Handbook of Pricing Management, Oxford University Press

Game Theory Models of Pricing

September 2010

Praveen Kopalle and Robert A. Shumsky

Tuck School of Business at Dartmouth

1. Introduction

In 1991, packaged-goods behemoth Procter & Gamble (P&G) initiated a "value pricing" scheme for sales to retailers. Value pricing was P&G's label for *everyday low pricing* (EDLP), a pricing strategy under which retailers are charged a consistent price rather than a high baseline price punctuated by sporadic, deep discounts. P&G had many reasons for converting to EDLP. Its sales force had grown to rely on discounting to drive sales and the use of deep discounts had spiraled out of control, cutting into earnings (Saporito, 1994). The discounts then created demand shocks, as both retailers and customers loaded up on products whenever the price dropped. This exacerbated the so-called bullwhip effect, as demand variability was amplified up the supply chain, reducing the utilization of distribution resources and factories and increasing costs (Lee, Padmanabhan, and Whang, 1997).

Despite the costs of its traditional discounting strategy, the move to value pricing was risky for P&G. How would retailers, consumers, and competing manufacturers respond to P&G's EDLP program? For example, would competitors deepen their own discounts to pull market share away from P&G? Or would they also move to a low, consistent price? To accurately predict competitors' behavior, P&G would have to consider each competitor's beliefs about P&G itself. For example, it was possible that Unilever, a competing manufacturer, might increase its discounting if it believed that P&G was not fully committed to EDLP but otherwise might itself adopt EDLP. Therefore, P&G's decision on whether and how to adopt EDLP depended on the responses of its competitors, which, in turn, depended upon their beliefs about what P&G was going to do.

P&G's actions depended on Uniliver's, which depended on P&G's—a chicken-and-egg problem that can make your head spin. There is, however, an approach to problems like this that can cut through the complexity and help managers make better pricing decisions. That approach, game theory, was first developed as a distinct body of knowledge by mathematicians and economists in the 1940s and '50s (see the pioneering work by von Neumann and Morgenstern, 1944, and Nash, 1950). Since then, it has been applied to a variety of fields, including computer science, political science, biology, operations management, and marketing science.

In this chapter, we first introduce the basic concepts of game theory by using simple pricing examples. We then link those examples with both the research literature and industry practice, including examples of how game theory has been used to understand P&G's value pricing initiative and the competitive response. In Section 2 we define the basic elements of a game and describe the fundamental assumptions that underlie game theory. In the remaining sections we examine models that provide insight into how games work (e.g., understanding various types of equlibria) as well as how competition affects pricing. The models may be categorized according to two attributes: the timing of actions and the number of periods (Table 1). In many of the models, competitors make simultaneous pricing decisions, so that neither player has more information about its competitor's actions than the other. In other games with sequential timing, one competitor sets its price first. Some games have a single period (we define a 'period' as a time unit in which each player takes a single action), while other games will have multiple time periods, and actions taken in one period may affect actions and rewards in later periods. Along the way, we will make a few side trips: a brief look at behavioral game theory in Section 3.2 and some thoughts on implications for the practice of pricing in Section 4.

		Timing of actions	
		Simultaneous	Sequential
Number of periods	One	Sections 2.1-2.4, 2.6	Section 2.5
	Multiple	Sections 3.1 – 3.4	Sections 3.5-3.6

Table 1: Game types in this chapter

This chapter does not contain a rigorous description of the mathematics behind game theory, nor does it cover all important areas of the field. For example, we will not discuss *cooperative* game theory, in which participants form coalitions to explicitly coordinate their prices (in many countries, such coalitions would be illegal). We will, however, say a bit more about cooperative games at the end of Section 2.6.

For a more complete and rigorous treatments of game theory, we recommend Fudenberg and Tirole (1991), Myerson (1991), and Gibbons (1992). For more examples of how practitioners can make use of game theory for many decisions in addition to pricing, we recommend Dixit and Nalebuff (1991). Moorthy (1985) focuses on marketing applications such as market-entry decisions and advertising programs, as well as pricing. Cachon and Netessine (2006) discuss applications of game theory to supply chain management. Nisan et. al. (2007) describe computational approaches to game theory for applications such as information security, distributed computation, prediction markets, communications networks, peer-to-peer systems, and the flow of information through social networks. The popular press and trade journals also report on how game theory can influence corporate strategy. For example Rappeport (2008) describes how game theory has been applied at Microsoft, Chevron, and other firms. Finally, information intermediaries such as Yahoo and Google employ research groups that develop game theory models to understand how consumers and firms use the web, and to develop new business models for the internet (see labs.yahoo.com/ and research.google.com/).

2. Introduction to Game Theory: Single-Period Pricing Games

A game is defined by three elements: players, strategies that may be used by each player, and payoffs.

A game consists of **players** (participants in the game), **strategies** (plans by each player that describe what action will be taken in any situation), and **payoffs (**rewards for each player for all combinations of strategies).

In the value-pricing example above, the players might be P&G and Unilever, their strategies are the prices they set over time under certain conditions, and the payoffs are the profits they make under various combinations of pricing schemes. Given that each player has chosen a strategy, the collection of players' strategies is called a *strategy profile*. If we know the strategy profile, then we know how every player will act for any situation in the game.

When describing a game, it is also important to specify the information that is available to each player. For example, we will sometimes consider situations in which each player must decide upon a strategy to use without knowing the strategy chosen by the other player. At other times, we will let one player observe the other player's choice before making a decision.

When analyzing a game, there are two basic steps. The first is to make clear the implications of the rules—to fully understand the relationships between the strategies and the payoffs. Once we understand how the game works, the second step is to determine which strategy each player will play under various circumstances. This analysis may begin with a description of each player's optimal strategy *given* any strategy profile for the other players. The analysis usually ends with a description of *equilibrium* strategies - a strategy profile that the players may settle into and will not want to change. One fascinating aspect of game theory is that these equilibrium strategies can often be quite different from the best overall strategy that might be achieved if the players were part of a single firm that coordinated their actions. We will see that competition often can make both players worse off.

A typical assumption in game theory, and an assumption we hold in much of this chapter, is that all players know the identities, available strategies, and payoffs of all other players, as in a game of tic-tac-toe. When no player has private information about the game that is unavailable to the other players, and when all players know that the information is available to all players, we say that the facts of the game are *common knowledge* (see Aumann, 1976, and Milgrom, 1981, for more formal descriptions of the common knowledge assumption). In some instances, however, a player may have private information. For example, one firm may have more accurate information than its competitor about how its own customers respond to price changes by the competitor. Such games of *incomplete information* are sometimes called *Bayesian games*, for uncertainty and probability are important elements of the game, and Bayes' rule is used to update the players' beliefs as the game is played. See Gibbons (1992) and Harsanyi (1967) for more background on Bayesian games.

We also assume that the players are rational, i.e., that firms make pricing decisions to maximize their profits. Finally, we assume that each player knows that the other players are rational. Therefore, each player can put itself in its competitors' shoes and anticipate its competitors' choices. Pushing further, each player knows that its competitors are anticipating its own choices, and making choices accordingly. In general, we assume that the players have an *infinite hierarchy of beliefs* about strategies in the game. This assumption is vital for understanding why competitors may achieve the Nash equilibrium described in the next Section. The assumption also highlights a key insight that can be developed by studying game theory – the importance of anticipating your competitors' actions.

2.1 A First Example: Pricing the UltraPhone and the Prisoner's Dilemma

Consider two players, firms A and B, that have simultaneously developed virtually identical versions of what we will call the UltraPhone, a hand-held device that includes technology for cheap, efficient 3-dimensional video-conferencing. Both firms will release their UltraPhones at the same time and, therefore, each has to set a price for the phone without knowing the price set by its competitor. For the purposes of this example, we assume both firms will charge either a High price or a Low price for all units (see Table 2). Therefore, the firms' possible actions are to price 'High' or 'Low,' and each firm's strategy is a plan to charge one of these two prices.

Price	Unit contribution margin at that price	Number in segment
High	$800	2,000,000
Low	$300	2,000,000

Table 2: Prices, margins, and segments for the UltraPhone

To complete the description of the game, we describe the payoffs under each strategy profile. At the High price, the firms' unit contribution margin is $800/phone; at the Low price, the firms' unit contribution margin is $300/phone. There are two segments of customers for the UltraPhone, and the maximum price each segment is willing to pay is very different. The two million people in the "High" segment are all willing to pay the High price (but would, of course, prefer the Low price), while the two million people in the "Low" segment are only willing to pay the Low price. We also assume that if both firms charge the same price, then sales are split equally between the two firms. And we assume both firms can always manufacture a sufficient number of phones to satisfy any number of customers. Note that in this example, there is no reliable way to discriminate between segments; if the phone is offered for the Low price by either firm, then both the two million Low customers and the two million High customers will buy it for that Low price.

Given the information in Table 1, both firms ask, "What should we charge for our phone?" Before looking at the competitive analysis, consider what the firms should do if they coordinate—e.g., if they were not competitors but instead were divisions of one firm that is a monopoly producer of this product. The total contribution margin given the High price is 2 million*$800 = $1.6 billion. The margin given the Low price is (2 million + 2 million)*$300 =

$1.2 billion. Therefore, the optimal solution for a monopoly is to set prices to exclusively target the high-end market.

		B Chooses…	
		High	Low
A Chooses...	High	8 8	12 0
	Low	0 12	6 6

Table 3: The *UltraPhone* game. Entries in each cell are payoffs to A (lower number in each cell) and to B (upper number in each cell) in $100 millions.

Now we turn to the competitive analysis—the game. The two competing firms, A and B, face the situation shown in Table 3, which represents a normal form of a game. The rows correspond to the two different strategies available to firm A: price High or Low. The columns correspond to the same strategies for B. The entries in the tables show the payoffs, the total contribution margins for firm A (the lower number in each cell of the table) and firm B (the upper number in each cell) in units of $100 million. If both price High (the upper left cell), then they split the $1.6 billion equally between them. If both price Low (bottom right cell), then they split the $1.2 billion. For the cell in the upper right, A prices High but B prices Low, so that B captures all of the demand at the Low price and gains $1.2 billion. The lower left cell shows the reverse.

Given this payoff structure, what should the firms do? First, consider firm A's decision, given each strategy chosen by B. If B chooses High (the first column), then it is better for A to choose Low. Therefore, we say that Low is player A's best response to B choosing High.

> A player's **best response** is the strategy, or set of strategies, that maximizes the player's payoff, given the strategies of the other players.

If B chooses Low (the second column), then firm A's best response is to choose Low as well. The fact that it is best for A to price Low, no matter what B does, makes pricing Low a particular type of strategy, a dominant one.

> A strategy is a **dominant strategy** for a firm if it is optimal, no matter what strategy is used by the other players.

Likewise, *dominated strategies* are never optimal, no matter what the competitors do.

Pricing Low is also a dominant strategy for B: no matter what A chooses, Low is better. In fact, the strategy profile [A Low, B Low] is the unique Nash Equilibrium of this game; if both firms choose Low, then neither has a reason to unilaterally change its mind.

> The firms are in a **Nash Equilibrium** if the strategy of each firm is the best response to the strategies of the other firms. Equivalently, in a Nash equilibrium, none of the firms have any incentive to unilaterally deviate from its strategy.

This is the only equilibrium in the game. For example, if A chooses High and B chooses Low, then A would have wished it had chosen Low and earned $600 million rather than $0. Note that for the Nash equilibrium of this game, both firms choose their dominant strategies, but that is not always the case (see, for example, the Nash equilibria in the Game of Chicken, Section 2.3).

Because we assume that both firms are rational and fully anticipate the logic used by the other player (the infinite belief hierarchy), we predict that both will charge Low, each will earn $600 million, and the total industry contribution margin will be $1.2 billion. This margin is less than the optimal result of $1.6 billion. This is a well-known result: under competition, both prices and industry profits decline. We will see in Section 3.1, however, that it may be possible to achieve higher profits when the game is played repeatedly.

This particular game is an illustration of a Prisoner's Dilemma. In the classic example of the Prisoner's Dilemma, two criminal suspects are apprehended and held in separate rooms. If neither confesses to a serious crime, then both are jailed for a short time (say, one year) under relatively minor charges. If both confess, they both receive a lengthy sentence (five years). But if one confesses and the other doesn't, then the snitch is freed and the silent prisoner is given a long prison term (10 years). The equilibrium in the Prisoner's Dilemma is for both prisoners to confess, and the resulting five-year term is substantially worse than the cooperative equilibrium they would have achieved if they had both remained silent. In the UltraPhone example, the Low-Low strategy is equivalent to mutual confession.

The strategy profile [A Low, B Low] is a *Nash equilibrium in pure strategies*. It is a *pure strategy* equilibrium because each player chooses a single price. The players might also consider a *mixed strategy*, in which each player chooses a probability distribution over a set of pure strategies (see Dixit and Skeath, 1999, chapter 5, for an accessible and thorough discussion of mixed strategies). The equilibrium concept is named after John Nash, a mathematician who showed that if each player has a finite number of pure strategies, then there must always exist at least one mixed strategy or pure strategy equilibrium (Nash, 1950). In certain games, however, there may not exist a pure strategy equilibrium and as we will see in Section 2.4, it is also possible to have multiple pure-strategy equilibria.

2.2 Bertrand and Cournot Competition

The UltraPhone game is similar to a *Bertrand competition*, in which two firms produce identical products and compete only on price, with no product differentiation. It can be shown that if demand is a continuous linear function of price and if both firms have the same constant marginal cost, then the only Nash equilibrium is for both firms to set a price equal to their marginal cost. As in the example above, competition drives price down, and in a Bertrand competition, prices are driven down to their lowest limit.

Both the UltraPhone example and the Bertrand model assume that it is always possible for the firms to produce a sufficient quantity to supply the market, no matter what price is charged. Therefore, the Bertrand model does not fit industries with capacity constraints that are expensive to adjust. An alternative model is *Cournot competition*, in which the firms' strategies are quantities rather than prices. The higher the quantities the firms choose, the lower the price. Under a Cournot competition between two firms, the resulting Nash equilibrium price is higher than the marginal cost (the Bertrand solution) but lower than the monopoly price. As the number of firms competing in the market increases, however, the Cournot equilibrium price falls. In the theoretical limit (an infinite number of competitors), the price drops down to the marginal cost. See Weber [this volume] for additional discussion of Bertrand and Cournot competition.

Besides the assumption that there are no capacity constraints, another important assumption of Bertrand competition is that the products are identical, so that the firm offering the lowest price attracts all demand. We will relax this assumption in the next section.

2.3 Continuous Prices, Product Differentiation, and Best Response Functions

In Section 2.1 we restricted our competitors to two choices: High and Low prices. In practice, firms may choose from a variety of prices, so now assume that each firm chooses a price over a continuous range. Let firm i's price be p_i (i=A or B). In this game, a player's strategy is the choice of a price, and we now must define the payoffs, given the strategies. First we describe how each competitor's price affects demand for the product. In the last section, all customers chose the lowest price. Here we assume that the products are differentiated, so that some customers receive different utilities from the products and are therefore loyal; they will purchase from firm A or B, even if the other firm offers a lower price. Specifically, let $d_i(p_i,p_j)$ be the demand for product i, given that the competitors charge p_i and p_j. We define the linear demand function,

$$d_i(p_i,p_j)=a-bp_i+cp_j.$$

If $c > 0$, we say that the products are substitutes: a lower price charged by firm A leads to more demand for A's product and less demand for B's (although if $p_B < p_A$, B does not steal all the demand, as in Section 2.1). If $c < 0$, then the products are complements: a lower price charged by A raises demand for both products. In the UltraPhone example, the products are substitutes; product A and a specialized accessory (e.g., a docking station) would be complements. If $c=0$

then the two products are independent, there is no interaction between the firms, and both firms choose their optimal monopoly price.

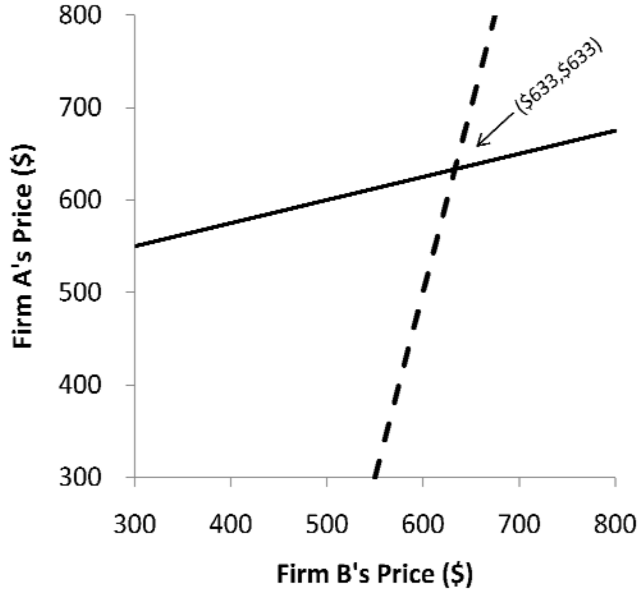

Figure 1: Best response functions in the UltraPhone pricing game with product differentiation

Now let m be the constant marginal cost to produce either product. Firm i's margin is therefore $(p_i - m)d_i(p_i,p_j)$. This function is concave with respect to p_i, and therefore the optimal price for firm i given that firm j charges p_j can be found by taking the derivative of the function with respect to p_i, setting that first derivative equal to 0, and solving algebraically for p_i. Let $p_i^*(p_j)$ be the resulting optimal price, and we find,

$$p_i^*(p_j) = \frac{1}{2b}(a + bm + cp_j). \tag{1}$$

This optimal price is a function of the competitor's price, and therefore we call $p_i^*(p_j)$ the *best response function* of i to j; this is the continuous equivalent of the *best response* defined in

Section 2.1. Figure 1 shows the two best response functions on the same graph, for a given set of parameters $a=3000$, $b=4$, $c=2$, and $m=\$200$/unit. For example, if Firm B sets a price of \$300, Firm A's best response is a price of \$550. From the function, we see that if the products are substitutes and $c > 0$ (as in this case), the best response of one firm rises with the price of the other firm. If the products are complements, the lines would slope downward, so that if one firm raises its prices then the best response of the competitor is to lower its price.

The plot also shows that if Firm A chooses a price of \$633, then Firm B's best response is also \$633. Likewise, if Firm B chooses \$633, then so does Firm A. Therefore, neither firm has any reason to deviate from choosing \$633, and [\$633, \$633] is the unique Nash equilibrium. In general, the equilibrium prices satisfy equation (1) when i is Firm A as well as when i is Firm B. By solving these two equations we find,

$$p_i^* = p_j^* = \frac{a + bm}{2b - c}.$$

Note that as c rises, and the products move from being complements to substitutes, the equilibrium price declines.

In this example, the firms are identical and therefore the equilibrium is symmetric (both choose the same price). With non-identical firms (e.g., if they have different parameter values in the demand function and/or different marginal costs), then there can be a unique equilibrium in which the firms choose different prices (Singh and Vives, 1984). We will discuss additional extensions of this model, including the role of reference prices, in Section 3.3.

2.4 Multiple Equilibria and a Game of Chicken

Now return to the UltraPhone pricing game with two price points, as in Section 2.1, although we also incorporate product differentiation, as in Section 2.3. Assume that 40 percent of the High-segment customers are loyal to firm A and will choose A's product even if A chooses a High price while B chooses a Low price. Likewise, 40 percent of the High-segment customers are loyal to firm B. The payoffs from this new game are shown in Table 4. If both firms price High or both price Low, then the payoffs do not change. If A prices High and B prices Low (the upper right cell), then A's margin is 40%*(2 million)*$800 = $640 million while B's margin is (60%*(2 million) + 2 million)*$300 = $960 million. The lower left cell shows the reverse.

		B Chooses…	
		High	Low
A Chooses...	High	8 8	9.6 6.4
	Low	6.4 9.6	6 6

Table 4: Payoffs for the UltraPhone game with loyal customers, in $100 millions

The presence of loyal customers significantly changes the behavior of the players. If B chooses High, then A's best response is to choose Low, while if B chooses Low, A's best response is to choose High. In this game, neither firm has a dominant strategy, and [A Low, B Low] is no longer an equilibrium; if the firms find themselves in that cell, both firms would have an incentive to move to High. But [A High, B High] is not an equilibrium either; both would then want to move to Low.

In fact, this game has *two* pure strategy Nash equilibria: [A High, B Low] and [A Low, B High]. Given either combination, neither firm has any incentive to unilaterally move to another strategy. Each of these two equilibria represents a split in the market between one firm that focuses on its own High-end loyal customers and the other that captures everyone else.

This example is analogous to a game of "chicken" between two automobile drivers. They race towards one another and if neither swerves, they crash, producing the worst outcome—in our case [A Low, B Low]. If both swerve, the payoff is, obviously, higher. If one swerves and the other doesn't, the one who did not swerve receives the highest award, while the one who "chickened out" has a payoff that is lower than if they had both swerved but higher than if the driver had crashed. As in Table 4, there are two pure strategy equilibria in the game of chicken, each with one driver swerving and the other driving straight. [1]

The existence of multiple equilibria poses a challenge for both theorists and practitioners who use game theory. Given the many possible equilibria, will players actually choose one? If so, which one? Behavioral explanations have been proposed to answer this question – see Section 3.2. In addition, game theorists have invented variations of the Nash equilibrium to address these questions. These variations are called *refinements*, and they lay out criteria that exclude certain types of strategies and equilibria. Many of these innovations depend upon the theory of dynamic games, in which the game is played in multiple stages over time (see Section

[1] There is also a mixed-strategy equilibrium for the game in Table 4. If each player randomly selects a High price with probability 0.2 and a Low price with probability 0.8, then neither player has any incentive to choose a different strategy.

3). The *subgame perfect equilibrium* described in the next section is one such refinement. Another is the *trembling-hand perfect equilibrium* in which players are assumed to make small mistakes (or "trembles") and all players take these mistakes into account when choosing strategies (Selten, 1975; Fudenberg and Tirole, 1991, Section 8.4). For descriptions of additional refinements, see Govindan and Wilson (2005).

2.5 Sequential Action and a Stackelberg Game

In the previous Section, we assumed that both firms set prices simultaneously without knowledge of the other's decision. We saw how this game has two equilibria and that it is difficult to predict which one will be played. In this Section we will see how this problem is resolved if one player moves first and chooses its preferred equilibrium; in the game of chicken, imagine that one driver rips out her steering wheel and throws it out her window, while the other driver watches.

For the UltraPhone competition, assume that firm A has a slight lead in its product development process and will release the product and announce its price first. The release and price announcement dates are set in stone—e.g., firm A has committed to announcing its price and sending its products to retailers on the first Monday of the next month, while firm B will announce its price and send its products to retailers a week later. It is important to note that these are the *actual* pricing dates. Either firm may publicize potential prices before those dates, but retailers will only accept the prices given by firm A on the first Monday and firm B on the second Monday. Also, assume that the one-week "lag" does not significantly affect the sales assumptions introduced in Sections 2.1 and 2.4.

This sequential game has a special name, In a **Stackelberg Game**, one player (the *leader*) moves first, while the other player (the *follower*) moves second after observing the action of the leader.

The left side of Figure 2 shows the extensive-form representation of the game, with A's choice as the first split at the top of the diagram and B's choice at the bottom.

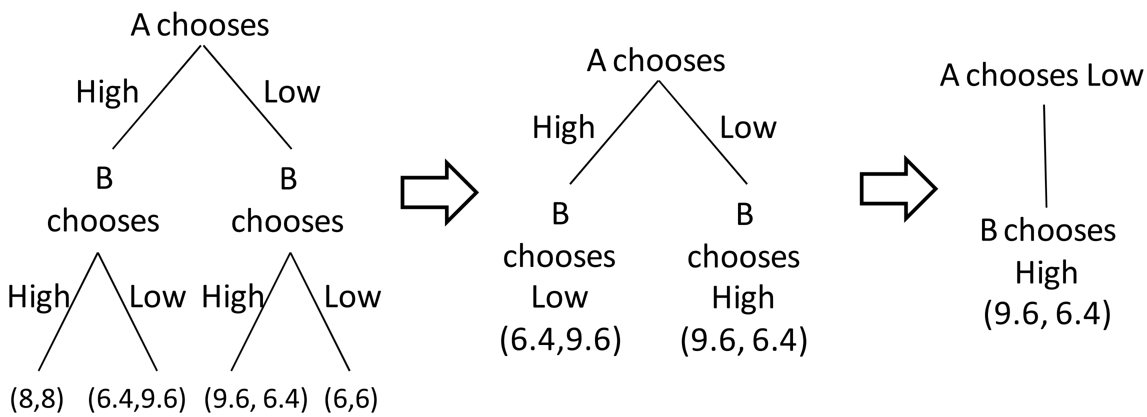

Figure 2: Analysis of the UltraPhone game with loyal customers and A as a first mover

One can identify an equilibrium in a Stackelberg Game (a Stackelberg equilibrium) by *backwards induction* - work up from the bottom of the diagram. First, we see what firm B would choose, given either action by firm A. As the second tree in Figure 2 shows, B would choose Low if A chooses High and would choose High if A chooses Low. Given the payoffs from these choices, it is optimal for A to choose low, and the Stackelberg equilibrium is [A Low, B High], with the highest payoff—$960 million—going to firm A. By moving first, firm A has guided the game towards its preferred equilibrium.

It is important to note, however, that [A High, B Low] is still a Nash equilibrium for the full game. Suppose that firm B states before the crucial month arrives that *no matter what* firm A decides, it will choose to price Low (again, this does not mean that firm B actually sets the price to Low before the second Monday of the month, just that it claims it will). If firm A believes this threat, then it will choose High in order to avoid the worst-case [A Low, B Low] outcome. Therefore, another equilibrium may be [A High, B Low], as in the simultaneous game of Section 2.4.

But can firm A believe B's threat? Is the threat credible? Remember that A moves first and that B observes what A has done. Once A has chosen Low, it is not optimal for B to choose Low as well; it would be better for B to abandon its threat and choose High. We say that B's choices (the second level in Figure 1) are subgames of the full game. When defining an equilibrium strategy, it seems natural to require that the strategies each player chooses also lead to an equilibrium in each subgame.

In a **subgame perfect equilibrium,** the strategies chosen for each subgame is a Nash equilibrium of that subgame.

In this case, B's strategy "price Low no matter what A does" does not seem to be part of a reasonable equilibrium because B will not price Low if A prices Low first; An equilibrium that includes B Low after A Low is not Nash. Therefore, we say that [A Low, B High] is a subgame perfect equilibrium but that [A High, B Low] is not.

2.6 An Airline Revenue Management Game

In the UltraPhone game, the two firms set prices, and we assume the firms can easily adjust production quantities to accommodate demand. The dynamics of airline pricing often are quite different. The capacity of seats on each flight is fixed over the short term, and prices for various types of tickets are set in advance. As the departure of a flight approaches, the number of seats available at each price is adjusted over time. This dynamic adjustment of seat inventory is sometimes called *revenue management*; see Talluri [this volume] for additional information of the fundamentals of revenue management as well as Barnes [this volume], Kimes et al. [this volume], and Lieberman [this volume] for descriptions of applications.

Here, we will describe a game in which two airlines simultaneously practice revenue management. As was true for the initial UltraPhone example, the game will be a simplified version of reality. For example, we will assume that the single flight operated by each airline has only one seat left for sale. As we will discuss below, however, the insights generated by the model hold for realistic scenarios.

Suppose two airlines, *X* and *Y*, offer direct flights between the same origin and destination, with departures and arrivals at similar times. We assume that each flight has just one seat available. Both airlines sell two types of tickets: a low-fare ticket with an advance-purchase requirement for $200 and a high-fare ticket for $500 that can be purchased at any time. A ticket purchased at either fare is for the same physical product: a coach-class seat on one flight leg. The advance-purchase restriction, however, segments the market and allows for price discrimination between customers with different valuations of purchase-time flexibility. Given the advance-purchase requirement, we assume that demand for low-fare tickets occurs before demand for high-fare tickets. Customers who prefer a low fare and are willing to accept the purchase restrictions will be called “low-fare customers.” Customers who prefer to purchase later, at the higher price, are called “high-fare customers.”

Assume marginal costs are negligible, so that the airlines seek to maximize expected revenue. To increase revenue, both airlines may establish a *booking limit* for low-fare tickets: a restriction on the number of low-fare tickets that may be sold. (In our case, the booking limit is either 0 or 1.) Once this booking limit is reached by ticket sales, the low fare is closed and only high-fare sales are allowed. If a customer is denied a ticket by one airline, we assume that customer will attempt to purchase a ticket from the other airline (we call these "overflow passengers"). Therefore, both airlines are faced with a random initial demand as well as demand from customers who are denied tickets by the other airline. Passengers denied a reservation by both airlines are lost. We assume that for each airline, there is a 100 percent chance that a low-fare customer will attempt to purchase a ticket followed by a 50 percent chance that a high-fare customer will attempt to purchase a ticket. This information is summarized in Table 5.

		Demand	
Fare Class	Ticket Price	Airline X	Airline Y
Low	$200	1 (guaranteed)	1 (guaranteed)
High	$500	Prob{0}=0.5 Prob{1}=0.5	Prob{0}=0.5 Prob{1}=0.5

Table 5: Customer segments for the revenue management game

The following is the order of events in the game:

1. Airlines establish booking limits. Each airline chooses to either close the low-fare class ("Close") and not sell a low-fare ticket, or keep it open ("Open"). In this revenue management game, the decision whether to Close or remain Open is the airline's strategy.

2. A low-fare passenger arrives at each airline and is accommodated (and $200 collected) if the low-fare class is Open.

3. Each airline either receives a high-fare passenger or does not see any more demand, where each outcome has probability 0.5. If a passenger arrives, and if the low-fare class has been closed, the passenger is accommodated (and $500 collected).

4. A high-fare passenger who is not accommodated on the first-choice airline spills to the alternate airline and is accommodated there if the alternate airline has closed its seat from a low-fare customer and has not sold a seat to its own high-fare customer.

		Y Chooses…	
		Open	Close
X Chooses...	Open	200, 200	200, 375
	Close	375, 200	250, 250

Table 6: Expected payoffs for the airline revenue management game.

Entries in each cell are payoffs to X, Y.

Table 6 shows the payoffs for each combination of strategies. If both airlines choose to keep the low-fare class Open, both book customers for $200. If both choose to Close, each has a 50 percent chance to book a high fare, so that each has an expected value of (0.5)*$500 = $250. In the upper right cell, if X chooses to Open and Y chooses to Close, then X books a low-fare customer for $200 while Y has a chance to book a high-fare customer. That high-fare customer may either be a customer who originally desired a flight on Y (who arrives with probability 0.5) or a spillover from X (who also arrives, independently, with probability 0.5). Therefore, Y's expected revenue is [1- (0.5)(0.5)]*$500 = $375. The lower left cell shows the reverse situation.

By examining Table 6, we see that the only Nash equilibrium is [X Close, Y Close]. If one airline chooses to Close, then the airline that chooses Open foregoes the opportunity to book a high-fare customer.

Note that this game is not precisely the same as the UltraPhone game/Prisoner's Dilemma discussed in Section 2.1; the airlines could have done even worse with [X Open, Y Open]. The competitive equilibrium, however, is inferior in terms of aggregate revenue to the cooperative solution, just as it is in the Prisoner's Dilemma. In Table 6, the highest total revenue across both airlines—$575—is obtained with strategies [X Open, Y Close] or [X Close, Y Open]. The Nash equilibrium has an expected total revenue of $500.

There is another interesting comparison to make with the UltraPhone game. In that game, competition drove prices down. In this game, competition leads the airlines from an optimal coordinated solution [X Open, Y Close] or [X Close, Y Open] to the Nash equilibrium [X Close, Y Close]. Therefore, under the coordinated solution, one low-fare ticket is sold and one high-fare ticket may be sold. Under the Nash equilibrium, only high-fare tickets may be sold. Netessine and Shumsky (2005) analyze a similar game with two customer classes and any number of available seats, and they confirm that this general insight continues to hold; under competition, more seats are reserved for high-fare customers, while the number of tickets sold (the load factor) declines, as compared to the monopoly solution. The use of seat inventory control by the airlines and the low-to-high pattern of customer valuations inverts the Bertrand and Cournot results; competition in the revenue management game raises the average price paid for a ticket as compared to the average price charged by a monopolist.

In practice, however, airlines may compete on multiple routes, offer more than two fare classes, and do not make a single revenue management decision simultaneously. Jiang and Pang (2007) extend this model to airline competition over a network. Dudey (1992) and Martinez-de-Albeniz and Talluri (2010) analyze the dynamic revenue management problem, in which the airlines adjust booking limits dynamically over time. d'Huart (2010) uses a computer simulation to compare monopoly and oligopoly airline revenue management behavior. He simulates up to 4 airlines with 3 competing flights offered by each airline and six fare classes on each flight. The simulation also incorporates many of the forecasting and dynamic revenue management heuristics used by the airlines. The results of the simulation correspond to our models' predictions: fewer passengers are booked under competition, but with a higher average fare.

Of course, the airlines may prefer to collaborate, agree to one of the monopoly solutions, and then split the larger pot so that both airlines would come out ahead of the Nash outcome. In general, such collusion can increase overall profits, while reducing consumer surplus. Such price-fixing is illegal in the United States, Canada, the European Union, and many other countries. Occasionally airlines do attempt to collude and get caught, e.g., British Airways and Virgin Atlantic conspired to simultaneously raise their prices on competing routes by adding fuel surcharges to their tickets. Virgin Atlantic reported on the plot to the U.S. Department of Justice, and British Airways paid a large fine (Associated Press, 2007).

In the absence of such legal risks, a fundamental question is whether such collusion is stable when airlines do not have enforceable agreements to cooperate as well as explicit methods for distributing the gains from cooperation. As we have seen from the examples above, each player has an incentive to deviate from the monopoly solution to increase its own profits. In Section 3 we will discuss how repeated play can produce outcomes other than the single-play Nash equilibrium describe above.

In certain cases in the United States, airlines have received permission from the Department of Justice to coordinate pricing and split the profits. Northwest and KLM have had such an agreement since 1993, while United, Lufthansa, Continental and Air Canada began coordinating pricing and revenue management on their transatlantic flights in 2009 (BTNonline.com, 2008; DOT, 2009). The ability of such joint ventures to sustain themselves depend upon whether each participant receives a net benefit from joining the coalition. Cooperative game theory focuses on how the value created by the coalition may be distributed among the participants, and how this distribution affects coalition stability. Such joint ventures, however, are the exception rather than the rule in the airline business and elsewhere, and therefore we will continue to focus on non-cooperative game theory.

3. Dynamic Pricing Games

Pricing competition can be viewed from either a static or a dynamic perspective. Dynamic models incorporate the dimension of time and recognize that competitive decisions do not necessarily remain fixed. Viewing competitive pricing strategies from a dynamic perspective can provide richer insights, and for some questions more accurate answers, than when using static models. In this Section we will first explain how dynamic models differ from the models presented in the last Section. We will then describe a variety of dynamic models, and we will see how the price equilibria predicted by the models can help firms to understand markets and set prices. Finally, and perhaps most significantly, we will show how pricing data collected from the field has been used to test the predictions made by game theory models.

First, we make clear what we mean by a dynamic game:

If a game is played exactly once, it is a **single-period**, or a **one-shot**, or a **static game**. If a pricing game is played more than once, it is a **dynamic** or a **multi-period game**.

In a dynamic game, decision variables in each period are sometimes called **control variables**. Therefore, a firm's strategy is a set of values for the control variables.

Dynamic competitive pricing situations can be analyzed in terms of either continuous or discrete time. Generally speaking, pricing data from the market is discrete in nature (e.g., weekly price changes at stores or the availability of weekly, store-level scanner data), and even continuous-time models, which are an abstraction, are discretized for estimation of model parameters from market data. Typically, dynamic models are treated as multi-period games because critical variables—sales, market share, and so forth—are assumed to change over time based on the dynamics in the marketplace. Any dynamic model of pricing competition must consider the impact of competitors' pricing strategies on the dynamic process governing changes in sales or market-share variables.

In this section we will describe a variety of dynamic pricing policies for competing firms. See Aviv and Vulcano [this volume] for more details on dynamic policies that do not take competitors' strategies directly into account. That chapter also describes models that examine, in more detail, the effects of strategic customers. Other related chapters include Ramakrishnan [this volume], Blattberg and Breisch [this volume], and Kaya and Ozer [this volume]. This last chapter is closely related to the material in Sections 2.5 and 3.5, for it also describes Stackelberg games, while examining specific contracts between wholesalers and retailers (e.g., buyback, revenue sharing, etc.), the effects of asymmetric forecast information, and the design of contracts to elicit accurate information from supply chain partners. For additional discussion of issues related to behavioral game theory (discussed in Section 3.2, below), see Ozer and Zheng [this volume].

3.1 Effective Competitive Strategies in Repeated Games

We have seen in Section 2.1 how a Prisoner's Dilemma can emerge as the non-cooperative Nash equilibrium strategy in a static, one-period game. In the real world, however, rarely do we see one-period interactions. For example, firms rarely consider competition for one week and set prices simply for that week, keeping them at the same level during all weeks into the future. Firms have a continuous interaction with each other for many weeks in a row, and this continuous interaction gives rise to repeated games.

A **repeated game** is a dynamic game in which past actions do not influence the payoffs or set of feasible actions in the current period, i.e., there is no explicit link between the periods.

In a repeated game setting, firms would, of course, be more profitable if they could avoid the Prisoner's Dilemma and move to a more cooperative equilibrium. Therefore, formulating the pricing game as a multi-period problem may identify strategies under which firms achieve a cooperative equilibrium, where competing firms may not lower prices as much as they would have otherwise in a one-shot Prisoner's Dilemma. As another example, in an information sharing context, Özer, Zheng, and Chen (2009) discuss the importance of trust and determine that repeated games enhance trust and cooperation among players. Such cooperative equilibria may also occur if the game is not guaranteed to be repeated, but may be repeated with some probability.

A robust approach in such contexts is the "tit-for-tat" strategy, where one firm begins with a cooperative pricing strategy and subsequently copies the competing firm's previous move. This strategy is probably the best known and most often discussed rule for playing the Prisoner's Dilemma in a repeated setting, as illustrated by Axelrod and Hamilton (1981), Axelrod (1984), and also discussed in Dixit and Nalebuff (1991). In his experiment, Axelrod set up a computer tournament in which pairs of contestants repeated the two-person prisoner's dilemma game 200 times. He invited contestants to submit strategies, and paired each entry with each other entry. For each move, players received three points for mutual cooperation and one point for mutual defection. In cases where one player defected while the other cooperated, the former received five points while the latter received nothing. Fourteen of the entries were from people who had published research articles on game theory or the Prisoner's Dilemma. The tit-for-tat strategy (submitted by Anatol Rapoport, a mathematics professor at the University of Toronto), the simplest of all submitted programs, won the tournament. Axelrod repeated the tournament with more participants and once again, Rapoport's tit-for-tat was the winning strategy (Dixit and Nalebuff 1991).

Axelrod's research emphasized the importance of minimizing echo effects in an environment of mutual power. Although a single defection may be successful when analyzed for its direct effect, it also can set off a long string of recriminations and counter-recriminations, and, when it does, both sides suffer. In such an analysis, tit-for-tat may be seen as a punishment strategy (sometimes called a *trigger strategy*) that can produce a supportive collusive pricing strategy in a non-cooperative game. Lal's (1990a) analysis of price promotions over time in a competitive environment suggests that price promotions between national firms (say Coke and Pepsi) can be interpreted as a long run strategy by these firms to defend market shares from possible encroachments by store brands. Some anecdotal evidence from the beverage industry in support of such collusive strategies is available. For example, a *Consumer Reports* article (Lal, 1990b) noted that in the soft-drink market, where Coca Cola and PepsiCo account for two-thirds of all sales, "[s]urveys have shown that many soda drinkers tend to buy whichever brand is on special. Under the Coke and Pepsi bottlers' calendar marketing arrangements, a store agrees to feature only one brand at any given time. Coke bottlers happened to have signed such agreements for 26 weeks and Pepsi bottlers the other 26 weeks!" In other markets, such as ketchup, cereals, and cookies (Lal, 1990b), similar arrangements can be found. Rao (1991) provides further support for the above finding and shows that when the competition is asymmetric, the national brand promotes to ensure that the private label does not try to attract consumers away from the national brand.

3.2 Behavioral Considerations in Games

We now step back and ask if our game theory models provide us with a sufficiently accurate representation of actual behavior. Behavioral considerations can be important for both static games and for repeated games, where relationships develop between players, leading to subtle and important consequences.

The assumptions we laid out at the beginning of Section 2 – rational players and the infinite belief hierarchy – can sometimes seem far-fetched when we consider real human interactions. Given that the underlying assumptions may not be satisfied, are the predicted results of the games reasonable? When describing research on human behavior in games, Camerer (2003) states,

> " …game theory students often ask: "This theory is interesting…but do people actually play this way?' The answer, not surprisingly, is mixed. There are no interesting games in which subjects reach a predicted equilibrium immediately. And there are no games so complicated that subjects do not converge in the direction of equilibrium (perhaps quite close to it) with enough experience in the lab."

Camerer and his colleagues have found that to accurately predict human behavior in games, the mathematics of game theory must often be merged with theories of behavior based on models of social utility (e.g., the concept of fairness), limitations on reasoning, and learning over time. The result of this merger is the fast-growing field of behavioral economics and, more specifically, behavioral game theory.

Behavioral game theory can shed light on a problem we saw in Section 2.4: when there is more than one equilibrium in a game, it is difficult to predict player behavior. Psychologists and behavioral game theorists have tested how player behavior evolves over many repetitions of such games. According to Binmore (2007), over repeated trials of games with multiple equilibria played in a laboratory, the players' behavior evolves towards particular equilibria because of a combination of chance and "historical events." These events include the particular roles ascribed to each player in the game (e.g., an employer and a worker, equal partners, etc.), the social norms that the subjects bring into the laboratory, and, most importantly, the competitive strategies that each player has experienced during previous trials of the game. Binmore also describes how, by manipulating these historical events, players can be directed to non-equilibrium "focal points," although when left to their own devices players inevitably move back to a Nash equilibrium.

Behavioral game theorists have also conducted laboratory tests of the specific pricing games that we describe in this chapter. For example, Duwfenberg and Gneezy (2000) simulated the Bertrand game of Section 2.1 by asking students in groups of 2, 3 or 4 to bid a number from 2 to 100. The student with the lowest bid in the group won a prize that was proportional to the size of that student's bid. Therefore, the experiment is equivalent to the one-period Bertrand pricing game of Section 2.1, in that the lowest price (or 'bid' in the experiment) captures the entire market. This bid-reward process was repeated 10 times, with the participants in each group randomized between each of the 10 rounds to prevent collusive behavior - as described in the previous Section of this Chapter - that may be generated in repeated games with the same players.

No matter how many competitors, the Nash equilibrium in this game is for all players to bid 2, which is equivalent to pricing at marginal cost in the Bertrand game. Duwfenberg and Gneezy found that when their experiments involved two players, however, the bids were consistently higher than 2 (over two rounds of ten trials each, the average bid was 38, and the average winning bid was 26). For larger groups of 3 and 4, the initial rounds also saw bids substantially higher than the Nash equilibrium value of 2, but by the tenth round the bids did converge towards 2.

Camerer (2003) describes hundreds of experiments that demonstrate how actual player behavior can deviate from the equilibrium behavior predicted by classical, mathematical game theory. He also shows how behavioral economists have gone a step further, to explicitly test the accuracy of behavioral explanations for these deviations, such as bounded rationality or a preference for fairness.

A great majority of experiments conducted by behavioral game theorists involve experimental subjects under controlled conditions. Of course, actual pricing decisions are not made by individuals in laboratories, but by individuals or groups of people in firms, who are subject to a variety of (sometimes conflicting) financial and social incentives. Armstrong and Huck (2010) discuss the extent to which these experiments may be used to extrapolate to firm behavior. An alternate to laboratory experiments is an empirical approach, to observe actual pricing behavior by firms and compare those prices with predictions made by game theory models. We will return to the question of game theory's applicability, and see an example of this empirical approach, in Section 3.6. In addition, Ozer and Zheng [this volume] discuss many issues related to behavioral game theory.

3.3 State-Dependent Dynamic Pricing Games

State-dependent dynamic pricing games are a generalization of repeated games, in which the current period payoff may be a function of the history of the game.

> In a dynamic game, a **state variable** is a quantity that (i) defines the state of nature, (ii) depends on the decision or control variables, and (iii) impacts future payoffs (e.g., sales or profitability).

In a **state-dependent dynamic game**, there is an explicit link between the periods, and past actions influence the payoffs in the current period.

In other words, pricing actions today will have an impact on future sales as well as sales today. We are taught early in life about the necessity to plan ahead. Present decisions affect future events by making certain opportunities available, precluding others, and altering the costs of still others. If present decisions do not affect future opportunities, the planning problem is trivial; one need only make the best decision for the present. The rest of this section deals with solving dynamic pricing problems where the periods are physically linked while considering the impact of competition.

In their pioneering work, Rao and Bass (1985) consider dynamic pricing under competition where the dynamics relate primarily to cost reductions through experience curve effects. This implies that firms realize lower product costs with increases in cumulative production. It is shown that with competition, dynamic pricing strategies dominate myopic pricing strategies.

The Table 7 illustrates the impact of dynamic pricing under competition relative to a static, two-period competitive game.

	Static	**Dynamic**
Period 1 Equilibrium Price	\$2.00	\$1.85
Period 2 Equilibrium Price	\$1.90	\$1.80

Table 7: Static vs. Dynamic Pricing Example

In other words, while a static game may suggest an equilibrium price of, say, \$2.00 per unit for a widget in the first period and \$1.90 in the second period, a dynamic game would recommend a price that is less than \$2.00 (say, \$1.85) in period 1 and even lower than \$1.90 (say, \$1.80) in the period 2. This is because a lower price in the first period would increase sales in that initial period, thus increasing cumulative production, which, in turn, would reduce the marginal cost in the second period by more than the amount realized under a static game. This allows for an even lower price in the second period. In such a setting, prices in each period are the control (or decision) variables, and cumulative sales at the beginning of each period would be considered a state variable that affects future marginal cost and, hence, future profitability.

Another pricing challenge involving state dependence is the management of *reference prices*. A reference price is an anchoring level formed by customers, based on the pricing environment. Consider P&G's value pricing strategy discussed earlier in the chapter. In such a setting, given the corresponding stable cost to a retailer from the manufacturer, should competing retailers also employ an everyday low pricing strategy or should they follow a High-Low pricing pattern? In addition, under what conditions would a High-Low pricing policy be optimal? The answer may depend upon the impact reference prices have on customer purchase behavior. One process for reference-price formation is the exponentially smoothed averaging of past observed prices, where more recent prices are weighted more than less recent prices (Winer 1986):

$$r_t = \alpha r_{t-1} + (1-\alpha) p_{t-1}, \quad 0 \leq \alpha < 1,$$

where r_t is the reference price in period t, r_{t-1} is the reference price in the previous period, and p_{t-1} is the retail price in the previous period.

If $\alpha = 0$, the reference price in period t equals the observed price in period t-1. As α increases, r_t becomes increasingly dependent on past prices. Thus, α can be regarded as a memory parameter, with $\alpha = 0$ corresponding to a one-period memory. Consider a group of frequently purchased consumer brands that are partial substitutes. Research (Winer 1986; Greenleaf 1995; Putler 1992; Briesch, Krishnamurthi, Mazumdar, and Raj 1997; Kalyanaram and Little 1994) suggests that demand for a brand depends on not only the brand price but also whether that brand price is greater than the reference price (a perceived loss) or is less than the reference price (a perceived gain). It turns out that the responses to gains and losses relative to the reference price are asymmetric. First, a gain would increase sales and a loss would decrease sales. Second, in some product categories, the absolute value of the impact of a gain is greater than that of a loss, and in yet other categories, it is the opposite. In this scenario, since past and present prices would impact future reference prices, prices can be considered control variables and consumer reference prices would be state variables, which then impacts future sales. In a competitive scenario, this might lead to a dynamic pricing policy, and this rationale differs from other possible explanations for dynamic pricing, such as heterogeneity in information among consumers (Varian, 1980), transfer of inventory holding costs from retailers to consumers (Blattberg, Eppen, and Lieberman, 1981), brands with lower brand loyalty having more to gain from dynamic pricing (Raju, Srinivasan, and Lal, 1990), and a mechanism for a punishment strategy (Lal, 1990b).

Consider a dynamic version of the one-period model of Section 2.3. In this version, a brand's sales is affected by its own price, the price of the competing brand, and the reference price. Demand for brand i, $i = 1,2$, in time t is given by:

$d_{it} = a_i - b_i p_{it} + c_i p_{jt} + g_i (r_{it} - p_{it})$

Where $g_i = \delta$ if $r_{it} > p_{it}$, else $g_i = \gamma$.

Ignoring fixed costs, the objective function for each competitor *i* is to maximize the discounted profit stream over a long-term horizon—i.e.,

$$\max_{p_{it}} \sum_{t=1}^{\infty} \beta^t (p_{it} - c) d_{it}$$

Where,

t = time

β = discount rate (between 0 and 1)

p_{it} = price

c = marginal cost

d_{it} = demand for competitor *i*'s product in time *t*

This problem can be solved using a dynamic programming approach (Bertsekas, 1987) in a competitive setting, where each competitor maximizes not just today's profits but today's profit *plus* all the future profits by taking into consideration how today's pricing would impact tomorrow's reference price and, hence, tomorrow's profitability (Kopalle et al., 1996).

Of course, each player cannot ignore the actions of competing players when solving this optimization problem, and we are interested in the players' equilibrium behavior. The specific type of equilibrium we use here is the subgame perfect equilibrium we defined for the

Stackelberg game in Section 2.5, but now applied to a more general dynamic game (also see Fudenberg and Tirole, 1991). That is, a strategy is a sub-game perfect equilibrium if it represents a Nash equilibrium of every sub-game of the original dynamic game.

Now we discuss the results of this model and distinguish two possible cases. One is where customers care more about gains than losses—i.e., the impact of a unit gain is greater than that of a unit loss ($\delta > \gamma$). Such was the case in Greenleaf's (1995) analysis of peanut butter scanner data. In such a situation, in a monopoly setting Greenleaf (1995) finds that the sub-game perfect Nash equilibrium solution is to have a cyclical (High-Low) pricing strategy. In a competitive setting, Kopalle, Rao, and Assunção (1996) extend this result and show that the dynamic equilibrium pricing strategy is High-Low as shown in Figure 3.

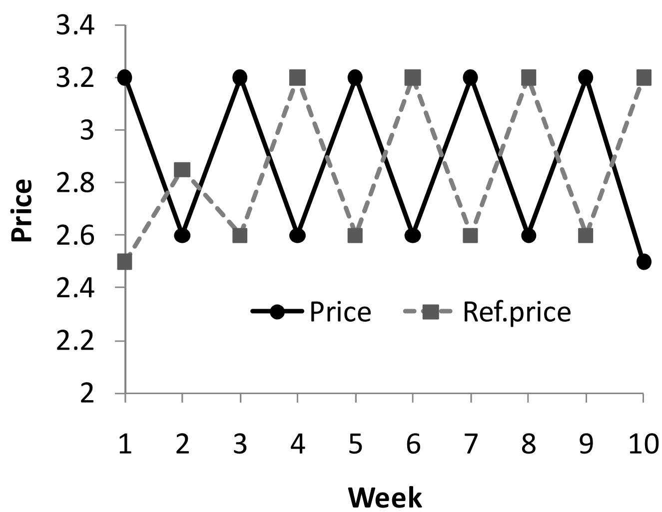

Figure 3: Equilibrium Prices Over Time

The other case is where customers are more loss averse (e.g., as found by Putler (1992) in the case of eggs)—i.e., the impact of a loss is greater than that of a corresponding gain ($\gamma > \delta$). It turns out that the Nash equilibrium solution is to maintain a constant pricing policy for both competitors. These results hold in an oligopoly consisting of *N* brands as well. Table 8 shows the two-period profitability of two firms in a duopoly. The first and second numbers in each cell refer to Firm 1's and Firm 2's profit over two periods. For both firms, constant pricing is the dominant strategy, and the Nash equilibrium is a constant price in both periods:

		Firm 2 Chooses…	
		Constant Pricing	**High-Low Pricing**
Firm 1 Chooses..	**Constant Pricing**	95 95	90 100
	High-Low Pricing	100 90	90 90

Table 8: Expected payoffs for the reference price game. Entries in each cell are payoffs to 1, 2.

In the more general case, where some consumers (say, segment 1) are loss averse ($\gamma_1>\delta_1$) and others (say, segment 2) are gain seeking ($\delta_2>\gamma_2$), Kopalle et al. (1996) find that in a competitive environment, a sufficient condition for a cyclical pricing policy to be a sub-game perfect Nash equilibrium is that the size of the gain seeking segment (i.e., non-loss-averse) segment is not too low. As seen in Figure 4, as the relative level of loss-aversion in segment 1 increases, i.e., $(\gamma_1-\delta_1)/(\delta_2-\gamma_2)$ increases, the size of the gain seeking segment needs to be larger for High-Low pricing to be an equilibrium.

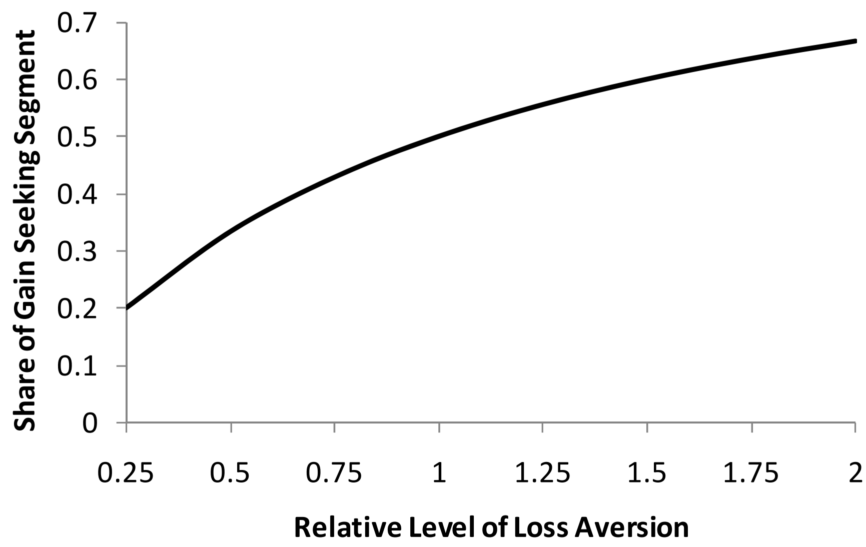

Figure 4: High-Low Equilibrium Pricing Policy and Share of Gain Seeking Segment

Finally, note that reference prices may also be influenced by the prices of competing brands. The equilibrium solutions of High-Low dynamic pricing policies hold under alternative reference-price formation processes, such as the reference brand context, where the current period's reference price is the current price of the brand bought in the last period.

3.4 Open- Versus Closed-Loop Equilibria

Two general kinds of Nash equilibria can be used to develop competing pricing strategies in dynamic games that are state-dependent: open loop and closed loop (Fudenberg and Tirole, 1991).

In a state-dependent dynamic game, an equilibrium is called an **open-loop equilibrium** if the strategies in the equilibrium strategy profile are functions of time alone.

In a state-dependent dynamic game, an equilibrium is called a **closed-loop** or a **feedback equilibrium** if strategies in the equilibrium strategy profile are functions of the history of the game, i.e., functions of the state variables.

The notation in dynamic models must be adapted to whether open or closed-loop strategies are used. Let S_{it} represent the i^{th} state variable at time t, and let P_c denotes the pricing strategy of competitor c. The notation $P_c(t)$ indicates an open-loop strategy, while a function $P_i(S_{it}, t)$ is a closed-loop strategy.

With open-loop strategies, competitors commit at the outset to particular time paths of pricing levels, and if asked at an intermediate point to reconsider their strategies, they will not change them. In an open-loop analysis of dynamic pricing in a duopoly, Chintagunta and Rao (1996) show that the brand with the higher preference level charges a higher price. In addition, they find that myopic price levels are higher than the corresponding dynamic prices. The authors also provide some empirical evidence in this regard by using longitudinal purchase data obtained from A. C. Nielsen scanner panel data on the purchases of yogurt in the Springfield, Missouri market over a two year period. Their analysis focuses on the two largest, competing brands, Yoplait (6 oz) and Dannon (8 oz). The empirical analysis shows that (i) the dynamic competitive model fits the data better than a static model, and (ii) the preference level for Dannon is noticeably higher than that for Yoplait. The equilibrium retail prices per unit are given in Table 9 under two different margin assumptions at the retailer.

Retail Margin	Brand	Dynamic Model	Static Model
40%	Yoplait	$0.46	$1.83
40%	Dannon	$0.61	$1.08
50%	Yoplait	$0.39	$1.72
50%	Dannon	$0.55	$1.05

Table 9: Open Loop Equilibrium Prices Versus Static Model Results

It turns out that the actual prices were fairly close to the equilibrium prices generated by the dynamic model, while the equilibrium static model prices are inflated relative to both the dynamic model as well as reality. Further, the higher preference level for Dannon is reflected in the higher pricing for the corresponding brand in the dynamic model results. The static model essentially ignores the evolution of consumer preference over time due to past purchases.

One drawback of open-loop Nash equilibria is that fixed, open-loop strategies cannot be modified on the basis of, say, current market share. Pricing managers, however, are unlikely to put their strategies on automatic pilot and ignore challenges that threaten their market positions. There are plenty of instances of companies and brands responding to competitive threats to their market shares—for example, the “leapfrog” pricing of Coke and Pepsi at stores, where Coke is on sale one week and Pepsi on sale the next. Closed-loop equilibria are more realistic than open-loop strategies because they do allow strategies to adjust to the current state of the market. Closed-loop equilibria, however, are harder to solve (Chintagunta and Rao 1996), and their analysis may require sophisticated mathematical techniques and approximations. For example, in an airline revenue management context, Gallego and Hu (2008) develop a game theoretic formulation of dynamically pricing perishable capacity (such as airplane seats) over a finite horizon. Since such problems are generally intractable, the authors provide sufficient conditions for the existence of open-loop and closed-loop Nash equilibria of the corresponding dynamic differential game by using a linear functional approximation approach to the solution. Dudey (1992) and Martinez-de-Albeniz and Talluri (2010) also describe dynamic revenue management games.

3.5 A Dynamic Stackelberg Pricing Game

The Stackelberg game defined in Section 2.5 has been applied to the sequential pricing process involving manufacturers and retailers. Initial research on this problem has focused on static Stackelberg games (Choi 1991; McGuire and Staelin 1983; Lee and Staelin 1997; Kim and Staelin 1999). In this game, the manufacturers maximize their respective brand profits, while the retailer maximizes its category profits, which include profits from sales of the manufacturer's brand as well as other brands in that category. The decision variables are the wholesale and retail prices, where the manufacturer decides the wholesale price first and then the retailer determines the retail price. The models described in the references above, however, do not include promotional dynamics over time.

To see how promotional dynamics can affect the game, consider a manufacturer that sells a single product (product 1) to a retailer. The retailer sells both product 1 and a competing brand, product 2, to consumers. Let $\bar{p}_k$ be the retailer's regular, undiscounted price of product k = 1, 2, and let p_{kt} be the retailer's chosen price in period *t*. A promotion, a temporary price reduction in product k, can lead to an immediate increase in sales, which we will capture in the demand function d_{kt} defined below. The promotion, however, may also have a negative impact on sales in future periods because consumers buy more than they need in period t (stockpiling), the discount may reduce brand equity, and/or the discount may reduce the product's reference price. We represent this effect of promotions in future periods with the term l_{kt}, k=1, 2:

$$l_{kt}=\alpha l_{k,t-1}+(1-\alpha)(\bar{p}_k-p_{kt}),\ 0\leq\alpha\leq 1.$$

We use l_{kt} in the following retail-level demand function, to capture this lagged effect of promotions on both baseline sales and price sensitivity. For (j,k) = (1,2) or (2,1),

$$d_{kt}=a_k-\left(\bar{b}_k+b_k l_{kt}\right)p_{kt}+\left(\bar{c}_k+c_k l_{jt}\right)p_{jt}-g_k l_{kt}. \quad (2)$$

The second term is the direct impact of product k's price on demand, where the price sensitivity is adjusted by the lagged effect of promotion, l_{kt}. The third term is the impact of product j's price on product k sales, adjusted by the effect of previous promotions for product j. The last term is the lagged effect of product k promotions on baseline sales of product k (baseline sales are sales that would occur when there are no promotions).

Now we describe the objective functions of the manufacturer and retailer. Within each period t the manufacturer first sets a wholesale price w_{1t} for product 1. Given marginal product cost c_1 and discount rate β (and ignoring fixed costs), the manufacturer's objective over a finite period t = 1…T is,

$$\max_{w_{1t}}\sum_{t=1}^{T}\beta^t(w_{1t}-c_1)d_{1t}.$$

Note that d_{1t} depends upon the prices p_{1t} and p_{2t} that are chosen by the retailer in each period *t* after the manufacturer sets its wholesale price. The retailer's objective is,

$$\max_{p_{1t},p_{2t}} \sum_{t=1}^{T} \beta^t \sum_{k=1}^{2} (p_{kt} - w_{kt}) d_{1t} .$$

Kopalle, Mela, and Marsh (1999) describe a more general version of this game, with any number of retailers, any number of competing brands, and controls for seasonality, features and displays. Their model also has an explicit term for the stockpiling effect, although when the model was fit to data, the direct stockpiling effect was not significant. Finally, the model in Kopalle et al. (1999) differs in a few technical details, e.g., the expressions above are in terms of the logarithm of demand and price.

This interaction between manufacturer and retailer is an example of a closed-loop, dynamic Stackelberg game.

> In a **dynamic Stackelberg game**, the leader and the follower make decisions in each period and over time based on the history of the game until the current period.

In the one-period Stackelberg game of Section 2.5, the leader (in this case, the manufacturer) only anticipates the single response from the follower (the retailer). In this dynamic game, the leader must anticipate the retailer's reactions in all future periods. Also, in the one-period game the follower simply responds to the leader's choice – the decision is not strategic. In this dynamic game, when the retailer sets its price in period *t*, the retailer must anticipate the manufacturer's response in future periods. As in our previous dynamic games, one can solve this game using Bellman's (1957) principle of optimality and backward induction. This procedure determines the manufacturer's and retailer's equilibrium pricing strategies over time.

Before solving for the manufacturer's and retailers' strategies, Kopalle et al. (1999) estimate the parameters of the demand model, using one hundred twenty-four weeks of A. C. Nielsen store-level data for liquid dishwashing detergent. The results suggest that when a brand increases use of promotions, it reduces its baseline sales; increases price sensitivity, thus making it harder to maintain margins; and diminishes its ability to use deals to take share from competing brands. Figure 5 shows a typical equilibrium solution for manufacturer's wholesale prices and the corresponding brand's retail prices over time in the game.

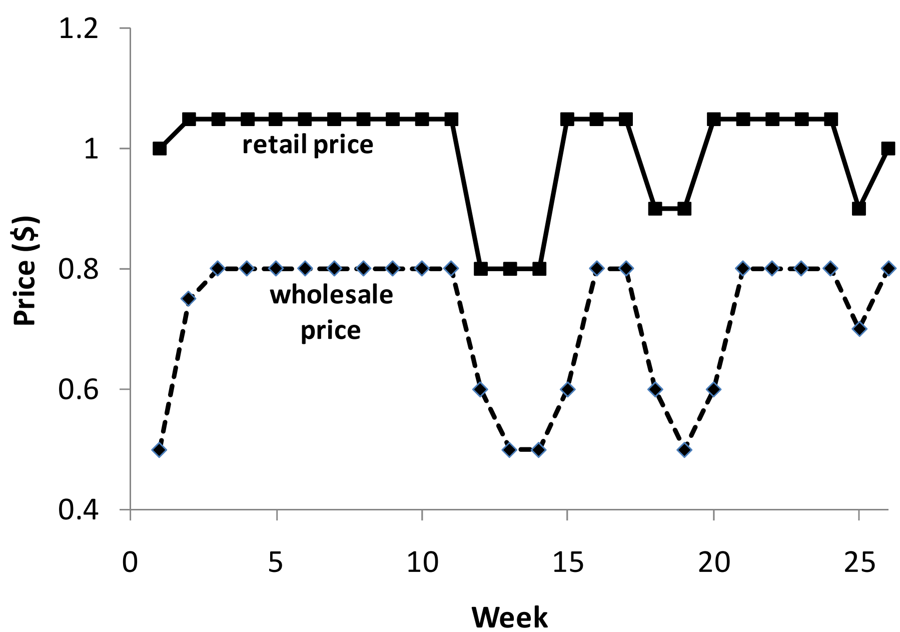

Figure 5: Equilibrium Wholesale and Retail Prices

Kopalle et al. (1999) find that many other brands have flat equilibrium price paths, indicating that for these brands, the manufacturer and retailer should curtail promotion because the immediate sales increase arising from the deal is more than eliminated by future losses due to the lagged effects of promotions.

In practice, estimates of baseline sales and responses to promotions often ignore past promotional activity. In other words, the lagged promotional effects l_{jt} and l_{kt} in demand function (2) are ignored. To determine the improvement in profits due to using a dynamic model rather than a static model, one must compute the profits a brand would have had, assuming that it had estimated a model with no dynamic effects. Comparing the static and dynamic Stackelberg solutions, Kopalle et al. (1999) find that the use of a dynamic model leads to a predicted increase in profits of 15.5 percent for one brand and 10.7 percent for another. These findings further suggest that managers can increase profits by as much as 7 percent to 31 percent over their current practices and indicate the importance of balancing the trade-off between increasing sales that are generated from a given discount in the *current period* and the corresponding effect of reducing sales in *future periods.*

Promotions may stem from skimming the demand curve (Farris and Quelch, 1987), from the retailer shifting inventory cost to the consumer (Blattberg et al., 1981), from competition between national and store brands (Lal, 1990a; Rao, 1991), or from asymmetry in price response about a reference price (Greenleaf 1995). Kopalle et al. (1999) show that another explanation for dynamic pricing exists: the trade-off between a promotion's contemporaneous and dynamic effects.

3.6 Testing Game—Theoretic Predictions in a Dynamic Pricing Context: How Useful Is a Game-Theory Model?

Prior research has illustrated some of the challenges of predicting competitive response. For example, a change by a firm in marketing instrument X may evoke a competitive change in marketing instrument Y (Kadiyali, Chintagunta, and Vilcassim, 2000; Putsis and Dhar, 1998). Another challenge is that firms may "over-react" or "under-react" to competitive moves (Leeflang and Wittink, 1996). Furthermore, it is unclear whether managers actually employ the strategic thinking that game theoretic models suggest. Research by Montgomery, Moore, and Urbany (2005) indicates that managers often do not consider competitors' reactions when deciding on their own moves, though they are more likely to do so for major, visible decisions, particularly those pertaining to pricing. Therefore, it is reasonable to ask whether game-theory models can be useful when planning pricing strategies and whether they can accurately predict competitor responses. Here we describe research that examines whether the pricing policies of retailers and manufacturers match the pricing policies that a game theory model predicts. The research also examines whether it is important to make decisions using the game theory model rather than a simpler model of competitor behavior, for example, one that simply extrapolates from previous behavior to predict competitor responses.

In one research study, Ailawadi, Kopalle, and Neslin (2005) use a dynamic game theory model to consider the response to P&G's "value pricing" strategy, mentioned at the beginning of this chapter. The high visibility of P&G's move to value pricing was an opportunity to see whether researchers can predict the response of competing national brands and retailers to a major policy change by a market leader. The researchers empirically estimate the demand function that drives the model, substitute the demand parameters into the game theoretic model, and generate predictions of competitor and retailer responses. To the extent their predictions correspond to reality, they will have identified a methodological approach to predicting competitive behavior that can be used to guide managerial decision making.

Ailawadi et al. (2005) base their game theoretic model on the Manufacturer-Retailer Stackelberg framework similar to Kopalle et al. (1999), described in Section 3.5. Their model is significantly more comprehensive in that it endogenizes (a) the national brand competitor's price and promotion decisions; (b) the retailer's price decision for P&G, the national brand competitor, and the private label; (c) the retailer's private-label promotion decision; and (d) the retailer's forward-buying decision for P&G and the national brand competitor. This model of the *process* by which profit-maximizing agents make decisions includes *temporal* response phenomena and decision making. In this sense, it is a "dynamic structural model." Such models are particularly well suited to cases of major policy change because they are able to predict how agents adapt their decisions to a new "regime" than are reduced-form models that attempt to extrapolate from the past (Keane, 1997). However, they are more complex and present researchers with the thorny question of how all-encompassing to make the model.

To develop their model, Ailawadi et al. (2005) use the scanner database of the Dominick's grocery chain in Chicago and wholesale-price and trade-deal data for the Chicago market from Leemis Marketing Inc. Data were compiled for nine product categories in which P&G is a player. Dominick's sold a private label brand in six of these nine categories throughout the period of analysis. The number of brands for which both wholesale and retail data are available varies across categories, ranging from three to six. In total, there were 43 brands across the nine categories (nine P&G brands, 28 national brand competitors, and six private label competitors).

	Dynamic Game Theoretic Model			**Statistical Model of Competition Dynamics**		
	Deal Amount	**Wholesale Price**	**Retail Price**	**Deal Amount**	**Wholesale Price**	**Retail Price**
% Directionally Correct Predictions	78%	71%	78%	52%	48%	55%
% Directionally Correct Predictions When Actual Values Increased	79%	50%	87%	46%	25%	64%
% Directionally Correct Predictions When Actual Value Decreased	77%	93%	65%	60%	67%	36%

Table 10: Predictive Ability of A Dynamic Game-Theoretic Model Versus A Statistical Model Competition Dynamics

As seen in Table 10, their results show that a game theoretic model combined with strong empirics does have predictive power in being able to predict directional changes in wholesale deal amount given to the retailer by the manufacturer, wholesale price, and retail price. The predictive ability of the directional changes ranges from 50% to 93%. The table shows that of the times when the game theoretic model predicted that competitors would increase wholesale deal amount, competitors actually increased it 79% of the time. This compares with correct predictions of 46% for the statistical model of competitor behavior, which is essentially is a statistical extrapolation of past behavior. The predictions of another benchmark model, where the retailer is assumed to be non-strategic, does no better with 46% correct predictions.

Their model's prescription of how competitors should change wholesale price and trade dealing and how the retailer should change retail price in response to P&G's value-pricing strategy is a significant predictor of the changes that competitors and Dominick's in the Chicago market actually made. This suggests that, in the context of a major pricing policy change, managers' actions are more consistent with strategic competitive reasoning than with an extrapolation of past reactions into the future or with ignoring retailer reaction.

In the wake of its "Value Pricing" strategy, it turns out that P&G suffered a 16% loss in market share across 24 categories (Ailawadi, Lehmann, and Neslin 2001). If P&G's managers had built a dynamic game theoretic model, they might have been able to predict (at least directionally) the categories in which the competitors would follow suit versus retaliate. This could have helped P&G to adopt a more surgical approach to its value pricing strategy instead of a blanket strategy that cut across many categories.

The above test was for a major policy change and the findings support the view of Montgomery et al. (2005) that when faced with highly visible and major decisions, especially with respect to pricing, managers are more likely to engage in strategic competitive thinking. It also supports the premise that when the agents involved are adjusting to a change in regime, structural models are preferable to reduced-form models. Other approaches such as the reaction function approach (Leeflang and Wittink 1996) or a simplified game-theoretic model (for example, assuming that the retailer is non-strategic and simply uses a constant mark-up rule for setting retail prices) may have better predictive ability for ongoing decisions and week-to-week or month-to-month reactions where managers are less likely to engage in competitive reasoning.

4. Implications for Practice

Firms have many opportunities to increase profitability by improving the way they manage price. Many need better systems, such as enterprise software to manage strategy effectively and optimize execution to support their decisions. Vendors provide tools that monitor customer demand and measure how price is perceived compared to other retailers, helping to optimize pricing. Assortment and space-optimization tools help retailers manage their merchandise and shelf space, but these tools lack sophisticated predictive analytics. Although retailers can access a range of data, such as point-of-sale, packaged goods product information, competitor prices, weather, demographics, and syndicated market data, they do not have the tools they need to integrate all these data sources to manage strategies and optimize price in a competitive setting. In particular, current pricing systems do not incorporate the strategic impact of pricing changes – the dynamics described by game theory. The flowchart of Figure 6 is adapted from Kopalle et al. (2008), and envisions a pricing system with modules that incorporate data-driven competitive analysis. In the diagram, the “Strategic Framework” module contains the underlying model of competition, and is driven by data on competitors.

The analysis conducted in the Strategic Framework module would focus on a variety of questions, guided by the models introduced in this chapter. In Table 11 we lay out the frameworks, using this chapter's taxonomy of models. For the simultaneous, one-shot pricing games in Box 1 (Nash games), the impact of product differentiation, and whether the products are substitutes or complements become important issues to consider. As discussed earlier in the chapter, there are advantages of moving first, thereby changing the game from a simultaneous to sequential game. For the one-period sequential (Stackelberg) game in Box 2, while the questions for Box 1 are still relevant, it is also important for a firm to recognize whether it is a leader or follower. If it is a leader, it should anticipate the followers' reactions to its moves. If a follower, a firm should determine if it is advantageous to become a leader (sometimes it is not), and consider how.

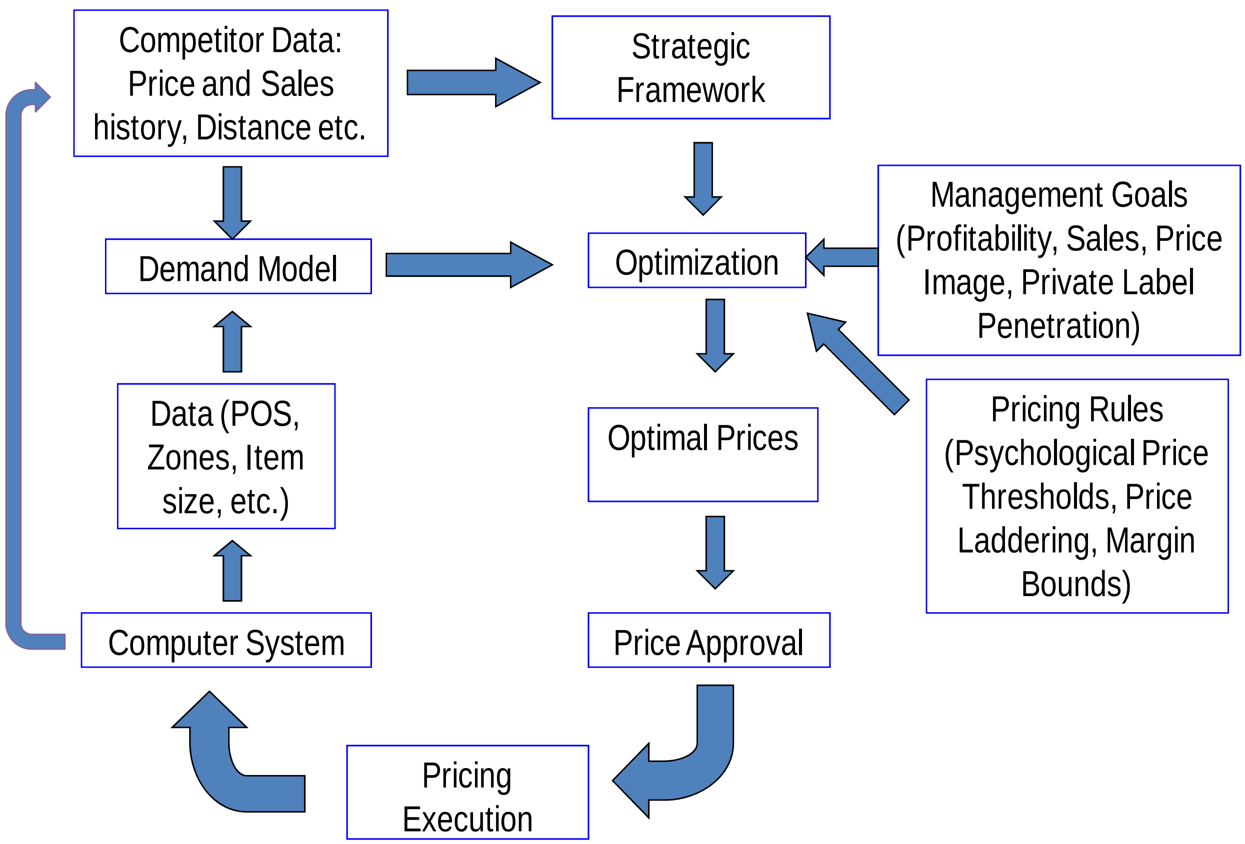

Figure 6: Pricing algorithm

Number of periods		Timing of actions: Simultaneous	Timing of actions: Sequential
	One	**Box 1:** Static (one-shot) Nash games	**2:** Stackelberg games
	Multiple	**3:** Repeated and dynamic games	**4:** Dynamic Stackelberg games

Table 11: Strategic frameworks

Box 1's questions also apply to Box 3, which refers to simultaneous games over multiple periods. Here firms should also consider the degree to which their competitors take a strategic, forward-looking view. In addition, recall that in a retail application, the optimal strategy is determined by the degree to which reference prices influence customer purchasing behavior, as well as the strength of loss aversion among customers. Finally, we saw that when players repeat a game over multiple periods, trigger strategies (such as tit-for-tat) may be important. In box 4, leaders in dynamic Stackelberg games should consider the impact of promotions on followers' actions and on future profits. In contrast to the single-period game, followers cannot ignore the future; they should consider the impact of their actions on leaders' actions in later periods.

For all of these environments (boxes 1-4), fruitful application of game theory to pricing requires that a firm understand its competitions' pricing strategy. For retailers in particular, the impact of competition has been well documented (Lal and Rao, 1997; Moorthy, 2005; Bolton and Shankar, 2003; Shankar and Bolton, 2004), but a lack of knowledge of their competitors' prices and strategies hinder retailers. Such knowledge is usually difficult to obtain, and that is one reason why firms either (i) implement pricing systems that do not attempt to take competitor strategies directly into account (perhaps by assuming that competitors' current actions will not change in the future), or (ii) implement simple, myopic strategies, given competitors' actions. If all firms in an industry choose (i), then they are still implicitly playing a game – they are responding to their competitors' actions indirectly, using information obtained by measuring changes in their own demand functions. Recent research has shown that strategies based on models that ignore competitors' actions can lead to unpredictable, and potentially suboptimal, pricing behavior (e.g., see Cooper et al. (2009)). If all firms in an industry choose (ii), however, the results depend upon the myopic strategy. If firms only match the lowest competitor price (and, say, ignore product differentiation), then firms simply drive prices down. There is a stream of research that shows that firms in an oligopoly who choose the "best response" to the current prices of their competitors will result in a sequence of prices that converges to the unique Nash equilibrium from the simultaneous game (Gallego et al., 2006).

The models and frameworks presented in this chapter are most useful when a firm can understand and predict its competitors' pricing strategies. With improving price visibility on the web and masses of data collected by retail scanners, there is much more data available about competitors. Market intelligence services and data scraping services collect and organize such information. Examples include QL2.com for the travel industries and Information Resources Inc. for the consumer packaged goods industry. In the retailing world, some firms use hand held devices to track competitors' prices of key items.

It is difficult, however, to distill raw data into an understanding of competitor strategy. In some industries, conferences can serve to disseminate price strategy information. In the airline industry, for example, revenue management scientists attend the AGIFORS conference (sponsored by the Airline Group of the International Federation of Operational Research Societies) and share information about their pricing systems (this is not collusion – participants do not discuss prices, pricing tactics, or pricing strategies but do discuss models and algorithms). Consider another example from outside the airline industry: when one major firm implemented a revenue management system, they were concerned that their competitors would interpret their targeted local discounting as a general price "drop" and that a damaging price war would ensue. A senior officer of the firm visited various industry conferences to describe revenue management at a high level to help prevent this misconception on the part of the competitors. In the abstract, such information may be used to design pricing systems that are truly strategic.

Of course, game theory motivates us to ask whether providing information about our pricing system to a competitor will improve our competitive position. There are certainly games where knowledge shared between competitors can improve performance for both players: compare, for example, the one-shot prisoner's dilemma (with its equilibrium in which both 'confess') with a repeated game in which competitors anticipate each others' strategies and settle into an equilibrium that is better for both.

In general, our hope is that this chapter will help nudge managers toward making analytics-based pricing decisions that take into account competitor behavior. Further, our view is that a significant barrier to the application of game theory to pricing is the extent to which mathematical and behavioral game theory, developed at the individual level, is not applicable to firm-level pricing decisions. In particular, a pricing decision is the outcome of group dynamics within the firm, which can be influenced by a range of factors that are different from factors that drive individual behavior. Recall from this chapter the research on individual behavior in laboratory pricing games (e.g., Duwfenberg and Gneezy, 2000; Camerer, 2003) as well as one attempt to empirically validate a game theory model of pricing at the firm level (Ailawadi et al., 2005). There has been little research in between these extremes. More work is needed to identify the general characteristics of firms' strategic pricing behavior, and we believe that there is a significant opportunity for research to bridge this gap.

Acknowledgements

We would like to thank Ozalp Ozer, Robert Phillips, and an anonymous referee for their helpful suggestions.

References

Ailawadi, K., P.K. Kopalle, and S.A. Neslin. (2005). 'Predicting Competitive Response to a Major Policy Change: Combining Game Theoretic and Empirical Analyses'. *Marketing Science*, 24/1: 12–24;

Ailawadi, K., D. R. Lehmann, S. A. Neslin. (2001). 'Market Response to a Major Policy Change in the Marketing Mix: Learning from Procter & Gamble's Value Pricing Strategy'. *Journal of Marketing*, 65/January: 44-61;

Associated Press. (2007). 'British Airways and Korean Airlines Fined in Fuel Collusion'. *The New York Times*, August 2, 2007;

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

Axelrod, R., and Hamilton, W.D. (1981). 'The Evolution of Cooperation'. *Science*, 211/4489: 1,390–1,396;

Armstrong, Mark and Steffen Huck. (2010). 'Behavioral Economics as Applied to Firms: A Primer'. CESifo Working Paper Series No. 2937. Accessed June 8, 2010, http://ssrn.com/abstract=1553645.

Aumann, R. (1976). 'Agreeing to Disagree'. *The Annals of Statistics*, 4/6: 1236-1239.

Bellman, Richard. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University Press.

Bertsekas, Dimitri P. (1987). *Dynamic Programming: Deterministic and Stochastic Models*. Englewood Cliffs, NJ: Prentice Hall;

Binmore, K. (2007). *Does Game Theory Work?* Cambridge, MA: The MIT Press;

Blattberg, R.C., G.D. Eppen, and J. Lieberman. (1981). 'Price-Induced Patterns of Competition'. *Marketing Science*, 8/Fall: 291–309;

Bolton, Ruth N. and Venkatesh Shankar. (2003). 'An Empirically Derived Taxonomy of Retailer Pricing and Promotion Strategies'. *Journal of Retailing*, 79/4: 213–224;

Briesch, Richard A., Lakshman Krishnamurthi, Tridib Mazumdar, and S. P. Raj. (1997). 'A Comparative Analysis of Reference Price Models'. *Journal of Consumer Research*. 24/2: 202-214;

BTNonline.com. (2008). 'Continental Plans UAL Joint Venture, Star Membership.' June 19, 2008. Accessed May 24, 2010, http://www.btnonline.com/businesstravelnews/headlines/article_display.jsp?vnu_content_id=1003818864;

Cachon, G.P., and S. Netessine. (2006). ‘Game Theory in Supply Chain Analysis’. INFORMS *TutORials in Operations Research*, Paul Gray, Series Editor: 200–233;

Camerer, Colin. (2003). *Behavioral Game Theory* Princeton, NJ: Princeton University Press.

Chintagunta, Pradeep K. and Vithala R. Rao. (1996). ‘Pricing Strategies in a Dynamic Duopoly: A Differential Game Model’. *Management Science*, 42/11: 1501-1514;

Choi, S. Chan. (1991). ‘Price Competition in a Channel Structure with a Common Retailer’. *Marketing Science*, 10/Fall: 271-296;

Cooper, W. L., Homem-de-Mello, T., and Kleywegt, A. J. (2009). ‘Learning and Pricing with Models that Do Not Explicitly Incorporate Competition’. Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia. Accessed June 11, 2010, http://www2.isye.gatech.edu/~anton/competitive.pdf;

DOT (Department of Transportation). (2009). ‘Joint Application of Air Canada, The Austrian Group, British Midland Airways LTD, Continental Airlines, Inc., Deutsche Lufthansa AG, Polskie Linie Lotniecze LOT S.A., Scandinavian Airlines System, Swiss International Air Lines Ltd., TAP Air Portugal, United Airlines Inc.to Amend Order 2007-2-16 under 49 U.S.C. 41308 and 41309 so as to Approve and Confer Antitrust Immunity.’ Order: 2009-7-10. Accessed May 24, 2010, http://www.regulations.gov/fdmspublic/ContentViewer?objectId=09000064809eda55&disposition=attachment&contentType=pdf;

d'Huart, Oliver. (2010). *A Competitive Approach to Airline Revenue Management.* Unpublished Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Dixit, Avinash K. and Barry J. Nalebuff. (1991). *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*. New York: W. W. Norton & Company;

Dixit, A. and S. Skeath. (1999). *Games of Strategy*. New York: W. W. Norton & Company;

Dudey, M. (1992). 'Dynamic Edgeworth-Bertrand competition'. *Quarterly Journal of Economics*, 107:1461-1477;

Duwfenberg, M. and U. Gneezy. (2000). 'Price Competition and Market Concentration: An Experimental Study'. *International Journal of Industrial Organization*, 18: 7-22;

Farris, Paul, and John A. Quelch. (1987). 'In Defense of Price Promotions'. MIT *Sloan Management Review*, 29/Fall: 63–69;

Fudenberg, D., and J. Tirole. (1991). *Game Theory*. Cambridge, MA: The MIT Press;

Gallego, Guillermo and Ming Hu. (2008). "Dynamic Pricing of Perishable Assets Under Competition." *SSRN Working Paper*, http://ssrn.com/abstract=1308848;

G. Gallego, T. Huh, W. Kang, and R. Phillips. (2006). 'Price competition with the attraction demand model: Existence of unique equilibrium and its stability'. *Manufacturing and Service Operations Management*, 8:359-375;

Gibbons, R. (1992). *Game Theory for Applied Economists*. Princeton, NJ: Princeton University Press;

Govindan S, and R. Wilson. (2005). ‘Refinements of Nash equilibrium’. In: Durlauf S, Blume L (eds) *The New Palgrave Dictionary of Economics 2*. London: Palgrave Macmillan.

Greenleaf, Eric A. (1995). ‘The Impact of Reference Price Effects on the Profitability of Price Promotions’. *Marketing Science*, 14/Winter: 82–104;

Harsanyi, J. (1967). ‘Games with Incomplete Information Played by Bayesian Players Parts I, II and III.’ *Management Science*, 14/3,5,7: 159-182, 320-34, 486-502;

Jiang, Houyuan and Zhan Pang (2007). ‘Network Capacity Management under Competition’. Judge Business School Working Papers, No.14/2007. University of Cambridge;

Kadiyali, Vrinda, Pradeep Chintagunta, and Naufel Vilcassim. (2000). ‘A Manufacturer-Retailer Channel Interaction and Implications for Channel Power: An Empirical Investigation of Pricing in a Local Market’. *Marketing Science*, 19/2: 127–148;

Kalyanaram, Gurumurthy and John D. C. Little. (1994). ‘An Empirical Analysis of Latitude of Price Acceptance in Consumer Packaged Goods’. *Journal of Consumer Research*, 21/3: 408-418;

Keane, M. (1997). ‘Modeling Heterogeneity and State Dependence in Consumer Choice Behavior’. *Journal of Business and Economic Statistics*, 15/3: 310–327;

Kim, Sang Yong and Richard Staelin. (1999). ‘Manufacturer Allowances and Retailer Pass-Through Rates in a Competitive Environment’. *Marketing Science*, 18 /1: 59-76;

Kopalle, Praveen. K., Dipayan Biswas, Pradeep K. Chintagunta, Jia Fan, Koen Pauwels, Brian Ratchford, and Jim Sills. (2008). 'Pricing and Competitive Effects'. *Journal of Retailing*, forthcoming;

Kopalle, Praveen K., Ambar G. Rao, and João L. Assunção. (1996). 'Asymmetric Reference Price Effects and Dynamic Pricing Policies'. *Marketing Science*, 15/1: 60–85;

Kopalle, P.K., C.F. Mela, and L. Marsh. (1999). 'The Dynamic Effect of Discounting on Sales: Empirical Analysis and Normative Pricing Implications'. *Marketing Science*, 18/3: 317–332;

Lal, Rajiv. (1990a). 'Price Promotions: Limiting Competitive Encroachment'. *Marketing Science*, 9 /3: 247-262;

Lal, Rajiv. (l990b). 'Manufacturer Trade Deals and Retail Price Promotions'. *Journal of Marketing Research*, 27/November: 428–444;

Lal, Rajiv, and Ram Rao. (1997). 'Supermarket Competition: The Case of Every Day Low Pricing'. *Marketing Science*, 16/1: 60–80;

Lee, Eunkyu and Richard Staelin. (1997). 'Vertical Strategic Interaction: Implications for Channel Pricing Strategy'. *Marketing Science*, 16 /3: 185-207;

Lee, H.L., V. Padmanabhan, and S. Whang. (1997). 'The Bullwhip Effect in Supply Chains'. *MIT Sloan Management Review*, 38/3: 93;

Leeflang, Peter S.H., and Dick R. Wittink. (1996). 'Competitive Reaction Versus Consumer Response: Do Managers Overreact?' *International Journal of Research in Marketing*, 13: 103–119;

Martinez-de-Albeniz, Victor and Kalyan Talluri. (2010). 'Dynamic Price Competition with Fixed Capacities'. Universitat Pompeu Fabra Department of Economics and Business, Working Paper 1205. Accessed June 10, 2010, http://www.econ.upf.edu/docs/papers/downloads/1205.pdf.

McGuire, Timothy W., and Richard Staelin. (1983). 'An Industry Equilibrium Analysis of Downstream Vertical Integration'. *Marketing Science*, 2/2: 161–192;

Milgrom, P. (1981). 'An Aximomatic Characterization of Common Knowledge'. *Econometrica*, 49/1: 419-422.

Montgomery, David B., Marian Chapman Moore, and Joel E. Urbany. (2005). 'Reasoning about Competitive Reactions: Evidence from Executives'. *Marketing Science*, 24/1: 138–149;

Moorthy, K.S. (1985). 'Using Game Theory to Model Competition'. *Journal of Marketing Research*, 22: 262–282;

Moorthy, Sridhar. (2005). 'A General Theory of Pass-Through in Channels with Category Management and Retail Competition'. *Marketing Science*, 24/1: 110–122;

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press;

Nash, J. (1950). ‘Equilibrium Points in N-Person Games’. *Proceedings of the National Academy of Sciences of the United States of America*, 36/1: 48–49;

Netessine, S., and R.A. Shumsky. (2005). ‘Revenue Management Games: Horizontal and Vertical Competition’. *Management Science*, 51/5: 813-831;

Neumann, John von and Oskar Morgenstern. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press;

Nisan, N., T. Roughgarden, E. Tardos, and V. V. Vazirani (eds.) (2007). *Algorithmic Game Theory.* Cambridge, MA: Cambridge University Press;

Özer, Özalp, Yanchong Zheng, and Kay-Yut Chen. (2009). ‘Trust in Forecast Information Sharing’. *SSRN Working paper*, http://ssrn.com/abstract=1477563;

Putler, D. (1992). ‘Incorporating Reference Effects into a Theory of Consumer Choice’. *Marketing Science*, 11/Summer: 287–309;

Putsis, William, and Ravi Dhar. (1998). ‘The Many Faces of Competition’. *Marketing Letters*, 9/3: 269–284;

Raju, Jagmohan S., V. Srinivasan, and Rajiv Lal. (1990). ‘The Effects of Brand Loyalty on Competitive Price Promotions’. *Management Science*, 36/3: 276–304;

Rao, Ram C. (1991). ‘Pricing and Promotions in Asymmetric Duopolies’. *Marketing Science*, 10/Spring: 131–144;

Rao, Ram C. and Frank M. Bass. (1985). ‘Competition, Strategy, and Price Dynamics: A Theoretical and Empirical Investigation’. *Journal of Marketing Research*, 22/August: 283-296;

Rappeport, Alan. (2008). ‘Game Theory Versus Practice’. CFO Magazine, July 15, 2008;

Saporito, B. (1994). ‘Behind the Tumult at P&G’. Fortune, 129/5: 74;

Schelling, T. (1960). The Strategy of Conflict. Cambridge, MA: Harvard University Press;

Selton, R. (1975). ‘Reexamination of the perfectness concept for equilibrium points in extensive games’. International Journal of Game Theory, 4: 25-55;

Shankar, Venkatesh, and Ruth N. Bolton. (2004). ‘An Empirical Analysis of Determinants of Retailer Pricing Strategy’. Marketing Science, 23/1: 28–49;

Singh, N. and X. Vives. (1984). ‘Price and Quantity Competition in a Differentiated Duopoly’. Rand Journal of Economics, 15/4: 546-554;

Varian, Hal. (1980). ‘A Model of Sales’. American Economic Review, 70/September: 651–659;

Winer, Russell S. (1986). ‘A Reference Price Model of Brand Choice for Frequently Purchased Products’. Journal of Consumer Research, 13/2: 250-256.