Intermediate Microeconomics: Cost Curves, Firm Supply and Industry Supply

Simon Naitram, PhD

November 20, 2020

Chapter 22: Cost Curves

READ SECTIONS 22.1 AND 22.2 in the 9th Edition or Section 21.1 and 21.2 in the 8th Edition.

The previous chapter ignored the firm's profit-maximising goals, and instead focused on minimising costs. We found that when we minimise costs, we find a *conditional* factor demand function. These conditional factor demand functions $x_1^*(w_1, w_2, y)$ and $x_2^*(w_1, w_2, y)$ depend on factor costs and output. When we plug those optimal values back into the cost function, we end up with a cost function

$$c(w_1, w_2, y) = w_1 x_1^*(w_1, w_2, y) + w_2 x_1^*(w_1, w_2, y). \tag{1}$$

This cost function now depends only on the factor prices and the output the owner chooses. Effectively then, this cost function tells us the minimum total cost the firm can produce any amount of output at, given factor prices. Think of it as a function that no matter what values you plug in, it always gives you the lowest possible cost you can produce that output at.

Cost Function Now let's make this simpler. Like before, we assume that factor costs are fixed. The firm cannot choose how much to pay its workers or the price of gas. If factor costs can't be changed, the the cost function is really only a function of output. This means we can rewrite this function as $c(y)$.

Fixed Cost v Variable Costs Fixed costs are those costs that must be paid whether or not the firm produces 1 or 10,000 units, or even none at all. Therefore, fixed cost does not depend on y. Instead, only variable costs depend on y. We can express this as:

$$c(y) = c_v(y) + F \tag{2}$$

where $c_v(y)$ are the variable costs which are dependent on the amount of output the firm produces (like electricity), and F are the fixed costs which the firm pays regardless (like rent).

Average Costs As before, we can simply divide total cost by y to get average cost. This gives:

$$AC(y) = \frac{c(y)}{y} = \frac{c_v(y)}{y} + \frac{F}{y}. \quad (3)$$

Look at these equations. When y goes up, F doesn't change. That means that as y increases, the last term F/y always falls. That is, average fixed costs will always decline as the firm produces more. On the other hand, variable costs are also determined by y. Whether variable costs are increasing or decreasing depends on the returns to scale of the short run production function. *Importantly, because there is some fixed factor of production creating fixed costs in the short run, then we must have a decreasing returns to scale production function.*[1] This means that as output increases (the firm expands), then average variable costs are expected to increase.

When we combine these two, the average total cost will typically fall as we begin increasing production, then begin to rise.

Marginal Costs Remember we always find the *marginal* by taking a derivative. We want to know how much cost will increase if we increase production by a little bit. This gives us:

$$MC(y) = \frac{\partial c(y)}{\partial y} = \frac{\partial c_v(y)}{\partial y}. \quad (4)$$

MC & AC Think of producing cricket bats. The first one cost you \$100 to produce; the second one cost you \$90 to produce; the third cost you \$80 to produce. After the first two, the average cost was \$95. The third one cost less to produce that this average cost. This means that when we take the average of the three, the average cost will have fallen to \$90. It means that if the cost of producing the *next unit* if lower than the *average of all the units* that have gone before, then the average cost will fall if you produce the next unit. The opposite hold too. Suppose the fourth unit costs \$100 to produce. What is the average now? It's \$92.50. The average cost has risen because the marginal cost of producing the next unit was higher than the average cost was before.

This is important. Whenever marginal cost is below average cost, then average cost is falling. Whenever marginal cost is above average cost, then average cost is rising.

Let's take a step further. Remember we said that average costs fall then rise. When average costs are falling, then marginal costs are below average costs. When average costs are rising, then marginal costs are above average costs. Then there must be a point where average costs are exactly equal to marginal costs. If marginal costs are a bit

[1] Remember I said that if we can double the entire world, then our new world must be able to produce at least as much as the original. The only way that doubling the entire world wouldn't be able to match the old one is if we forgot to double something. This means that the 'something' is fixed, and we can't change it. This gives us decreasing returns to scale. Mathematically, this says that with a Cobb-Douglas function, we're missing a factor so that the exponents don't add up to 1.

above average costs, then average costs are rising. If marginal costs are a bit below average costs then they're falling. Just in between there, average cost must change from falling to rising. And just in between there, that must be where marginal costs are equal to average costs. **Marginal costs are equal to average costs when average costs are at a minimum.**

Marginal cost is a good concept to understand. How much will your company's cost increase if you produce one more unit? This then begins to allow you to think of what that one unit is worth to your company!

Average Cost This is simple. Average cost is always just the cost function divided by y. That is, total cost divided by the number of units you're producing. If you've got your factor demands, then you can plug them into the original cost function and get $c(w_1, w_2, y)$, which tells you the lowest cost you can produce y units of output. Divide this by y and you get the average cost function:

$$\text{AC}(y) = \frac{c(w_1, w_2, y)}{y} \tag{5}$$

The average cost is an intuitive measure of how cost-efficient you are.

Constant Returns to Scale Let's think about this for a moment. If you have constant returns to scale, it means that doubling inputs will lead to exactly double the output. Remember that cost is given by $c(x_1, x_2) = w_1x_1 + w_2x_2$. If factor costs remain constant, the total cost will exactly double when we double inputs, obviously. If y doubles as well, then we get:

$$\text{AC}(y) = \frac{2(w_1x_1 + w_2x_2)}{2y} = \frac{c(w_1, w_2, y)}{y} \tag{6}$$

This gives you back exactly the same average cost as you had before! This says that for a constant returns to scale function, the average cost is *always* the same for all levels of output!

Increasing Returns to Scale For an increasing returns to scale function, when I double inputs, I get *more* than double the output. What does this imply about average cost? Well total cost doubles, but output more than doubles. So the average cost is now:

$$\text{AC}(y) = \frac{2(w_1x_1 + w_2x_2)}{3y} < \frac{w_1x_1 + w_2x_2}{y} \tag{7}$$

This means that, for an increasing returns to scale function, average cost falls as we increase output!

Decreasing Returns to Scale For a decreasing returns to scale function, when I double inputs, I get *less* than double the output. What does this imply about average cost? Well total cost doubles, but output less than doubles. So the average cost is now:

$$\text{AC}(y) = \frac{2(w_1x_1 + w_2x_2)}{1\frac{1}{2}y} > \frac{w_1x_1 + w_2x_2}{y} \tag{8}$$

This means that, for a decreasing returns to scale function, average cost increases as we increase output.

Sunk Costs The text provides a really nice example of sunk costs. Sunk cost is an idea that you can actually quite broadly apply to your life. Read it, and think about how often you come to a point where you say: "but I've put too much into it already to give up now".

Chapter 23: Firm Supply

READ SECTION 23.3 CAREFULLY, or Section 22.3 in the 8th Edition. Read Section 23.2 (22.2) to get a nice background on what is competitive market.

Let's work from this point. How much is the extra unit of output worth to your firm? Well, given that you face a fixed selling price[2], then it is worth however much you can make from selling it—the selling price. Let's call the extra income you get from selling the additional unit your marginal revenue.

[2] Remember we said we are operating in a competitive market? There are lots of firms selling exactly the same thing, so you just have to take the market price as given. You can't change it yourself.

Optimal So if we know how much the extra unit of output costs, and we know how much the extra unit of output will sell for, we can figure out exactly how much we should be producing. If the extra income you get from producing the extra unit is higher than the extra cost of producing it, you should probably produce it since it's a *net benefit* to your firm. When you reach the point where producing an extra unit costs more than the extra income you get, then it's not a good idea to produce that extra unit. In more technical terms, we should continue increasing producing as long as marginal revenue is higher than marginal cost. **The point where we stop producing is exactly where marginal revenue is equal to marginal cost!**

Since we know that marginal revenue is equal to price[3], then we can write this condition for optimality as:

[3] The extra money you get on an extra unit is exactly the price it is sold for.

$$p = MC(y) \tag{9}$$

Optimisation Let's do that mental exercise as an optimisation problem now. Remember in cost minimisation, we said we would split the problem into two? The owner chooses how much output he wants to produce and the manager takes that as given and minimises cost? Well the manager has done his job, and has given the owner a function that tells him what the lowest cost would be for any level of output he chooses. The owner, equipped with this minimum cost function $c(y)$, now has to choose the amount of output that would maximise profit. So we're back at a profit maximising problem. The owner's problem is therefore:

$$\max_{y} \quad \pi = py - c(y) \tag{10}$$

This looks a lot easier than the original profit maximising problem we did, right? Well that's because the manager has already done half the job for us by finding the optimal minimum cost function $c(y)$. All we need to do now, is choose the optimal y that maximises profits. Let's take the derivative and set it equal to zero to get our first order condition:

$$\frac{\partial \pi}{\partial y} = p - \frac{\partial c(y)}{\partial y} = 0 \tag{11}$$

or, we can rewrite this more simply as:

$$p = \frac{\partial c(y)}{\partial y}. \tag{12}$$

Once again we find that the optimal condition for profit maximisation is to produce the amount of output y^* where price is equal to marginal cost.

Chapter 24: Industry Supply

READ SECTIONS 24.1 TO 24.3, AND 24.5 TO 24.6 in the 9th Edition (or the corresponding Sections for Chapter 23 in the 8th Edition).

Industry supply is just as easy as market demand. All we're doing is summing the total supply of all the firms. If all firms take the output price as given, they will each choose an optimal amount of output y_i^*. Let's call the firm's choice of output it's *supply* given price p. Let's write this as $S_i(p)$. The industry supply is simply:

$$S(p) = \sum_{i=1}^{n} S_i(p) \tag{13}$$

Profits A firm optimally produces where marginal cost is equal to marginal revenue (the price). At this point of production, the firm's short run average cost might be higher than, lower than, or equal to price. If average cost is higher than price, the firm is making a loss in the short run. If average cost is lower than price, the firm is making a profit in the short run.

Loss A firm that makes a loss in the short run but makes a profit in the long run will stay in business. A firm that makes a loss in both the short run and the long run will exit the industry (go out of business). By exiting (not producing), the firm can limit its losses to zero in the long run.

Eliminating Profit If firms can enter, what about exit? If firms in the industry are making a profit in the long run, then other firms will be attracted to enter the industry. New firms will get all the factors of production the others are using and replicate their production processes and start their own gig. This entry will eventually eliminate the profits that existing firms were making. New firms will harm the economic profits of old firms.

Economic Profits Remember we're talking about economic profits here. So we're eliminating the opportunity cost of shareholders investing their money somewhere else, the opportunity cost of entrepreneurs starting other businesses, and the like. In equilibrium, zero economic profits implies that *all factors of production are being paid their market price.* It implies that shareholders aren't earning *excess profits* or *supernormal profits*. Do you think firms should earn supernormal profits?[4]

[4] Profits beyond the opportunity cost of their original investment.

Fixed Factors Fixed factors of production might mean that firms can't enter the market as they please. Entry requires that new firms can get all the factors they need. What if they can't? What if they're only owned by certain people?[5] What if ownership is regulated by law? This means that new firms might be unable to get the factors they need to enter the market.

[5] Think of limestone quarries. Or oil deposits. Or even beaches!

Zero Profit Identity But this should mean that the fixed factor increases in value. Because ownership of this factor gives the owners economic profits, then its value should increase by exactly the amount of economic profit. Why? Well because the amount of economic profit is exactly the amount of additional money someone else would be willing to pay to get that fixed factor of production, because they'd get to earn those economic profits. This is exactly the definition of opportunity cost. This means that the economic profit is really always zero if you measure the definition correctly, even if entry into the market isn't free.[6]

[6] Note that this talk about zero profits has nothing to do with returns to scale, and all to do with entry and exit.