INTRODUCTION TO GAME THEORY & THE BERTRAND TRAP

Benjamin E. Hermalin
University of California, Berkeley

Contents

1 Introduction **1**

2 Introduction to Game Theory **1**
2.1 The Prisoners' Dilemma 1
2.2 Best Responses and Dominant Strategies 3
2.3 Nash Equilibrium 4

3 The Bertrand Model of Competition **8**
3.1 Assumptions Underlying the Bertrand Model 8
3.2 Equilibrium in the Bertrand Model 10
3.3 Two Ghosts 12

4 More Game Theory **12**
4.1 The Monitoring Game 12
4.2 An Entry Game 15

5 Summary **16**

6 Exercises **17**

1 Introduction

Most firms face competition selling their goods and services. Coca-Cola duels with Pepsi (and a handful of small firms) in the soft-drink market. Anheuser-Busch battles Miller and Coors, as well as a slew of smaller brewers, in the beer market. In Japan, Kajima competes against Shimizu, Obayashi, Taisei, and Kumagi Gumi (and a number of other builders) in the construction industry. Locally, area supermarkets and groceries are engaged in some level of competition.

Competition among firms—indeed, just the potential for competition—can have a serious impact on firm behavior. For instance, if Coca-Cola lowers it price, it is probably *un*reasonable for it to assume that Pepsi won't react: By lowering its price, Coca-Cola will not only attract new cola drinkers, but it will also lure away some of Pepsi's customers. Unless Pepsi is willing to put up with a loss of revenue—an unlikely scenario—it will find itself obligated to respond to Coca-Cola's price cut. Indeed, whatever it might do, its *reaction* to Coca-Cola's price cut will affect the value of the price cut to Coca-Cola. Analyzing and anticipating its rival's reaction must, therefore, play a large role in Coca-Cola's decision making.

In this reading, we start our study of how managers make decisions when their firms face rivals or potential rivals. We begin with a set of tools—game theory—that will allow us to study strategic situations in greater depth and with more insight. Next we introduce a model of competition in which firms compete solely on the basis of price—go "head-to-head" on price. In economics, this model is called the Bertrand model of competition. As we will see, *if* the Bertrand model describes your competitive situation, you're not going to do too well. For this reason, we will describe the Bertrand model as the "Bertrand trap"—a situation that you wish to avoid being caught in. We'll wrap up this chapter by considering some more game theory.

2 Introduction to Game Theory

To better analyze strategic situations, economists and strategists have come to use a set of tools known as *game theory*. Game theory is an extension of decision theory to situations in which there is more than one decision maker (now called a *player*).

game theory

player

2.1 The Prisoners' Dilemma

The best way to introduce game theory is through an example. Figure 1 illustrates a "game." Two firms, Row Inc. and ColumnCo compete to sell a product. Each firm has a choice of strategies: It can be "aggressive" or "passive." Aggressive corresponds to behaviors such as big advertising campaigns, strong discounting, etc. Passive corresponds to more conventional and less extreme behaviors. The way the game is played is that both firms simultaneously select

		ColumnCo Aggressive	ColumnCo Passive
Row Inc.	Aggressive	$3 $3	$2 $6
	Passive	$6 $2	$4 $4

Figure 1: Game between Row Inc. and ColumnCo (payoffs in millions)

their strategies. Their profits—*payoffs*—as a function of the chosen strategies are shown in Figure 1. The number in the lower left corner of each cell is the payoff to Row Inc. and the number in the upper right corner of each cell is the payoff to ColumnCo. Hence, if Row Inc. is aggressive and ColumnCo is passive, then Row Inc.'s profit is $6 million and ColumnCo's profit is $2 million. The question, now, is what strategy will each firm choose? PAYOFFS

To answer this question, imagine that you're the CEO of Row Inc. What should you do if ColumnCo will be passive? If you're aggressive, then your firm earns $6 million; but if you're passive, then your firm earns just $4 million. Hence, if you think ColumnCo will be passive, you should be aggressive. But what if you think it will be aggressive? If you're also aggressive, then your firm earns $3 million; but if you're passive, then your firm earns just $2 million. Again, you should be aggressive. We've just seen that, no matter what you anticipate ColumnCo's strategy to be, your best strategy as Row Inc.'s CEO is to be aggressive. Hence, we would predict that Row Inc. will be aggressive. What should we anticipate ColumnCo will do? The analysis that we did as Row Inc.'s CEO can be repeated for ColumnCo. Given the symmetry of the game, it should be clear that such an analysis would reveal that ColumnCo should be aggressive no matter what it anticipates Row Inc.'s strategy will be. In summary, then, we've concluded that both firms will choose the aggressive strategy.

Essentially, this what game theory is about.[1] A multiple-decision-maker situation is described in terms of the strategies available to each decision maker (here, "aggressive" and "passive") and the payoffs for each decision maker as a function of the strategies (here, the eight monetary amounts in Figure 1). Having described the situation—or game—the next step is to solve it; that is, make predictions as to how the players will play it (what strategies they will choose).

Before proceeding with a more general analysis of games, we should first analyze the outcome predicted for the game in Figure 1. Observe, first, that if both firms are aggressive—as we've predicted they will be—then each makes only $3 million in profit. If, in contrast, they were each passive, they would make $4 million. Hence, we see that *un*like single decision-maker problems, rational play does *not* necessarily produce desirable outcomes with multiple

[1] To be 100% accurate, we should add "in *normal form*."

decision makers. Indeed, if somehow the firms could cooperate (which might violate antitrust laws!), they would do better agreeing to be passive than being aggressive. This illustrates the

Fundamental Rule of Competitive Strategy

> **Fundamental Rule of Competitive Strategy**: *Competitive strategy is like driving, not football—head-on collisions are to be avoided.*

The dilemma here is that what is in each firm's private interest is not in their collective interest. This dilemma is so common to strategic situations that it has a name: It's the *prisoners' dilemma*.[2,3]

Prisoners' dilemma: *A game in which the strategies that are* privately *optimal for the players lead to an outcome that is* not mutually *optimal.*

2.2 Best Responses and Dominant Strategies

Turning, now, to a more systematic approach to game theory, let's review how we solved the version of the prisoners' dilemma shown in Figure 1. For each possible strategy of the opponent (ColumnCo in the preceding analysis), we asked what was *our* best strategy if the opponent played that strategy. So we first asked what was our best strategy if ColumnCo played passive. Since \$6 million is greater than \$4 million, we concluded our best strategy was aggressive. In the parlance of game theory, we concluded that aggressive was our *best response* to passive. That is, in a game between two players—#1 and #2—a given strategy for #1 is a best response to a specific strategy of #2's if #1's payoff from playing the given strategy is at least great as her payoff from playing any other of her strategies if #2 plays that specific strategy. Similarly, when we presumed ColumnCo would play aggressive(ly), we observed that \$3 million is greater than \$2 million, so that aggressive was again Row Inc.'s best response to aggressive by ColumnCo. We can extend this notion of best response to any number of strategies and players:

best response

> **Best Response:** *Consider a game with N players. A* best response *for player n to the $N-1$ strategies of her opponent is a strategy for her that maximizes her payoff against these $N-1$ strategies of her opponents.*

In the Row Inc. vs. ColumnCo game, aggressive was Row Inc.'s best response to *both* strategies of its opponent. When a single strategy for a player is a best

[2] The original "story" behind the Prisoners' Dilemma is that two prisoners are suspected in some crime. Each is held in isolation and each is visited by the prosecutor, who informs each prisoner of the following: If he *confesses* and testifies against his partner, while his partner remains *silent* then he will go free and his partner will serve 10 years in jail. If they both confess, then they will each get 6 years in jail. If he remains silent and his partner confesses, then he gets 10 years, while his partner goes free. Finally, there is enough circumstantial evidence to send each to jail for 2 years if they both remain silent. It can be shown that the outcome is both confess (see Exercise 1), even though the jointly optimal outcome occurs if they both remain silent.

[3] Another example is "rat racing" with fellow employees: You and a co-worker have equal odds of being promoted if you're both relaxed or if you're both "hyper." The problem is that if one of you is hyper while the other is relaxed, the hyper one will be certain to get the promotion. The outcome is that you're both hyper, although *jointly* you'd be better of both being relaxed. See Exercise 2.

		Ms. C	
		Northside	Southside
Mr. R	Northside	$10,000 (Ms. C) $10,000 (Mr. R)	$15,000 (Ms. C) $15,000 (Mr. R)
	Southside	$15,000 (Ms. C) $15,000 (Mr. R)	$10,000 (Ms. C) $10,000 (Mr. R)

Figure 2: Game between Mr. R and Ms. C

response to *all* the strategies of her opponents, then that strategy is called a dominant strategy; formally,

DOMINANT STRATEGY

> **Dominant strategy:** *A single strategy for a player that is a best response to* each *and* every *strategy of her opponents.*

As the name suggests, you always want to play a dominant strategy if you have one. In the Row Inc. vs. ColumnCo game, each player—as we've seen—has a dominant strategy, namely to be aggressive.

We can now give our first method of solving games for that subset of games in which all players have a dominant strategy: If all players have a dominant strategy, then the solution of the game is the outcome corresponding to each player playing her dominant strategy. This solution is known as a *solution in dominant strategies* or *an equilibrium in dominant strategies*. Observe that prisoners' dilemma games always have a solution in dominant strategies.

SOLUTION IN DOMINANT STRATEGIES

2.3 Nash Equilibrium

Although a solution in dominant strategies seems an eminently reasonable solution concept—what else would a rational player do except play the strategy that is best regardless of his opponents' choices of strategy—it has the drawback that not all games have dominant strategies for all their players. As an example of such a game, consider the game in Figure 2. Two entrepreneurs, Mr. R and Ms. C, plan to open competing firms. There are two possible places to open, Northside and Southside. Knowing the fundamental rule of competitive strategy, they know that if they open at the same location competition between them will be fiercer and, thus, profits lower. Specifically, monthly profits are $10,000 each if they open at the same location, but $15,000 each if they open at different locations. Observe that Mr. R's best response to Ms. C's strategy of Northside is to play Southside ($\$15,000 > \$10,000$); but his best response to Ms. C's strategy of Southside is to play Northside. Mr. R, therefore, has no dominant strategy since the best strategy for him to play *varies* with the strategy he anticipates Ms. C will play. Given the symmetry of the game, it's readily seen that Ms. C doesn't have a dominant strategy either.

Since we want to have a prediction of what will happen in the game depicted in Figure 2, we need some method of solving games when players don't have dominant strategies. To this end, we define the concept of Nash equilibrium:

Nash equilibrium

> **Nash Equilibrium:** *Consider a game with N players. A set of strategies, one for each player, is a* Nash equilibrium *if, for* each *player, her strategy is a best response to the* $N-1$ *other strategies in this set (i.e., to the strategies her opponents will play in this equilibrium).*

That is, the strategies chosen by the players of a game constitute a Nash equilibrium when each player's strategy is a best response to the strategies of the other players. In other words, when the strategies are mutual best responses.

For example, in the game shown in Figure 2, a Nash equilibrium is Mr. R plays Northside and Ms. C plays Southside: As we concluded above, Northside is the best response to Ms. C's strategy of Southside, and it is readily seen that Southside is Ms. C's best response to Mr. R's strategy of Northside. As a second example, the outcome in which both Row Inc. and ColumnCo play aggressive(ly) is also a Nash equilibrium: The best response to aggressive is aggressive, as we saw above. In fact, this last example is a general result:

> **Result:** *A solution in dominant strategies is also a Nash equilibrium.*

There is, however, a drawback to using the Nash-equilibrium concept to solve games: It sometimes admits more than one solution (prediction). To see this, return to the game between Mr. R and Ms. C and observe that *another* Nash equilibrium is for Mr. R to play Southside and Ms. C to play Northside.[4] [Exercise: demonstrate this.] Fortunately, in many games of interest there is a unique Nash equilibrium; that is, this solution concept yields a single prediction. Moreover, for many games in which it yields multiple solutions—known as *multiple equilibria*—there may be other factors that allow us to make more precise predictions. For instance, if Mr. R lives on the northside of town, while Ms. C lives on the southside, the prediction—equilibrium—in which Mr. R opens on the northside and Ms. C opens on the southside would likely be the better one.

multiple equilibria

Another *potential* drawback to any solution concept is that it may fail to make a prediction (as, for example, solution in dominant strategies would do for the game between Mr. R and Ms. C). Fortunately, this drawback is absent with Nash equilibrium: Every game has at least one Nash equilibrium.[5]

To review, finding a Nash equilibrium is a two-step process:

1. For each player and each set of strategies of her opponents, identify that player's best response to that set of strategies.

2. Find a set of strategies for all players that are mutual best responses.

[4]Later, when we consider *mixed* strategies, we will see that this game also has a third Nash equilibrium!

[5]To be 100% accurate, we should add "in this course." Although one can construct weird games that fail to have Nash equilibria, these games are primarily of mathematical, rather than practical, interest. For more on this, consult Roger B. Myerson: *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Mass., 1991, particularly Sections 3.12 and 3.13.

Row Player		Column Player: Left	Column Player: Right
	Up	0, $\underline{5}$	6, 4
	Middle	3, 4	$\underline{7}$, $\underline{8}$
	Down	$\underline{6}$, 2	5, $\underline{6}$

Figure 3: Another Game

To illustrate this process, consider a third game, shown in Figure 3 (since it's for illustrative purposes only, we won't bother giving it a business interpretation). There are two players, Row and Column. Column has two possible strategies: left and right. Row has three possible strategies: up, middle, and down. Let's find Row's best responses first. If Column played left, then Row's best response is down since 6 is greater than 0 (up) or 3 (middle). To help us keep track of best responses, we've underlined the 6. If Column instead played right, then Row's best response is middle since 7 is greater than 6 (up) or 5 (down). Again, we've underlined the 7 to help us keep track of the best responses. Now let's find Column's best responses. If Row played up, then Column's best response is left since 5 is greater than 4. If Row played middle, then Column's best response is right since 8 is greater than 4. Finally, if Row played down, then Column's best response is right since 6 is greater than 2. That completes step one. Turning to step two, we need to find mutual best responses. Given our convention of underlining the payoff corresponding to each best response, a set of mutual best responses will correspond to the strategies that yield a cell with *both* numbers underlined. In the game in Figure 3, the one cell with both numbers underlined corresponds to Row playing middle and Column playing right. That is, the Nash equilibrium of this game is Row plays middle and Column plays right.

Before going on, it's worth asking why we think the Nash equilibrium of a game is a good prediction of what will happen in the game. There are a number of justifications we can give for this view:

1. **A Nash equilibrium is rationalizable.** By this we mean that the players could come to play the Nash equilibrium through an introspective process. Consider, for example, the preceding game. Column might reason as follows: "No matter what Row thinks I'm going to do, it can *never* be a best response for him to play up (observe neither of his payoffs in the up row are underlined).[6] Hence, I can rule out Row playing up. But if

[6] As a technical matter, we should also consider whether up could be a best response against a *mixed* strategy for Column. Since we haven't yet introduced mixed strategies—see Section 4—we can't yet do that. If we could, however, it would turn out that up also fails to be a best response against a mixed strategy for Column (see Exercise 3).

Row would never play up—that is, I'm sure he'll play middle or down—then it's best for me to play right, since right is my best response to both middle and down.[7] Therefore, I should play right." Since Row knows that Column is rational, Row would also reason in this way if he were to put himself in Column's shoes. Hence, he would conclude that Column would play right. But if Column will play right, he wants to play middle and we've, thus, arrived at the Nash equilibrium.

2. **A Nash equilibrium is a stable solution to the game.** By this we mean that *if* the players knew they were supposed to play the Nash equilibrium, then no player would want to deviate from playing his part of the equilibrium provided he assumed that the other players would be playing their strategies for this equilibrium (their *equilibrium strategies*): By construction, if they're playing their equilibrium strategies, then his best response to these strategies is to play the strategy he's supposed to play in the equilibrium. Thus, for example, if Row and Column were told the equilibrium of the preceding game is Row plays middle and Column plays right and if they each believed the other would play his part of the equilibrium, then neither Row nor Column would wish to deviate: As shown, middle is the best response to right and right is the best response to middle. Hence, if there's any communication among the players before they play, they could agree to play a Nash equilibrium since that agreement is *self-enforcing*. Moreover, a Nash equilibrium is the *only* agreement that can be self-enforcing. SELF-ENFORCING

3. **A Nash equilibrium is often reachable by adaptive play.** By this we mean that if the players learn in a myopic fashion, they will often come to learn to play the Nash equilibrium. For instance, in the preceding game imagine that, today, the players play their best responses to the strategy their opponents' played yesterday. So, for example, if Row played up yesterday, Column would play left today. To see how this process could converge to a Nash equilibrium, suppose the players start with (up, right). The next day, under this adaptive process, they would play (middle, left). Then (down, right). Then, finally, (middle, right), the Nash equilibrium. Since, as discussed, the Nash equilibrium is stable, once they reach it, they'll never leave it; that is, once the players come to learn to play the Nash equilibrium, they will never deviate from it.

4. **What else?** Finally, even if Nash equilibrium isn't a *perfect* prediction, no one's yet devised a universally *better* prediction. Moreover, many of the predictions reached by employing Nash equilibrium have been borne out either experimentally or in actual business practice.

We have so far considered only games in which the players have quite small strategy spaces (the largest so far is Row's in the last game, which has only

[7] In the parlance of game theory, right becomes a dominant strategy in the reduced game in which we've deleted the up row.

three strategies in it). Fortunately, our analysis extends to games in which the players have larger strategy spaces. For instance, let's consider a "race" between two companies to develop a new product. The first firm on the market with the product will earn V millions of dollars from it. Assume that this new product can be patented, so that the runner-up can't produce it at all. Hence, the runner-up earns 0. Suppose, somewhat unrealistically, that each firm is limited to spending no more than \$1 million on research and development (R&D) and that a firm's probability of winning the race is

$$\Pr\{\text{win}\} = \frac{1}{2} + \frac{e_o - e_r}{2},$$

where $e_o \times \$1$ million is the firm's OWN expenditure on R&D and $e_r \times \$1$ million is its RIVAL's expenditure (by assumption, $0 \leq e_0 \leq 1$ and $0 \leq e_r \leq 1$). Each firm's *expected* profit from investing is

$$\begin{aligned} \mathbb{E}\pi &= (\Pr\{\text{win}\} V + \Pr\{\text{lose}\} 0 - e_o) \times \$1 \text{ million} \\ &= \left(\left[\frac{1}{2} + \frac{e_o - e_r}{2}\right] V - e_o\right) \times \$1 \text{ million}. \end{aligned}$$

In maximizing this expression, we're free to ignore the $\times \$1$ million part. Observe that the marginal benefit of an expenditure on R&D is $V/2$ *regardless* of your rival's expenditure. The marginal cost is 1 *regardless* of your rival's expenditure. Hence, if $V/2 > 1$—marginal benefit exceeds marginal cost—you want to invest as much as possible (*i.e.,* $e_o = 1$). If $V/2 < 1$—marginal benefit is less than marginal cost—you want to invest as little as possible (*i.e.,* $e_o = 0$). Since these conclusions hold true *regardless* of your rival's expenditure, they must constitute dominant strategies for you. We can then solve the game in dominant strategies (which is also, recall, the Nash equilibrium): If $V/2 > 1$, then the solution is both firms expend \$1 million on R&D. If $V/2 < 1$, then the solution is that neither firm expends anything on R&D. [Were $V/2 = 1$, then any pair (e_o, e_r) would be consistent with rational play.] Note, if $V/2 > 1$, then this game is very much a prisoners' dilemma: In equilibrium, each firm's expected profit is $\left(\frac{V}{2} - 1\right) \times \$1$ million; whereas, if they could somehow agree that neither would make R&D expenditures, each firm's expected profit would be $\frac{V}{2} \times \$1$ million.

3 The Bertrand Model of Competition

In this section, we introduce a simple, yet important, model of competition among two or more firms known as the *Bertrand model*. Because it will help motivate our later analyses of competitive strategy, it is worth being very explicit about the assumptions that underlie this model. BERTRAND MODEL

3.1 Assumptions Underlying the Bertrand Model

1. Firms produce a homogenous product; that is, the product made by any one firm is absolutely equivalent in the consumers' minds with that made by any other firm.

2. Customers know the price being charged by each firm.
3. Customers incur *no* costs switching from one firm to another.
4. No cost advantages: Regardless of the units produced by one firm versus another, they each have precisely the same marginal cost for the next unit produced.
5. No capacity constraints: *Each* firm is capable of handling the *entire* industry on its own.
6. No future considerations: Firms do not consider possible future interactions among themselves.

The consequence of the first assumption is that the only dimension on which firms can differentiate themselves in consumers' eyes is price. Moreover, since customers know all prices and face no switching costs, customers will buy from the firm charging the lowest price. If more than one firm is charging the lowest price then we may assume that customers divide themselves equally, on average, among those firms charging this lowest price.

In terms of notation, let $D_{\min}$ denote the *market* demand at the lowest or minimum price, $p_{\min}$, being charged by the N firms in the industry. And let the units demanded from firm n if it charges price p_n be denoted d_n. Then assumptions 1–3 imply

$$d_n = \begin{cases} 0 \text{ if } p_n > p_{\min} \\ \frac{1}{M} D_{\min} \text{ if } p_n = p_{\min} \end{cases},$$

where M is the number of firms (including n) charging the lowest price. That is, firm n sells no units if its price is not equal to the lowest price being charged; but, if it is, then it splits the market demand at that price equally with the other firms (if any) charging that price.

Assumption 4 implies a number of points. First, the only way the number of units produced can have no effect on the marginal cost is if each firm has a constant marginal cost. Moreover, this assumption requires that they each have the same marginal cost. Let c denote this constant marginal cost. Assumption 5 implies that we don't have to worry about a given firm reaching a capacity limit—each firm can service every customer who comes in its door.

Finally, assumption 6 allows us to ignore the possibility that firms need to worry about retribution from their competitors. This is either because there is no tomorrow (firms compete once) or because firms discount future profits too heavily for future retribution to affect their behavior today.

You might well ask at this point, "what industries are described by these assumptions?" The answer, admittedly, is not many—these assumptions are too extreme to be absolutely realistic. Two points, however, are worth making. First, few industries look like this because it's not in the interest of firms to be in an industry such as we've described by these six assumptions. That is—as we will see—firms want to avoid being in a Bertrand industry. Hence, we need to analyze this hypothetical industry to see what is they're trying to avoid and what strategies are available to them to avoid it. Second, a number of industries come close enough to these assumptions that the conclusions and insights we will develop will help us to understand these industries and the problems managers face in them.

3.2 Equilibrium in the Bertrand Model

We now solve for the Nash equilibrium of the Bertrand game by calculating each firm's best response to various strategies of its rivals. First, observe that a given firm, n, only cares about the minimum price being charged by its competitors, p_* (observe p_* is the minimum price charged by the $N-1$ competitors of firm n, while $p_{\min}$ is the minimum price charged by all N firms; the two quantities can, however, be the same). Let D_* denote the market demand at this price. Hence, this minimum price, p_*, is sufficient for summarizing the strategies of the competitors. Assume $M-1$ rivals are charging p_*. Firm n's profits are, then,

$$\pi_n = \begin{cases} 0 \text{ if } p_n > p_* \\ \frac{1}{M} D_* \times (p_* - c) \text{ if } p_n = p_* \\ D \times (p_n - c) \text{ if } p_n < p_* \end{cases},$$

where D is market demand at p_n, $p_n < p_*$. That is, if firm n charges more than the minimum being charged by its rivals, then it gets no customers and earns zero profit. If it matches this minimum price, then it gets one-Mth of the market demand at that price and its profit is that share of demand times the profit per unit. Finally, if it undercuts everybody, then it gets all the demand and its profit is that demand times profit per unit. We'll hypothesize three scenarios for p_*:

1. $p_* > c$ (where, recall, c is the constant marginal cost common to all firms);
2. $p_* < c$; and
3. $p_* = c$.

In the first scenario, matching p_* yields more profits than *exceeding* p_* (*i.e.*, $\frac{1}{M} D_* \times (p_* - c) > 0$). But suppose, instead, that firm n *undercut* p_* by a little bit. Then $p_n - c$ would be only slightly less than $p_* - c$ while D would be at least as great as D_* (demand curves, recall, slope down) and, most importantly, firm n would get all the demand, so its profits would be greater; that is, undercutting slightly means UNDERCUT

$$D \times (p_n - c) > \frac{1}{M} D_* \times (p_* - c).$$

For example, suppose

$$D = 100{,}000 - 25{,}000p,$$

$c = 1$, $p_* = 2$, and $M - 1 = 3$. Then matching would yield firm n a profit of

$$\frac{1}{3+1}(100{,}000 - 25{,}000 \times 2) \times (2-1) = 12{,}500.$$

Suppose, however, that firm n undercut to $p_n = 1.95$. Profit would be

$$(100,000 - 25,000 \times 1.95) \times (1.95 - 1) = 48,688;$$

almost four-times as great. Hence, in the first scenario, the best response to a $p_* > c$ is to undercut slightly (*i.e.,* charge price $p_* - \varepsilon$, where ε is some small amount).

In the second scenario, matching or undercutting means selling to a positive demand at a price less than cost. Hence, matching or undercutting means losing money. On the other hand, if firm n charges more than its competitors, it gets no customers, but it doesn't lose money either. The best response, therefore, to a $p_* < c$ is to charge any price greater than p_*.

Finally, we come to the third scenario. Since charging p_* yields zero profits—you're pricing at cost—matching or charging more than your competitors are equally good strategies; they both yield you zero profits. Were you, however, to undercut p_*, then you would get all the demand, but since you'd then be charging *less than* cost, you'd be losing money. Hence, we see that the best response to $p_* = c$ is either to match or charge a higher price.

Having derived the best responses, we can now look for a situation of mutual best responses. Observe, first, that all firms charging a price equal to marginal cost (*i.e.,* $p = c$) is a situation of mutual best responses: If all other firms are charging c, then $p_* = c$ and, as we just saw, matching is a best response in this situation. Conclusion:

> **Equilibrium of the Bertrand Model:** *A Nash equilibrium of the Bertrand model is for all firms to charge marginal cost.*

Could there be an equilibrium at another price? Well we can't have an equilibrium in which $p_{\min} > c$. To see why, observe first that at least one firm must be charging $p_{\min}$. Hence, for any *other* firm,

$$p_* = p_{\min} > c.$$

But we know that other firm's price must be at least $p_{\min} = p_*$, which means this firm is *not* playing its best response to p_*—remember the best response is to *undercut* $p_* > c$, not charge a price greater or equal to it! Since this firm (at least) is not playing a best response, it can't be a Nash equilibrium to have $p_{\min} > c$. It can also be readily shown that we can't have a Nash equilibrium in which $p_{\min} < c$; one firm, at least, can't be playing a best response since it could make a greater (non-negative) profit by raising its price above $p_{\min}$. Conclusion:

> **Equilibrium Price in the Bertrand Model:** *The market price in a Nash equilibrium of the Bertrand model is equal to marginal cost.*

In light of this last conclusion, we can conclude that

> **Equilibrium Profit in the Bertrand Model:** *All firms in a Bertrand industry must make zero economic profits.*

This is why being in a Bertrand industry is undesirable and why we will refer to being in the Bertrand model as being in the *Bertrand trap*. Bertrand trap

3.3 Two Ghosts

In the tale *A Christmas Carol* by Charles Dickens, the main character, Ebenezer Scrooge, is visited by three ghosts. One is the Ghost of Christmas Past; one is the Ghost of Christmas Present; and one is the Ghost of Christmas Future. The Ghost of Christmas Present shows Scrooge *what is*. The Ghost of Christmas Future shows him *not* what will be, but rather *what could be*. In our tale of competitive strategy, game theory gets double-billing: It is both the Ghost of Christmas Present—explaining what is and what will happen now—and the Ghost of Christmas Future—warning us of what *could* transpire at some future date if we are not careful today.

Once your firm is in the Bertrand trap, it's too late—the only equilibrium payoff possible for your firm is zero profit. Hence, one of the keys to competitive strategy is to view the Bertrand model like the Ghost of Christmas Future—a warning of what *could* transpire unless your firm unilaterally, or in tacit cooperation with its rivals, takes steps to keep your industry from being a Bertrand industry. Much of our focus in this course will be on avoiding the trap; means and methods of preventing competition being on the single dimension of price.

4 More Game Theory

Periodically in our analysis of strategic situations we will need to consider *mixed strategies*; that is, rather than play a particular strategy with certainty, as we've so far considered, players could randomize systematically over their strategies (their *pure strategies*). Systematic randomization admittedly sounds like an oxymoron: What we mean is that yes, the player randomizes over what she does, but she uses a distribution over strategies that she's chosen systematically. Such a systematic randomization is a mixed strategy. For example, in the game of Figure 2 (page 4), a mixed strategy for Ms. C might be to choose the Northside with probability $\frac{3}{4}$ and the Southside with probability $\frac{1}{4}$. At first glance, such strategies might seem weird and you might question whether they have any use. To motivate their use, consider the following model.

mixed strategies

pure strategies

4.1 The Monitoring Game

Have you ever wondered why a single security guard "watches" a whole bank of security-camera screens and why that bank might, at any moment in time, show the images from only a subset of the security cameras on the premises? Why not have enough guards to watch the image from every camera all the

		Company: Watch	Company: Other Activity
Thief	Steal	$\underline{\$0}$ $-\$10,000$	$-\$975$ $\underline{\$1000}$
	Don't Steal	$\$0$ $\underline{\$0}$	$\underline{\$25}$ $\$0$

Figure 4: Monitoring Game

time? The answer to this question can be found through an application of game theory with mixed strategies.

To simplify the problem somewhat, imagine that there is a single security camera and that an employee, who has other duties as well, can either watch the screen or do his other tasks. Let the value of his other tasks be $25 per hour. Imagine that if he is watching when a crime is committed, then he calls the police and your company loses nothing. If he's not watching when a crime is committed, then your company loses $1000 worth of property. Consider also a potential thief. If he is successful (not caught), he gets away with property worth $1000 to him. If he's caught, he must give back the property plus serve time in jail. The opportunity cost of his time in jail is $10,000. Figure 4 summarizes this game.

Observe that in this game that the best response for the thief to watch is don't steal and the best response to other activity (*i.e.,* don't watch) is steal. For the company, the best response to steal is to have the employee watch and the best response to don't steal is have him do his other activity. Note that we don't have a pair of mutual best responses (a point underscored by the fact that no cell in Figure 4 has both numbers underlined). Does this mean that this game doesn't have a Nash equilibrium? No, it means that the game doesn't have a Nash equilibrium in *pure* strategies. If we allow for mixed strategies, however, the game does have a Nash equilibrium as we will see.

Let s be the probability that the thief steals and let w be the probability that the employee watches. Then the company's expected payoff from watching is $0 and its expected payoff from the other activity is

$$-\$975s + \$25\,(1-s)\,.$$

If the company is willing to mix (*i.e.,* use a $w < 1$ but > 0), then it must like the two possible outcomes equally well: If it didn't, then why would it risk getting the one it liked less? Hence, if $0 < w < 1$, we know that the expected payoff from watching, $0, and the expected payoff from the other activity, $-\$975s + \$25\,(1-s)$, must be equal. That is,

$$\begin{aligned} -\$975s + \$25\,(1-s) &= \$0\text{; or} \\ s &= \frac{1}{40}. \end{aligned}$$

In other words, for mixing to be a best response for the company, the thief must steal with probability 1/40.

What about the thief. His expected payoff if he steals is

$$-\$10,000w + \$1000\,(1-w)\,,$$

while his expected payoff from not stealing is \$0. Again, if he is willing to mix (*i.e.,* choose $s = 1/40$), then he must like the two possible alternatives equally well. This implies

$$\begin{aligned} -\$10,000w + \$1000\,(1-w) &= \$0\text{; or} \\ w &= \frac{1}{11}. \end{aligned}$$

In other words, for mixing to be a best response for the thief, the company must watch with probability 1/11.

Upon first examination, you might not care for this equilibrium. Ninety-one percent of the time that the thief steals, he gets away with it. So why not completely deter the thief by raising the probability of monitoring to, say, 1/10? The answer is that if the thief is completely deterred from stealing why monitor at all? That is, the best response to a deterred thief is not to monitor. Of course, if you don't monitor, the thief will always steal from you and he'll always get away with it.

This same logic also explains surprise inspections. If you always inspect, then whomever you're inspecting will always be performing up to standards. But in that case, your inspections are rather wasteful. If you announce your inspections in advance, then whomever you're inspecting will have his or her act together when you arrive, but will be performing sub-par every other day. Again, your inspections are not accomplishing much. But with surprise (*i.e.,* random) inspections, whomever you're inspecting will perform up to standards most days (which is what you want); but without you having to inspect most days.

We can also illustrate this equilibrium graphically by plotting the best responses for each player on a common graph. Specifically, let S_R be the thief's best response and let W_R be the company's best response, where a best response is the probability of stealing or watching, respectively. That is, for example, the best response $S_R = 0$ corresponds to the pure strategy don't steal, $S_R = 1$ corresponds to the pure strategy steal, and $0 < S_R < 1$ corresponds to a mixed strategy. From our previous analysis we have

$$S_R = \begin{cases} 0 \text{ if } -\$10,000w + \$1000\,(1-w) < \$0 \;(i.e.,\ w > \frac{1}{11}) \\ s,\ 0 \le s \le 1, \text{ if } -\$10,000w + \$1000\,(1-w) = \$0 \;(i.e.,\ w = \frac{1}{11}) \\ 1 \text{ if } -\$10,000w + \$1000\,(1-w) > \$0 \;(i.e.,\ w < \frac{1}{11}) \end{cases}$$

and

$$W_R = \begin{cases} 0 \text{ if } -\$975s + \$25\,(1-s) > 0 \;(i.e.,\ s > \frac{1}{40}) \\ w,\ 0 \le w \le 1, \text{ if } -\$975s + \$25\,(1-s) = \$0 \;(i.e.,\ s = \frac{1}{40}) \\ 1 \text{ if } -\$975s + \$25\,(1-s) < \$0 \;(i.e.,\ s > \frac{1}{40}) \end{cases} .$$

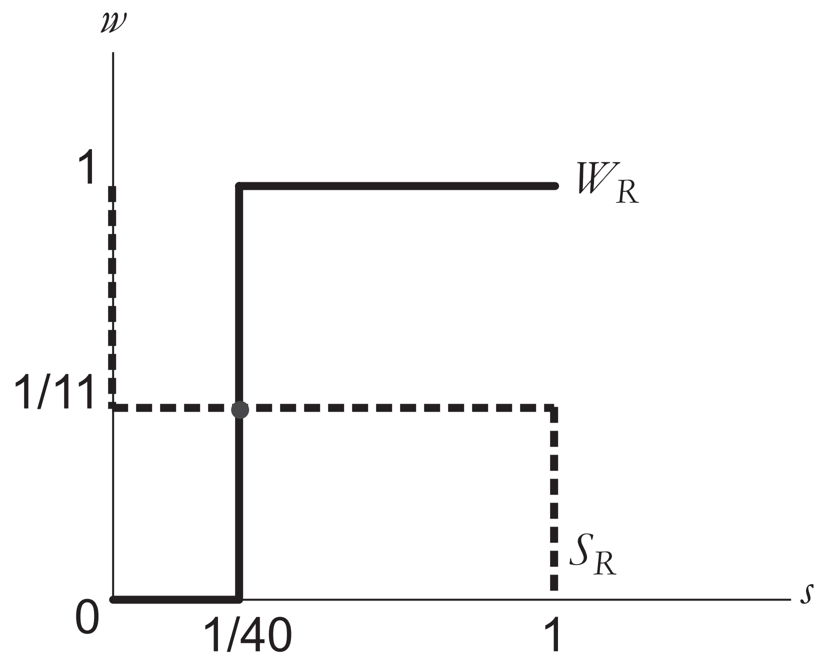

Figure 5: The company's best response, W_R, is the thick solid curve and the thief's best response, S_R, is the thick dotted curve. Observe they cross at $\left(\frac{1}{40}, \frac{1}{11}\right)$, which is the Nash equilibrium. [Note: scale has been altered for illustrative purposes.]

Figure 5 illustrates these two best responses. A Nash equilibrium occurs when we have mutual best responses; which, graphically, means where the two best responses cross. Hence the Nash equilibrium is the crossing point, which corresponds to $w = 1/11$ and $s = 1/40$.

4.2 An Entry Game

Consider the following game between two firms. Each can enter a market by investing I. If both enter the market, then they will be Bertrand competitors and, thus, will make zero profit (as we saw above). If just one enters, it has a monopoly. Assume its monopoly profit is π_M, where $\pi_M > I$. Not entering is free, but yields no profit. Figure 6 illustrates.

As can readily be seen, there are two Nash equilibria in *pure* strategies. In one, Firm 1 enters and Firm 2 stays out. In the other, these roles are reversed. [Exercise: Can you see why these are both Nash equilibria?] Yet these equilibria might strike you as odd: How is it that the firms coordinate so that only one enters? Indeed, there is a third equilibrium, in *mixed* strategies, that might

		Firm 2	
		Enter	Stay Out
Firm 1	Enter	$-\$I$ $-\$I$	$\underline{\$0}$ $\underline{\$\pi_M - \$I}$
	Stay Out	$\underline{\$\pi_M - \$I}$ $\underline{\$0}$	$\$0$ $\$0$

Figure 6: Entry Game

seem more likely.

To find this third equilibrium, let e_n be the probability that firm n decides to enter. Firm 1's expected profit from entering is

$$e_2 \times (-\$I) + (1 - e_2) \times (\$\pi_M - \$I)\,,$$

while its expected profit from staying out is \$0. Mixing—playing an e_1 between 0 and 1—is a best response only if Firm 1 likes the two outcomes equally well; that is, only if

$$\begin{aligned} e_2 \times (-\$I) + (1 - e_2) \times (\$\pi_M - \$I) &= 0\text{; or} \\ e_2 &= 1 - \frac{I}{\pi_M}. \end{aligned}$$

Reasoning in a similar fashion, it is readily shown that

$$e_1 = 1 - \frac{I}{\pi_M}.$$

Hence, the third Nash equilibrium is a mixed-strategy equilibrium in which each firm enters with probability $1 - I/\pi_M$. If we think about it, this equilibrium makes the most sense of the three. Each firm gambles on whether to enter or not. The probability it gambles on entering is smaller the greater is the cost of entering, $\$I$, *ceteris paribus*; while the probability it enters is greater the greater are the monopoly profits, $\$\pi_M$.

We can also illustrate these equilibria by plotting the best responses (see Exercise 7). Unlike Figure 5, this time the best responses cross three times—at $(0, 1)$, $\left(1 - \frac{I}{\pi_M}, 1 - \frac{I}{\pi_M}\right)$, and at $(1, 0)$—corresponding to the three Nash equilibria.

5 Summary

We introduced a set of tools, game theory, with which to more deeply understand competitive strategy. In particular, we focused attention on the Bertrand model—the *Bertrand trap*—that illustrates what we should expect to happen if firms compete myopically solely on the dimension of price.

Key Points

- Fundamental rule of competitive strategy: *Competitive strategy is like driving, not football: Head-on collisions are to be avoided.*
- Best responses and dominant strategies.
- Nash equilibrium and solution in dominant strategies.
- The Bertrand trap: *Firms in a Bertrand industry make zero economic profits.*
- Game theory as the Ghosts of Christmas Present & Future.
- Mixed strategies.

6 Exercises

1. Consider the original prisoners' dilemma (see footnote 2 on page 3). Write the game, similar in form to Figure 1, letting x years in jail have a payoff of $-x$. Let each prisoner have two strategies "confess" and remain "silent." Show that confess is a dominant strategy for each player. Is the outcome of the prisoners' dilemma a prisoners' dilemma?

2. Consider the following game. You and a co-worker are in competition for a promotion. Prior to the decision of who will be promoted, each of you have two possible strategies: You can work in a relaxed fashion ($w = 0$) or you can work in a hyper-active fashion ($w = 1$). Imagine the cost of working (measured, *e.g.,* in cost of time) is $H \times w$. Imagine that the benefit (in, *e.g.,* net present value of increased salary) is $P > H$. Assume that your probability of getting promoted is

$$\Pr\{\text{promotion}\} = \frac{1}{2} + \frac{w_y - w_r}{2},$$

where w_y is how you work and w_r is how your rival (co-worker) works. The probability that your rival is promoted is one minus the probability you're promoted.

   (a) Write this game in a form similar to Figure 1.

   (b) Show that if $P/2 < H$, then relaxed work ($w = 0$) is a dominant strategy for both workers.

   (c) Show that if $P/2 > H$, then hyper-active work ($w = 1$) is a dominant strategy for both workers.

   (d) Under what circumstances are you in a prisoners' dilemma ("rat racing") with your co-worker?

		Column Player	
		Left	Right
	Up	5 8	4 2
Row Player	Middle	4 5	3 6
	Down	1 4	8 4

Figure 7: Another Game

(e) What would the Nash equilibria be if we made the level of work, w, a continuous variable between 0 and 1? [Hint: consider the entry game on page 8.]

3. Show that "up" in the game in Figure 3 is never a best response for Row even against mixed strategies by Column. That is, show that his (expected) payoff from up must be less than his expected payoff from either middle or down for all of Column's strategies (included her mixed strategies).

4. In a game, a strategy, s, for a player is said to be *strictly dominated* if that player has another strategy, $\hat{s}$, such that this player's payoff is greater playing $\hat{s}$ than playing s regardless of his opponents' strategies. Consider the game in Figure 7. Show that "down" is a strictly dominated strategy for Row.

5. A solution concept that lies between solution in dominant strategies and Nash equilibrium is iterated deletion of strictly dominated strategies (IDSDS). The way this works is, first, each player's strictly dominated strategies are deleted. In a two-player game, this corresponds to crossing out the row or column corresponding to that strategy. What is left is considered a new game that is like the original game except it lacks the crossed out strategies. If there are any strictly dominated strategies in this game, they are crossed out. This procedure continues until no strictly dominated strategies are left. The remaining (not-crossed-out) cells are called the solution in IDSDS. If there is only one remaining cell, then IDSDS has yielded a unique solution.

   (a) For the game in 7 show that it does *not* have a solution in dominant strategies.

   (b) The game does, however, have a unique solution in IDSDS—find it.

   (c) The solution you found in the last part is also a Nash equilibrium—verify this (*i.e.,* check that it corresponds to mutual best responses).

(d) [Hard] Show that a solution in IDSDS must always be a Nash equilibrium.

6. Consider a variation of the Bertrand model in which market demand is 100 units if price does not exceed \$5 and is 0 if it does. Consider two competitors each of whom has a constant marginal cost of \$1. Each firm has a capacity of 50 units. Each firm can add 50 units of capacity by paying \$25.

   (a) Suppose both firms have left their capacity at 50 units. Show that the equilibrium price is \$5.

   (b) Suppose both firms have expanded their capacity to 100 units. Show that the equilibrium price is \$1 (hint: this is now a Bertrand trap).

   (c) Suppose one firm has a capacity of 50 and the other has a capacity of 100. Show that there can*not* be a *pure-strategy* Nash equilibrium in prices.

   (d) What does this problem suggest about limitations on capacity as a means of avoiding the Bertrand trap?

7. For the entry game on page 16, draw the best responses for the two firms and show that they intersect three times. Verify that the points at which they cross correspond to Nash equilibria of the game.