

Lecture Notes

on

Pricing

(Revised: August 2015)

These lecture notes cover a number of topics related to strategic pricing. Some of these are topics already presented in 15.013, and some are new. The objective is to provide you with a pricing “toolbox,” i.e., a set of pricing techniques, each of which might apply in some situations but not in others.

I begin with a discussion of *markup pricing*. In 15.010, you saw how the profit-maximizing price-cost margin is inversely related to the firm’s price elasticity of demand. (If you have forgotten this, go back and reread Chapter 10 of Pindyck & Rubinfeld, *Microeconomics*.) But how does one obtain an estimate of the *firm’s* price elasticity of demand? We will see how this can be done using the simple Cournot model. We will also see how — if determining output is not important — a linear approximation to a general demand curve can be used to simplify the entire pricing problem.

Next, I turn to *product line pricing*. We discussed this briefly in the context of pricing and bundling complementary products; here, we will go into it in more detail. Afterwards, I turn to pricing in markets with strong network externalities. First, we will see how the markup pricing rule can be adjusted to account for a positive or negative network externality. Second, we will examine dynamic pricing policies. In particular, how should one set price when competing with a firm in a “winner-take-all” market? If the firm currently has a small market share, should it lower its price aggressively to try to capture more market share from the leading firm, or maintain a high price and let its market share dwindle further.

I then turn to a complex but important problem: dynamic pricing with capacity constraints. This is sometimes called dynamic yield management, and it applies most directly to airlines (the airplane might have only 150 seats, which can be sold at different prices with different kinds of advanced purchase rules), cruise lines, and hotels. Finally, I will discuss some of the issues that arise with long-term contracting. What are the advantages and disadvantages – for both a seller and a buyer – of transacting under a long-term contract versus selling and buying on the spot market?

1 Markup Pricing

In 15.010, you saw that profit maximization implies that marginal revenue should equal marginal cost, which in turn implies:

$$\frac{P - MC}{P} = -\frac{1}{E_d} \quad (1)$$

Here, E_d , is the *firm's* price elasticity of demand. Note that this equation can be rewritten as:

$$P = \frac{MC}{1 + (1/E_d)} \quad (2)$$

If the firm is a monopolist, then the relevant elasticity is the *market* elasticity of demand, which I will denote by E_D . Obtaining an estimate of this market elasticity of demand may or may not be difficult, but for the moment, let's assume that we have such an estimate. Given this estimate of E_D , how can we obtain an estimate of the *firm's* price elasticity of demand, E_d ?

1.1 Cournot Competition

If we assume that the firms in the market compete à la Cournot (which may or may not be a reasonable assumption), obtaining the elasticity of demand for an individual firm is fairly straightforward. We will do this first for the simple case of equal-sized firms, and then for the more common case of firms with different sizes.

Equal-Sized Firms. Suppose that there are n equal-sized firms in the market, and that they all have the same marginal cost, c . In the Cournot model, each firm chooses its profit-maximizing output taking the outputs of its competitors as fixed.

For Firm i , profit is given by:

$$\Pi_i = [P(Q) - c]Q_i \quad (3)$$

where Q is the total output of the industry. We want to maximize this profit with respect to Q_i , treating the Q_j 's for the other firms as fixed:

$$\frac{\partial \Pi_i}{\partial Q_i} = P(Q) - c + Q_i \frac{dP}{dQ} \frac{\partial Q}{\partial Q_i} = 0 \quad (4)$$

Note that $dQ/dQ_i = 1$, and $dQ_j/dQ_i = 0$. Thus, the reaction curve for Firm i is:

$$P(Q) + Q_i \frac{dP}{dQ} = c \quad (5)$$

Because there are n equal-sized firms, we know that $Q_i = Q/n$. Thus eqn. (5) becomes:

$$P(Q) + \frac{Q}{n} \frac{dP}{dQ} = c \quad (6)$$

If we divide both sides of this equation by P and rearrange the terms, we get:

$$\frac{P - c}{P} = -\frac{1}{n} \frac{Q}{P} \frac{dP}{dQ} = -\frac{1}{nE_D} \quad (7)$$

Compare this result to eqn. (1). Note that:

$$E_d = nE_D \quad (8)$$

Thus, with n equal sized firms, going from the market elasticity of demand to the firm's elasticity of demand is quite easy: just multiply the market elasticity by n .

Firms of Different Sizes. There are very few markets in which all of the firms are the same size. Suppose that instead we have m firms of *unequal size*. As long as the assumption of Cournot competition holds, obtaining the elasticity of demand for each firm is still straightforward. I will not go through the details (you can try to derive this as an exercise), but the procedure is as follows:

First, calculate the Herfindahl-Hirschman index (HHI) for the industry:

$$\text{HHI} = \sum_{i=1}^m S_i^2, \quad (9)$$

where S_i is the market share of Firm i .

Next, find n^* , the equivalent number of equal-sized firms that yields the same value of the HHI calculated above. In other words, find n^* such that:

$$\sum_{i=1}^{n^*} (1/n^*)^2 = n^*(1/n^*)^2 = 1/n^* = \text{HHI} \quad (10)$$

Thus, n^* is simply the reciprocal of the HHI, i.e., $n^* = 1/\text{HHI}$. Now, just use this n^* to obtain E_d :

$$E_d = n^* E_D \quad (11)$$

Suppose, for example, that there are four firms with market shares of 40%, 30%, 15%, and 15% respectively. In that case, the HHI is given by: $\text{HHI} = (.4)^2 + (.3)^2 + (.15)^2 + (.15)^2 = .295$. Thus, $n^* = 1/.295 = 3.39$, and $E_d = 3.39E_D$.

Applications. Earlier in the course, we applied these results to the market for beer. You might recall that for mass market beers (Budweiser, Miller, etc.), the market elasticity of demand has been shown from statistical studies to be about -0.8. Recall that competition in attribute space is local, i.e., a particular brand only competes with its nearby neighbors in the attribute space. For mass market beer, this local competition implies a value of n (or n^*) of about four or five. Thus, for a typical mass market brand of beer, the elasticity of demand is approximately $E_d = (5)(-.8) = -4.0$. Using eqn. (2), this means that the profit maximizing price is roughly:

$$P = \frac{\text{MC}}{1 - (1/4)} = 1.33\text{MC}$$

(Remember that this is a wholesale price.) We can see from this that profit-maximizing price-cost markups for mass market beer are likely to be fairly small. This means that aggressive price competition can be extremely damaging, which, as we discussed, is one reason that firms prefer to compete aggressively using advertising rather than price.

We know that in many industries firms use price as a strategic variable, rather than quantity. Recall that the use of price rather than quantity as a strategic variable intensifies

competition and reduces profits. (If you don't remember this, go back and review the Lecture Notes on Game Theory.) This means that the effective elasticity of demand for a particular brand is likely to be larger in magnitude than the Cournot model would predict. But how much larger? That depends on the cross-price elasticities among the different brands.

If the cross-price elasticities are large (i.e., the brands are fairly close substitutes), the Cournot analysis won't work: E_d is likely to be much larger than nE_D . If, on the other hand, substitutability among brands is limited, our Cournot result will likely still hold.

In what kinds of markets is our oligopoly markup rule likely to work, and in what kinds of markets will it not work? It will work in markets where competition is fairly stable, and involves choices of output and/or capacity. One example that we have studied in depth is beer, where the major firms try to avoid price competition, and focus instead on advertising. It will also work in markets for automobiles (where firms choose *capacity*, which makes the outcome similar to Cournot). Likewise, it is likely to work in markets for mineral resources (copper, aluminum) and some service industries (such as shipping and insurance).

The oligopoly markup rule is unlikely to work, however, in markets where competition and pricing are dynamic in nature, and where strategic multi-period gaming is important. We discuss this below.

1.2 Strategic Considerations

Remember that the Cournot model is essentially static; it assumes that each firm takes the output of its competitors as fixed. As explained above, this assumption is quite reasonable for the large number of markets where competition is stable and dynamic price wars are generally avoided. As we have seen throughout this course, this kind of stable pricing could arise if the same firms having been competing for a long time (and there is little prospect of entry by new firms), production cost and market demand conditions are relatively stable, and the firms face a repeated Prisoners' Dilemma with no end point (and thus little or no likelihood of "unravelling"). But the Cournot assumption will not be reasonable for markets where pricing and output decisions are dominated by dynamic gaming considerations.

Airlines. One example is airlines, where very low short-run marginal costs result in intense price competition, particularly for certain fare categories. As you have seen, airline pricing is also complicated by the fact that prices are linked to yield management (i.e., the allocation, which changes from day to day, of the number of seats for each fare category). What can an airline do to avoid the intense price competition that is so prevalent in the industry? As we have discussed, one strategy is seek and maintain “monopoly routes,” where it is the only airline to offer non-stop service.

Retail Store Pricing. Retail store pricing is in many ways similar to the Strategic Oligopoly Game that you play every week — it is repeated Prisoners’ Dilemma with a fixed end point. You know that in this kind of repeated game, “unravelling” is likely to occur. That is exactly what happens to retail store pricing each year during November and December. The last play of the retail store repeated game occurs on December 24. Price competition becomes intense well before that date as stores face the unhappy prospect of being stuck with large amounts of unsold inventories that they will have to dump, often at below cost.

Over the years, managers of large retail chains and department stores have become increasingly sophisticated in their understanding of game theory. What has this done for them? The result is that they are more aware of the unravelling problem, and they are aware that their competitors are also aware of the problem. Thus they anticipate that the unravelling will start earlier. Not wanting to be the “sucker” who gets undercut and winds up with lots of unsold inventories, each store tries to get a jump on the unravelling. And as you would expect, the unravelling starts earlier. Ten years ago the unravelling generally began on or just after Thanksgiving, five years ago in mid-November, and in more recent years unravelling has started by the beginning of November.

If you ran a retail store chain, how should you go about setting prices? First of all, you are stuck in a bad place and deserve a lot of sympathy! Beyond that, the best you can do try to predict when and how the unravelling is likely to occur, make sure you are not on vacation or asleep when it starts, and take it into account when you determine the quantities that you will purchase at wholesale.

Commercial Aircraft. Finally, one more more example of an industry in which pricing

is dominated by gaming and bargaining is commercial aircraft. As we have discussed at length in this course, the durable goods monopoly problem, combined with the ability of airlines to play Boeing and Airbus off against each other, can drive prices down toward marginal cost.

2 The Linear Demand Approximation

A nice thing about industries such as beer, soft drinks, breakfast cereals, airlines, automobiles, etc., is that they have been around for a long time, and a good deal of data are available so that we can estimate market demand curves. Often, however, firms must set the prices of new products for which there is little or no history of demand or consumer response to price changes. You saw an example of this early in this course when we studied the emerging market for Internet music downloads and the pricing problem facing Apple. How might managers think about pricing in such situations?

Suppose that you were trying to set the price of a prescription drug such as Prilosec (an antiulcer medication). You would know that Prilosec competes with other proton-pump inhibitor antiulcer drugs such as Prevacid and Nexium, as well as the earlier generation H_2 -antagonist antiulcer drugs such as Zantac, Pepcid, and Axid. You would also know that marginal production cost is quite low. However, you might find it difficult to estimate the elasticity of demand for the drug. Rather than try to estimate the elasticity of demand directly, an alternative approach is to approximate the demand curve as linear. In particular, suppose that the demand curve can be approximated by the following equation:

$$P = P_{max} - bQ \tag{12}$$

where P_{max} is the *maximum* price you think consumers would pay for this drug. This linear approximation is illustrated in Figure 1, which shows the actual demand curve for a drug like Prilosec, the corresponding marginal revenue curve, and the linear approximation to the demand curve.

You can check that equating marginal revenue with marginal cost gives the following




Figure 1: Linear Approximation of Demand Curve

profit-maximizing quantity and profit-maximizing price:

$$Q^* = (P_{max} - c)/2b \quad (13)$$

$$P^* = (P_{max} + c)/2 \quad (14)$$

where c is marginal cost. From the demand curve, note that the elasticity of demand is:

$$E_d = -\frac{1}{b} \frac{P}{Q} = -\frac{P_{max} + c}{P_{max} - c} \quad (15)$$




Figure 2: Linear Demand Curves

Observe that P^* does not depend on b . (Q^* does, but our concern right now is with pricing.) This is a property of linear demand curves, as shown in Figure 2. What is nice about this linear curve is that we can measure c , so to get the optimal price, P^* , all we need is an estimate of the maximum price, P_{max} . Estimating P_{max} is often much easier to do than measuring the elasticity of demand.

This is a remarkably simple rule for pricing: given the maximum price P_{max} and marginal cost c , just set a price $P^* = (P_{max} + c)/2$. But how well can the firm expect to do if it uses this pricing rule? Suppose that if the firm knew its true demand curve, it would set a price P^{**} and earn a profit Π^{**} . Using the price P^* it earns a profit Π^* , which is presumably lower than Π^{**} . But how much lower?




Figure 3: Randomly Generated Demand Curve

A recent paper addressed this question.¹ The authors showed that if the “true” demand curve is one of the widely used demand functions, e.g., quadratic, log-log, or semi-log, the ratio Π^{**}/Π^* will be quite close to 1. They then generated random demand curves, an example of which is shown in Figure 3. For each random demand curve, they calculated the price the firm would charge if it knew the curve as well as the resulting profit Π^{**} , and also the profit that would result if the firm used the simple linear pricing rule and set a price $P^* = (P_{max} + c)/2$. They then calculated the ratio Π^{**}/Π^* . They repeated this exercise 100,000 times, which resulted in 100,000 values of the ratio Π^{**}/Π^* . Figure 4 shows a histogram of the profit ratio, first for $c = 0$ and then for $c = .5P_{max}$.

As you can see from the histograms, for the vast majority of cases, $\Pi^{**}/\Pi^* < 1.2$, i.e., Π^* is within 20% of Π^{**} . Thus this pricing rule can be expected to perform extremely well, most of the time.

¹See M. Cohen, G. Perakis, and R. Pindyck, “Pricing with Limited Knowledge of Demand,” Sloan School Working Paper, July 2015.




Figure 4: Histograms of Π^{**}/Π^* for $c = 0$ and $c = 0.5P_{max}$

Now, let's use some examples to see how this pricing rule can be applied in practice. We'll begin with our example of Prilosec. Marginal cost (including packaging and distribution) is around \$0.30 to \$0.40 per daily dose, so set $c = \$0.40$. As a proton pump inhibitor, Prilosec is much more effective than the earlier H_2 -antagonist drugs such as Zantac, etc. The prices of Zantac, etc. were about \$3.00 (per daily dose). We might expect P_{max} to be double or triple this number.

If $P_{max} = \$9.00$, $P^* = \$9.40/2 = \4.70 , which implies that the elasticity of demand is $E_d = -(9.40)/(8.60) = -1.09$. On the other hand, if $P_{max} = \$6.00$, this implies that $P^* = \$3.20$, so that the elasticity of demand is -1.14 . Taking an average of these two demand elasticities suggests that $E_d = -1.12$ is a good estimate. This, in turn, implies that the optimal price should be roughly $P^* = 0.40/(1 - 1/1.12) = \3.73 . This is indeed close to the actual price that was charged for Prilosec.²

To take one more example, consider Gillette's pricing of the Mach 3 razor. Gillette has a large market share, but still faces some competition from Schick and others. What might be

²Prilosec was introduced by Astra-Merck in 1995. The pricing of Prilosec is discussed in Example 10.1 of Pindyck and Rubinfeld, *Microeconomics*.

a reasonable estimate of the maximum price that people would pay for this razor? Given the prices of other razors (including Gillette's Sensor), a good estimate would probably be in the range of \$8 to \$12. Given that the marginal cost of production and distribution is \$4, this implies a profit-maximizing price in the range of \$6.00 to \$8.00. When Gillette introduced the Mach 3 in 1998, they priced it to sell at retail for \$6.49 to \$6.99.

3 Product Line Pricing

When a firm produces two or more products that have *interrelated demands* (i.e., they are substitutes or complements), to maximize profits the firm must price the products *jointly*, not individually. The reason is that this enables the firm to internalize the effects of sales of each product on the sales of the other products.

You saw this earlier when we studied the bundling of complementary products. (Recall that complements are goods that tend to be used together, e.g., a computer operating system and applications software that runs on that operating system.) You saw that when two products are complements, a firm that produces *both* of them (e.g., the operating system and the applications software) should set lower prices than would two independent firms, each of which sells just one of the products. We saw that, in some cases, the firm would want to sell one of the products at a price that was below its marginal cost, or even a price that was zero or negative. (If all of this seems a bit hazy, go back and review the Lecture Notes on Bundling and Brand Proliferation, Section 1.)

Now, let's consider a case in which a firm sells two products that are *substitutes*. For example: (i) a supermarket sells gourmet foods and staples, (ii) Embraer sells the 80-seat Model 175 and the 100-seat Model 190 passenger jets. For simplicity, we will assume that the demand curves are linear:

$$Q_1 = a_0 - P_1 + .5P_2 \tag{16}$$

$$Q_2 = b_0 + .5P_1 - P_2 \tag{17}$$

We will assume that the marginal cost for both products is the same and equal to the constant c .

3.1 Pricing the Products Individually

Suppose the firm sets prices *individually*, i.e., suppose it sets P_1 to maximize Π_1 , and sets P_2 to maximize Π_2 . The resulting prices are found as follows.

$$\Pi_1 = (P_1 - c)Q_1 \quad (18)$$

set $\partial\Pi_1/\partial P_1 = 0$ and get:

$$P_1 = (a_0 + c)/2 + \frac{1}{4}P_2 \quad (19)$$

Likewise, setting $\partial\Pi_2/\partial P_2 = 0$, we get

$$P_2 = (b_0 + c)/2 + \frac{1}{4}P_1 \quad (20)$$

Combining these equations for P_1 and P_2 , we find that the optimal prices are:

$$P_1 = (8a_0 + 2b_0 + 10c)/15 \quad (21)$$

$$P_2 = (8b_0 + 2a_0 + 10c)/15 \quad (22)$$

These prices are the same as those that would result if the two products were produced by separate firms, each taking its competitor's price into account, yielding a Nash equilibrium in prices. The reason is that the single firm producing both products is treating each product individually and ignoring the fact that it can internalize the demand spillover from one product to another.

Suppose that $a_0 = 100$, $b_0 = 50$, $c = 10$. Then $P_1 = \$66.70$, $P_2 = \$46.70$, $Q_1 = 56.60$, $Q_2 = 36.6$, and $\Pi_T = \$4,552.40$.

3.2 Pricing the Products Jointly

Now suppose that instead, the firm sets the prices of the two products *jointly*, i.e., it sets P_1 and P_2 to maximize total profit, Π_T :

$$\Pi_T = (P_1 - c)Q_1 + (P_2 - c)Q_2 \quad (23)$$

To maximize this, set $\partial\Pi_T/\partial P_1 = 0$ to get:

$$P_1 = (a_0 + .5c)/2 + \frac{1}{2}P_2 \quad (24)$$

Likewise, set $\partial\Pi_T/\partial P_2 = 0$ and get:

$$P_2 = (b_0 + .5c)/2 + \frac{1}{2}P_1 \quad (25)$$

Combining these two equations, we find that the optimal prices are:

$$P_1 = (2a_0 + b_0 + 1.5c)/3 \quad (26)$$

$$P_2 = (2b_0 + a_0 + 1.5c)/3 \quad (27)$$

Using the same parameter values, i.e., $a_0 = 100$, $b_0 = 50$, and $c = 10$, we have $P_1 = \$88.30$, $P_2 = \$71.70$, $Q_1 = 47.3$, $Q_2 = 22.5$, and $\Pi_T = \$5,088.80$. Note that now the prices are higher, and total profit Π_T is also higher. Why? The reason is that the higher price for each product stimulates the sale of the other product, and the firm is taking this into account.

3.3 More General Demand Curves

Obtaining the profit-maximizing prices is straightforward when the demand curves are linear, as in our example above, but how can those prices be found for more general demand curves? Consider the general demand curves $Q_1 = Q_1(P_1, P_2)$ and $Q_2 = Q_2(P_1, P_2)$. Once again, write down the expression for total profit, and then maximize this with respect to P_1 and P_2 by setting the derivatives with respect to each price equal to 0. You can check that doing that yields the following two equations:

$$P_1 + (P_1 - c)E_{11} + (P_2 - c)(Q_2/Q_1)E_{21} = 0 \quad (28)$$

$$P_2 + (P_2 - c)E_{22} + (P_1 - c)(Q_1/Q_2)E_{12} = 0 \quad (29)$$

where E_{11} and E_{22} are own-price elasticities (< 0), and E_{21} and E_{12} are cross-price elasticities (> 0 if the products are substitutes, < 0 if the products are complements). In principle, these two equations can be solved (probably numerically) for the two prices P_1 and P_2 .

We can get some insight into how the prices should be set without actually solving these two equations numerically. Consider Product 1. Rearranging the first equation yields:

$$P_1 = \frac{c}{1 + 1/E_{11}} - \frac{(P_2 - c)(Q_2/Q_1)E_{21}/E_{11}}{1 + 1/E_{11}} \quad (30)$$

The first term on the right-hand side of this equation is the standard markup. The second term is an adjustment for the impact of P_1 on the sales of Product 2. We can determine this adjustment using estimates of our expected profit margin for Product 2 (as a substitute for $P_2 - c$), the relative sales volumes (Q_2/Q_1), and the “relative elasticity” E_{21}/E_{11} . Note that if $E_{21} > 0$, i.e., the products are substitutes, the adjustment *raises* P_1 , but if $E_{21} < 0$, it *lowers* P_1 .

4 Pricing with Network Externalities

When the demand for a product is subject to network externalities, pricing becomes somewhat more complicated. First, in a static context, the presence of a network externality affects the elasticity of demand for the product. As you have seen before, a positive network externality makes demand more elastic, and a negative network externality makes it less elastic. At issue is how one can take this into account when setting price.

In a dynamic context, a strong positive network externality can complicate price setting. Suppose you face a single competitor in what is likely to be a “winner-take-all” market. If you and your competitor are just starting out, should you set your price very low in order to become the “winner,” even if this means losing a great deal of money in the short run. And if you do, how would you expect your competitor to respond? Alternatively, suppose that you have been in the market for some time, but have only a 30% market share, compared to your competitor’s 70% share. Should you sharply undercut your competitor in an attempt to increase your market share (and eventually overtake your competitor), or should you simply maintain a high price, knowing that your market share will eventually dwindle toward zero? We will examine both the static and dynamic aspects of pricing. We will begin by considering a simple adjustment to the markup pricing rule that can account for the presence of a network externality, at least in static terms.

4.1 Static Pricing

Recall how the presence of a network externality affects the elasticity of demand of a product. Figures 5 and 6 show how positive and negative network externalities, respectively, affect demand. (These figures came from Section 4.5 of Pindyck and Rubinfeld, *Microeconomics*.) Observe how a positive network externality makes demand more elastic than it would be otherwise. Looking at Figure 3, suppose we started at a price of \$30, so that 40,000 units per month were being sold. If we lower the price to \$20, what happens? If there were no network externality, we would move down the demand curve labeled D_{40} , and the quantity demanded would increase to 48,000 units. (In the figure, this is called a “pure price effect.”) However, the increase in sales from 40,000 to 48,000 units creates more value for those people purchasing, or thinking of purchasing the good. This increases sales further, perhaps to 55,000 or 60,000 units. But this increases the value of the good still more, increasing demand further. As shown in Figure 5, the new equilibrium is reached at a level of sales of 80,000 units. The increase from 48,000 to 80,000 units is the network effect. As you can see in Figure 6, just the opposite happens when the network effect is negative.

The issue now is *how much* larger or smaller does the elasticity of demand become as a result of the network externality, and how should the price be adjusted accordingly? The simplest way to answer these questions is to work with a demand curve that is either linear or isoelastic. I will consider an isoelastic demand curve; you can do the analysis yourself for a linear demand curve.

Suppose, then, that the demand curve is given by:

$$\log Q = a - b \log P + \beta \log q \quad (31)$$

where q = “perceived demand” (or sales), and Q is actual demand. If the network effect is positive, $\beta > 0$; if the network effect is negative, $\beta < 0$. In equilibrium, $q = Q$, and therefore:

$$\log Q = \frac{a}{1 - \beta} - \frac{b}{1 - \beta} \log P \quad (32)$$

(Note that we must have $\beta < 1$ for this isoelastic demand curve to make any sense.) The elasticity of demand is the coefficient multiplying the log of price. Thus, the elasticity of




Figure 5: Positive Network Externality

demand is $E_d = -b/(1-\beta)$. In the absence of a network externality, we would have $E_d = -b$. So, all we have to do is “mark up” E_d using an estimate of the elasticity of Q with respect to q .

Let’s consider a couple of examples. The first example is the Apple iPod. Without any network externalities, we might expect the price elasticity of demand for this product to be around -1.5. However, there is a positive network externality, and β is probably around 0.5. This would imply a doubling of the price elasticity of demand.

As a second example, consider the Rolex watch. Suppose there are five competing brands. Without any network externality, the elasticity of *market* demand is about -2 , so $b = -2$. However, exclusiveness is important for luxury watches, i.e., there is a strong negative




Figure 6: Negative Network Externality

network externality. Perhaps, $\beta = -10$ (so that a 1% increase in q results in a 10% decrease in Q). Then, the elasticity of market demand is $E_D = -2/11$, or about -0.2 . However, $n = 5$, so the brand elasticity is $E_d = (5)(-0.2) = -1$. Thus, P should be much, much higher than marginal cost, which, in fact, it is.

4.2 Dynamic Pricing

Dynamic pricing applies to situations in which firms have competing “networks,” and there is a process of saturation. Each firm wants its consumers to use its “network,” rather than the “network” of its competitor. Because of the positive network externality, the greater the

share of consumers that a firm has, the easier it is to attract additional consumers. Each firm can affect the rate at which it attracts consumers by altering its price and/or its level of advertising. However, as the market saturates, a firm with a small market share will find it very difficult to attract additional consumers, even if it lowers its price substantially.

A much-cited example of this kind of competition was the battle over VHS versus Beta in videocassette recorders, in which the VHS “network” eventually won. Other examples include Microsoft Word versus Word Perfect in word processing software, Microsoft’s Windows versus IBM’s OS/2 in desktop operating systems, the Apple iPod (which stores music in a proprietary format) versus competing MP3 music players, and PDAs using the Palm operating system versus those using Windows CE.

Suppose that you are competing against another firm, and both of you are starting with small but equal numbers of consumers using your “networks.” Should you price aggressively in order to “win” the market, or should you generally match your competitor’s price and advertising levels in the hope of eventually sharing the market? We examined this problem qualitatively in our discussion of markets with network externalities. (If you do not recall that discussion, reread Section 2 of the Lecture Notes on Network Externalities.) As we saw, this is a complicated problem, because you and your competitor are locked in a dynamic game. In general, there is no analytical solution to the problem.

Exercise 9 provides an opportunity to study this problem quantitatively. The exercise portrays competition between DOS-based PCs (DOS) and Macintosh systems (MAC), during 1984-1992. The growth of each network follows a diffusion equation subject to market saturation, and those equations are presented and explained in the exercise. You have an opportunity to examine alternative pricing rules for DOS by using the model to perform simulations. You can also use the model as a “flight simulator,” in which you and a classmate repeatedly set DOS and MAC prices, each of you attempting to maximize your cumulative profits. In class, we will discuss some of the basic lessons that came out of that exercise.

5 Yield Management

Yield management refers to the strategy used in a number of industries to manage the allocation of a firm's capacity to different price categories over time in order to maximize profits. (It is also referred to as *revenue management*.) This form of dynamic pricing is complicated because the prices for several categories of capacity must be determined, the total capacity must be divided up among these categories, and the prices and category allocations are continually modified and updated over time.

Yield management is applicable under the following conditions:

- The firm is selling a fixed stock of *perishable* capacity.
- Customers book capacity prior to use.
- The firm controls a set of price categories (“buckets”), each with a specific price.
- The firm can change the availability of various price categories over time.

The most obvious example of yield management is airline pricing. Consider a flight by United Airlines that is scheduled to go from Boston to Chicago six months from now. The plane has 150 seats (in coach). Customers vary in their willingness to pay, and in their ability to book in advance. Furthermore, United is likely to face competition on this route from one or more other airlines (e.g., American).

Suppose we can identify three categories of customers that differ in their willingness to pay: business travelers, high-demand leisure travelers, and low-demand leisure travelers (who will only book a flight if they can get a rock-bottom price). Fortunately for the airline, these groups can be sorted based on their ability to book in advance: business travelers are least able to book in advance, high-demand leisure travelers are somewhat willing to do so, and low-demand leisure travelers are very willing if they can get a low fare in return.

The airline will therefore create three *fare categories*, which are sometimes referred to as “fare buckets.” Now the airline has to decide *what prices to set for each of the three fare categories, and how many seats to allocate to each category*. Furthermore, the airline will

want to modify these prices and seat allocations as the departure date gets closer, as it updates its demand estimates for this particular flight based on how many tickets in each category have been sold.

If the seat allocations were fixed, i.e., if the airline had no control over how many seats went to each category, this would be a problem of pure price discrimination. What makes this difficult is allocating seats and deciding how to change those allocations over time.

Let's see how yield management might be done in practice. Once again, suppose we have three segments, which I will call (1) business (2) middle, and (3) leisure. We can safely assume that segment (1) competes with segment (2) but not with segment (3). Likewise for segment (3). Segment (2) competes with both segment (1) and segment (3). We thus have three substitute products, and we can write the demand curves for the products as follows:

$$Q_1 = a_0 - a_1P_1 + a_2P_2$$

$$Q_2 = b_0 + b_1P_1 - b_2P_2 + b_3P_3$$

$$Q_3 = c_0 + c_2P_2 - c_3P_3$$

As long as the plane is not full, marginal cost is zero for all three segments. This looks like a problem in product-line pricing, except that we face the additional constraint that every passenger must have a seat:

$$Q_1 + Q_2 + Q_3 \leq 150$$

The flight is 6 months away, but some passengers will book early, so we need to set the fares and allocate seats. We will do this in stages. First, we will use historical data (the last three years) to estimate the own-price and cross-price demand elasticities for the three segments. (For United's Boston-Chicago flights, there is a great deal of data available, so these elasticities can be estimated reasonably well. For other routes, there may not be as much data, so we will have to look at comparable routes, and our estimates will be less precise.) Given these initial estimates:

- Calculate the prices P_1 , P_2 , and P_3 that maximize total revenue (= profit). At this point we have imposed no seat restrictions. Now check:

- Given these prices, do the implied quantities add up to more than 150? If yes, raise all of the prices until the total quantity demanded is just 150. We start with with these prices, and no seat restrictions, 6 months before the flight.
- Two or three months before the flight, check on how many seats have been sold in each fare category. Are seats selling faster than expected?
 - If yes, maybe “close” bucket #3, or raise the fare on that bucket.
 - If no, don’t impose restrictions.
- One month before the flight, check once again. Are seats selling faster than expected? If yes, close bucket #3, i.e., do not allocate any more seats to that category. Also, consider raising the fare on bucket #2.
- Now, repeat this process *every day*, right up to the departure of the flight. If, as the departure date nears, we have already closed bucket #3 but seats in buckets #1 and #2 are selling faster than expected, we might close bucket #2, so that a sufficient number of seats will be available for business travelers willing to pay the highest fare. But if seats are selling more slowly than expected, we might re-open bucket #3, and lower the fares on both that bucket and on bucket #2.
- If seats are selling much more slowly than usual, *or competitors have lowered fares*, reduce fares on some or all buckets, perhaps by introducing “special offers.”

Does this sound complicated? Well, it is. And keep in mind that for most airlines, there are likely to be 15 to 20 fare buckets in coach alone! The major airlines have spent a great deal of money over the past 20 years to develop and refine methods of estimating demands for different segments, for continually updating those estimates, and for setting fares and allocating seats accordingly. The actual algorithms are extremely sophisticated and complex, and represent an important asset for an airline.

If you have somehow concluded that yield management is only for airlines, you are mistaken. Dynamic pricing of this kind is now used in a wide range of industries:

- **Cruise Lines:** A typical cruise line will start selling cabins six months before the departure date of the cruise, with about five or six “fare buckets.” (In the case of a cruise, each fare bucket applies to a different size cabin, and thus capacity allocations are fixed.) As with an airline, the cruise company will monitor sales and change prices as the departure date nears.
- **Sporting events and concerts:** A baseball stadium or concert hall has a fixed number of seats. Usually the high-price “buckets” are for better seats, but the stadium or concert hall can decide in advance how to divide the available seats up into different categories. Thus capacity allocations are somewhat flexible initially, but cannot be easily changed as the date of the ball game or performance approaches.
- **Hotels:** Each hotel has a fixed number of rooms available each night. As with an airplane flight, customers differ in willingness to pay and ability to book in advance. Thus yield management for a hotel looks very similar to yield management for an airline. The hotel will allocate lower-price buckets for corporate discounts, AAA discounts, and various kinds of “specials.” If hotel is filling up quickly, it will close some of the low-priced buckets.
- **Broadway shows:** A new discovery for the theater industry! Broadway theaters are now using sophisticated yield management systems to price their seats, and change those prices as the performance date approaches.
- **Casino Hotels:** Harrah’s has developed very sophisticated dynamic pricing algorithms for hotel rooms in its Las Vegas casinos. The firm price discriminates based on public and private customer data. (Gary Loveman, CEO of Harrah’s, has a PhD in Economics from MIT.)
- **Retail chains.** This is new, but now consulting firms are helping retail chains use dynamic pricing during the November and December holiday sales period. As of December 26, whatever hasn’t sold “perishes,” in that it will be marked down drastically.

Do you still think that yield management is only about airlines?

6 Long-Term Contracting

Firms must often decide whether to transact on the spot market, or instead buy and sell via a long-term contract. The owner of a coal-fired electric plant, for example, could purchase coal month-to-month on the spot market, or could sign a contract for the delivery of coal over the next ten years. Most electricity producers would prefer to buy coal on a long-term contract basis, on the grounds that this allows them to reduce risk related to fluctuations in the price of coal. (As you learned in your finance courses, however, investors can just as easily diversify away that risk.)

Long-term contracting is often done as a means of reducing risk, even though the same risk reduction can often be achieved by taking positions in the futures market, or by diversification on the part of investors. Our concern will be with a different motivation for long-term contracting: the ability to get a strategic advantage, either as a buyer or a seller.


Consider the case of commercial aircraft. As we have seen, selling aircraft via long-term contracts enables Boeing or Airbus to plan and smooth production over time. But we have seen that long-term contracting also creates a strategic advantage for *buyers*, and lets them gain monopsony power. The reason is that an airline (even when in bankruptcy) negotiating the purchase of twenty or thirty airplanes over the next ten years can play the two manufacturers off against each other. Both Boeing and Airbus have much to lose if the competitor wins the sale. That, combined with the durable goods monopoly problem, gives the airline considerable power in the negotiations.

Long-term contracting can also provide a strategic advantage to sellers. This is particularly true of natural resources that are crucial to the buyer and that may be in limited supply in the future.

An example of this is the market for uranium, or more specifically, uranium oxide (U_3O_8), commonly referred to as “yellowcake,” because of its color and consistency. Companies in the United States, Canada, Australia, and South Africa are the main sellers of yellowcake, and the main buyers are the owners of nuclear power plants (in the U.S. and around the world). Naturally occurring uranium contains less than 1% of the fissile isotope U_{235} ; the remainder

On the Road to Uranium Enrichment

Iran resumed its conversion of yellowcake uranium to uranium tetrafluoride yesterday.



The New York Times

Figure 7: Uranium Fuel Cycle

is the non-fissile heavier isotope, U_{238} . Power plants and other buyers contract to have their yellowcake processed and enriched to yield uranium metal that is about 50% U_{235} for use in fuel rods. (Figure 7 shows the uranium fuel cycle. Enrichment is done by converting U_3O_8 into uranium fluoride, which is then spun in high-speed centrifuges to separate the heavier U_{238} . Enrichment to 90% U_{235} yields bomb-grade uranium.)

Beginning in 1973, an international uranium cartel formed, and pushed the price of yellowcake from about \$5.00 per pound to about \$45.00 per pound over a period of two or three years. Long-term contracting facilitated their ability to push prices up sharply. Electric utilities had spent billions of dollars on nuclear power plants, and those plants would not run without uranium. Power plant owners wanted to line up supplies of uranium over fifteen or twenty years, and as prices rose, they became frantic over the availability of future supplies. As a result, uranium producers were able to lock in high prices over periods of fifteen or more years, even though (after the cartel was discovered), the spot price fell back to about \$10.00 per pound.

The uranium cartel is a somewhat extreme example. It is often the case, however, that a natural resource is a critical input for an industrial producer. Electrical equipment manufacturers, for example, cannot operate without a steady supply of copper. That dependence

enables copper refiners to sell under long-term contracts at prices well above average spot prices.