GEOG590-002/EVPP741-002 Fall 2006: Theory of the Firm: Profit Maximization and Cost Minimization*

Dawn C. Parker, Dept. of Environmental Science and Policy

September 13, 2006

1 Learning Objectives

- Understand the components of firm profits;
- Graphically identify total revenue, total costs, and firm profits;
- Understand why the firm will choose to produce where marginal revenue equals marginal cost;
- Be able to calculate the firm's profit maximizing output;
- Understand the firm's cost function, and why the firm would adjust its input mix to minimize production costs;
- Understand how changes in production technology and input costs will change the supply curve.

2 Profit Maximization

We are going to cover the topics covered in Nicholson in a slightly different order than he does. We are going to look at a firm that produces a single output, using two inputs, labor (L) and capital (K). Note that when this firm maximizes profits, it implicitly makes two choices:

1. What quantity of the good to produce,
2. What combination of capital and labor to use to produce this output.

The process of profit maximization solves both problems. But, we are going to look at the broad problem first, that of maximizing profits, and then look at the corollary problem of minimizing the cost of producing the profit-maximizing level of output. *Profits* in general are simply the difference between total revenues and total costs.

2.1 Total Revenue

Total Revenue is simply the total number of dollars that the firm takes in from sales: how much they get per unit, times how many units they sell. Right now, we are going to look at firms which are *price takers*: their decisions about how much to produce don't change the prevailing market price of the good. So, if the market price is equal to p, for each unit they sell, they gain *marginal revenue* equal to p.

The second part of revenue is the quantity produced. The quantity produced is determined by the *production function*, which describes the relationship between production inputs and the level of output of the good.[1] Nicholson has a nice discussion of production functions and a series of examples. Today, we are going to assume that we produce our good using a combination of capital and labor. (For example, when I mow my lawn, I use a combination of capital (shears, push mower, power mower, or riding mower) and my own labor.) We call capital (K) and labor (L) *factors of production*. So, our production function can be expressed generically as:

$$q = f(L, K) \tag{1}$$

and this week we will also use a specific example:

$$f(L, K) = L^{\frac{1}{3}} K^{\frac{1}{3}} \tag{2}$$

So, total revenue is just:

$$pq = pf(L, K)$$

2.2 Total and Marginal Costs

Total costs are simply the sum of the *marginal cost* for each factor of production times the amount that we employ. (This is because we are assuming that a firm can hire as much of an input as it needs at a fixed price.) Note that each factor of production has to be paid at least its opportunity cost, in theory. Following Nicholson, we will call the wage rate for labor w and the rental rate for capital v. For a given level of L and K that produce quantity q of output, our cost function is:

$$c(q) = c(L, K) = wL + vK \tag{3}$$

Keep in mind that since these firms are maximizing profits, they are going to choose their combination of capital and labor very carefully, so as to make as much profit as they can. They may substitute one factor for another as costs of each factor change. But for now, take the quantities of those factors of production, or inputs, as given, and remember that this function will tell us how much it will cost to produce any given amount of output.

[1] Remember last week how we had a utility function that described how much satisfaction we got from consuming a combination of two goods? You could almost think of the utility function as a production function for well-being.

It makes sense that as output increases, costs will always increase. However, there is an interesting concept in economics that also deals with *how* costs of output will increase as output increases. It depends on the production function, which translates inputs into outputs. This concept is *returns to scale*. We can't spend much time on it. But, for our purposes, we will assume that as this firm produces more, it costs it a little bit more to produce each additional unit, because it requires a bit more of each input to produce another unit of output. This is called *diminishing returns to scale*, and implies that the total cost curve is convex, and that the marginal cost curve slopes upward. (See figures 7.1 and 7.4 in Nicholson, and Figure **??**.)

2.3 Conditions for Profit Maximization

For some unknown reason, economists like to use the symbol Π for profit. Note that we can express our profit function, total revenue less total cost, in two different ways:

$$\Pi = pq - c(q) \tag{4}$$

$$\Pi = pf(L, K) - c(L, K) \tag{5}$$

But for now, let's stick with Equation 4, and look at our decision in terms of output.

If we were to use calculus to solve this problem, the first-order condition that defines the solution would be:

$$p - \frac{dC(q)}{dq} = 0 \Rightarrow \underbrace{p}_{\text{marginal revenue}} = \underbrace{\frac{dC(q)}{dq}}_{\text{marginal cost}} \tag{6}$$

But let's forget about that for now, and take a graphical approach to understanding why this condition should hold. (Looking at modified versions of Nicholson's figures 7.1 and 7.4) First, the 7.1-like graph. Notice that profits are just the vertical distance between the revenue and cost curves. Where the distance between these two curves is greatest, profits are at a maximum. This is also where marginal revenue equals marginal cost. The profit line tells us how much we make above our costs. (*Question:* why would we want to produce at q^*, where marginal revenue equals marginal cost, instead of where total revenue equals total cost?)

We can learn even more by looking at the second graph, a plot of marginal revenue (a straight line, since no matter how much we produce we get the same price for it) and our marginal cost curve (which slopes up because of diminishing returns to scale). A very important point about this graph: the area below the marginal cost curve and above the q axis represents the total costs for the firm. The rectangle between p and q^* represents total revenue for the firm. So the difference, the triangular area above the marginal cost curve and below the price line, represents profits for the firm. (We also call this *producer surplus.*) Note also that the height of this distance represents the *marginal profit*, or the additional profit the firm makes from producing another unit. Some questions for you:

1. If the firm produced less than q^*, how much would its profits be? How does this compare to producing at q^*?

2. If the firm produced more than q^*, how much would its profits be? How does this compare to producing at q^*?

3. So, do you think that producing at q^*makes sense?

Solving for the profit-maximizing level of output is quite simple. For example, what if the firm's marginal cost/supply curve is:

$$MC = \frac{1}{2}Q$$

and the firm receives fixed marginal revenue (the market price) of \$2 per unit. Simply set $MR = MC$ and solve for Q:

$$\frac{1}{2}Q = 2 \Rightarrow Q^2 = 4$$

3 Cost minimization

Now (hopefully) we all agree that the firm should produce q^*. It makes sense that a profit-maximizing firm would try to produce this at minimum cost. But, it also makes sense that firms that aren't profit maximizers might also try to economize on costs.[2] To be minimizing costs, the firm must be getting all the output possible from a given combination of inputs. More formally, it must be on the frontier of the production possibilities set, or, in mathematical terms, it must be that $q^* = f(L, K)$. (Remember, we talked about the production possibilities frontier on the first day.) Second, it must be making the best choices of inputs given their market prices. Even if one input contributes relatively more to production than another, if it is very expensive, it may be economically efficient to use less of that input and more of the other.[3]

In words, the firm will choose input levels in order to minimize the total cost of producing q^*, subject to the fact that only certain combinations of inputs will produce that level of output. Notice, we begin with a cost function that is in term of levels of inputs. We want to get back to the cost function that we used before, which is in terms of output, q.

It turns out that we can set up this problem quite similarly to how we set up the utility maximization problem. *Question:* What is the major difference between the two problems?

[2]For example, many people accept jobs for reasons other than making money. But, they still probably want to do their jobs as efficiently as they can, so that they have time for other things.

[3]For example, what if you could hire either a professor or a econ major to tutor you? The professor might be very good, but very expensive, so you might hire the undergrad instead, since you could afford more hours of tutoring.

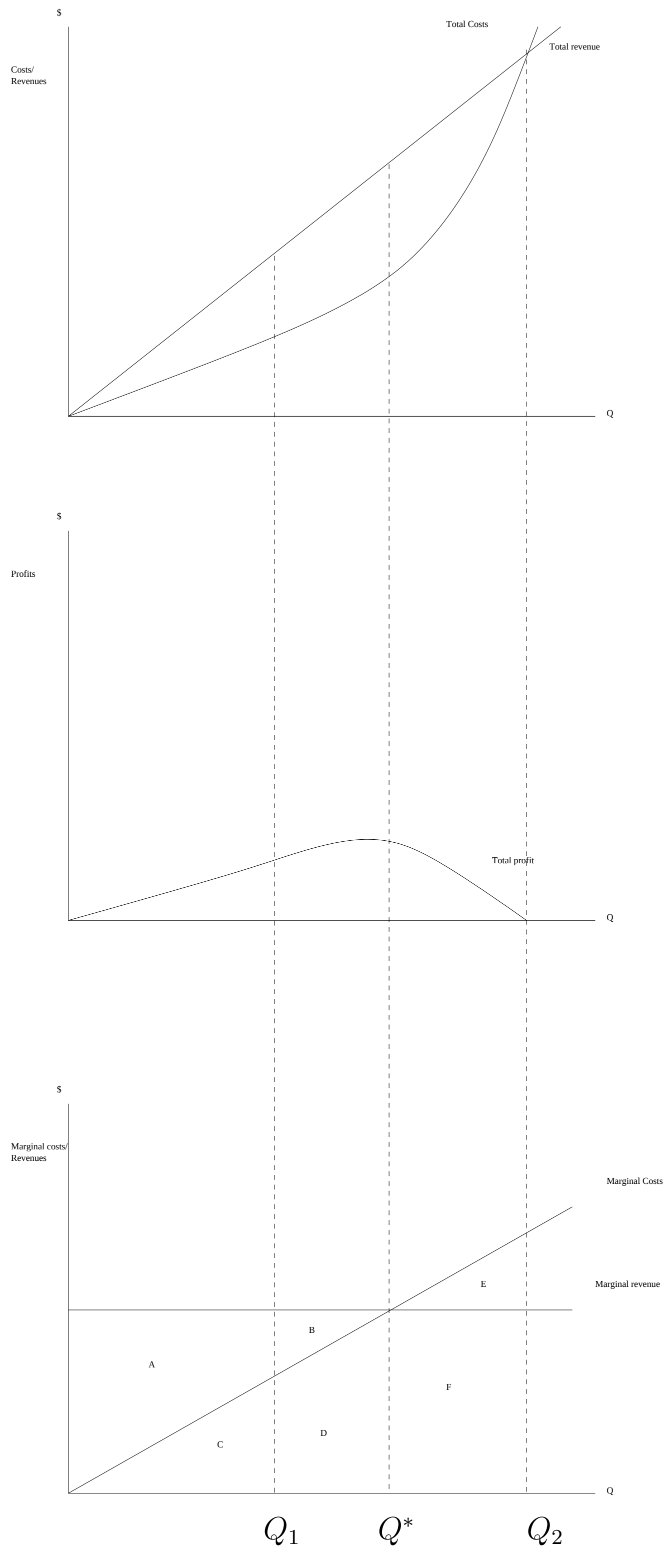

Figure 1: Profit maximization at Q^*

3.1 Production Isoquants

First, we can map out the relationship between levels of our two inputs and levels out output using an *isoquant*. This curve shows all combinations of our factor inputs that produce the same level of output. (Note similarities with indifference curves). Just as the indifference curve that we looked at exhibited diminishing marginal utility for each good, the isoquant we will look at exhibits *diminishing marginal productivity* for each factor of production. This means that, holding the other factor fixed, as we use more of one factor, the increase in total production from that factor is less than before. The lawn-mowing case that I mentioned before is an extreme example. With only one lawn mower, an additional person to help will not get the job done any faster. A more intermediate example is several people working together to prepare a meal. With only one stove and counter, the second person may not contribute as much to the production of the meal as the first, and the third may contribute even less. Diminishing marginal productivity for each input implies that the isoquants are convex to the origin. Note also that isoquants farther to the Northwest represent higher levels of output. (*Question:* why does this make sense intuitively?)

The slope of the iso-cost line is called the *marginal rate of technical substitution*. It tells us how much more of one factor we will need to produce a given amount of output, given that we reduce our use of another factor by one unit. Similar to the slope of the indifference curve, the slope of the iso-cost line is given by the ratio of the marginal products of the two factors.

3.2 Iso-cost lines

We can also draw a series of lines that represent combinations of factors of production that have the same total cost. The equation for an iso-cost line is:

$$wL + vK = \bar{C} \tag{7}$$

Question: How would you graph this line? How is it similar to the budget constraint from utility maximization? How is is different?

3.3 Conditions for Cost Minimization

The firm's goal is to produce a given output at lowest possible cost. Graphically, this means that the firm wants to reach the lowest possible iso-cost line that it can, given that it still produces the profit-maximizing output, q^*. (Nicholson Figure 6.1) This point will occur where the iso-cost line is tangent to the isoquant:

$$\underbrace{\frac{w}{v}}_{\text{slope of isocost line}} = \underbrace{\frac{MP_L}{MP_K}}_{\text{slope of isoquant } = MRTS} \tag{8}$$

We can also rearrange this equation in another form:

$$\underbrace{\frac{w}{MP_L}}_{\text{value added from labor}} = \underbrace{\frac{v}{MP_K}}_{\text{value added from capital}} \tag{9}$$

This is perhaps a more intuitive way of looking at the condition. It says that the contribution to output per dollar that you get for each input should be equal. As Nicholson says, you should be getting the same "bang for the buck".

3.4 Getting to the Cost Function

Somewhat similarly to the way that we derived a demand function by changing the price of one good and mapping out how the quantity of the other good chosen changed, we can map out a total cost function by shifting the firm's isoquant upwards. (Nicholson Figure 6.2) On the capital/labor graph, these points map out the firm's *expansion path*, which shows how cost-minimizing input use increases as output increases. Each of these points has associated with it a total cost of production. (*Question:* Graphically, how is this total cost of production for a particular level of output found?)

Since each of these points provides a quantity/total cost pair, we can also graph the set of points on a q/$ graph, such as Nicholson Figure 7.1. Now we are back to where we started, but hopefully with a deeper understanding of where the total cost function comes from.

Once again, the marginal cost curve is the slope of the total cost curve. Intuitively, it represents the cost to the firm of producing one additional unit. In order for the firm to make this decision, it needs to make at least as much on that unit as it costs the firm to produce it. So, it will supply that unit only if price is at least as high as its marginal cost. This is the sense in which the marginal cost curve also represents the firm's supply function.