Producer Theory - Monopoly

Mark Dean

Lecture Notes for Fall 2009 Introductory Microeconomics - Brown University

1 Introduction

Up until now, we have assumed that all the agents in our economies are **price takers**. Consumers accept that there are set prices at which they can buy goods and sell their services. Firms accept that there are set prices at which they can sell goods and sell factors of production. Of course, it should be easy enough to think of cases in which this is not a natural assumption - in fact, it strikes most people as an extremely *unnatural* assumption. Surely most firms get to choose at what price they sell their stuff?

We are now going to explicitly analyze such a situation. We are going to start off by thinking of a market in which there is only one firm - the monopolist, and they get to choose the price at which they sell their goods. What is to stop such a firm making an infinite amount of money (and so making our problem boring)? Well, we are of course going to have to also drop the assumption that the monopolist can sell as much as they want at the price they choose. Instead, the monopolist will be constrained by the demand curve of the product they are selling. This tells them how much they can sell at any given price: The problem of the monopolist can therefore be seen as the problem of selecting the profit maximizing location on the demand curve.

Before we analyze such a situation, it is worth taking a minute to think about how such a situation might arise: How could we end up with an industry in which there is only one firm, making (possibly very large) positive profits? Given our previous discussion of market entry, shouldn't these profits attract other firms into the industry? Usually, there are two reasons given for the existence of monopolies

1. **Barriers to entry** In some cases, it may be impossible for other firms to enter a market. Consider, for example, a firm who holds the patent for a particular drug. Other firms may want to start selling the same drug, but they simply cannot by law. States may also set up monopolies, by granting the rights to sell goods to only one firm (for example liquor stores in British Columbia). Note that whether a situation can be considered a monopoly in this sense depends a bit on how we define a particular market - for example Apple has a monopoly in the market for iMacs, but not in the market for personal computers

2. **Increasing returns to scale.** If the cost structure in a particular industry exhibits increasing returns to scale (decreasing marginal costs), then it may be impossible for the industry to support more than one firm: If there were two firms, then each would always be trying to grow, and therefore become more efficient, and force the other out of business. Typical examples of such industries are ones in which there are very high fixed costs - for example electricity supply - it would seem like a bad idea for there to be 7 different electricity firms, all with their own distribution system running to each house.

For the most part, we are going to concentrate on the first case: the firms that we will look at are identical to the ones we considered in the perfect competition case, apart from that we will drop the assumption of price taking. However, we will also show quickly how, in this framework, we can study the behavior of firms that have increasing returns to scale.

2 The Behavior of Monopolist

In order to think about how we are going to model the monopolist, let's first remember how we modelled the perfectly competitive firm. Being the evil capitalist swine that they were, we assumed that our perfect competitors maximized profit, of

$$\pi(q) = pq - c(q)$$

Where q is the output the firm, p is the price that they sell at and $c(q)$ is the cost of making q.

The assumption that defines the firm as perfectly competitive is that the price p is fixed. The firm can sell as much as they like at price p, but only at that price. We now want to do is relax that assumption and instead assume that the firm can select it's own price. Of course, thy cannot sell as much as they want at any price: the relationship between price and quantity is given by the demand curve in the market, or $D(p)$. Remember that this tells us the demand (or the amount that the firm can sell, at any given price. We can therefore rewrite the profit function as

$$\pi(p) = pD(p) - c(D(p))$$

Of course as the demand curve tells us that the amount the firm can sell at any given price, it also tells us the price that the firm can charge if it wants to sell a particular quantity q (i.e. the inverse of the demand function). We can write that as $p(q)$ where

$$p(q) = D^{-1}(q)$$

and rewrite the profit function as

$$\pi(q) = p(q)q - c(q)$$

We will usually work with this formulation, as it makes it easier to compare to the perfectly competitive case.

Now we have the profit function, we know what it is that the monopolist will do. Assuming that the second order conditions are satisfied, they will maximize profits by setting the derivative of this function equal to zero

$$\begin{aligned} \frac{\partial \pi}{\partial q} &= p(q) + \frac{\partial p}{\partial q}q - \frac{\partial c}{\partial q} = 0 \\ &\Rightarrow p(q) + \frac{\partial p}{\partial q}q = \frac{\partial c}{\partial q} \end{aligned}$$

What is this equation telling us? Well, exactly as in the case of the perfect competitor, the monopolist is setting marginal revenue equal to marginal cost. However, now the marginal revenue term is more complicated that it was before. In the case of the perfectly competitive firm, the marginal revenue of the firm was just the price, p, because the firm could sell as much as it wanted at that price. Now, the marginal revenue has to take into account that producing and selling one more item will decrease the price at which the firm can sell all it's output - that is the $\frac{\partial p}{\partial q}q$ term . Note that, if this term is equal to 0, then the monopolist looks exactly like the perfect competitor. Effectively, we can think of a perfectly competitive firm as one that 'forgets' to take into account the fact that an increase in output will push down prices

What if the $\frac{\partial p}{\partial q}q$ term is not zero? How does the monopolist compare to the perfect competitor? Well, assuming that the demand curve is downward sloping, then $\frac{\partial p}{\partial q} < 0$, and so the marginal revenue curve is below the marginal revenue curve, as shown in figure 1. As we have shown above, the monopolist will chose to produce at the point where the marginal revenue curve crosses the marginal cost curve. In contrast, we showed that for the perfectly competitive industry equilibrium occurred when the marginal cost curve crossed the *demand* curve. The difference is shown in figure 2. The monopolistic chooses to produce an amount q_m, which is below the perfectly competitive equilibrium output q_e. In order to sell this output, they charge the price p_m, which is *above* the perfect competitive price p_e.

How does this effect the relative surpluses of the actors in our economy? Remember that consumer surplus is still the area between the price line and the demand curve, while producer surplus (i.e. profit) is the area between the price line and the marginal cost curve. Figure 3 shows the producer and consumer surplus for the monopolistic case. The consumer surplus is given by the triangle D, while the produced surplus is given by the rectangle A plus the triangle B. How does this compare to the perfectly competitive case? Well, two things have happened: first some consumer surplus has shifted to the producer (rectangle A), but there is also a deadweight loss (triangle D). Producers are unambiguously better off, and consumers are ambiguously worse off. Moreover, on this measure (and in fact, in general) monopolies are inefficient.

One interesting question is how the profits of the monopolist relate to the properties of the demand function. To see the link, remember that

$$\pi(q) = p(q)q - c(q)$$

and so

$$\begin{aligned} p(q) + \frac{\partial p}{\partial q}q &= \frac{\partial c}{\partial q} \\ p(q)\left(1 + \frac{\partial p}{\partial q}\frac{q}{p}\right) &= MC(q) \\ p(q) &= \frac{MC(q)}{\left(1 + \frac{\partial p}{\partial q}\frac{q}{p}\right)} \end{aligned}$$

but remember that $\varepsilon(q) = -\frac{\partial q}{\partial p}\frac{p}{q}$ is just the price elasticity of demand, or the percentage response of demand to a percentage change in price, so

$$p(q) = \frac{MC(q)}{\left(1 - \frac{1}{\varepsilon(q)}\right)}$$

Thus, the 'mark up' that a monopolist charges over marginal cost is dependent on the elasticity of demand with respect to prices: The more elastic that demand is, the lower the mark up.

In order to fix ideas, let's consider the case in which the inverse demand and cost functions are linear (we will be using this case a lot). So

$$\begin{aligned} p(q) &= a - bq \\ c(q) &= cq \end{aligned}$$

Then profits are given by

$$\begin{aligned} \pi(q) &= (a - bq)q - cq \\ \frac{\partial \pi}{\partial q} &= a - 2bq - c = 0 \\ &\Rightarrow q = \frac{a-c}{2b} \end{aligned}$$

And price is given by

$$\begin{aligned} p &= a - bq \\ &= a - b\frac{a-c}{2b} \\ &= \frac{a+c}{2} \end{aligned}$$

2.1 Increasing Marginal Cost

So far, we have considered firms that have looked 'standard' in the sense that the have had increasing marginal costs. However, we can also use the same analysis in the case of decreasing marginal costs. Remember that

$$\pi(q) = p(q)q - c(q)$$

and so

$$\frac{\partial \pi}{\partial q} = p(q) + \frac{\partial p}{\partial q}q - \frac{\partial c}{\partial q}$$

Note that, even, if $\frac{\partial c}{\partial q}$ decreases as q increases, in may still be the case that profits will start falling at some point as $\frac{\partial p}{\partial q}q$ is negative, and may be increasing in q (i.e. for linear demand functions). This is the situation shown in figure 4. In such a case, the monopolist will maximize profits by producing a positive, but non-infinite level of output.

3 Price Discrimination

So far, we have restricted attention to the case in which the monopolist can only charge a fixed per-unit price to all its customers. In fact, firms have all sorts of clever strategies that can allow them to extract more of the surplus from a market . Intuitively, the best that the monopolist can do is to extract the total surplus from the market - i.e. all the producer surplus and all the consumer surplus. There is no way that the monopolist can extract more than that, as the consumer will refuse to enter the market if it costs them more than their total consumer surplus. We are now going to examine some of the ways that that the monopolist can boost their share of the

3.1 Two-Part Tariff.

One thing that some firms do is to charge what is called a **two part tariff**. The two part tariff consists of two fees: A fixed fee for joining the service, then an additional fee every time you use the service - for example zipcar, or pay-as-you-go mobile phone tarrifs

Why might this be a good idea? To see this, lets think about the person with quasi linear preferences. Remember that this is someone who has preferences over two goods x and y given by

$$u(x, y) = v(x) + y$$

Lets think about x as being the movies of Adam Sandler, and y being money. If you recall, for such a consumer , we showed that the area under the demand curve (the consumer surplus) was equal to the utility they got from being allowed to buy the videos of Adam Sandler.

Let's imagine the demand curve for our consumer is given by $a - bx$, and the cost curve for the monopolist given by cx. If the firm was only allowed to set one price, then they would do so by setting

$$\begin{aligned} a - 2bx - c &= 0 \\ &\Rightarrow \quad x = \frac{a-c}{2b} \end{aligned}$$

and making profits equal to

$$\begin{aligned} &\frac{a-c}{2b}\left(a - \frac{a-c}{2}\right) - c\left(\frac{a-c}{2b}\right) \\ = &\frac{(a-c)^2}{4b} \end{aligned}$$

The profits are shown in figure 5. Now Consider the following strategy for the monopolist

1. Charge the consumer a fixed amount

$$\frac{(a-c)^2}{2b} - \alpha$$

   For the right to buy Adam Sandler videos where α is a small positive number

2. Charge c per movie for those that have chosen to buy the right to buy moves.

What will happen in this case? Well, if the consumer chooses to purchase the rights to but movies, then number of movies they will buy is given by

$$\begin{aligned} x &= \frac{a-p}{b} \\ &= \frac{a-c}{b} \end{aligned}$$

Their consumer surplus is therefore given by $\frac{bx^2}{2} = \left(\frac{a-c}{b}\right)^2 \frac{b}{2} = \frac{(a-c)^2}{2b}$. The utility they get from being allowed to buy in this market is therefore the consumer surplus *minus* the fixed cost of being allowed to buy Adam Sandler films, or

$$\begin{aligned} &\frac{(a-c)^2}{2b} - \frac{(a-c)^2}{2b} - \alpha \\ = \ &\alpha \end{aligned}$$

Thus, it is (just about worth) the consumer paying the fixed cost, and yet the profits of the firm are given by $\frac{(a-c)^2}{2b} - \alpha$, which is much larger than it was in the one tariff case.

So what is going on here? The answer is shown in figure 6: The firm is maximizing *total* surplus, by setting price equal to marginal cost, then extracting the whole surplus with the fixed cost it charges the consumer. This outcome is therefore efficient (though not great for the consumer.)

One question you might ask is: why do we not see more of this type of behavior, given it seems so great for the firm? In fact, if anything we seem to see more of the opposite - firms subsidizing fixed costs in order to attract customers. One possible explanation for this is hyperbolic discounting. We will come back to this in later lectures, but the basic idea is the people may be very averse to paying costs here and now, and so firms that charge large up front costs may not do well.

3.2 Perfect Price Discrimination

So far, we have concentrated on the case in which there is only one consumer, or at least all consumers are identical. Remember that another interpretation we had of the demand curve is that it represented the aggregate behavior of a bunch of different consumers, each with a different willingness to pay. If this is the case, what can the monopolist do to extract more of the surplus for themselves?

It turns out the answer is: If they can charge difference prices to different people, quite a lot! Imagine that, for each quantity q, there is an individual that is prepared to pay $p(q)$, and that these prices fall as quantities increase giving the demand curve. Imagine also that the firm can charge different prices to different people (a situation called perfect price discrimination). The best thing the firm can do is charge an amount $p(q)$ to each person. Thus (and assuming constant marginal cost), the profit from selling an amount q is given by

$$\pi(q) = \int_0^q p(x)dx - cq$$

Optimal output is given by

$$\begin{aligned} \frac{\partial \pi(q)}{\partial q} &= p(q) - c = 0 \\ p(q) &= c \end{aligned}$$

Thus, the firm keeps producing up to the point where price is equal to marginal costs, and extract all the consumer surplus. This is therefore another outcome which is efficient (in that total surplus is maximized), but is not great for consumers.

This is, however, considered to be something of an extreme case, as it requires three strong assumptions

1. The monopolist can identify all different classes of consumer, and charge them different price
2. The monopolist knows the maximum price that each person will pay
3. There is no re-selling of goods (otherwise the person being charged the lowest price could just buy enough for everyone, then resell. it on.

3.3 Second Degree Price Discrimination

What about the case in which the firm cannot discriminate between different purchasers - or in other words cannot charge different prices to different customers? Consider the following example. Say that there are two consumers in the market. One has a demand curve given by

$$D_1(q) = a - bq$$

and the other by

$$D_2(q) = a - dq$$

with $b < d$. The two demand curves are shown in figure 7. Let's also assume that the consumers have quasi-linear preferences, so the area under the demand curve is equal to additional utility the get from being allowed to buy at a particular price.

If the firm could perfectly discriminate between the two consumers, they would like to make consumer 1 buy an amount q_1, and charge him an amount equal to their total surplus $(A + B + C + D + E + F)$. At this price, consumer 1 is indifferent between buying q_1 or not buying anything. Similarly, the firm would like to sell consumer 2 an amount q_2 and change them an amount equal to their total surplus $A + D$. Thus, a monopolist might try and offer individuals two possibilities:

1. Buy an amount q_1 for a total price A+B+C+D+E+F
2. Buy an amount q_2 for the total price A+D

However, what would agent 2 do in such a situation? Well, the utility they get from buying amount 1 for A+B+C+D+E+F is zero. However, the surplus they get from buying q_2 at $A + D$ is equal to B. Therefore, they would choose to buy the q_2 - or pretend to be the other type of consumer. In other words these prices are not *incentive compatible*.

In order to get consumer 2 to buy the correct bundle, they would have to charge them $A + C + D + E + F$, which would give them the same surplus as if they bought the amount q_2 at price $A + D$. Is this the best that the consumer can do? The answer is no, as you will see for homework.