10. On the other hand, if the coin comes up tails, then the individual loses \$10. The reader should verify that the expected value of the wager is \$0. Suppose, on the other hand the payoffs were \$1,000 and –\$1,000 for a head and tail, respectively. Once again, the reader will verify that the expected value of the wager is \$0. While the expected values of the two wagers are the same, clearly the wagers themselves are different. While the potential payoff is much greater than in the second scenario, so too is the potential loss. While the individual may be prepared to accept the first bet, that person may not be willing to accept the second because the possibility of such a large loss may be unacceptable. For this individual, the second wager may simply be too risky.

The second wager is riskier because the spread, or dispersion, of the possible payoffs is greater. Each has the same expected value, but the swing between a gain and a loss is considerably greater. It is this dispersion in the possible payoffs that is the distinguishing characteristic of risk. The most commonly used measure of the dispersion of a set of random outcomes is the variance. The variance is the weighed average of the squared deviations of all possible random outcomes from its mean, with the weights being the probability of each outcome. The variance of a set of random payoffs may be calculated by using Equation (14.3).

$$E\left[(x-\mu)^2\right] = \sigma^2 = \sum_{i=1\to n} (x_i - \mu)^2 p_i \tag{14.3}$$

When the probability of each outcome is the same, then the variance is simply the sum of the squared deviations divided by the number of outcomes.

$$E\left[(x-\mu)^2\right] = \sigma^2 = \left(\frac{1}{n}\right)\sum_{i=1\to n} (x_i - \mu)^2 \tag{14.4}$$

Definition: The variance of a set of random outcomes is the expected value of the squared deviations of an outcome from its mean. The variance is a measure of the dispersion of a data series around its expected value. The greater this dispersion, the greater the value of the variance. The variance is the sum of the products of the square of the deviation of each outcome from its mean and the probability of the occurrence of the outcome. When the probability of the occurrence of each outcome is the same as the probability of the occurrence of every other outcome, the mean is the sum of the squared deviations divided by the number of outcomes.

Denoting a win and a loss as x_1 and x_2, respectively, the variances of the two wagers, σ_1^2 and σ_2^2 are

$$\begin{aligned}\sigma_1^2 &= \sum_{i=1\to n} (x_i - \mu)^2 p_i = 0.5(10-0)^2 + 0.5(-10-0)^2 \\ &= 0.5(100) + 0.5(100) = 50 + 50 = 100\end{aligned}$$

$$\begin{aligned}\sigma_1^2 &= \sum_{i=1\to n} (x_i - \mu)^2 p_i = 0.5(1{,}000-0)^2 + 0.5(-1{,}000-0)^2 \\ &= 0.5(1{,}000{,}000) + 0.5(1{,}000{,}000) = 500{,}000 + 500{,}000 = 1{,}000{,}000\end{aligned}$$

Since $\sigma_2^2 > \sigma_1^2$, then the second wager is riskier than the first.

An alternative way to express the riskiness of a set of random outcomes is the *standard deviation*. The standard deviation is simply the square root of the variance, σ.

$$\sigma = \sqrt{\sigma^2} \tag{14.5}$$

Definition: The standard deviation is the square root of the variance.

For the foregoing wagers the standard deviations are $\sigma_1 = \sqrt{\sigma_1^2} = \sqrt{100} = 10$ and $\sigma_2 = \sqrt{\sigma_2^2} = \sqrt{1{,}000{,}000} = 1{,}000$. Since the standard deviation is a monotonic transformation of the variance, the ordering of relative risks of the wagers is preserved. Thus, since $\sigma_2 > \sigma_1$ the second wager is riskier than the first.

Problem 14.3. Using the information provided in Problem 14.1, calculate the variance and the standard deviation of Silver Zephyr's expected profits.

Solution. From Problem 14.1, expected profits are \$640. The variance of Silver Zephyr's expected profits is

$$\begin{aligned}
E\left\{[\pi - E(\pi)]^2\right\} &= \sum_{i=1\to 2} [\pi_i - E(\pi)]^2 p_i = \sigma^2 \\
&= [\pi_1 - E(\pi)]^2 p_1 + [\pi_2 - E(\pi)]^2 p_2 \\
&= 0.4(100 - 640)^2 + 0.6(1{,}000 - 640)^2 \\
&= 0.4(-540)^2 + 0.6(360)^2 = 0.4(291{,}600) + 0.6(129{,}600) \\
&= 116{,}640 + 77{,}760 = \$194{,}400
\end{aligned}$$

The standard deviation is

$$\sigma = \sqrt{\sigma^2} = \sqrt{194{,}400} = \$440.91$$

Problem 14.4. From Problem 14.2, calculate the variance and standard deviation of Bob's expected payout.

Solution. Since the probability of any number between 1 and 6 is 1/6, then Bob's expected payout is

$$\begin{aligned}
E\left\{[v - E(v)]^2\right\} &= \left(\frac{1}{n}\right) \sum_{i=1\to 2} [v - E(v)]^2 = \sigma^2 \\
&= \left(\frac{1}{n}\right)\left\{[v_1 - E(v)]^2 + [v_2 - E(v)]^2 + \ldots + [v_6 - E(v)]^2\right. \\
&= \left(\frac{1}{6}\right)\left[(1-3.5)^2 + (2-3.5)^2 + (3-3.5)^2 + (4-3.5)^2\right. \\
&\quad \left. + (5-3.5)^2 + (6-3.5)^2\right] \\
&= \left(\frac{1}{6}\right)\left[(-2.5)^2 + (-1.5)^2 + (-0.5)^2 + (0.5)^2 + (-1.5)^2 + (-2.5)^2\right] \\
&= \left(\frac{1}{6}\right)[6.25 + 2.25 + 0.25 + 0.25 + 2.25 + 6.25] = \left(\frac{1}{6}\right)(17.5) \\
&= \$2.92
\end{aligned}$$

The standard deviation is

$$\sigma = \sqrt{\sigma^2} = \sqrt{2.92} = \$1.71$$

COEFFICIENT OF VARIATION

Unfortunately, neither the variance nor the standard deviation can be used to compare the riskiness involving two or more risky situations with different expected values. The reason for this is that neither measure is independent of the units of measurement. To measure the relative riskiness of two or more outcomes, we may use the coefficient of variation, which may be calculated by using Equation (14.6). The coefficient of variation allows us to compare the riskiness of alternative projects by "normalizing" the standard deviation of each by its expected value.

$$CV = \frac{\sigma}{\mu} \tag{14.6}$$

Definition: The coefficient of variation is a dimensionless number that is used to compare risk involving two or more outcomes involving different expected values. It is calculated as the ratio of the standard deviation to the mean.

Problem 14.5. Suppose that capital investment project *A* has an expected value of μ_A = \$100,000 and a standard deviation of σ_A = \$30,000. Additionally, suppose that project *B* has an expected value μ_B = \$150,000 and a standard deviation of σ_B = \$40,000. Which is the relatively riskier project?

Solution. From Equation (14.6) the relative riskiness of projects *A* and *B* are

$$CV_A = \frac{\sigma_A}{\mu_A} = \frac{30{,}000}{100{,}000} = 0.300$$

$$CV_B = \frac{\sigma_B}{\mu_B} = \frac{40{,}000}{150{,}000} = 0.267$$

Thus, although project *B* has the larger standard deviation, it is the relatively less risky project.

CONSUMER BEHAVIOR AND RISK AVERSION

Suppose that a manager is confronted with the choice of two investment projects with the same expected rate of return. Which project will the manager choose? Most managers will select the project with the lowest risk, that is, the one with the smallest standard deviation. These managers are said to be *risk averse*. On the other hand, *risk-loving* managers would choose the riskier project. Managers who are indifferent to risk are said to be *risk neutral*. The reason for these differences in managers' behavior toward risk may be explained in terms of the *marginal utility of money*.

In Figure 14.1, which illustrates three total utility of money functions, money income or wealth is measured along the horizontal axis, and a car-

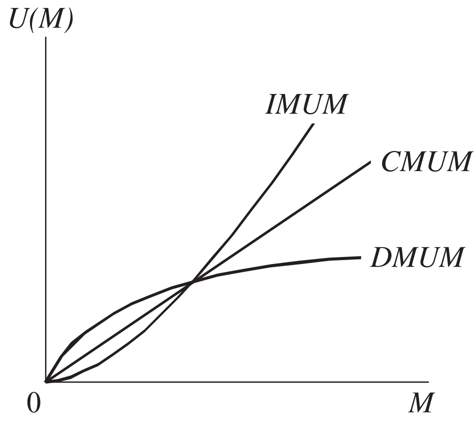

FIGURE 14.1 Increasing, constant, and diminishing marginal utility of money.

dinal index of the utility (satisfaction) of money is measured along the vertical axis. The three total utility of money functions in Figure 14.1 illustrate the concepts of constant marginal utility of money (*CMUM*), increasing marginal utility of money (*IMUM*), and diminishing marginal utility of money (*DMUM*). When conditions of increasing marginal utility of money exist, as more money income is received, the total utility of money increases at an increasing rate. Similarly, constant marginal utility of money means that the total utility of money increases at a constant rate. Finally, decreasing marginal utility of money means that the total utility of money increases at a decreasing rate.

Most individuals are risk averse because their total utility of money function exhibits decreasing marginal utility. To see this, consider an individual who offers the following wager. If the individual flips a coin that comes up "heads," then the individual wins $1,000. On the other hand, if the individual flips a coin that comes up "tails," then the individual loses $1,000. The coin is assumed to be "fair" so there is an even chance of flipping either "heads" or "tails." If we denote the value of the wager as M, the expected value of this wager is $E(M) = 0.5(-\$1{,}000) + 0.5(\$1{,}000) = -\$5{,}000 + \$5{,}000 = 0$. This wager is sometimes referred to as a *fair gamble* because the expected value of the payoff is zero.

Definition: A fair gamble is one in which the expected value of the payoff is zero.

Problem 14.6. Lugg Hammerhands has been offered the following wager (M). Blindfolded, Lugg may draw a single marble from an urn containing 10 marbles that are perfectly identical in terms of size, shape, and weight. Nine of the marbles in the urn are green and one marble is red. If Lugg draws a green marble, then he loses $50. If Lugg draws the red marble, wins $450. Is this a fair gamble?

Solution. The expected value of the wager is

$$E(M) = 0.9(-\$50) + 0.1(\$450) = -\$45 + \$45 = 0$$

Since the expected value of the wager is zero, then this is a fair gamble.

Problem 14.7. In the United States, many state governments sponsor lotteries to support of public education. In New York State, for example, \$1 purchases two games of Lotto. Each game involves selecting six of 59 numbers. The New York State Lottery Commission randomly draws six numbers, and whoever has selected the correct combination wins, or shares, the top prize, which is in the millions of dollars. According to the New York State Lottery Commission, the odds of winning the top prize on a \$1 bet are 1 in 22,528,737. Suppose, for example, that the top prize is \$20 million. Is this a fair gamble?

Solution. Denoting a \$1 wager as M, the expected value of one player winning the top prize is

$$E(M) = \left(\frac{1}{22{,}528{,}737}\right)(20{,}000{,}000) + \left(\frac{22{,}528{,}736}{22{,}528{,}737}\right)(-1) = -0.11$$

Since the expected value is negative, then this game of Lotto is an unfair gamble.

More formally, the utility of money function may be written as

$$U = U(M) \tag{14.7}$$

Utility is assumed to be an increasing function of money, that is, $dU/dM > 0$. Constant marginal utility of money requires that $d^2U/dM^2 = 0$. Increasing marginal utility of money requires that $d^2U/dM^2 > 0$. Diminishing marginal utility of money requires that $d^2U/dM^2 < 0$. The utility of money function of risk-averse individuals exhibits diminishing marginal utility of money (i.e., the total utility of money increases at a decreasing rate). The reason for this is that a risk-averse individual will experience a greater loss of utility by losing \$1,000 than he or she would gain by winning \$1,000. To see this, consider once again the fair gamble of winning or losing \$1,000 on the flip of a coin.

Suppose that the individual's utility of money function is $U = 100M^{0.5}$. The reader will readily verify that this total utility of money function exhibits diminishing marginal utility of money, since $dU/dM = 50M^{-0.5} > 0$ and $d^2U/dM^2 = -25M^{-1.5} < 0$. As we saw earlier, this is a fair gamble because $E(M) = (1{,}000)0.5 + (-1{,}000)0.5 = 0$. Even though this is a fair gamble, a risk-averse individual would not accept the wager. To see this, suppose that the individual's initial money wealth is $M = \$50{,}000$. The utility of money for this individual is $U = 100(50{,}000)^{0.5} = 22{,}361$ units. Now, suppose that the individual flips "heads" and wins \$1,000. The individual's new money wealth is $M' = \$51{,}000$. The individual's total utility of money is $U = 100(51{,}000)^{0.5}$

= 22,583 units. That is, the individual gains 222 utility units. Suppose, on the other hand, that the individual flips "tails" and loses \$1,000. The individual's new money wealth is now \$49,000. The individual's total utility of money is $U = 100(49{,}000)^{0.5} = 22{,}136$ units. In this case, the individual's utility index falls by 225. The expected change in utility from the bet is

$$E(\Delta U) = \sum_{i=1\to n} (\Delta U_i)p_i = (\Delta U_1)p_1 + (\Delta U_2)p_2$$
$$= (222)0.5 + (-225)0.5 = -15 \text{ units}$$

Since the expected utility change is negative, this individual will not accept this fair gamble.

There is yet another way to interpret risk-averse behavior. In the example just given, a risk-averse individual would prefer not to bet, a sure amount of zero than to wager \$1,000 with an expected value of zero. In other words, a risk-averse individual will not accept a fair gamble. Denoting the sure amount as M, a risk-averse individual will prefer a sure amount to its expected value, that is, $M > E(M)$. An individual is said to be risk loving if the reverse is true; that is, the expected value of a payoff is preferred to its certainty equivalent, or $E(M) > M$. Finally, an individual who is indifferent between a certain payoff and its expected value, that is, $M \approx E(M)$, is said to be risk neutral.

It should be noted that while most individuals are risk averse most of the time, under certain circumstances they may be risk loving. In particular, many risk-averse individuals are risk loving for small gambles. An individual who, for example, is willing to wager \$1.00 on the flip of a coin may would not be willing to wager \$1,000. In the first instance $E(M) > M$, but in the second instance $M > E(M)$. In Problem 14.7, we saw that playing Lotto is an unfair gamble. Yet, risk-averse individuals frequently play Lotto because it involves a very small wager and a potentially very large payoff.

Definition: An individual is risk averse if he or she prefers a sure amount to a risky payoff with the same expected value.

Definition: An individual is risk loving when the expected value of a risky payoff is preferred to a sure amount of the same value.

Definition: An individual is risk neutral when the individual is indifferent between a sure amount and a risky payoff with the same expected value.

Problem 14.8. Suppose that an individual is offered the fair gamble of receiving \$1,000 on the flip of a coin showing heads and losing \$1,000 on the flip of a fair coin showing tails. Suppose further that the individual's utility of money function is

$$U = M^{1.1}$$

a. For positive money income, what is this individual's attitude toward risk?
b. If the individual's initial money income is \$50,000, will he or she accept this bet? Explain.

Solution

a. The first derivative of the utility function with respect to money income is

$$\frac{dU}{dM} = (1.1)M^{0.1} > 0$$

That is, the individual's utility is an increasing function of money income. The second derivative of the utility function with respect to money income is

$$\frac{d^2U}{dM^2} = (0.1)(1.1)M^{-0.9} > 0$$

Since the second derivative of the utility function with respect to money income is positive, this individual is a risk lover.

b. Suppose, for example, that the individual's initial money income is M = \$50,000. At this level of money income the index of the utility of money is

$$U = (50{,}000)^{1.1} = 147{,}525.47$$

If the individual wins \$1,000, then the corresponding utility index is

$$U = (51{,}000)^{1.1} = 150{,}774.26$$

That is, ΔU_1 = 3,248.79.
If the individual's loses \$1,000, then the corresponding utility index is

$$U = (49{,}000)^{1.1} = 144{,}283.17$$

That is, ΔU_2 = –3,242.30.
The expected utility of the bet is given as

$$E(\Delta U) = \sum_{i=1 \to n} (\Delta U_i)p_i = (\Delta U_1)p_1 + (\Delta U_2)p_2$$
$$= (3{,}248.79)0.5 + (-3{,}241.70)0.5 = 3.55$$

Since the expected utility change from the bet is positive, this risk-loving individual will accept this fair bet.

EXAMPLES OF RISK-AVERSE CONSUMER BEHAVIOR

Knowledge of risk-averse behavior by consumers has a wide range of applications in managerial decision making. Suppose, for example, that a firm plans to introduce a new brand of coffee. Suppose further that the new brand has only one competitor. Will knowledge of risk-averse behavior by consumers influence the firm's marketing strategy? The challenge to the firm is to persuade consumers to give the new brand of coffee a try. If both brands cost the same, then a risk-averse consumer will tend to stay with the old brand rather than switch to the new brand with an uncertain outcome. This, of course, suggests two possible marketing strategies. Either the firm can offer the product, at least initially, at a lower price to compensate the consumer for the risk of trying the new brand, or the firm can adopt an advertising campaign designed to convince the consumer that the new brand is superior. Either marketing strategy will raise the expected value to the consumer of trying the new brand.

Another example of the consequences of risk-averse behavior relates to the benefits enjoyed by chain stores and franchise operations over independently owned and operated retail operations. A risk-averse American tourist visiting, say, Athens, Greece, for the first time is more likely to have his or her first meal at McDonald's or Burger King rather than sample native victuals at a neighborhood bistro. The reason for this is that the risk-averse tourist may initially prefer a familiar meal of predictable quality to exotic menus of unpredictable quality. Of course, this will very likely change as the tourist over time becomes familiar with the indigenous cuisine and the reputation of local dining establishments. It is left as an exercise for the student to explain why large retail chain stores or franchise operations are typically found in areas in which there are a relatively large number of out-of-town visitors.

Perhaps the most familiar example of risk-averse behavior relates to the purchase of insurance. People purchase insurance, which typically involves small premium payments (relative to the potential loss), to protect themselves against the possibility of catastrophic financial loss. Many homeowners, for example, purchase fire insurance in the unlikely event that their house will burn down. If the insurance premiums for given level of financial protection are equal to the expected value of financial loss resulting from a fire, then this may be viewed as a fair gamble. For a fair gamble, a risk-averse homeowner will purchase fire insurance because he or she prefers a sure outcome to a risky prospect of equal expected value. Because of the difficulties associated with estimating the probability of catastrophic loss, it should not be surprising that insurance companies employ actuaries to determine insurance premiums.

FIRM BEHAVIOR AND RISK AVERSION

As the examples thus for illustrate, an understanding of consumer behavior in situations of risk and uncertainty is an important element in the pricing and output decisions of firms. Risk and uncertainty also have important implications for the firm's investment and production decisions. The concepts introduced in the foregoing discussion of decision making by consumers under conditions of risk are directly applicable to decision making by managers.

RISK-ADJUSTED DISCOUNT RATES

Chapter 12 introduced the concept of the net present value of a capital investment project. The reader will recall that the net present value of a capital project is the difference between the net present value of cash inflows and cash outflows. If the net present value of a project is negative, then it should be rejected. If the net present value of a project is positive, then the project should be considered for adoption. It was demonstrated that the net present value method could be used to evaluate projects of equal or equivalently equal lives. In general, when this method is used, projects with higher net present values are preferred to projects with lower net present values. Equation (12.26) summarizes the net present value of a project as the difference between cash inflows (revenues), R_t, and cash outflows, O_t.

$$NPV = \frac{\sum_{t=1\to n} R_t}{(1+K)^t} - \frac{\sum_{t=1\to n} O_t}{(1+k)^t} \tag{12.26}$$

where k is the appropriate discount rate. Recall from Chapter 12 that the rate of interest used to discount a cash flow is called the *discount rate*. The reader will immediately recognize that calculations of net present value by means of Equation (12.26) are made under conditions of certainty. No mention was made of the potential the riskiness of alternative capital investment projects under consideration.

The use of risk-adjusted discount rates introduces the investor's attitude toward risk directly into the manager's net present value calculations. In Figure 14.2, which illustrates three possible risk–return trade-off functions, the riskiness of a capital investment project, measured as the standard deviation of the expected rate of return, is measured along the horizontal axis, and the discount rate (k), interpreted as the expected rate of return on an investment, or portfolio of investments, is measured along the vertical axis.

The risk–return trade-offs illustrated in Figure 14.2 are called investor *indifference curves*. These indifference curves summarize the expected rates of return that an investor must receive in excess of the expected rate of return from risk-free investment to compensate for the risk associated with a particular investment project. Risk–return indifference curves also reflect the investor's different attitudes toward risk. To see this, consider a risk-free investment where $\sigma = 0$. Here, the risk-free rate of return for each investor is k_{rf}. For a risky investment in which $\sigma > 0$, however, the investor must be compensated with a *risk premium*.

Definition: Investor indifference curves summarize the combinations of risk and expected rate of return for which the investor will be indifferent between a risky and a risk-free investment.

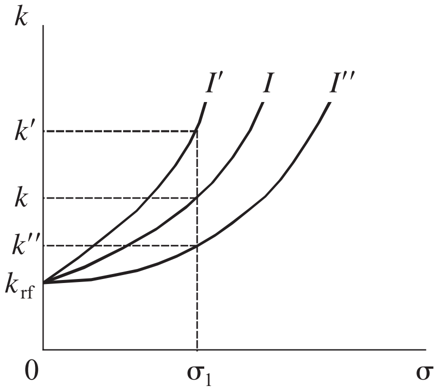

FIGURE 14.2 Investor risk–return indifference curves.

The risk premium on an investment is the difference between the expected rate of return on a risky investment and the expected rate of return on a risk-free investment. The size of the risk premium will depend on the investor's attitude toward risk. Consider, for example, the investor's indifference curve I in Figure 14.2. In this case, a risk premium of $k - k_{rf}$ is required to make this investor indifferent between an investment with risk $\sigma_1 > 0$ and a risk-free investment, $\sigma = 0$. On the other hand, the indifference curve labeled I' illustrates the risk–return trade-offs of a more risk-averse investor. In this case, the investor will require a larger risk premium, $(k' - k_{rf}) > (k - k_{rf})$ as compensation for the same level of risk incurred. Similarly, the indifference curve labeled I'' summarizes the risk–return trade-offs for a less risk-averse investor. Here, a risk premium of $(k'' - k_{rf}) < (k - k_{rf})$ is required to make this investor indifferent to a risk-free investment.

Definition: A risk premium is the difference between the expected rate of return on a risky investment and the expected rate of return on a risk-free investment.

The risk–return indifference curves may be used to evaluate *mutually exclusive* and *independent* investment projects. The reader may recall from Chapter 12 that projects are mutually exclusive if acceptance of one project means rejection of all other projects. Projects are said to be independent if the cash flows from alternative projects are unrelated to each other. Figure 14.3 illustrates management's risk–return indifference curve and three mutually exclusive investment opportunities. As measured by the standard deviation of the expected rates of return, projects A, B, and C are assumed to be equally risky.

The reader may well question the usefulness of proceeding in this manner. After all, when confronted with alternative, mutually exclusive investment projects of equivalent risk, is it not logical to presume that management would choose the project with the highest rate of return? The

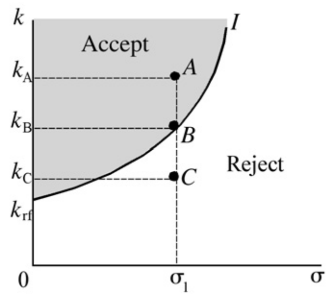

FIGURE 14.3 Risk–return indifference curve and alternative investment projects.

answer to this question is yes, but only if the investment project with the highest expected rate of return is acceptable. It should be readily apparent from Figure 14.3 that only risk–return combinations in the shaded region, such as project *A*, represent acceptable investments. The reason for this is that expected rate of return from project *A* (k_A) is greater than the expected rate of return required to make management indifferent between accepting or rejecting the project (k_B). By contrast, the rate of return from project *B* (k_B) is just sufficient to compensate management for the risk incurred. Clearly, if the projects are mutually exclusive, the investor will prefer project *A* to project *B*. On the other hand, project *C* is unacceptable and will be rejected outright because the rate of return (k_C) is not sufficient to compensate the investor for the risk incurred. Thus, any risk–return combination in the unshaded region will be rejected by a risk-averse investor. By contrast, if the projects under consideration are independent, then the investor will choose project *A* and may choose project *B*, but will reject project *C*.

We have suggested that knowledge of the expected rate of return and of the riskiness of a project, as measured by the standard deviation, is not sufficient to identify the manager's optimal investment strategy. It is also necessary to know the investor's attitude toward risk, which is summarized in the investor's risk–return indifference curve. To amplify this point, consider the situation depicted in Figure 14.4.

The reader will visually verify from Figure 14.4 that the expected rate of return from project *C* is greater than that from project *D*, which is greater than the expected rate of return from project *A*. Finally, project *B* has the lowest expected rate of return. On the other hand, as measured by the standard deviation of the expected rates of return, project *C* is the riskiest of the four projects, while project *B* is the least risky. It should be clear from Figure 14.4 that if projects *A*, *B*, and *C* are independent, then management will accept projects *B* and *D*, but will reject project *C*. Since point *A* lies on the risk–return indifference curve, the investor is indifferent between accepting or rejecting project *A*. On the other hand, if the projects are mutu-

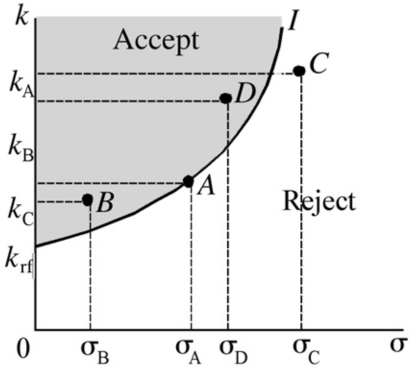

FIGURE 14.4 Risk-indifference curve and alternative investment projects.

ally exclusive, should management will accept project *B* or project *D*? The answer to this question will be addressed in the next section.

RISK–RETURN INDIFFERENCE MAP

We saw in the preceding section that knowledge of the expected rates of return and standard deviations of alternative investment projects is not sufficient to determine the investor's optimal investment strategy. An understanding of the individual's attitude toward risk is absolutely essential for determining the optimal investment strategy. We also saw that an investor's indifference curve summarizes the combinations of risk and return at which the investor will be indifferent between a risky and a risk-free investment. Each investor has a "map" of such risk–return indifference curves. Consider the investor's indifference map in Figure 14.5.

The concept of an investor indifference curve is similar to that of the consumer's indifference curve of utility theory (see Appendix 3′A). The higher the risk–return indifference curve, the greater the investor's level of utility (satisfaction). In Figure 14.5, for example, the risk–return combinations summarized by indifference curve I_3 are preferred to those of I_2 because for any given level of risk, the investor receives a higher expected rate of return. Each investor has an infinite number of such risk–return indifference curves, and each investor has a unique indifference map. Return again to Figure 14.4. Will the investor choose project *A* or project *B*? If project *B* lies on a higher risk–return indifference curve, which seems likely, then project *B* will be preferred to project *A*.

EQUILIBRIUM

Suppose that an investor is considering investing a certain amount in both a risky asset and a risk-free asset. If the investor invests the entire

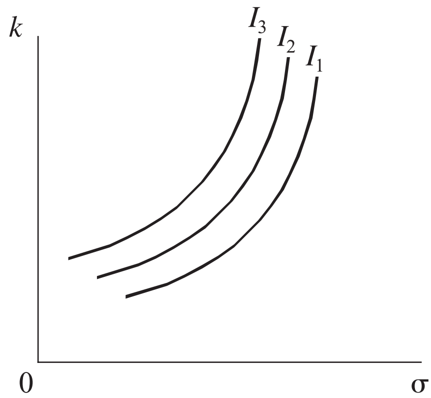

FIGURE 14.5 Investor's indifference map.

amount in the risk-free asset, he or she will earn the expected risk-free rate of return, k_{rf}. If the investor's portfolio includes a combination of the risky and risk-free assets, the expected rate of return from the combination of risky and risk-free assets is

$$k_p = \theta k_{rf} + (1-\theta)k_r \tag{14.8}$$

where k_p is expected return from the portfolio of assets, θ is the percentage of the portfolio comprising the risk-free asset, and $(1 - \theta)$ is the percentage of the portfolio made up of the risky asset, k_r. Equation (14.8) is simply a weighted average of the expected return from the individual assets in the portfolio.

The relationship between the expected rate of return and riskiness of the portfolio consisting of risky and risk-free assets is called the *capital market line*. Figure 14.6 illustrates two capital market lines, M_0 and M_1. The equation of the capital market line is

$$k_p = k_{rf} + \left(\frac{k_r - k_{rf}}{\sigma_r}\right)\sigma_p \tag{14.9}$$

where k_p is the expected rate of return, σ_p is the standard deviation of returns on the portfolio, and σ_r is the standard deviation of returns on the risky assets.[1]

Definition: The capital market line summarizes the market opportunities available to an investor from a portfolio consisting of alternative combinations of risky and risk-free investments.

The slope of the capital market line is the difference between the expected rate of return from the risky asset and the expected rate of return

[1] The slope of the linear capital market line is $\Delta k_p/\Delta\sigma_p$. Starting at the vertical axis, $\Delta\sigma_p = \sigma_r - \sigma_p = \sigma_r - 0 = \sigma_r$ and $\Delta k_p = k^r - k^{rf}$. Thus, the slope of the capital market line is $(k^r - k^{rf})/\sigma_r$.

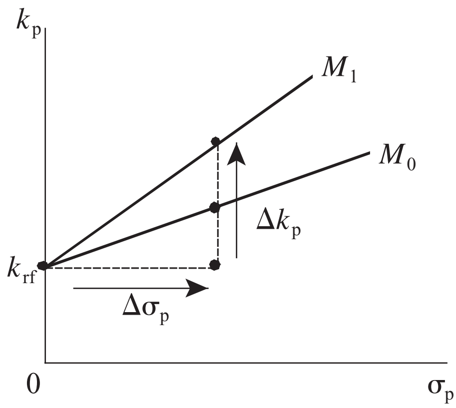

FIGURE 14.6 Investor's opportunity set.

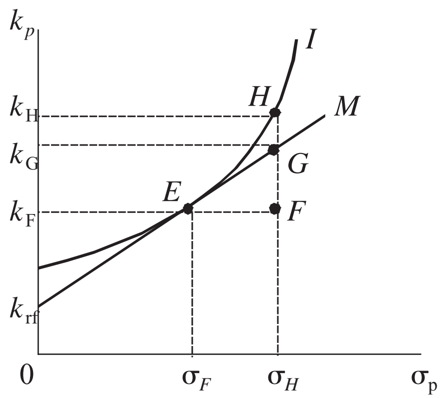

FIGURE 14.7 Investor equilibrium.

from the risk-free asset, divided by the standard deviation of the risky asset, $(k_r - k_{rf})/\sigma_r$, which is called the *market risk premium*. The slope of the capital market line is the expected return on a portfolio of risky and risk-free assets. The steeper the slope of the capital market line, the greater the additional expected rate of return from higher levels of risk associated with holding a greater percentage of the risky asset. In Figure 14.6, the capital market line M_1 represents the higher expected rate of return that is required to compensate the investor for any given level of additional risk incurred.

In Figure 14.7, point *E* represents the highest risk–return indifference curve that this investor can attain given the capital market line *M*. At point *E*, the slope of the risk–return indifference curve is equal to the slope of the capital market line. Equilibrium requires that the expected rate of return on an efficient portfolio of risky and risk-free assets be equal to the expected rate of return at which the investor is indifferent between a risky and a risk-free investment.

There are, of course, other portfolios with equivalent risk–return trade-offs, such as point *H* in Figure 14.7, that are equally preferred to portfolio *E*, but these portfolios are not attainable given the amount of the investment and the expected rates of return on risky and risk-free assets. Portfolio *H*, for example, has a higher level of risk, σ_H, than portfolio *E*. For the investor to assume this higher level of risk, he or she would have to be compensated with the additional expected rate of return, $k_H - k_F$, which is not possible. An increase in the expected risk-free rate of return or an increase in the market risk premium, $(k_r - k_{rf})/\sigma_r$, would enable the investor to purchase portfolio *H* by moving the investor to a higher risk–return indifference curve.

Problem 14.9. Webb Ungoliant has just won \$1,000,000 in the state lottery. Webb has decided to invest his winnings in either U.S. Treasury bills that yield a risk-free expected rate of return of 5%, or risky equity shares in the Lugburz Corporation, which has an expected rate of return of 11%. Webb has analyzed the company's past performance and has determined that the standard deviation of returns is \$3 per share. Suppose that Webb's investment utility function is

$$U = k_p - \sigma_p^2$$

where k_p and σ_p are the portfolio's expected return and standard deviation on the portfolio, respectively. How should Webb's investment be divided between risk-free U.S. Treasury bonds and risky Lugburz shares?

Solution. From Equation (14.9), the capital market line is

$$k_p = k_{rf} + \left(\frac{k_r - k_{rf}}{\sigma_r}\right)\sigma_p = 5 + \left(\frac{11-5}{3}\right)\sigma_p = 5 + 2\sigma_p$$

The problem confronting Webb Ungoliant is to create a portfolio of U.S. Treasury bonds and Lugburz shares that maximizes his investment utility subject to a fixed investment of \$1,000,000. This problem may be conveniently expressed as the constrained maximization problem

$$\text{Maximize:} \quad U(k_p, \sigma_p) = k_p - \sigma_p^2$$

$$\text{Subject to:} \quad k_p = 5 + 2\sigma_p$$

There are at least two solution methods to this problem, which were discussed in Chapter 2. One approach to this constrained profit maximization problem is the substitution method. Substituting the capital market line into the objective function yields

$$U = k_p - \sigma_p^2 = 5 + 2\sigma_p - \sigma_p^2$$

The first-order condition for maximizing this equation with respect to σ_p is

$$\frac{dU}{d\sigma_p} = 2 - 2\sigma_p = 0$$

which yields the solution value

$$2\sigma_p = 2$$

$$\sigma_p^* = 1\%$$

Note that $d^2U/d\sigma_p^2 = -2 < 0$, which is the second-order condition for utility maximization. Substituting this solution value into the equation for the capital market line results in the expected return from the portfolio:

$$k_p^* = 5 + 2(1) = 7\%$$

The composition of Webb's portfolio is

$$k_p = \theta k_{rf} + (1 - \theta)k_r$$

where k_p is expected return from the portfolio of the risky and risk-free assets, p is the percentage of the portfolio consisting of the risk-free asset, and $(1 - p)$ is the percentage of the portfolio consisting of the risky asset. Substituting yields

$$7 = \theta(5) + (1 - \theta)11$$

$$\theta = 0.6667$$

or 66.67% of Webb Ungoliant's portfolio consists of U.S. Treasury bills and $(1 - \theta) = 0.3333$, or 33.33% consists of Lugburz shares.

An alternative solution to this problem is the Lagrange multiplier method. The first step in the Lagrange multiplier method is to bring all terms to left side of the constraint.

$$k_p - 5 - 2\sigma_p = 0$$

The resulting Lagrange function may be written as

$$\mathcal{L}(k_p, \sigma_p, \lambda) = k_p - \sigma_p^2 + \lambda(k_p - 5 - 2\sigma_p)$$

The first-order conditions for a maximum are

$$\frac{\partial \mathcal{L}}{\partial k_p} = 1 + \lambda = 0$$

$$\frac{\partial \mathcal{L}}{\partial \sigma_p} = -2\sigma_p - 2\lambda = 0$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = k_p - 5 - 2\sigma_p = 0$$

Assuming that the second-order conditions for a maximum are satisfied, the simultaneous solution of this system of equations yields the solution values

$$K_p^* = 7\%; \sigma_p^* = 1\%; \lambda^* = -1$$

As demonstrated in Chapter 2, the Lagrange multiplier is the marginal change in the maximum value of the objective function with respect to a parametric change in the value of the constraint. In the present example, the constraint is the expected rate of return on the portfolio. The value of the Lagrange multiplier is

$$\lambda^* = \frac{\partial \mathcal{L}}{\partial k_p} = \frac{\partial U^*}{\partial k_p} = -1$$

From the value of the Lagrange multiplier, we can say that in the limit, an increase of 1% in expected portfolio rate or return will reduce the maximum value Webb's investment utility index by 1 unit. By construction, the optimization procedure guarantees that Webb's investment utility function will always be maximized subject to the capital market line. Altering the expected rate of return on the portfolio will simply change the maximum value of U^*.

Our discussion of the selection of an optimal portfolio was very simple. It was restricted to the choice of the optimal combination of a single risk-free and a single risky asset. A more realistic discussion of the selection of an optimal portfolio must include the possibility of multiple risk-free and risky asset holdings. Although a more extensive treatment of optimal portfolio selection is beyond the scope of this book, the preceding discussion underscores the importance of identifying the investor's behavior toward risk when constraints on alternative investment opportunities are under consideration. For a more detailed treatment of this and related topics dealing with portfolio optimization, the reader is encouraged to consult a text on financial management, such as Brigham, Gapenski, and Ehrhardt (1998).

In Chapter 12 we considered the net present value method for evaluating alternative capital investment projects. We are now in a position to adjust the net present value method to explicitly consider the relative riskiness of the alternative capital projects. We will begin the discussion by rewriting Equation (12.26) as

$$NPV = \frac{\sum_{t=1 \to n} (R_t - O_t)}{(1+k)^t} \tag{14.10}$$

where $(R_t - O_t)$ represents the firm's net revenues on an investment and k is the discount rate. Suppose, for example, that a firm is considering investing in a project that promises annual net revenues of $100,000 for the next

5 years. If the risk-free discount rate is 10%, then the net present value of this investment project is

$$NPV = \frac{\sum_{t=1\to 5} (\$100,000)}{(1.10)^t} = \$100,000 \sum_{t=1\to 5} \left(\frac{1}{1.10}\right)^t$$

$$= \$100,000\left[\left(\frac{1}{1.10}\right)^1 + \left(\frac{1}{1.10}\right)^2 + \left(\frac{1}{1.10}\right)^3 + \left(\frac{1}{1.10}\right)^4 + \left(\frac{1}{1.10}\right)^5\right]$$

$$= \$100,000(0.91 + 0.83 + 0.75 + 0.68 + 0.62) = \$379,000$$

Suppose, on the other hand, the investor perceives the project as risky and uses a risk-adjusted discount rate of 20%. The net present value of the project is

$$NPV = \frac{\sum_{t=1\to 5} (\$100,000)}{(1.20)^t} = \$100,000 \sum_{t=1\to 5} \left(\frac{1}{1.10}\right)^t$$

$$= \$100,000\left[\left(\frac{1}{1.20}\right)^1 + \left(\frac{1}{1.20}\right)^2 + \left(\frac{1}{1.20}\right)^3 + \left(\frac{1}{1.20}\right)^4 + \left(\frac{1}{1.20}\right)^5\right]$$

$$= \$100,000(0.83 + 0.69 + 0.58 + 0.48 + 0.40) = \$298,000$$

Since the net present value of the investment is positive, the project will be considered for adoption. On the other hand, it should be clear from the example that, *ceteris paribus*, riskier investment projects are less preferred.

Definition: Risk-adjusted discount rates are used when one is calculating net present values to compensate for the perceived riskiness of an investment. The greater the perceived risk, the higher will be the discount rate used to calculate the net present value.

Problem 14.10. Suppose that the Orcrist Sword and Blade Company is considering an expansion of its production capability by purchasing a new grinding and polishing machine. While the cost of the investment is known with certainty, the cash inflows are not. The cost of the machine is \$100,000 and the expected annual cash inflows are \$40,000 annually for 5 years.

a. Should Orcrist consider the investment if the discount rate is 10%?
b. Suppose that the riskiness of expected cash inflows is such that management requires a 30% rate of return. Should Orcrist consider this investment?

Solution

a. At a discount rate of 10%, the net present value of this investment project is

$$NPV = \sum_{t=1\to n}\left[\frac{R_t}{(1+k)^t}\right] - O_0 = \frac{\sum_{t=1\to 5}(\$40,000)}{(1.10)^t} - \$100,000$$

$$= \$400,000\sum_{t=1\to 5}\left(\frac{1}{1.10}\right)^t - \$100,000$$

$$= \$400,000\left[\left(\frac{1}{1.10}\right)^1 + \left(\frac{1}{1.10}\right)^2 + \left(\frac{1}{1.10}\right)^3 + \left(\frac{1}{1.10}\right)^4 + \left(\frac{1}{1.10}\right)^5\right]$$

$$- \$100,000$$

$$= \$40,000(0.83 + 0.69 + 0.58 + 0.48 + 0.40) - \$100,000$$

$$= \$51,600$$

Since the net present value of the investment project is positive, Orcrist should undertake the investment.

b. At a risk-adjusted discount rate of 30%, the net present value of this investment project is

$$NPV = \sum_{t=1\to n}\left[\frac{R_t}{(1+k)^t}\right] - O_0$$

$$= \frac{\sum_{t=1\to 5}(\$40,000)}{(1.30)^t} - \$100,000 = \$400,000\sum_{t=1\to 5}\left(\frac{1}{1.30}\right)^t - \$100,000$$

$$= \$40,000\left[\left(\frac{1}{1.30}\right)^1 + \left(\frac{1}{1.30}\right)^2 + \left(\frac{1}{1.30}\right)^3 + \left(\frac{1}{1.30}\right)^4 + \left(\frac{1}{1.30}\right)^5\right]$$

$$- \$100,000$$

$$= \$40,000(0.77 + 0.59 + 0.46 + 0.35 + 0.27) - \$100,000$$

$$= \$2,400$$

Since the net present value of the risk-adjusted investment project is negative, Orcrist should not undertake the expansion.

Although the foregoing examples highlight the importance of considering the potential riskiness of investment projects, they suffer from at least two shortcomings. The first deals with the subjective selection of the risk-adjusted discount rate. In general, it will not be possible to consistently and objectively determine the risk-adjusted discount rate, especially in the absence of historical data. Another, more conceptual, shortcoming is that the risk-adjusted discount rate approach does not explicitly consider the investor's attitude toward risk. While the first shortcoming is problematic, it is possible to compensate for the second shortcoming by evaluating potential investments by using the *certainty-equivalent approach*.

CERTAINTY-EQUIVALENT APPROACH

Earlier, the perceived riskiness of an investment was incorporated into the net present value method by raising the discount rate. Recall from Chapter 12 that the net present value of a project may be calculated as

$$\begin{aligned} NPV &= CF_0 + \frac{CF_1}{(1+k)^1} + \frac{CF_2}{(1+k)^2} + \ldots + \frac{CF_n}{(1+k)^n} \\ &= \frac{\sum_{t=0 \to n} CF_t}{(1+k)^t} \end{aligned} \tag{12.25}$$

where CF_t is the expected net cash flow in period t, k is the cost of capital, and n is the life of the project. Net cash flows are defined as the difference between cash inflows (revenues), R_t, and cash outflows, O_t. Thus, Equation (12.25) was rewritten as

$$\begin{aligned} NPV &= \frac{\sum_{t=0 \to n} R_t}{(1+k)^t} - \frac{\sum_{t=0 \to n} O_t}{(1+k)^t} \\ &= \frac{\sum_{t=1 \to n} (R_t - O_t)}{(1+k)^t} \end{aligned} \tag{12.26}$$

By contrast, the certainty-equivalent approach incorporates risk into the net present value method by multiplying expected net cash flows by the term α_t, which is called the *certainty-equivalent coefficient*. The new expression is

$$NPV = \frac{\sum_{t=1 \to n} \alpha_t (R_t - O_t)}{(1+k)^t} = \frac{\sum_{t=1 \to n} \alpha_t CF_t}{(1+k)^t} \tag{14.11}$$

The certainty-equivalent coefficient is defined as the ratio of a risk-free net cash flow to its equivalent risky cash flow

$$\alpha_t = \frac{CF_t^*}{CF_t} \tag{14.12}$$

where CF_t^* is the risk-free cash flow and CF_t is the actual, risky cash flow that is considered to be equivalent to CF_t^*. Since $CF_t^* \leq CF_t$, then $0 \leq \alpha_t \leq 1$. When $CF_t^* = CF_t$, then $\alpha_t = 1$, in which case the investment is considered to be risk free. On the other hand, when the risky cash flow that is considered to be equivalent to the risk-free cash flow is infinitely large (i.e., $CF_t = \infty$), then $\alpha_t = 0$, in which case the project should be rejected out of hand.

Definition: The certainty-equivalent approach modifies the net present value approach to evaluating capital investment projects by incorporating risk directly into expected cash flows by means of a certainty-equivalent coefficient.

Definition: The certainty-equivalent coefficient is the ratio of a risk-free net cash flow to its equivalent risky cash flow. The smaller the coefficient, the greater the perceived riskiness of an investment.

To illustrate the certainty-equivalent approach, suppose that a firm is considering an investment with expected net revenues of $100,000 for the next 5 years. Suppose, further, that management believes that the risk of future cash-flow receipts will increase over time. Thus, management has subjectively determined that $\alpha_1 = 0.9$, $\alpha_2 = 0.8$, $\alpha_3 = 0.7$, $\alpha_4 = 0.6$, and $\alpha_5 = 0.5$. The risk-free discount rate is assumed to be 10%. Using the certainty-equivalent approach, the risk-adjusted net present value of this investment project is

$$NPV = \frac{\sum_{t=1\to n} \alpha_t CF_t}{(1+k)^t}$$

$$= \left[0.9(\$100{,}000)\left(\frac{1}{1.10}\right)^1 + 0.8(\$100{,}000)\left(\frac{1}{1.10}\right)^2 + 0.7(\$100{,}000)\left(\frac{1}{1.10}\right)^3 + 0.6(\$100{,}000)\left(\frac{1}{1.10}\right)^4 + 0.5(\$100{,}000)\left(\frac{1}{1.10}\right)^5\right] = \$272{,}600$$

Of course, since the net present value of the investment is positive, the firm will consider the project for adoption. The reader should verify that riskier projects with lower certainty-equivalent coefficients will result in lower risk-adjusted net present values.

As with the risk-adjusted discount rate approach, the certainty-equivalent method suffers from the shortcoming of the subjective determination of the certainty equivalent cash flow (CF_t) in Equations (14.11) and (14.12). On the other hand, the certainty-equivalent approach is conceptually superior to the risk-adjusted discount rate approach in that it explicitly considers the investor's attitude toward risk.

COST OF EQUITY CAPITAL: CAPITAL ASSET PRICING MODEL (CAPM)

Earlier, we defined the cost of common stock, k_c, as the stockholder's required rate of return on common stock. It was also noted that there are two sources of equity capital: retained earnings and capital financing obtained by issuing new shares of common stock. What is the *required rate* *of return*? It may be argued that investors expect to receive the stock's dividend yield plus its expected growth rate, g. If we denote the dividend payout at time t as D_t and the share price at the time of purchase as S_0, we may write the expected rate of return as $D_t/S_0 + g$. Alternatively, the required rate of return may be interpreted as the rate of return on a risk-free investment, say the short-term (30-day) U.S. Treasury bill rate, k_{rf}, plus a risk premium, rp. The risk premium may be taken to be the difference between the expected rate of return from a "market" investment, k_{m} (i.e., the rate of return on an "average" investment) and k_{rf}. Thus, the required rate of return may be defined as

$$k_c = k_{\text{rf}} + rp = \frac{D_t}{S_0} + g \tag{14.13}$$

One approach for estimating the cost of common stock by means of the concepts just introduced is to use the *capital asset pricing model*, or *CAPM*. The *CAPM* model is summarized in Equation (14.14).

$$k_{\text{c}} = k_{\text{rf}} + \beta_i(k_{\text{m}} - k_{\text{rf}}) \tag{14.14}$$

The capital asset pricing model is used to analyze the relationship between the risk associated with the purchase of a stock and its rate of return. The *beta coefficient*, β_i, is one measure of this risk. The beta coefficient is a measure of the tendency of stock prices for company i to move up and down with "average" stock prices as measured by some market index, such as the S&P 500, the New York Stock Exchange Index, or the Dow Jones Industrial Average. Stock prices that move in step with the market average have a beta coefficient equal to unity ($\beta_i = 1$). Stocks that exaggerate fluctuations in the market—that is, stock prices exhibiting more than the average price fluctuations—have beta coefficients greater than unity ($\beta_i > 1$), while stock prices that are less volatile have positive beta coefficients less than unity ($0 < \beta_i < 1$).

Definition: The capital asset pricing model (CAPM) establishes a relationship between the risk associated with the purchase of a stock and its rate of return. CAPM asserts that the required return on a company's stock is equal to the risk-free rate of return plus a risk premium.

Definition: The beta coefficient measures the price volatility of a given stock with the price volatility of "average" stock prices.

Equation (14.14) admits to several interesting relationships. Taking the first partial derivatives of k_{c} with respect to k_{rf}, k_{m}, and β_i yields

$$\frac{\partial k_{\text{c}}}{\partial k_{\text{rf}}} = 1 - \beta_i \tag{14.15}$$

$$\frac{\partial k_{\text{c}}}{\partial k_{\text{m}}} = \beta_i \tag{14.16}$$

$$\frac{\partial k_c}{\partial \beta_i} = k_m - k_{rf} \tag{14.17}$$

Equation (14.15) says that the required rate of return varies directly or inversely with the risk-free rate of return depending upon the value of β_i. When the rate of return on a company's stock is more volatile than the market average ($\beta_i > 1$), then an increase in the risk-free rate results in a decline in the required rate of return, and vice versa. On the other hand, when the rate of return on a company's stock price is less volatile than the market average ($0 < \beta_i < 1$), an increase in the risk-free rate will result in an increase in the required rate of return.

Equation (14.16) says that the required rate of return varies directly with the market rate of return, which is the value of the beta coefficient. Finally, Equation (14.17) says that as long as the market rate of return is greater than the risk-free rate, the more volatile the stock, the greater the required rate of return.

The procedure for estimating the cost of common stock using the CAPM approach is straightforward. The procedure begins with selecting some risk-free rate, k_{rf}, which is usually taken to be the return on a U.S. government security. Next, estimate β_i, which may be accomplished by "regressing" the historic realized returns on a company's stock, k_c, against the historic realized returns on average stock prices, k_m, using some market index as a proxy (see, e.g., Greene, 1997). These values may be substituted into Equation (14.14) to determine a "risk-adjusted" cost of capital stock.

Problem 14.11. Suppose that the rate of return on 3-month U.S. Treasury bills is 7% and the market rate of return is 10%.

a. If $\beta_i = 1.5$, what is a company's required rate of return? What is the required rate of return if $\beta_i = 0.75$?
b. Suppose that $\beta_i = 1.5$. If the rate of return on 3-month U.S. Treasury bills rises to 8%, what is the company's required rate of return?

Solution

a. According to the capital asset pricing model, the required rate of return on an individual stock is calculated as

$$k_c = k_{rf} + \beta_i (k_m - k_{rf})$$

where k_c is the required rate of return on stock i, k_{rf} is the risk-free rate of return, and k_m is the market rate of return. Substituting into this expression the values given in the problem yields

$$k_c = 7 + 1.5(10 - 7) = 11.5\%$$

Thus, the required rate of return is 11.5%.

When $\beta_i = 0.75$, then

$$k_{\mathrm{c}} = 7 + 0.75(10 - 7) = 9.25\%$$

b. When $k_{\mathrm{rf}} = 8$ and $\beta_i = 1.5$, then

$$k_{\mathrm{c}} = 8 + 1.5(10 - 8) = 11.0\%$$

GAME THEORY AND UNCERTAINTY

In Chapter 13 we introduced non cooperative, simultaneous, one-shot, finitely and infinitely repeated games with known payoffs. The focus of those discussions was decision making that involves the strategic interaction (move and countermove) of players. How might the implications of game theory be modified in circumstances in which the payoffs are uncertain? To illustrate, consider the following variation of the Prisoners' Dilemma, called the Slumlords' Dilemma, which was first discussed by Davis and Whinston (1962).

SLUMLORDS' DILEMMA

Suppose that there are two owners of adjacent slum tenements: Slumlord Larry and Slumlady Sally. Both Slumlord Larry and Slumlady Sally are considering investing $100,000 to renovate their apartment buildings. If both individuals invest in their properties, they will have the most appealing low-rent apartments in the area and can expect higher profits. If Slumlady Sally invests but Slumlord Larry does not, then Slumlord Sally will lose money while Slumlord Larry will earn positive profits.

The reason for these outcomes is that this type of investment involves externalities. Slumlady Sally will experience only a small increase in the demand for her apartments because of the negative externality of being located near Slumlord Larry's run-down tenement. Slumlord Larry, on the other hand, will find a sharp increase in the demand for his apartments because of the positive effect on the neighborhood from Slumlady Sally's investment. The opposite result would occur if Larry invested and Sally did not. Finally, if neither invests, the economic profits of both will be zero. This normal-form game is depicted in Figure 14.8.

The reader will readily verify that if Slumlord Larry and Slumlady Sally do not cooperate, the solution to the Slumlords' Dilemma will be the strategy profile {*Don't invest*, *Don't invest*}. That is, in the absence of an agreement between Larry and Sally, it will be in both individuals' interest not to invest in their properties. The reason for this is that both players have a strictly dominant strategy. For example, regardless of whether Slumlady Sally invests, it will be in Slumlord Larry's best interest not to invest because not investing will result in the highest profit. The same is true

		Slumlady Sally: Invest	Slumlady Sally: Don't invest
Slumlord Larry	Invest	($5,000, $5,000)	(–$3,000, $7,500)
	Don't invest	($7,500, –$3,000)	($0, $0)

Payoffs: (Slumlord Larry, Slumlady Sally)

FIGURE 14.8 Slumlords' Dilemma.

for Slumlady Larry regardless of whether Slumlord Larry invests or does not invest. Moreover, this solution constitutes a Nash equilibrium, since neither player will be able to increase profits by unilaterally switching strategies.

It is obvious that it will be in the best interest of both Slumlord Larry and Slumlady Sally to cooperate and agree to invest in their properties. If both tenement owners trust each other to keep the agreement, then the strategy profile {*Invest*, *Invest*} will result in a mutually beneficial outcome for both players. But since there is an obvious incentive for Slumlord Larry and Slumlady Sally to mislead each other, neither one will be certain whether the other will actually invest in spite of the agreement to do so. Now, magnify this scenario into a more realistic situation involving three or more tenement owners in which the decision by one owner to invest depends on the investment decisions of each of several others. As difficult as it may be for Slumlord Larry and Slumlady Sally to trust each other, imagine the implausibility of simultaneous trust among three or more tenement owners, even though investment in their properties would be in the best interest of owners and residents alike. This inability of tenement owners to trust one another will result in the perpetuation of uneconomic slum conditions.

In this and other cases involving externalities, the solution to the problem is found by internalizing these third-party effects. In this case, the problem is the uncertainty arising from a lack of trust among the participants. This problem would be eliminated if there was only one owner of all the tenements in the area. If there are only two owners, as in the situation depicted in Figure 14.8, it might be possible to, say, persuade Slumlord Larry to sell out to Slumlady Sally. When there are three or more owners, however, the likelihood that one owner will be able to buy out the others diminishes as the number of owners increases. One public policy solution to this problem is for local government to exercise the power of eminent domain and purchase the run-down properties. The government might redevelop the properties on its own, or sell them to a single developer who will agree to do so. This process is commonly known as urban renewal.

		State A	State B
Slumlord Larry	Invest	$5,000	–$1,000
	Don't invest	$3,000	$0

FIGURE 14.9 The Slumlords' Dilemma under alternative states of nature.

The Slumlords' Dilemma is typical of the simultaneous move, one-shot games discussed in Chapter 13 in which the outcomes of the combined strategies are known with certainty. The nature of the game fundamentally changes, however, if uncertainty is explicitly introduced into the deliberations. To see this, consider Figure 14.9, which summarizes the revised possible outcomes confronting Slumlord Larry under alternative states of nature.

The reader will immediately recognize from Figure 14.9 that Slumlord Larry no longer has a dominant strategy. The payoff to Slumlord Larry is greater by investing in state A and by not investing in state B. Now, suppose that Slumlord Larry believes that the state of nature is determined by Slumlady Sally, and that state A is to invest and state B is not to invest. In this case, Slumlord Larry recognizes that Slumlady Sally has a dominant strategy not to invest. If Slumlady Sally invests, then it is in Slumlord Larry's best interest to invest, and if Slumlady Sally does not invest, then it is in Slumlord Larry's best interest not to invest. Thus, if Slumlord Larry and Slumlady Sally do not cooperate, the strategy profile {*Don't invest*, *Don't invest*}, which still constitutes a Nash equilibrium.

Suppose, now, that the states of nature depicted in Figure 14.9 no longer represent Slumlady Sally's investment decisions. Instead, suppose the states of nature represent circumstances outside either player's control, which nevertheless affect the likelihood of the outcomes. For example, suppose that Slumlord Larry's tenement is located in an area of the city known for its high arson rate. Suppose that it is commonly known that state *B* represents a 20% probability that the tenement will be "torched." The expected payoffs of investing and not investing given these states of nature are 0.8($5,000) + 0.2(–$1,000) = $3,800 and 0.8($3,000) – 0.2($0) = $2,400, respectively.

With expected values of the outcomes, Slumlord Larry's optimal strategy is to invest because of the greater expected value, regardless of what Slumlady Sally decides to do. A shortcoming of this solution is that in the presence of uncertain outcomes it does not explicitly consider Slumlord Larry's attitudes toward risk. In other cases, the probabilities of the different states of nature are not known, nor can they be inferred, in which case some other decision rule must be used to determine the player's optimal strategy. These situations will be considered in the next section.

GAME TREES

In Chapter 13 we introduced the concept of the extensive form of multistage games. Multistage games differ from simultaneous-move games in that the players' moves are sequential. An extensive-form game summarizes the players, game stages, information set, strategies, order of the moves, and payoffs from alternative strategies. The collection of decision nodes and branches that characterize the extensive for of multistage games was referred to as a *game tree*. The game trees presented in Chapter 13 were used to analyze certain outcomes from alternative strategies. In this section, we will use game trees to analyze expected (risky) outcomes from alternative strategies.

Definition: Game trees are used to analyze outcomes from alternative strategies in multistage (sequential move) games. Game trees are made up of decision nodes and branches. At each decision node a decision maker must choose from among alternative moves. The first decision node is called the root of the game tree. Each move is represented by a branch. A game tree is used to determine a decision maker's optimal strategy. A strategy is the decision maker's game plan under all eventualities.

To illustrate how game trees may be used in the decision-making process, consider the following multistage game involving the management of two firms deciding whether to adopt a high-price or a low-price strategy. This is a sequential-move game in that firm *A* must first decide whether to charge a high price or a low price without having the luxury of knowing how firm *B* will respond. The players, strategies, and profits of this game are summarized in Figure 14.10.

According to the backward induction solution methodology discussed in Chapter 13, if firm *A* adopts a "high-price" strategy, then firm *B*'s best move

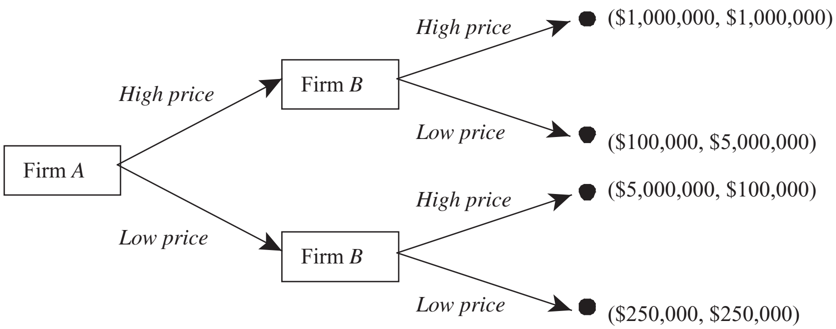

Payoffs: (Firm *A*, Firm *B*)

FIGURE 14.10 Game tree for a multistage game.

is to adopt a "low-price" strategy, since the profit is $5,000,000 compared with a profit of $1,000,000 by charging a high price. Thus, the "high-price" branch from the firm *B* decision node should be disregarded. On the other hand, if firm *A* adopts a "low-price" strategy, firm *B*'s optimal strategy is to charge a low price because the resulting profit is $250,000 compared with a profit of $100,000.

An examination of Figure 14.10 readily reveals that the optimal strategy for firm *B* is to charge a low price if firm *A* charges a high price, or to charge a low price if firm *A* charges a low price. Firm *B*'s optimal strategy may be summarized as *(Low price, Low price)*. It is clear from the extensive-form game in Figure 14.10 that the optimal strategy for firm *A* is to charge a low price, which will result in a profit of $250,000, compared with a profit of $100,000 if it adopts a high-price strategy. Thus, the final optimal strategy profile for this game is {*Low price*, (*Low price, Low price*)}, which yields profits of $250,000 for both firm *A* and firm *B*.

We will now modify the game depicted in Figure 14.10 by introducing risk into the decision-making process. Suppose, for example, that firm *A* believes that by adopting a high-price strategy there is a 40% probability that firm *B* will charge a high price and a 60% probability that its rival will change a low price. Similarly, if firm *A* believes that by adopting a low strategy there is an 80% probability that firm *B* will charge a high price and a 20% probability of a low price. The resulting extensive form for this multi-stage game is illustrated in Figure 14.11.

In Figure 14.11, the first entry in the parentheses at the terminal nodes indicates the expected profit to firm *A*. The second entry indicates a certain profit to firm *B*. The reason for this is that firm *A* is uncertain whether firm *B* will adopt a "high-price" or a "low-price" strategy. On the other hand, once firm *A* has decided on its strategy, firm *B*'s strategy is certain. Once again, using the technique of backward induction introduced in Chapter 13,

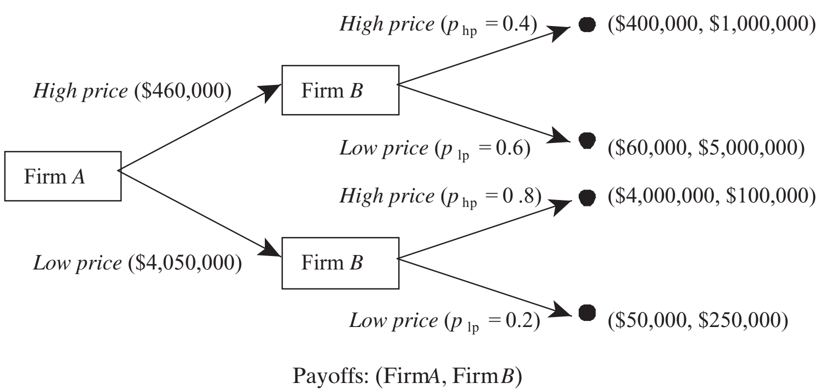

FIGURE 14.11 Game tree for an extensive-form game.

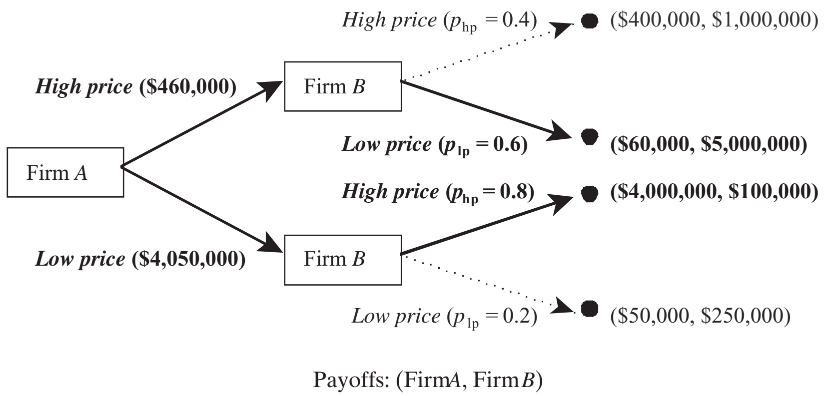

FIGURE 14.12 Using backward induction with probabilistic payoffs to find a Nash equilibrium.

if firm *A* adopts a "high-price" strategy, then firm *B*'s best move is to adopt a "low-price" strategy. If firm *A* adopts a "low-price" strategy, then firm *B*'s best move is to charge a "high price." The resulting extensive-form game is illustrated in Figure 14.12.

Once again, an examination of Figure 14.10 reveals that the optimal strategy for firm *B* is to charge a low price if firm *A* charges a high price, or to charge a low price if firm *A* charges a low price. As before, the optimal strategy for firm *B* is to charge a low price. In the situation depicted in Figure 14.12, however, the optimal strategy of firm *A* is based on the expected profit from a "high-price" strategy versus the expected profit from a "low-price" strategy. If firm *A* adopts a "high-price" strategy, the expected profit is π_{hp} = 0.4($1,000,000) + 0.6($100,000) = $460,000. If firm *A* adopts a "low-price" strategy, the expected profit is π_{lp} = 0.8($5,000,000) + 0.2($250,000) = $4,050,000. But should firm *A* adopt a "low-price" strategy simply because of its higher expected profit?

Whether firm *A* adopts a "high-price" strategy or a "low-price" strategy will depend on management's attitude toward risk. To see this, let us calculate the standard deviation of expected profit from each strategy. The standard deviation of firm *A*'s expected profit from a "high-price" strategy is

$$\sigma_{hp} = (\sigma^2)^{0.5} = \left\{[\pi_{hp} - E(\pi)]^2 p_{hp} + [\pi_{lp} - E(\pi)]^2 p_{lp}\right\}^{0.5}$$
$$= \left\{0.4(1{,}000{,}000 - 460{,}000)^2 + 0.6(100{,}000 - 460{,}000)^2\right\}^{0.5}$$
$$= \left\{0.4(540{,}000)^2 + 0.6(-360{,}000)^2\right\}^{0.5} = 440{,}908$$

The standard deviation of firm *A*'s expected profit from a "low-price" strategy is

$$\sigma_{\mathrm{lp}} = (\sigma^2)^{0.5} = \left\{[\pi_{\mathrm{hp}} - E(\pi)]^2 p_{\mathrm{hp}} + [\pi_{\mathrm{lp}} - E(\pi)]^2 p_{\mathrm{lp}}\right\}^{0.5}$$
$$= \left\{0.8(5,000,000 - 4,050,000)^2 + 0.2(250,000 - 4,050,000)^2\right\}^{0.5}$$
$$= \left\{0.8(950,000)^2 + 0.2(-3,800,000)^2\right\}^{0.5} = 1,900,000$$

These calculations indicate that the higher expected profit from a "low-price" strategy is relatively less risky than the lower expected profit from a "high-price" strategy. These results suggest that the optimal pricing strategy for firm *A* is to charge a low price. Before concluding that the optimal strategy profile for this game is {*Low price*, (*Low price*, *Low price*)}, it is important to consider management's attitude toward risk. Although a "low-price" strategy has the highest expected return and the lowest risk, if the expected rate of return is insufficient to compensate the investor for the associated riskiness of the project (i.e., the risk–return combination is in the unshaded region of Figure 14.3), then even this investment opportunity will be rejected.

Problem 14.12. Consider the extensive-form game in Figure 14.10. Suppose that firm *A* believes that adopting a high-price strategy will result in a 95% probability that firm *B* will charge a high price and a 5% probability that it will charge a low price. Similarly, firm *A* believes that a low-price strategy will result in a 2% probability that firm *B* will charge a high price and a 98% probability of a low price. What is the optimal strategy profile for this game?

Solution. The revised extensive form of this game is illustrated in Figure 14.13, which indicates that the optimal strategy for firm *B* is to charge a low price if firm *A* charges a high price, or to charge a low price if firm *A* charges a low price. Thus, firm *B*'s optimal strategy may be summarized as (*Low price, Low price*). Firm *A*'s optimal strategy is based on the expected profit from charging a high price versus the expected profit from charging a low price. If firm *A* adopts a "high-price" strategy, the expected profit is 0.95($1,000,000) + 0.05($100,000) = $955,000. If firm *A* adopts a "low-price" strategy, the expected profit is 0.02($5,000,000) + 0.98($250,000) = $345,000. Although firm *A*'s expected profit is greater for a "high-price" strategy, the firm's strategy will depend on management's attitude toward risk. To determine the riskiness of each strategy, calculate the standard deviations of the expected payoffs.

$$\sigma_{\mathrm{hp}} = (\sigma^2)^{0.5} = \left\{[\pi_{\mathrm{hp}} - E(\pi)]^2 p_{\mathrm{hp}} + [\pi_{\mathrm{lp}} - E(\pi)]^2 p_{\mathrm{lp}}\right\}^{0.5}$$
$$= \left\{0.95(1,000,000 - 955,000)^2 + 0.05(100,000 - 955,000)^2\right\}^{0.5}$$
$$= \left\{0.95(45,000)^2 + 0.05(-855,000)^2\right\}^{0.5} = 196,150$$

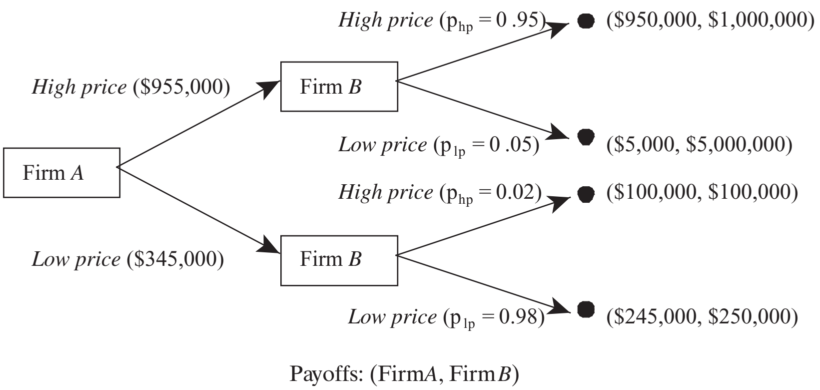

FIGURE 14.13 Game tree for problem 14.12.

The standard deviation of firm *A*'s expected profits from a "low-price" strategy is

$$\sigma_{lp} = (\sigma^2)^{0.5} = \left\{[\pi_{hp} - E(\pi)]^2 p_{hp} + [\pi_{lp} - E(\pi)]^2 p_{lp}\right\}^{0.5}$$

$$= \left\{0.02(5{,}000{,}000 - 345{,}000)^2 + 0.98(250{,}000 - 345{,}000)^2\right\}^{0.5}$$

$$= \left\{0.02(4655{,}000)^2 + 0.98(-95{,}000)^2\right\}^{0.5} = 665{,}000$$

The standard deviations of expected profits suggest that the "high-price" strategy is also the less risky. Since the 'high-price" strategy also has the higher expected profit, then the optimal pricing strategy for firm *A* is to charge a high price. Thus, the optimal strategy profile for this game is {*High price*, (*Low price, low price*)}.

In the preceding examples, we assumed that firm *A* would select that strategy with the highest expected profit and the lowest risk. Although this may appear quite logical, it fails to explicitly consider management's attitude toward risk. This is important because it might very well be that neither strategy is acceptable, in which case the firm might consider a third, "no change," price strategy. To amplify this point, suppose that a "high-price" strategy had the highest expected profit and was most risky, while the "low-price" strategy had the lowest expected profit and was the least risky. In this case, a risk-loving management might prefer a "high-price" strategy, whereas a risk-averse management might prefer the "low-price" strategy. It is important to recognize that while game trees are useful analytical devices for evaluating sequential managerial decisions, the selection of optimal strategies should also take into account management attitudes toward risk.

DECISION MAKING UNDER UNCERTAINTY WITH COMPLETE IGNORANCE

It was mentioned earlier that whenever the decision maker is able to use personal knowledge, intuition, and experience to assign subjective probabilities to outcomes, decision making under uncertainty is transformed into decision making under risk. These situations were described as decision making under conditions of uncertainty with partial ignorance. When managers are unable to assign probabilities to alternative outcomes, some other rational decision-making criteria must be used. As mentioned earlier, this is referred to as decision making under conditions of uncertainty with complete ignorance. In this section we will examine four such rational decision criteria: the *Laplace criterion*, the *Wald* (*maximin*) *criterion*, the *Hurwicz criterion*, and the *Savage* (*minimax regret*) *criterion*. No single decision rule is appropriate for all decision-making situations. The choice of the criterion should be appropriate to the circumstances and consistent with organizational objectives and philosophy.

LAPLACE DECISION CRITERION

Before examining in detail the Laplace decision criterion for selecting among alternative strategies under conditions of complete ignorance, consider the situation depicted in Figure 14.14. This figure summarizes the payoffs from three possible pricing strategies given three different states of the economy: economic expansion, stability, and contraction. The payoffs in the matrix represent the firm's expected rates of return.

Under conditions of risk or partial ignorance, however, the decision maker may be able to assign objective or subjective probabilities to the different states of the economy. These probabilities (in parentheses), and the expected values of the payoffs from each strategy, $E(S_i) = \mu_i$, are summarized in Figure 14.15. As in the Slumlords' Dilemma, if management decides to adopt the pricing strategy with the highest expected rate of return, then the best strategy is to "raise price." The most significant draw-

		Economy		
		Expansion	Stability	Contraction
	Raise price	25	15	–10
Strategy	No change	15	20	–5
	Lower price	15	0	5

FIGURE 14.14 Payoff matrix for pricing strategies under alternative states of nature.

		Economy				
		Expansion	Stability	Contraction	μ_i	σ_i
	Raise price	25(0.35)	15(0.5)	–10(0.15)	14.75	11.34
Strategy	No change	15(0.35)	20(0.5)	–5(0.15)	9.5	9.86
	Lower price	15(0.35)	0(0.5)	5(0.15)	6	6.8

FIGURE 14.15 Decision making under risk: expected values and standard deviations of returns for each pricing strategy with different probabilistic outcomes.

back of this decision is that it fails to consider management's attitude toward risk.

A more complete examination of the alternative strategies under different states of nature requires an examination of the risk associated with each strategy. In addition to the expected rates of return from each strategy, Figure 14.15 summarizes the standard deviations of the expected rates of return as a measure of the riskiness of each strategy. An examination of the payoff matrix reveals that while a "raise price" strategy has the greatest expected rate of return, it is also the most risky as measured by the standard deviation. By contrast, a "lower price" strategy is the least risky, but it also has the lowest expected rate of return. Clearly, the selection of the optimal strategy cannot be determined on the basis of a comparison of the expected rates of return and risk alone. In this and many similar situations, it is necessary to examine management's attitude toward risk before determining management's optimal strategy, as the following problem i llustrates.

Problem 14.13. Consider again the payoff matrix from three possible pricing strategies given three different states of nature in Figure 14.15. The payoffs in the matrix represent the firm's expected rate of return from each strategy. Suppose that management requires at least the equivalent of a 7% risk-free rate of return. Management's attitude toward risk is summarized by the risk–return indifference curve

$$\mu_i = 6 + 2^{\sigma_i}$$

where μ_i represents the expected rate of return and σ_i is the standard deviation of the expected rates or return. In this case, management's risk–return indifference curve summarizes the different combinations of expected rates of return and standard deviations for which management is indifferent between accepting and rejecting a particular price strategy.

a. Verify that management's risk-free rate of return is 7%.
b. On the basis management's risk–return indifference curve, determine the firm's most preferred pricing strategy.

Solution

a. A strategy with zero risk requires that $\sigma_i = 0$. Thus, the risk-free rate of return is

$$\mu_i = 6 + 1.2^{\sigma_i} = 6 + 1.2^0 = 6 + 1 = 7\%$$

For a risk-free strategy, management will accept any strategy with an expected rate of return greater than 7% and will reject any strategy with an expected rate of return less than 7%. Management will be indifferent between accepting and rejecting any pricing strategy that promises an expected rate of return of exactly 7%.

b. To determine the optimal pricing strategy for this firm, consider the rates of return for which the firm would be indifferent accepting or rejecting a particular pricing strategy.

$$\mu_{\text{Raise price}} = 6 + 1.2^{\sigma_i} = 6 + 1.2^{11.34} = 6 + 7.91 = 13.91\%$$

$$\mu_{\text{No change}} = 6 + 1.2^{\sigma_i} = 6 + 1.2^{9.86} = 6 + 6.04 = 12.04\%$$

$$\mu_{\text{Lower price}} = 6 + 1.2^{\sigma_i} = 6 + 1.2^{6.8} = 6 + 3.45 = 9.45\%$$

Management will reject both the "no change" and "lower price" strategies because the respective risk–indifferent rates of return (12.04 and 9.45%) are greater than the respective expected rate of return (9.5 and 6%). On the other hand, management will accept the "raise price" strategy because the risk-indifferent rate of return (13.91%) is less than the expected rate of return (14.75%). These solutions are illustrated in Figure 14.16.

The shaded area in Figure 14.16 indicates the combinations of expected rates of return at alternative standard deviations for which the pricing strategy is acceptable. Strategies with expected rates of return in the unshaded region are rejected. The reader will verify that only the

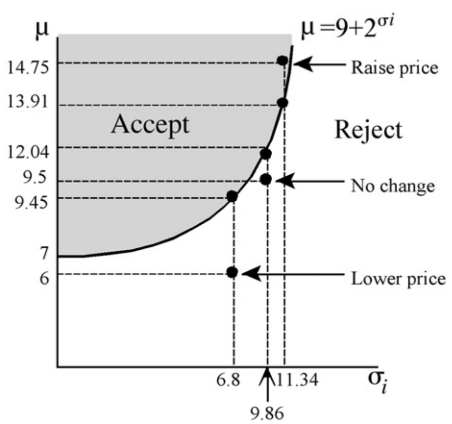

FIGURE 14.16 Diagrammatic solution to problem 14.13.

expected rate of return from the "raise price" strategy is located in the shaded region.

The situations considered thus far have assumed that probabilities could be objectively or subjectively assigned to each outcome from alternative strategies under alternative states of nature. It was also demonstrated that it is not sufficient to determine an optimal strategy based solely on expected values of the outcomes. It is necessary to consider not only the relative risk of each strategy, but also the decision maker's attitude toward risk. But what if it is not possible to assign probabilities to alternative outcomes? What other rational decision criteria are available to the manager in the presence of complete ignorance?

The Laplace decision criterion requires only a minor conceptual modification to the procedure just outlined. Decision making under conditions of complete ignorance assumes that the probabilities of the possible outcomes are not known, nor can they be inferred. The Laplace decision criterion asserts that since the probabilities of the outcomes are unknown, all possible outcomes must be assumed to be equally likely. This assumption effectively transforms the problem of decision making under uncertainty into decision making under risk. Consider again the situation depicted in Figure 14.15. The Laplace decision criterion assumes that the probability of each of these payoffs is one-third. The revised expected values and coefficients of variation from each strategy are summarized in Figure 14.17.

Of course, assuming equal probabilities for each outcome under the alternative states of the economy is equivalent to asserting that the expected value of the payoffs is a simple average of the outcomes. The information summarized in Figure 14.17 indicates that a "no change" strategy has the same expected return as a "raise price" strategy, but is less risky. A "lower price" strategy is the least risky, but also has the lowest expected rate of return. To determine which strategy will be adopted, it is necessary to consider the manager's attitude toward risk.

Definition: The Laplace decision criterion transforms decision making under complete ignorance to decision making under risk by assuming that all possible outcomes are equally likely.

		Economy				
		Expansion	Stability	Contraction	μ_i	σ_i
	Raise price	25(0.333)	15(0.333)	−10(0.333)	10	14.71
Strategy	No change	15(0.333)	20(0.333)	−5(0.333)	10	10.79
	Lower price	15(0.333)	0(0.333)	5(0.333)	6.67	6.23

FIGURE 14.17 Laplace decision criterion: expected values and coefficients of variation for each pricing strategy assuming equal probabilistic outcomes.

The primary deficiency of the Laplace decision criterion is the rather arbitrary manner in which the outcomes are assumed to be equally likely. The assumption of equiprobability of the different states of nature is especially problematic in the short run because it is oblivious to current conditions and known circumstances. Because of this, the Laplace decision criterion is more appropriate for strategic decisions, especially by larger firms that are better able to afford the cost of selecting a nonoptimal strategy.

WALD DECISION CRITERION

The Wald decision criterion is yet another rational approach to decision making under conditions of complete ignorance. The Wald decision criterion is analogous to the maximin, or secure, strategy for one-shot, simultaneous-move games as discussed in Chapter 13. The Wald decision criterion has also been called the strategy of extreme pessimism. A manager who employs the Wald (maximin) decision criterion will examine all possible payoffs associated with alternative strategies under different states of nature and will choose the strategy that results in the largest payoff from among the worst possible payoffs.

Definition: The Wald (maximin) decision criterion is a decision-making approach in the presence of complete ignorance that involves the selection of the largest payoff from among the worst possible payoffs.

Application of the Wald decision criterion is illustrated in Figure 14.18, which replicates the strategies and payoffs summarized in Figure 14.12. Figure 14.15 also summarizes the minimum (m) and maximum (M) payoffs from each strategy. The best of the worst (maximin) payoffs and the best of the best (maximax) payoffs are identified with asterisks (*). The maximax payoff is also identified because it will serve as a useful counterpoint to the maximin payoff in the subsequent discussion of the *Hurwicz decision criterion*.

The Wald decision criterion represents an extremely risk-averse approach to decision making in the presence of complete ignorance. In essence, the Wald decision criterion attempts to maximize management's

		Economy				
		Expansion	Stability	Contraction	**m**	**M**
Strategy	Raise price	25	15	–10	–10	25*
	No change	15	20	–5	–5	20
	Lower price	15	0	5	0*	15

FIGURE 14.18 Wald (maximin) decision criterion.

feelings of security; in Figure 14.18 the indicated solution is a "lower price" strategy with a maximin payoff of 0, which stands in contrast to the selection of a "no change" pricing strategy obtained by using the Laplace decision criterion.

How are we to assess the wisdom of a maximin approach to the decision-making process? What are the benefits and drawbacks of adopting such a conservative approach? As with other decision criteria considered, the answer ultimately turns on management's attitudes toward risk. In the situation depicted in Figure 14.18, only a "lower price" strategy avoids an economic loss in the event of an economic contraction. For any other state of the economy, however, a "lower price" strategy results in the lowest of all possible payoffs. In the end, only management can decide whether a low rate of return under the best economic conditions is worth the feeling of financial security that comes with the knowledge that the firm will be able to weather an economic downturn. The managements of what types of firm are most likely to adopt the Wald decision criterion? Most likely, the firms that are least able to weather financial reversals, such as small, start-up companies whose very survival will often depend on management's ability to avoid losses.

While the maximin strategy represents an extremely pessimistic approach to the decision-making process, a maximax strategy by contrast is extremely optimistic. Managers who use this approach will select as optimal that strategy that promises the best of the best of all possible outcomes. In the situation depicted in Figure 14.18, the decision to raise price represents one such maximax strategy. But, how likely is it that this, or any, firm would knowingly adopt such a strategy? The selection of a maximax strategy suggests that managers are risk lovers who are willing to gamble with the firm's assets in the hope of a big payoff, which in Figure 14.18 occurs with economic expansion. Under the other two phases of the business cycle, this firm will earn the lowest possible payoff. Since managers are ultimately responsible to the shareholders, it is very unlikely that such a strategy would ever be adopted. So why is it presented here? Minimax and maximax decision criteria are two extreme examples of the Hurwicz decision criteria.

HURWICZ DECISION CRITERION

We introduced the Wald decision criterion as a rather mechanistic approach to the selection of an optimal strategy by extremely risk-averse managers under conditions of complete ignorance. The manager using the maximin approach will select the strategy that results in the best of the worst possible outcomes. Unfortunately, the Wald approach makes no effort to explicitly incorporate management's attitude toward risk when it is not possible to assign probabilities to each outcome.

Although it is not possible to estimate the probabilities of all possible outcomes, the Hurwicz decision criterion is an attempt to incorporate the decision maker's attitude toward risk into the Wald decision criterion by creating a decision index for each strategy. This index is a weighted average of the maximum and minimum payoff from each strategy. These weights are called *coefficients of optimism*. The equation for estimating the Hurwicz decision index for each strategy is

$$D_i = \alpha M_i + (1-\alpha)m_i \tag{14.18}$$

where D_i is the decision index, M_i is the maximum payoff from each strategy, m_i is the minimum payoff from each strategy, and α is the coefficient of optimism. The optimal strategy using the Hurwicz decision criterion has the highest value for D_i.

Definition: The Hurwicz decision criterion is a decision-making approach in the presence of complete ignorance in which the optimal strategy is selected based on a decision index calculated from a weighted average of the maximum and minimum payoff of each strategy. The weights, which are called coefficients of optimism, are measures of the decision maker's attitude toward risk.

The value of the coefficient of optimism, which ranges in value from 0 to 1, represents management's subjective attitude toward risk. When $\alpha = 0$, the decision maker is completely pessimistic about the outcomes. When $\alpha = 1$, the decision maker is completely optimistic about the outcomes. Figure 14.19 summarizes the estimated values of the Hurwicz indices for selected values of α between 0 and 1. Consider, for example, a relatively pessimistic manager with a coefficient of optimism of $\alpha = 0.3$. From the maximum and minimum payoffs summarized in Figure 14.12, the Hurwicz decision index for a "raise price" strategy is

$$\begin{aligned} D_i &= \alpha M_i + (1-\alpha)m_i \\ &= 0.3(25) + (1-0.3)(-10) = 0.5 \end{aligned}$$

The reader should verify that when $\alpha = 0$ the optimal strategy under the Hurwicz decision is identical to the optimal strategy that would be selected by using the extremely pessimistic Wald (maximin) decision criterion. Moreover, when $\alpha = 1$, the optimal strategy under the Hurwicz decision criterion is identical to the optimal strategy obtained by using the maximax decision criterion. Figure 14.19 identifies the optimal strategies from the highest values for D_i with an asterisk. For values for $\alpha < 0.5$, the optimal (risk-averse) decision criterion is the "lower price" strategy. For values of $\alpha > 0.5$, the optimal (risk-loving) decision criterion is a "raise price" strategy. When $\alpha = 0.5$, the decision maker is indifferent to the different pricing strategies.

The Hurwicz decision criterion is superior to the Wald decision criterion because it forces managers to confront their attitudes toward risk. More-

α =	**0.0**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
Raise price	−10	−6.5	!3	0.5	4	7.5*	11*	14.5*	18*	21.5*	25*
No change	−5	−2.5	0	2.5	5	7.5*	10	12.5	15	17.5	20
Lower price	0*	1.5*	3*	4.5*	6*	7.5*	9	10.5	12	13.5	15

FIGURE 14.19 Estimated Hurwicz *D* values for selected values of α, the coefficient of optimism.

over, it forces managers to be consistent when they are considering the relative merits of alternative strategies. Of course, one drawback to this approach is the possible negative impact on company earnings should management's sense of optimism prove to be misplaced. Of course, this criticism might be leveled at any decision criterion that involves the subjective determination of probabilistic outcomes. In spite of this, the Hurwicz decision criterion does represent a conceptual improvement over the somewhat arbitrary Wald decision criterion.

SAVAGE DECISION CRITERION

The Savage decision criterion, which is sometimes referred to as the *minimax regret criterion*, is based on the opportunity cost (or regret) of selecting an incorrect strategy. In this instance, opportunity costs are measured as the absolute difference between the payoff for each strategy and the strategy that yields the highest payoff from each state of nature. Once these opportunity costs have been estimated, the manager will select the strategy that results in the minimum of all maximum opportunity costs.

Definition: The Savage decision criterion is used to determine the strategy that results in the minimum of all maximum opportunity costs associated with the selection of an incorrect strategy.

Figure 14.20 illustrates the calculations of the opportunity costs for the payoffs summarized in Figure 14.12. For example, the maximum possible payoff during an economic expansion is 25 for a "raise price" strategy. The absolute difference between the maximum payoff and the payoffs from each strategy during an economic expansion are calculated and summarized in each cell of the matrix. Figure 14.20 summarizes the maximum regret (opportunity cost) from each strategy. The minimum of these maximum opportunity costs, which is identified with an asterisk, is the strategy that will be selected by means of the Savage decision criterion.

Neither overly optimistic nor overly pessimistic, the Savage decision criterion is most appropriate when management is interested in earning a satisfactory rate of return with moderate levels of risk over the long term. Thus, the Savage decision criterion may be more appropriate for long-term capital investment projects.

		Economy			
		Expansion	Stability	Contraction	***Maximum regret***
Strategy	Raise price	|25 –25| = 0	|15 –20| = 5	|–10 –5| = 15	15
	No change	|15 –25| = 10	|20 –20| = 0	|–5 –5| = 10	10*
	Lower price	|15 –25| = 10	|0 –20| = 20	|5 –5| = 0	20

FIGURE 14.20 Savage regret matrix.

MARKET UNCERTAINTY AND INSURANCE

Markets operate best when all parties have equal access to all information regarding the potential costs and benefits associated with an exchange of goods or services. When this condition is not satisfied, then uncertainty exists and either the buyer or the seller may be harmed, which will result in an inefficient allocation of resources. In this section, we will examine some of the problems that arise in the presence of market uncertainty.

ASYMMETRIC INFORMATION

For markets to operate efficiently, both the buyer and the seller must have complete and accurate information about the quantity, quality, and price of the good or service being exchanged. When uncertainty is present, market participants can, and often do, make mistakes. An important cause of market uncertainty is *asymmetric information*. Asymmetric information exists when some market participants have more and better information than others about the goods and services being exchanged. An extreme example of the problems that might arise in the presence of asymmetric information is fraud. The reader will recall from Chapter 13 the discussion of the "snake oil" salesman, who traveled from frontier town to frontier town in the American West selling bottles of elixirs promising everything from a cure for toothaches to a remedy for baldness. Of course, these claims were bogus, but by the time customers realized that they had been "had" the snake oil salesman was long gone. Had the customer known that the elixir was worthless, the transaction would never have taken place.

In the extreme case, the knowledge that, some market participants had improperly exploited their access to privileged information could result in a complete breakdown of the market. In insider trading, for example, some market participants have access to classified information about a firm whose shares are publicly traded. Thus an executive who discovers that senior management of his firm plans to merge with a competitor, which will result in an increase in the firm's stock price, might act on this information by buying shares of stock in his own company. This person is guilty of insider trading. When insider trading is pervasive, rational investors who are not privy to privileged information may choose not to participate at all, rather than to put themselves at risk of buying or selling shares at the wrong price.

The uncertainty arising from asymmetric information affects managerial decisions as well. The reader will recall from Chapter 7, for example, that a profit-maximizing competitive firm will hire additional workers as long as the additional revenue generated from sale of the increased output (the marginal revenue product of labor) is greater than the wage rate. The marginal revenue product of labor is defined as the price of the product times the marginal product of labor, $P \times MP_L$. But how is the manager to know the potential productivity of a prospective job applicant? This is a classic example of asymmetric information. The prospective job applicant has much better information than the manager about his or her skills, capabilities, integrity, and attitude toward work. Since the potential cost to the firm of hiring an unproductive worker may be very high, managers will take whatever reasonable measures are necessary to rectify this asymmetry. This is why firms require job applicants to submit résumés, college transcripts, letters of recommendations, and so on. The firm's human resources officer may require job applicants to be interviewed by responsible professionals within the firm. Firms may also conduct background and credit checks, require applicants to sit for examinations to evaluate job skills, mandate probationary periods prior to full employment, and so forth.

ADVERSE SELECTION

The problem of *adverse selection* arises whenever there is asymmetric information. The classic example of adverse selection is the used-car market (Akerlof, 1970). A person with a used car to sell has the option of selling the vehicle to a used-car dealer or selling it privately. For simplicity, assume that all the used cars for sale are similar in every respect (age, features, etc.) except that half are "lemons" (bad cars) and the others are plums (good cars). Finally, suppose that potential buyers are willing to pay $5,000 for a plum and only $1,000 for a lemon.

Potential buyers have no way of distinguishing between lemons and plums. Since there is a fifty-fifty chance of getting a lemon, the expected market price of the used car is $3,000. Since only the sellers know whether their cars are lemons, there is a problem of asymmetric information. The seller has the option of selling to a used-car dealer or selling privately. If a lemon is sold to the used-car dealer for $3,000, then the seller will extract $2,000 at the expense of the buyer, while if a plum sells for $3,000, then the buyer will extract $2,000 at the expense of the seller. Thus, it is in the best interest of lemon owners to sell to used-car dealers, while it is in the best interest of plum owners to sell privately.

Buyers of used cars have the choice of buying from a used-car dealer or buying directly from an owner. Of course, buyers come to realize that probability of buying a lemon from a used-car dealer is greater than from buying from the owner directly. Thus, the used-car dealer price will fall. This will further exacerbate matters, since it will create an even greater incentive for plum owners to avoid the used-car market and sell privately. In the end, only lemons will be available from used-car dealers. In this case, the lemons drive the plums out of the market. This is an example of adverse selection. Here, the market has adversely selected the product of inferior quality because of the presence of asymmetric information.

Definition: In the presence of asymmetric information, adverse selection refers to the process in which goods, services, and individuals with economically undesirable characteristics tend to drive out of the market goods, services, and individuals with economically desirable characteristics.

The problem of adverse selection is particularly problematic in the market for insurance. As discussed earlier, risk-averse individuals purchase insurance to eliminate the risk of catastrophic financial loss in exchange for premium payments that are small relative to the potential loss. The problem confronting an insurance company is that it is difficult to distinguish high-risk from low-risk individuals. One possible solution would be for insurance companies to charge an insurance premium that is a weighted average of the premiums charged to individuals falling into different risk categories. In this case, high-risk individuals will purchase insurance policies while low-risk individuals will not. As a result, the insurance company will have to revise upward its insurance premium just to break even.

As an illustration of adverse selection in the insurance market, consider a firm that sells automobile collision insurance to residents of a particular area. The insurance company has identified two, equal-sized groups of high-risk and low-risk individuals. The insurance company has decided that the probability of an automobile accident is $p = 0.1$ for a member of the high-risk group and only $p = 0.01$ for a member of the low-risk group. If there are 100 people in each group, this is tantamount to an average of 10 automobile accidents per year for the high-risk group compared with one for the low-risk group. Suppose that the average repair bill per automobile accident is $1,000. If the insurance premium charged is the expected average repair bill loss, then the firm should charge the high-risk group 0.1($1,000) = $100 per year and the low-risk group 0.01($1,000) = $10 per year. If it is not possible for the insurance company to identify the members of each group, then the insurance company could decide to charge a premium based on the average risk, that is, 0.5($100) + 0.5($10) = $55.

The situation just described gives rise to the problem of adverse selection. If the insurance company charges a premium of $55, then some members of the low-risk group will opt not to purchase insurance. If 50 members of the low-risk group decide to withdraw from the insurance

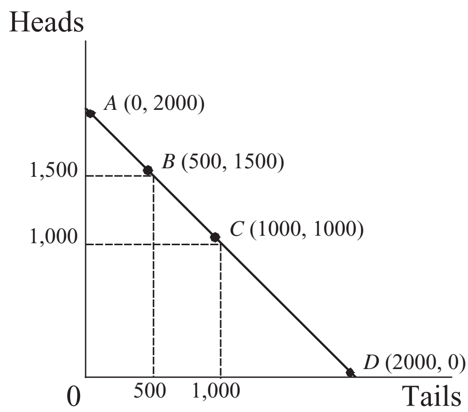

FIGURE 14.21 Fair-odds line for different states of nature.

market, then the total pool of individuals buying insurance falls from 200 to 150. As a result, the premium charged will increase to 0.67($100) + 0.33($10) = $70.3. Of course, some of the remaining individuals in the low-risk group will find that this premium is too high and will, in turn, withdraw from the insurance market. This process will continue until, in the end, only the most risk-averse individuals continue to buy insurance or, which is more likely, only members of the high-risk group remain.

FAIR-ODDS LINE

It is possible to analyze the problem of adverse selection by recasting individuals' attitudes toward risk within the framework of state-dependent indifference curves.[1] Consider again the situation in which an individual is offered a fair gamble on the flip of a coin. Suppose that the individual has $1,000. The person can bet all or part of this amount on the flip of a coin. If the coin comes up "heads," then the individual wins $1 for every $1 wagered. If the coin comes up "tails," then the individual loses $1 for every $1 wagered. Figure 14.21 illustrates the results of alternative wagers from this fair gamble. The horizontal axis represents the individual's money holdings if the coin comes up tails, while the vertical axis represents the individual's money holdings if the coin comes up heads. In a broader sense, the horizontal and vertical axes of Figure 14.21 may be thought of as the outcomes of two probabilistic states of nature. Point *C* in Figure 14.21 identifies the individual's money holdings on a decision not to bet. That is, regardless of the results of the flip of the coin, the individual will still have a cash "endowment" of $1,000, since no amount was placed at risk.

[1] For a detailed discussion of indifference curves see, for example, Walter Nicholson, Microeconomic Theory: Basic Principles and Extensions, 6th ed. (Font Worth: The Dryden Press, 1995), Chapter 3.

Suppose that the individual decides to wager \$500 on the flip of the coin. If the coin comes up heads, then the individual wins \$500. If the coin comes up tails, then the individual loses \$500. Point B in Figure 14.21 illustrates the possible outcomes of this bet. If the individual loses the wager, then his or her endowment is reduced to \$500. On the other hand, if the individual wins the wager, his or her endowment is increased to \$1,500. This combination of outcomes is identified in the parentheses at point B. Alternatively, if the individual wagers the entire \$1,000, then the possible combination of outcomes corresponds to point A, where the individual is left penniless if the coin comes up tail but has an endowment of \$2,000 if the coin comes up heads. What about the points in Figure 14.21 below C, such as point D? Points below point C represent a reversal of the terms of the wager (i.e., tails wins and heads loses).

The situation depicted in Figure 14.21 is analogous to the budget constraint introduced in Chapter 7 in that the endowments define the individual's consumption possibilities. Figure 14.21 is referred to as the individual's *fair-odds line*. In general, whenever the expected value of a wager is zero, then the gamble is said to be actuarially fair. A gamble is said to be fair if its expected value is zero. In the foregoing example, if the individual decides not to wager any amount, he or she is left with the initial endowment of \$1,000. If the individual decides to wager some amount, the expected value of the bet is zero, in which case the expected value of the endowment is still \$1,000.

The fair-odds line in Figure 14.21 is summarized in Equation (14.19), which represents an actually fair gamble where p is the probability of a monetary gain if the individual wins the bet and $(1 - p)$ is the probability of a monetary loss if the individual loses the bet.

$$pW + (1 - p)L = 0 \tag{14.19}$$

The slope of the fair-odds line is given as the monetary gain divided by the monetary loss from a fair gamble. Suppose, for example, that the individual places a wager of \$500. If the individual wins the bet, his or her endowment will increase to \$1,500 (i.e., the amount of the gain is $W = \$500$). On the other hand, if the individual loses the bet, his or her endowment is reduced to \$500 (i.e., $L = -\$500$). This is illustrated as a move from point C to point B in Figure 14.21. Solving Equation (14.19), we obtain

$$\frac{W}{L} = \frac{1 - p}{p} \tag{14.20}$$

The reader should verify that the budget constraint depicted in Figure 14.21 had a slope of –1. The reader should also verify that, in general, an increase in the probability of winning means that for the gamble to remain fair, the amount of the win will have to decrease. For example, when $p = 0.5$, then $W/L = -(1 - 0.5)/0.5 = -1$. If the probability of winning increases

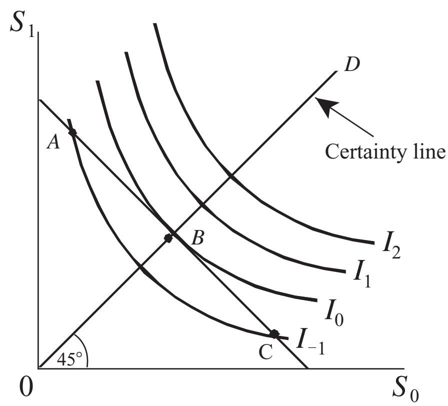

FIGURE 14.22 Indifference map of risk-averse preferences.

to $p = 0.75$, then $W/L = -(1 - 0.75)/0.75 = -0.25/0.75 = -0.33$. Similarly, if the probability of losing increases, the amount of the win will have to increase for the gamble to remain fair. These three situations are illustrated in Figure 14.22.

STATE PREFERENCES

The indifference curve framework can also be used to identify an individual's attitudes toward risk. In this case, however, the two goods that are normally identified along the horizontal and vertical axes are replaced with different combinations of state-dependent consumption levels that yield equal levels of utility. The shapes of these indifference curves reflect the individual's behavior when confronted with risky situations.

In Figure 14.22, which illustrates the case of an individual with risk-averse preferences, S_1 and S_0 represent two different states of nature. It will be recalled that an individual with risk-averse preferences will never accept a fair gamble with an expected value equal to zero. This is because a risk-averse individual will always prefer a certain sum to an uncertain sum with the same expected value. Thus, the indifference curves of an individual with risk-averse preferences are convex with respect to the origin.

The individual described in Figure 14.22 will prefer a consumption level corresponding to point B to any other point on the fair-odds line. Consumption levels that correspond to points A and C are found on an indifference curve that is closer to the origin, which yields a lower level of utility. The point of tangency between the fair-odds line and the indifference curve I_0 at point B represents the highest level of utility that this individual can attain with a given endowment. At point B the slope of the indifference curve is $-(1 - p)/p$. Line $0D$, which represents the locus of all such fair-odds tangency points at fair odds, is called the *certainty line*, which is analytically equivalent to the income consumption curve in utility theory and the expansion path in production theory. The certainty line represents equal consumption in either state of nature.

The choices confronting a person with risk-neutral preferences are illustrated in Figure 14.23. Points A, B, and C all yield the same level of utility, since the indifference curve I_0 corresponds to the fair-odds line. A risk-neutral individual is indifferent between a certain sum and an uncertain sum with the same expected value. Finally Figure 14.24 illustrates the case of a risk-loving individual. A risk lover will always accept a fair gamble with an expected value equal to zero. Risk lovers have indifference curves that are concave with respect to the origin. Accepting a fair gamble will move the individual away from point B and result in a higher level of utility. In fact, concave indifference curves will invariably result in a corner solution, such as points A and C, in which the individual will gamble the total amount of his or her endowment.

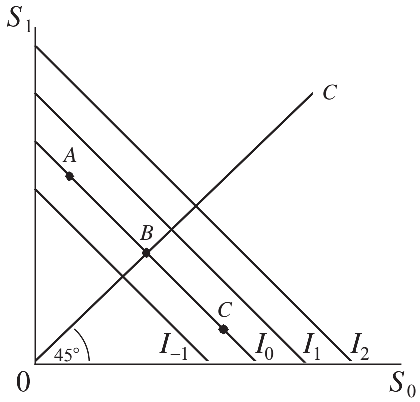

FIGURE 14.23 Indifference map of risk-neutral preferences.

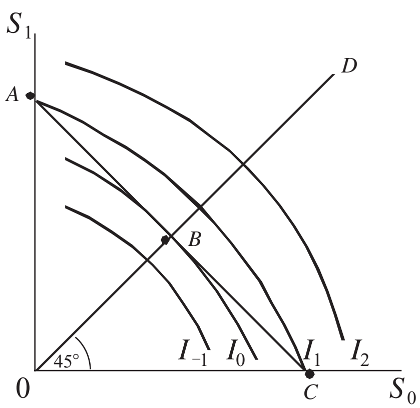

FIGURE 14.24 Indifference map of risk-loving preferences.

INSURANCE PREMIUMS

The state preferences model just presented can be used to analyze the demand for insurance. We will initially assume that insurance is provided at zero administrative cost. We will also assume that insurance is offered at actuarially fair terms. In the event of an adverse state of nature, the insurance company agrees to pay out the full amount of the loss, while in a favorable state of nature the insurance company pays nothing. The insurance premium is equal to the expected value of the payout, that is,

$$P = pL \tag{14.21}$$

where $(1 - p)$ is the probability of an adverse state of nature, such as the financial loss arising from an accident (L), and p is the probability of a favorable state of nature. In our example of automobile collision insurance, if the insurance policy provides \$1,000 annual coverage and the probability of an automobile is 10%, then an actuarially fair premium is \$100 per year. For each additional \$100 of coverage the additional premium will be \$10. Figure 14.25 illustrates the situation of an individual buying fair insurance.

In Figure 14.25 the individual's endowment is at point A. Suppose that the individual wishes to equalize his or her consumption in either state of nature. This will involve moving along the fair-odds line from point A to point B on the full insurance line $0D$. This will involve the payment of an insurance premium AC in exchange for an insurance payout of CB should the adverse event occur. In general, risk-averse individuals will purchase full insurance offered at fair odds. But what if insurance is offered at unfair odds? This situation is depicted in Figure 14.26.

Thus far we have assumed that insurance companies operate at zero cost. This assumption allowed us to assume that insurance companies are able to provide insurance at actuarially fair terms. This assumption is obviously

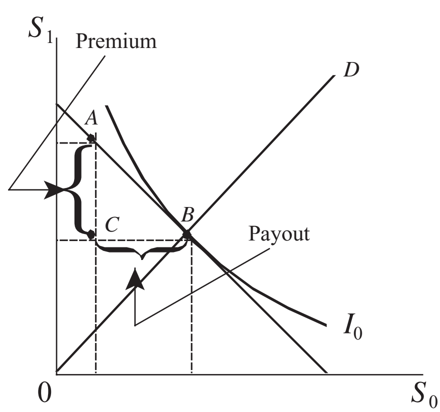

FIGURE 14.25 Full insurance at fair odds.

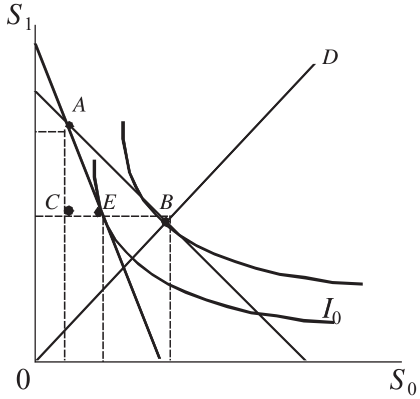

FIGURE 14.26 Partial insurance at unfair odds.

unrealistic, since insurance companies are analytically subject to the same long-run and short-run production considerations faced by any other firm. Thus, since the provision of insurance, or any other good or service, is not free, we must modify our analysis to recognize that the premium charged is not equal to the expected payout. When insurance is not provided at fair odds, the fair-odds line will pivot in a clockwise direction around the individual's initial endowment. In Figure 14.26, this is illustrated by the individual's new budget line that passes through points *A* and *E*.

Inspection of Figure 14.26 reveals that when insurance is not provided at actuarially fair terms, the individual will purchase partial insurance *CB* – *EB* for the same insurance premium *AC*. In other words, when insurance is not provided at actuarially fair terms, a risk-averse individual will nonetheless purchase partial coverage even though the premium payments are greater than the expected loss. It is evident from Figure 14.26 that insurance provided at unfair odds will move the individual's consumption level in either state of nature to a lower indifference curve than would be the case if insurance were provided at fair odds. As before, in equilibrium the individual's marginal rate of substitution between the state-dependent consumption levels is equal to the slope of the fair-odds (budget) constraint, although consumption levels will obviously be less than in a favorable state of nature.

We are now in a position to formally analyze the problem of adverse selection arising from asymmetric information. Recall from the automobile collision insurance example that the problem of adverse selection arises when the insurance company is unable to distinguish individuals belonging to the high- and low-risk groups. In terms of the state preference model, Figure 14.27 illustrates the fair-odds lines of the high-risk group, the low-risk group, and the average market risk.

In Figure 14.27, the fair-odds lines of the high- and low-risk groups are F_H and F_L, respectively. The average-market fair-odds line is F_M. Figure 14.27

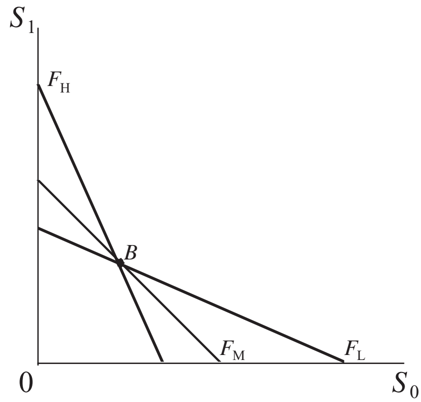

FIGURE 14.27 High-risk, low-risk, and average-market fair-odds lines.

assumes that both the high- and low-risk groups have the same initial endowment, with is indicated at point B. The different risks associated with each group are reflected in the slopes of the fair-odds lines, that is, $[-(1 - p_H)/p_H] < [-(1 - p_L)/p_L]$. This is because the probability that an individual in the high-risk group will have an accident $(1 - p_H)$ is greater than the probability that an individual in the low-risk group will have an accident $(1 - p_L)$.

The different risks faced by individuals in both groups are also reflected in the slopes of the indifference curves. Figure 14.28 illustrates the indifference curves for the high- and low-risk groups. Individuals belonging to the low-risk group are less likely to make a claim under an insurance policy than individuals belonging to the high-risk group. Thus, low-risk individuals will require greater compensation for a given reduction in consumption in a favorable state of nature. In Figure 14.28, low-risk individuals making a claim will require an additional amount AE in state of nature S_0, while high-risk individuals will require AC < AE. Thus, the indifference curve for the low-risk individual (I_L) is flatter than the indifference curve for the high-risk group (I_H).

The problem of adverse selection is illustrated in Figure 14.29. Note that the slope of the low-risk individual's indifference curve is flatter than the market-average fair-odds line, F_M, at the initial endowment point B. In exchange for a sure amount in a favorable state of nature, AB, the individual is able to obtain only AC coverage in an adverse state of nature. But, to be as well off as at point B, the low-risk individual would require an additional amount CE in an adverse state of nature. Thus, the low-risk individual would be better off with no insurance at all.

In general, adverse selection is more likely to be a problem when the market consists of a high proportion of high-risk individuals, which has the effect of moving the average-market fair-odds line closer to the fair-odds

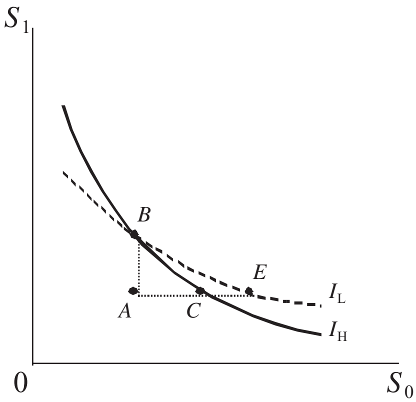

FIGURE 14.28 The indifference curve of a low-risk individual is flatter than the indifference curve of a high-risk individual.

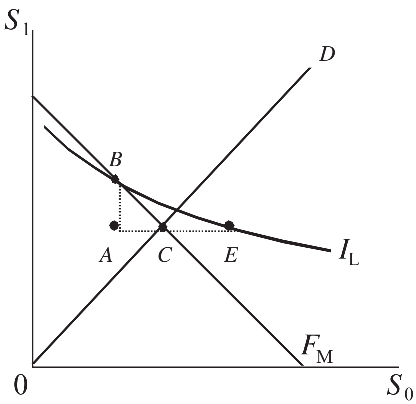

FIGURE 14.29 Adverse selection: low-risk individuals choose not to purchase insurance.

line for the high-risk group (F_H) in Figure 14.27. Adverse selection is also more likely to be a problem if there is a large gap in the perceptions toward risk of the high- and low-risk groups. Adverse selection will be less a problematic if some individuals are extremely risk averse. In practice, it is common for insurance companies to differentiate candidates for insurance to capture different attitudes toward risk. Thus, differential premiums based on age, sex, occupation, lifestyle, and domicile are a commonly found in the insurance industry.

MORAL HAZARD

Another problem that arises in the presence of asymmetric information is the problem of *moral hazard*. We saw earlier that risk-averse individuals will purchase insurance to protect themselves against catastrophic financial losses. Of course, the probability that such catastrophic losses will occur is inversely related to individual efforts to avoid such losses. For example, the probability of having an automobile accident depends on how carefully one drives. Other things being equal, individuals tend to be more careful behind the wheel if they are not insured than when they are fully insured. The reason for this is that the insured knows that he or she will be fully compensated for damages incurred as a result of an accident. If an insured individual has a reduced incentive to be careful, a moral hazard is said to exist. Other examples of moral hazard include individuals who lead less-than-healthy lifestyles after obtaining health insurance, or doctors who are less than conscientious about administering medical care after obtaining medical malpractice insurance.

Definition: A moral hazard exists when insurance coverage causes an individual to behave in such a way that changes the probability of incurring a loss.

In general, a moral hazard exists when an individual can determine the probability of an undesirable outcome. To see this, consider the case of an insurance company that has estimated that the probability p that an automobile will be stolen. Ignoring administrative costs, the insurance company will provide coverage against automobile theft for the premium payment P in Equation (14.21). Now, suppose that an insured individual can determine the probability that his or her car will be stolen. Suppose, for example, that the insured is able to set $p = 1$. In this case, the insured individual is effectively attempting to use the insurance policy to obtain the price of a new car. Of course, if the insurance company knows this, automobile theft insurance will not be offered. In this case, a moral hazard exists because the insurance company does not, indeed cannot, know the probability that the insured will submit a claim.

The problem of moral hazard may be represented diagrammatically by means of the state preferences model. Figure 14.30 illustrates the amount of care that an individual exercises to avoid the probability of an adverse state of nature. The flatter the indifference curve, the greater the care an individual takes to avoid a loss. The indifference curves in Figure 14.30 associated with low and high probabilities of an adverse state of nature are identified as I_L and I_H, respectively. To understand why this is the case, we can ask ourselves the following question: How much will an individual be willing to sacrifice in an adverse state of nature to obtain a given amount in a favorable state of nature?

The answer to this question depends on how likely it is that the individual will experience the adverse state of nature, which, of course, depends on the actions of the individual. In Figure 14.30, for an extra amount of consumption in a favorable state of nature, AB, the careful individual is willing to sacrifice a larger amount in the adverse state of nature than would the careless individual. The reason for this is that the probability that an adverse

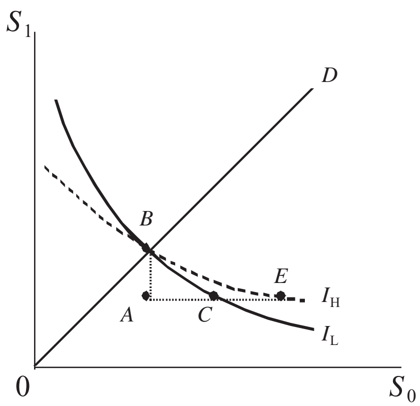

FIGURE 14.30 The slope of the state preference indifference curve is flatter when more care is taken to avoid an adverse state of nature.

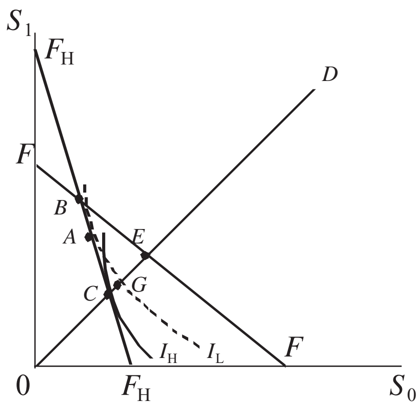

FIGURE 14.31 Moral hazard and partial insurance.

state of nature will occur is less because of the greater care exercised. The additional amount that the high-care individual is willing to sacrifice is given by the distance *CE*. Thus, the indifference curve I_H reflects the greater care that an individual takes to avoid a loss, compared with individuals who are less careful and are willing to sacrifice only *AC*.

Figure 14.31 illustrates the situation in which the individual's initial endowment is given at point *B* and the fair-odds line is given as *FF*. If an insured individual is able to increase the probability of an adverse state of nature by exercising less care, then the fair-odds line will pivot clockwise around point *B*. This is illustrated in Figure 14.31 as F_HF_H. Point *E* on the fair-odds line *FF* is no longer an equilibrium in the presence of a moral hazard, since no insurance company would offer such coverage at the premium *pL*.

In Figure 14.31 the new equilibrium at point *C* represents the individual's behavior along the new fair-odd line F_HF_H associated with the higher probability that the adverse state of nature will occur. An individual offered insurance along the new fair-odds line might obtain a higher level of utility by exercising greater care and purchasing partial insurance coverage. This situation is depicted at point A because I_L passes through the certainty equivalent (full insurance) line $0D$ at point G, which is above point C. The insured individual will be better off paying the amount CG, provided it does not represent a cost greater than the cost associated with exercising greater care to avoid the adverse state of nature.

Insurance companies attempt to reduce the problem of moral hazard by requiring insured individuals to share the losses that arise from an adverse state of nature by applying a deductible on all insurance claims. To be effective, the amount of the deductible should be no greater than the distance CG in Figure 14.31. Provided the deductible is not too large, an insured individual is likely to drive more carefully or choose a more healthy lifestyle when he or she is required to share the cost of an accident or illness.

CHAPTER REVIEW

Most economic decisions are made with something less than perfect information, and the consequences of these decisions cannot be known with any degree of precision. Moreover, the uncertainty of outcomes associated with those decisions increases with time. Most economic decisions are made under conditions of *risk* and *uncertainty*.

Risk involves choices with multiple outcomes in which the probability of each outcome is known or can be estimated. Uncertainty, on the other hand, involves multiple outcomes in which the probability of each one is unknown or cannot be estimated.

There are two sources of uncertainty. Uncertainty with *complete ignorance* refers to situations in which no assumptions can be made about the probabilities of alternative outcomes under different states of nature. Uncertainty with *partial ignorance* refers to situations in which the decision maker is able to assign subjective probabilities to possible outcomes. These subjective probabilities may be based on personal knowledge, intuition, or experience. Decision making under conditions of partial ignorance is effectively the same as decision making under risk. Uncertainty with complete ignorance requires alternative approaches to the decision-making process.

The most commonly used summary measures of uncertain, random outcomes are the *mean* and the *variance*. The expected value of random outcomes, such as profits, capital gains, prices, and unit sales, is called the *mean*. The mean is the weighted average of all possible random outcomes, where the weights are the probabilities of each outcome.

Risk may be measured as the dispersion of all possible payoffs. The most commonly used measure of the dispersion of possible outcomes is the *vari-*

ance. The variance is the weighed average of the squared deviations of all possible random outcomes from its mean, where the weights are the probabilities of each outcome. An alternative way to express the riskiness of a set of random outcomes is the *standard deviation*, which is the square root of the variance.

Neither the variance nor the standard deviation can be used to compare risk when there are two or more risky situations involving different expected values. The *coefficient of variation* is used to compare the relative riskiness of alternative outcomes. The project with the lowest coefficient of variation is the least risky.

Whether an individual undertakes a risky project will depend on the individual's attitude toward risk. An individual who prefers a certain payoff to a risky prospect with the same expected value is said to be *risk averse*. An individual who prefers the expected value of a risky prospect to its certainty equivalent is said to be a *risk lover*. Finally, an individual who is indifferent between a certain payoff and its expected value is *risk neutral*.

Generally speaking, most individuals are risk averse in accordance with the principle of the *diminishing marginal utility of money*. Most individuals, however, are not risk averse under all circumstances. It is not unusual to find that even extremely risk-averse individuals become risk lovers for "small" gambles, such as buying a lottery that costs far less than the expected value of winning.

Managers often evaluate equal or, equivalently, equal-lived capital investment projects, by calculating the net present values of net cash flows. *Risk-adjusted discount rates* are used in the calculation of net present values to compensate for the perceived riskiness of alternative capital investment projects. The greater the perceived risk, the higher will be the discount rate that will be used to calculate the net present value. The difference between the risk-free discount rate and the risk-adjusted discount rate is called the *risk premium*. The size of the risk premium will depend on the investor's attitude toward risk.

An alternative to the use of risk-adjusted discount rates for assessing capital investment projects is the *certainty-equivalent approach*. The certainty-equivalent approach incorporates risk directly into the net present value method by using the *certainty-equivalent coefficient* to modify expected net cash flows. As with the risk-adjusted discount rate approach, however, the certainty-equivalent method suffers from the shortcoming of the subjective determination of the certainty-equivalent cash flow. It is conceptually superior to the risk-adjusted discount rate approach, however, in that it explicitly considers the investor's attitude toward risk.

Decision making under conditions of uncertainty with complete ignorance requires rational decision-making criteria that do not rely on probabilistic outcomes. Four such rational decision criteria include the *Laplace* *criterion*, the *Wald* (*maximin*) *criterion*, the *Hurwicz criterion*, and the *Savage* (*minimax regret*) *criterion*.

The Laplace decision criterion transforms decision making under complete ignorance to decision making under risk by assuming that all possible outcomes are equally likely. The Wald (maximin) decision criterion selects the largest of the worst possible payoffs. The Hurwicz decision criterion involves the selection of an optimal strategy based on a decision index calculated from a weighted average of the maximum and minimum payoffs of each strategy. The weights, which are called coefficients of optimism, are measures of the decision maker's attitude toward risk. Finally, the Savage decision criterion is used to select a strategy that results in the minimum of all maximum opportunity costs associated with the selection of an incorrect strategy.

For markets to operate efficiently, both buyers and sellers must have complete and accurate information about the quantity, quality, and price of the good or service being exchanged. When uncertainty is present, market participants can, and often do, make mistakes. An important cause of market uncertainty is *asymmetric information*. Asymmetric information exists when some market participants have more and better information about the goods and services being exchanged. The problem of *adverse selection* arises whenever there is asymmetric information. In adverse selection, the interaction of buyers and sellers results in the market provision of goods and services with undesirable characteristics.

Another problem that arises in the presence of asymmetric information is called *moral hazard*. When obtaining information is costly, monitoring the behavior of the parties to a transaction becomes difficult. When the parties to a contract have an incentive alter their behavior from what was anticipated when the contract was entered into, a moral hazard exists.

KEY TERMS AND CONCEPTS

Adverse selection The process whereby, in the presence of asymmetric information, goods, services, and individuals with economically undesirable characteristics tend to drive out of the market goods, services, and individuals having economically desirable characteristics.

Beta coefficient (β) A measure of the price volatility of a given stock versus the price volatility of "average" stock prices.

Capital asset pricing model (CAPM) Establishes a relationship between the risk associated with the purchase of a stock and its rate of return. CAPM asserts that the required return on a company's stock is equal to the risk-free rate of return plus a risk premium.

Capital market line Summarizes the market opportunities available to an investor from a portfolio consisting of alternative combinations of risky and risk-free investments.

Certainty-equivalent approach Modifies the net present value approach to evaluating capital investment projects by incorporating risk directly into expected cash flows by means of a certainty-equivalent coefficient.

Certainty-equivalent coefficient The ratio of a risk-free net cash flow to its equivalent risky cash flow. The smaller the coefficient, the greater the perceived riskiness of an investment.

Coefficient of variation A measure used to compare risk of two or more outcomes when there are different expected values. It is calculated as the ratio of the standard deviation to the mean.

Fair gamble A gamble in which the expected value of the payoff is zero.

Hurwicz decision criterion A decision-making approach in the presence of complete ignorance an optimal strategy in which is selected based on a decision index calculated from a weighted average of the maximum and minimum payoff of each strategy. The weights, which are called coefficients of optimism, are measures of the decision maker's attitude toward risk.

Investor indifference curve Summarizes the combinations of risk and expected return in which the investor will be indifferent between a risky and a risk-free investment.

Laplace decision criterion A decision-making approach that transforms decision making under complete ignorance to decision making under risk by assuming that all possible outcomes are equally likely.

Mean The expected value of a set of random outcomes. The mean is the sum of the products of each outcome and the probability of its occurrence.

Moral hazard Exists when insurance coverage causes an individual to behave in such a way that change the probability of incurring a loss.

Risk The existence of choices involving multiple possible outcomes in which the probability of each outcome is known or may be estimated.

Risk-adjusted discount rate The discount rate used to calculate net present values to compensate for the perceived riskiness of an investment. The greater the perceived risk, the higher will be the discount rate that is used to calculate the net present value.

Risk aversion An individual who prefers a certain payoff to a risky prospect with the same expected value is said to be risk averse.

Risk loving Preferring the expected value of a payoff to its certainty equivalent.

Risk neutrality Indifference between a certain payoff and its expected value.

Savage decision criterion A decision-making approach in the presence of complete ignorance that involves the selection of the strategy that results in the minimum of all maximum opportunity costs. Opportunity costs are measured as the absolute difference between the payoff for each strategy and the strategy that yields the highest payoff for each state of nature.

Standard deviation The square root of the variance.

Uncertainty The existence of choices involving multiple possible outcomes in which the probability of each outcome is unknown and cannot be estimated.

Variance A measure of the dispersion of a set of random outcomes. It is the sum of the products of the squared deviations of each outcome from its mean and the probability of each outcome.

Wald (maximin) decision criterion A decision-making approach in the presence of complete ignorance in which one selects the largest from among the worst possible payoffs.

CHAPTER QUESTIONS

14.1 What is the difference between risk and uncertainty?

14.2 What are the most commonly used measures of risk?

14.3 Can uncertainty be estimated? If not, then why not? Explain.

14.4 When is the process of decision making under conditions of risk the same as the process of decision making under conditions of uncertainty?

14.5 Decision making under conditions of uncertainty with complete ignorance is never the same as decision making under conditions of uncertainty under partial ignorance. Do you agree? Explain.

14.6 What is the difference between the standard deviation and the coefficient of variation as a measure of risk? When would it be appropriate to use each one?

14.7 Risk-averse individuals will always reject a fair gamble. Do you agree? Explain.

14.8 Can the internal rate of return method discussed in Chapter 12 be used to determine the risk-adjusted discount rate?

14.9 Explain why many life insurance policies contain clauses stipulating that benefits will not to the heirs of a policyholder who commits suicide.

14.10 Explain why insurance companies charge higher premiums to male drivers between 18 and 25 years of age than for all other drivers.

14.11 What risk preferences are described by L-shaped indifference curves?

14.12 An individual with L-shaped indifference curves is indifferent to insurance offered at fair or unfair odds. Do you agree with this statement? Explain.

14.13 Briefly explain the following decision criteria and the conditions under which each might be used:

a. Laplace criterion
b. Wald (maximin) criterion
c. Hurwicz criterion
d. Savage criterion

14.14 Insurance companies require a deductible on all insurance claims to reduce costs and bolster profits. Do you agree? Explain.

14.15 Define adverse selection. Give an example.

14.16 Define moral hazard. Give an example.

14.17 How do deductibles on insurance claims address the problem of moral hazard?

CHAPTER EXERCISES

14.1 Illustrate, with the use of investor indifference curves, that project *A* is the most preferred project when the expected rates of return from the investment projects are $k_C > k_A > k_B$ and the risks associated with each project are $\sigma_C > \sigma_A > \sigma_B$.

14.2 Illustrate, with the use of investor indifference curves, that project *A* is the most preferred project when the expected rates of return from the investment projects are $k_A > k_B > k_C$ and the risks associated with each project are $\sigma_C > \sigma_A > \sigma_B$.

14.3 Rosie Hemlock offers Robin Nightshade the following wager. For a payment of $10, Rosie will pay Robin the dollar value of any card drawn from a standard deck of 52 cards. For example, for an ace of any suit Rosie will pay Robin $1. For an 8 of any suit Rosie will pay Robin $8. A ten or picture card of any suit is worth $10.

a. What is the expected value of Rosie's offer?
b. Should Robin accept Rosie's offer?

14.4 Suppose that capital investment project *X* has an expected value of $\mu_X = \$1{,}000$ and a standard deviation of $\sigma_X = \$500$. Suppose, also, that project *Y* has an expected value $\mu_Y = \$1{,}500$ and a standard deviation of $\sigma_Y = \$750$. Which is the relatively riskier project?

14.5 The management of Rubicon & Styx is trying to decide whether to advertise its world-famous hot sauce Sergeant Garcia's Revenge on television (campaign *A*) or in magazines (campaign *B*). The marketing department of Rubicon & Styx has estimated the probabilities of alternative sales revenues (net of advertising costs) using each of the two media outlets, summarized in Table E14.5.

a. Calculate the expected revenues from sales of Sergeant Garcia's Revenge from each advertising campaign.
b. What is the standard deviation of the distribution of profits from each advertising campaign?
c. Which advertising campaign appears relatively riskier?
d. Which advertising campaign should Rubicon & Styx select?

TABLE E14.5 Probabilities of alternative sales revenues for chapter exercise 14.5

Campaign *A* (television)		Campaign *B* (magazines)	
Sales, S_i	Probability	Sales, S_i	Probability
$5,000	0.20	$6,000	0.15
$8,000	0.30	$8,000	0.35
$11,000	0.30	$10,000	0.35
$14,000	0.20	$12,000	0.15

14.6 Suppose that Ted Sillywalk offers Will Wobble the fair gamble of receiving $500 on the flip of a coin showing heads and losing $500 on the flip of a fair coin showing tails. Suppose further that Will's utility of money function is

$$U = M^{1.2}$$

a. For positive money income, what is Will's attitude toward risk?
b. If Will's current income is $5,000, will he accept Ted's offer? Explain.

14.7 Mat Heathertoes has just inherited $10,000 from his Aunt Lobelia. Mat has decided to invest his inheritance either in 3-month Treasury bills, which yield a risk-free expected rate of return of 8%, or in shares of Hardbottle Company, which have an expected rate of return of 15%. Mat has analyzed Hardbottle's past performance and has determined that the standard deviation of returns is $3.50 per share. Mat's investment utility equation is

$$U = k_p - 100\sigma_p^2$$

where k_p and σ_p are the portfolio's expected return and standard deviation, respectively. How should Mat's investment be divided between 3-month Treasury bills and Hardbottle shares?

14.8 Harry Frogfoot is the proprietor of The Floating Log restaurant, which is located on the Delaware River near Frenchtown. Harry is considering expanding the dining area of his restaurant. The $150,000 cost of the investment is known with certainty. Harry has estimated that the expected cash inflows are $50,000 per year for the next 5 years.

a. Should Harry consider the investment if the discount rate is 8%?
b. Suppose that the riskiness of expected cash inflows was such that management requires a 25% rate of return. Should Harry consider this investment?

14.9 Suppose that you are given the information in Table E14.9 on cash flows and their probabilities for a proposed project.

If the discount rate is 0.0%, what is the expected value of the cash flows?

14.10 Suppose that the discount rate in Exercise 14.9 is 10.0%.

TABLE E14.9 Cash flows and probabilities for chapter exercise 14.9

Period 1		Period 2	
Probability	Cash flow	Probability	Cash flow
0.20	500	0.15	250
0.60	750	0.70	500
0.20	1,000	0.15	750

TABLE E14.11 Sales revenue expectations and probabilities for chapter exercise 14.11

Sales ($000s)	Probabilities
100	0.05
120	0.15
140	0.30
160	0.30
180	0.15
200	0.05

a. What is the expected value of the project?
b. If the initial investment was $1,000, what is net present value of this project?

14.11 Consider the sales revenue expectations and probabilities given in Table E14.11.

a. Calculate expected sales revenues.
b. Calculate the standard deviation of expected sales revenues.
c. Calculate the coefficient of variation.

14.12 Suppose that the equation for the risk–return indifference curve in Exercise 14.13 is

$$\mu_i = 3 + 2^{\sigma_i}$$

a. What is the new required risk-free rate of return?
b. What is the firm's optimal pricing strategy?

14.13 Suppose that the senior management of Red Wraith Enterprises is provided with the data for a proposed capital investment project given in Table E14.13.

a. Calculate the net present value of the proposed capital investment project if the risk-free discount rate is 10%.
b. On the basis of your answer to part a, should senior management of Red Wraith invest in this project?

TABLE E14.13 Data for proposed capital investment project for chapter exercise 14.13

Year	Cash flow	Certainty-equivalent coefficient
0	–$65,000	1.00
1	10,000	0.95
2	15,000	0.90
3	20,000	0.85
4	25,000	0.80
5	30,000	0.75

SELECTED READINGS

Akerlof, G. "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84 (1970), pp. 488–500.

Baumol, W. J. *Economic Theory and Operations Analysis*, 4th ed. Englewood Cliffs, NJ: Prentice Hall, 1977.

Bierman, H. S., and L. Fernandez. *Game Theory with Economic Applications*, 2nd ed. New York: Addison-Wesley, 1998.

Brigham, E. F., L. C. Gapenski, and M. C. Erhardt. *Financial Management: Theory and Practice*, 9th ed. New York: Dryden Press, 1998.

Davis, O., and A. Whinston. "Externalities, Welfare, and the Theory of Games." *Journal of Political Economy*, 70 (June 1962), pp. 241–262.

Dreze, J. "Axiomatic Theories of Choice, Cardinal Utility and Subjective Utility: A Review." In P. Diamond and M. Rothschild, eds., *Uncertainty in Economics*. New York: Academic Press, 1978, pp. 37–57.

Friedman, L. *Microeconomic Policy Analysis*. New York: McGraw-Hill, 1984.

Friedman, M., and L. Savage. "The Utility Analysis of Choices Involving Risk." *Journal of Political Economy*, 56 (August 1948), pp. 279–304.

Greene, W. H. *Econometric Analysis*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 1997.

Hirshleifer, J., and J. Riley. "The Analytics of Uncertainty and Information—An Expository Survey." *Journal of Economic Literature*, 57(4) (December 1979), pp. 1375–1421.

Hope, S. *Applied Microeconomics*. New York: John Wiley & Sons, 1999.

Knight, F. H. *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin, 1921.

Kunreuther, H. "Limited Knowledge and Insurance Protection." *Public Policy*, 24(2) (Spring 1976), pp. 227–261.

Pauly, M. "The Economics of Moral Hazard." *American Economic Review*, 58 (1968), pp. 531–537.

Schotter, A. *Free Market Economics: A Critical Appraisal*. (New York: St. Martin's Press, 1985).

Silberberg, E. *The Structure of Economics: A Mathematical Analysis*, 2nd ed. New York: McGraw-Hill, 1990.

Simon, H. "Theories of Decision-Making in Economics and Behavioral Science." American Economic Review, 49 (1959), pp. 253–283.

Varian, H. *Microeconomic Analysis*, 2nd ed. New York: W. W. Norton, 1984.

This Page Intentionally Left Blank

15

Market Failure and Government Intervention

Thus far, we have generally assumed that market transactions were, for the most part, free of government interference. In reality, however, government intervention in private transactions in the form of taxes and regulations is pervasive. Why? The cynical response to this question might be that legislators are more concerned with generating tax revenues to subsidize pork-barrel projects to curry the favor of one particular group of voters over another, or to attract the political and monetary support of special interest groups. While these explanations may be valid, the fact is that government intervention is often motivated by the failure of free markets to provide a socially optimal mix of goods and services. A socially optimal mix of goods and services may be defined as one in which the collective welfare of society has been maximized. Maximizing social welfare requires not only that the economy be producing efficiently given the productive resources available to it, but also that it be consuming efficiently. By this we mean that economy needs to be producing the goods and services that are most in demand by society.

Definition: Market failure occurs when private transactions result in a socially inefficient allocation of goods, services, and productive resources.

This chapter will examine three sources of market failure: market power, externalities, and public goods. Another source of market failure, asymmetric information, was discussed in Chapter 14. While the problems and potential solutions to the problems of market failure were touched on in earlier chapters, this chapter will focus on specific government remedies to problems arising from production and allocation inefficiencies.