Intermediate Microeconomics: Profit Maximisation and Cost Minimisation

Simon Naitram, PhD

November 17, 2020

Chapter 20: Profit Maximisation

READ SECTIONS 20.1 AND 20.5-20.12 (9th Edition). These are Sections 19.1 and 19.5-19.12 in the 8th Edition. Read and work through the Appendix to this chapter.

Now that we've described how much the firm can produce using its technology, we move onto the main thing that firms care about—profits. This leads us to our main *behavioural theory* about the firm: **the firm chooses production to maximise profits**.

Profits Let's first define profits:

$$\pi = py - wx \tag{1}$$

where p is the price of the good, $y = f(x)$ is the number of units produced, w is the cost of the input, and x is the number of inputs. The firm's aim is really to maximise profits *given* the technology it has available. Technically, the firm's technology is the firm's constraint. But when we deal with profit maximisation, we do not use the Lagrange method. Instead, we substitute for y and write this as:

$$\pi = p \cdot f(x) - wx. \tag{2}$$

This is much easier in this case, and very intuitive.[1]

[1] In the following chapter, we'll return to the Lagrange method.

Multiple Inputs Frequently we use production function that requires more than one input. Let's generalise this to two inputs:

$$\pi = p \cdot f(x_1, x_2) - w_1x_1 - w_2x_2. \tag{3}$$

We need to subtract the cost of both inputs when calculating profits.

Economic Costs Some of you may have done accounts, and know that in real life, we only subtract the costs that the firm actually pays. So if the firm uses an input (like the owner's garage) but doesn't pay for it, accountants don't count it as a cost.

In economics, however, we try to measure the market price of that input even if the firm doesn't pay for it. That is, we imagine what the firm *would have* had to pay for it. We can estimate this by considering what would have been the price the input would've received in it's next best usage. Maybe, for example, what would've been the amount someone else would pay for the owner's garage to be a workshop?

One further difference is that while accountants frequently use historical costs (the price the firm actually paid for it), economists use economic costs (they price the firm would have to pay to buy it now).

Economic Profit Economic profit is when we subtract all the 'economic costs'[2] from revenue (py). As we said before, it's important to include all the factors of production when carrying out this exercise. This means we eliminate all possible costs from revenue, coming to as narrow a definition of economic profit as possible.

[2] Including both things we didn't pay for at all, and things valued at market rather than historical cost.

The concept of economic profit is a bit weird. You'll almost never see it in reality, and firms almost never know (or care) what it is. Why do we use it? Well as economists we care about efficiency. We need to measure the returns to all factors of production, and it allows us to think about whether everything is being used in its best possible place. So think of this chapter less as a true description of how firms maximise profits, but more as a way for us to understand how the economy chooses where factors are used and how much to produce.

Competitive Output Market Much like with the consumer, we assume that the firm cannot set price. Why do we assume this? We think that because there are so many firms out there selling the same thing, consumers simply won't buy your products if you increase the price. If you're a single small player in a huge market, you can't change the entire market price. So you just take the price as given.

Competitive Factor Market Similarly, we assume that the firm takes factor prices as being fixed. This doesn't always hold in reality. For now, however, we say that the firm is one small buyer. This means, for example, that they can't change the wage being offered to a worker. Or they can't try to bargain with the gas station for a lower price on gas.

Competitive Demand You may be wondering where the customers are. Well, a competitive market assumes that there's always demand *once you sell at the market price*. That is, there are infinite consumers who want to buy your product at the price p, and no higher. So effectively we ignore the issue of demand for now.

Maximising Profit Now that we've defined profit, let's look at how we figure out what is the firm's optimal choice. The firm chooses inputs x_1 and x_2 to maximise profits, given its technology. This means the problem is simply:

$$\max_{x_1, x_2} \quad \pi(x_1, x_2) = p \cdot f(x_1, x_2) - w_1 x_1 - w_2 x_2 \tag{4}$$

Since we've substituted in for $y = f(x_1, x_2)$, there's no need for the Lagrangian method. This is now an unconstrained maximisation problem. This means we only take partial derivatives with respect to x_1 and x_2. This gives:

$$p \frac{\partial f(x_1, x_2)}{\partial x_1} - w_1 = 0 \tag{5}$$

$$p \frac{\partial f(x_1, x_2)}{\partial x_2} - w_2 = 0 \tag{6}$$

These are the first order conditions of the firm. Rearranging them says that:

$$p\text{MP}_1 = w_1 \tag{7}$$

$$p\text{MP}_2 = w_2 \tag{8}$$

where the marginal product of x_1 is the partial derivative of the production function with respect to x_1, so that $\text{MP}_1 = \partial f(x_1, x_2)/\partial x_1$.

Optimal Conditions The term MP_1 is the additional output generated by increasing x_1 by a little bit. The term $p\text{MP}_1$ is therefore the increase in total revenue that occurs when the firm increases the input x_1 by a little bit. The condition $p\text{MP}_1 = w_1$ therefore says that the firm should increase input x_1 up to the point where the extra revenue from increasing x_1 is exactly equal to the extra cost of increasing x_1. More simply, at the profit maximising point, **the revenue from the marginal product of x_1 should equal the marginal cost of x_1.**[3]

[3] Notice that the text is a bit tricky with terminology regarding the 'marginal product'. We've defined the marginal product as the increase in *output* from a small increase in x_1: $\partial y/\partial x_1$. However, Varian starts calling $p\partial y/\partial x_1$ the *value* of the marginal product. That is effectively how much the extra output generate is worth *in dollar terms*. Let me try to call it the *revenue from the marginal product* instead.

Marginal Increases The firm asks itself hypothetically what would happen if it adds a bit more of x_1. If the firm increases x_1 and the extra revenue the firm would generate is greater than the extra cost of x_1, the firm would decide that this is a good move. The firm keeps increasing x_1 as long as the extra revenue it gets is more than the extra cost. Since the marginal product of x_1 is decreasing as x_1 increases, there comes a point where the revenue from adding a bit more of x_1 doesn't cover the cost of increasing x_1. That is, when the firm hypothetically asks itself what would happen if adding a bit more x_1 would generate more revenue than it costs, there comes a point when the answer is **no**. At that point, the firm stops. It doesn't add the extra unit of x_1. It is at that point that the marginal revenue from a bit more of x_1 is equal to the marginal cost of x_1, so the firm is producing optimally!

Factor Demand When we solve the first order conditions for x_1 and x_2, we get **factor demand functions**. They are typically a function of output and factor prices (p, w_1, w_2) and the parameters of the production function:

$$x_1^*(p, w_1, w_2) \tag{9}$$
$$x_2^*(p, w_1, w_2) \tag{10}$$

These tell us how much of each factor the firm wants to buy at the optimum, taking output and factor prices as given.

Example

LET'S work through the main example. We'll use a simple Cobb-Douglas production function with two inputs. We just want to know what the factor demand functions will look like so we can interpret how a firm would choose optimally.

Problem The firm's production function is:

$$f(x_1, x_2) = x_1^{\alpha} x_2^{\beta} \tag{11}$$

Plugging this into the profit function, we get the firm's problem:

$$\max_{x_1, x_2} \quad \pi(x_1, x_2) = p(x_1^{\alpha} x_2^{\beta}) - w_1 x_1 - w_2 x_2 \tag{12}$$

This says that the firm wants to maximise profits by choosing optimal values of x_1 and x_2, given the technology it has available.

First Order Conditions The first order conditions describe the firm's optimal choices of x_1 and x_2. To get these conditions, all we do is take the partial derivative of the profit function with respect to each variable in turn.

$$\frac{\partial \pi(x_1, x_2)}{\partial x_1} : \quad \alpha p x_1^{\alpha-1} x_2^{\beta} - w_1 = 0 \tag{13}$$
$$\frac{\partial \pi(x_1, x_2)}{\partial x_2} : \quad \beta p x_1^{\alpha} x_2^{\beta-1} - w_2 = 0 \tag{14}$$

Solving Let's simplify these expressions a bit into something that makes some sense. Take the first order condition with respect to x_1

first:

$$
\begin{aligned}
\alpha p x_1^{\alpha-1} x_2^{\beta} - w_1 &= 0 \\
\alpha p x_1^{\alpha-1} x_2^{\beta} &= w_1 \\
\alpha p x_1^{\alpha-1} x_2^{\beta} &= w_1 \\
\text{using } x_1^{\alpha-1} = x_1^{\alpha} x_1^{-1}: \quad \alpha p x_1^{\alpha} x_1^{-1} x_2^{\beta} &= w_1 \\
\alpha p x_1^{\alpha} x_2^{\beta} &= \frac{w_1}{x_1^{-1}} \\
\alpha p x_1^{\alpha} x_2^{\beta} &= w_1 x_1 \\
\text{using } y = x_1^{\alpha} x_2^{\beta}: \quad \alpha p y &= w_1 x_1 \\
\text{factor demand function}: \quad \frac{\alpha p y}{w_1} &= x_1^* \qquad (15)
\end{aligned}
$$

which gives us our factor demand function. Notice that x_1 is also a function of the amount we want to produce y.

Profits

Short Run In the short run, we can assume that one of the factors of production is fixed (or unchangeable). This means that we'd only take partial derivatives with respect to the variables that can change, and treat the fixed factor as a constant. We have to continue to include it, we just treat it as not changing. For example, if x_2 is fixed, then we denote it $\bar{x}_2$. Then we simply take the derivative with respect to x_1, since this is the only thing the firm can change.

Zero Profits Consider a constant returns to scale profit function. We said that a constant returns to scale function means that when we double inputs, output would double as well. Well if the cost of the inputs are constant, then doubling inputs also means doubling profits. So with a constant returns to scale function you could always double inputs and double profits. But that would mean you could always increase profits simply by producing more. That doesn't make any intuitive sense, since it would mean that $x^* = \infty$ and $y^* = \infty$. **A constant returns to scale function should optimally generate zero profits.**

What is the idea of zero profits? Well remember we're talking about *economic profits*. Normally accounting profits are exactly the returns to entrepreneurship or returns to equity capital. But in an economic model of the firm, we consider those to be costs.[4] So if you take out a reasonable amount of 'accounting profits' that go to the owners of the firm, what 'economic profits' would be left?[5]

[4] These are valued at the market price of these factors in their next best alternative use.

[5] Also check out the reasons given in the text in Section 20.10. I find this explanation to be far more compelling, however.

Revealed Profitability

If you thought revealed preference was interesting, you'll similarly find this compelling. The analogy is simple. It isn't easy to directly observe the firm's production technology. It's probably hard to look at a firm and simply write down a production function. We can use the firm's observed choices to understand their optimal behaviour.

Simply put, **for any given prices the firm's chosen production plan was more profitable than other feasible production plans.** We can use this simple observation to approximate the firm's production technology.

Let's consider a production function with a single input $f(x)$. The firm makes profit $\pi = p \cdot f(x) - wx$. Now let's add superscripts that capture the time period. So at time period 1, the firm faces prices p^1 and w^1. Let's call its optimal (observed) choice x^1. This means the firm makes a profit:

$$\pi^1 = p^1 \cdot f(x^1) - w^1 x_1 \tag{16}$$

The simple observation coming from the assumption that the firm is profit maximising, is that there's no way the firm could have made *more profits*. So this is the upper bound of the firm's profitability given prices (p^1, w^1). We know that the observed choice was feasible, and anything that made *more profit* was obviously an unfeasible production plan. We we are putting an upper boundary on the firm's production function. Can we shrink this boundary a bit more to get closer to the firm's production function?

When prices change (either factor prices or output prices or both) from (p^1, w^1) to (p^2, w^2), we observe the firm's new choice of inputs (x^2) and its new level of profitability (π^2). Once again, we know that this production plan is feasible, and anything that makes more profits at these prices must have been unfeasible, otherwise the firm would have chosen it—obviously. We are now saying that we can at least put an upper limit on where the firm's production technology might lie. As we observe more price changes and more optimal production plans, we can make this bound tighter and tighter, and we get closer and closer to observing the firm's true production technology.

Chapter 21: Cost Minimisation

THIS CHAPTER isn't easy reading. We will cover Section 21.1 thoroughly, but go through only the main points of 21.2 to 21.6 briefly. We'll go through the appendix in some detail. Let's focus on getting the *main idea* in this chapter. This is Chapter 20 in the 8th Edition.

Main Idea The main idea in this chapter is that you can flip the question we asked last week. Instead of saying that the firm maximises profits, we make this equivalent to saying that the firm minimises costs. This is, in fact, a more natural way to think about how businesses work. They frequently decide *how much they want to produce* first, then they figure out what is the most efficient way to produce them.

For Given Output In this version of the firm's problem, we take the amount of output as given. What we're trying to do now is to derive a decision rule that tells us: *for any level of output, what's the best way to produce that amount of output given input prices?*

It's a slightly different question. How does it relate to what we were asking before? Before, we were asking "given prices and technology, how much output should we produce so as to maximise profits?" This time, we are now ignoring the part about maximising profits. So we've split off a chunk of the question, imagining for a week that we magically knew exactly how much to produce.[6] Instead, we focus on taking that number and organising our business so as to produce that *most efficiently*.

[6] Or maybe the firm's resident economist told them how much to produce!

Hotel Think of a hotel. Consider a hotel that is built with 100 rooms. They have decided that this is the optimal amount they are going to produce. Now, all that is left to do on a day-to-day basis is try to run the hotel as efficiently as possible. That is, keep all 100 rooms on the market while being as cost-effective as possible.

Combos This type of analysis is obviously only useful when there are a bunch of different *ways* to produce a given amount of output. If there were only one way to produce 100 pairs of shoes, then why are we discussing the *most efficient* way?

Cost The cost that we're seeking to minimise is easy. This is simply the sum of the price of each input multiplied by the quantity of the input:

$$c(x_1, x_2) = w_1 x_1 + w_2 x_2. \tag{17}$$

Output Since we are saying that the firm must produce a specific amount of output, this amount of output becomes the constraint! What is exactly is the constraint? The constraint is that your choices of x_1 and x_2, when put through the black box that is the firm's production function, must give you exactly the amount of output that you want. So we're *constraining* ourselves to *only* combinations of

x_1 and x_2 that give us exactly the amount of output we want. This constraint is simply:

$$f(x_1, x_2) = y. \tag{18}$$

As in profit maximisation, the firm's technology is always its constraint. In cost minimisation, however, we keep the constraint.[7]

[7] Unlike in profit maximisation, where we substitute in for y.

Cost Minimisation Problem The firm's cost minimisation problem therefore takes on a familiar constrained optimisation format. It is minimisation this time, however. The cost min problem is:

$$\min_{x_1, x_2} \quad w_1 x_1 + w_2 x_2$$
$$\text{s.t.} \quad f(x_1, x_2) = y \tag{19}$$

Min v Max Remember that minimisation is the same idea as maximisation, just in reverse. We now want to reach the bottom of the valley instead of the top of the hill. From a calculus perspective, the approach is basically the same. Why? The ground still needs to be flat, whether at the peak of the mountain or the bottom of a valley.

Optimal Choices In this case, we end up with a **conditional factor demand function**. Why are they *conditional*? Well these demand functions are *for a given amount of output.* So they are *conditional* on the amount of output we are trying to produce. These demand functions are a function of the *input* prices and output. They take the form:

$$x_1^*(w_1, w_2, y) \tag{20}$$
$$x_2^*(w_1, w_2, y) \tag{21}$$

Since we're not concerning ourselves with profit, the price of the final good is of no interest to us. Think of a manager whose only job is to keep the hotel running efficiently—he has no control over price setting. That's way above his pay grade.

Cost Function These conditional demand functions give us optimal quantities of the inputs x_1^* and x_2^*. When we take these optimal values, we can figure out what the total cost of producing is. How? Just plug those values back into the original cost function. This would give us:

$$c^*(w_1, w_2, y) = w_1 x_1^*(w_1, w_2, y) + w_2 x_2^*(w_1, w_2, y) \tag{22}$$

This is an *optimal cost function*. Since the firm is cost minimising we can replace x_1 and x_2 with their decision rules (conditional factor demand functions). Since it contains decision rules (conditional factor demand functions), optimal cost is now a function of input prices and output rather than $c(x_1, x_2)$. Why? Those demand functions depend on (w_1, w_2, y). So now the cost function depends on the exact same things.

Optimisation The process of optimisation is pretty easy with the Lagrangian method. What do you do? Set up your Lagrangian first:

$$\min_{x_1, x_2, \lambda} \quad \mathcal{L} = w_1 x_1 + w_2 x_2 - \lambda(f(x_1, x_2) - y). \tag{23}$$

First Order Conditions The first order conditions are easy. Jut take the first derivative of the Lagrangian with respect to x_1, x_2, and λ.

$$\frac{\partial \mathcal{L}}{\partial x_1}: \quad w_1 - \lambda \frac{\partial f(x_1, x_2)}{\partial x_1} = 0 \tag{24}$$

$$\frac{\partial \mathcal{L}}{\partial x_2}: \quad w_2 - \lambda \frac{\partial f(x_1, x_2)}{\partial x_2} = 0 \tag{25}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda}: \quad f(x_1, x_2) - y = 0 \tag{26}$$

Condition for Optimisation Once more, let's combine the first two first order conditions to find a single condition that describes the *cost minimising* point. Let's solve both equations for λ and put them together.

$$w_1 - \lambda \frac{\partial f(x_1, x_2)}{\partial x_1} = 0$$

$$w_1 = \lambda \frac{\partial f(x_1, x_2)}{\partial x_1}$$

$$\frac{w_1}{\partial f(x_1, x_2) / \partial x_1} = \lambda \tag{27}$$

You'd get a similar equation for w_2 and x_2. Putting them together:

$$\frac{w_1}{\partial f(x_1, x_2) / \partial x_1} = \frac{w_2}{\partial f(x_1, x_2) / \partial x_2}$$

$$\frac{w_1}{w2} = \frac{\partial f(x_1, x_2) / \partial x_1}{\partial f(x_1, x_2) / \partial x_2}$$

$$-\frac{w_1}{w2} = -\frac{\partial f(x_1, x_2) / \partial x_1}{\partial f(x_1, x_2) / \partial x_2} \tag{28}$$

Once again, we find out that the technical rate of substitution must be equal to negative of the ratio of factor prices. That is:

$$-\frac{w_1}{w2} = -\frac{\partial f(x_1, x_2) / \partial x_1}{\partial f(x_1, x_2) / \partial x_2}$$

$$-\frac{w_1}{w2} = \text{TRS} \tag{29}$$