Chapter 13
The Firm as a Price Taker

Quotes

Mr. Scorpio says productivity is up 2% and it's all because of my motivational techniques, like donuts, and the possibility of more donuts to come.

--Homer Simpson

Introduction

In many industries firms are price-takers, meaning they have little over control over the price they receive for their output and the price they pay for their inputs. Profits for these firms depend crucially on buying just the right amount of inputs and producing just the right amount of output. This chapter first discusses the production process itself. This production process is then extended to evaluate input decisions by the firm. Basically, we will use information on prices and the production process to answer the question: how much of an input should a firm purchase? The discussion then leads to production costs. It turns out that a firm can be described simply as a set of cost curves. In some ways this is unrealistic, as firms consist of many people and complexities, not just a few curves on a graph. Yet, those curves capture many important features of a firm in an easy-to-understand picture. We give up a little bit of reality for a large amount of understandability. We use these curves and price information to answer the all important question to a firm: how much should the firm produce? Finally, the cost of production to a firm as they differ across the short-run and long-run is considered.

To reiterative, the goals of this chapter is to

1. Develop a model of production for a firm.
2. Use this model to study input use by a firm.
3. Extend these concepts to understand costs of production.
4. Describe how a price-taking firm maximizes profits.
5. Distinguish between firm behavior in the short- and long-run.

It is common for students to consist the material in this chapter " a bunch of irrelevant junk." It is easy to think we are oversimplifying agribusiness management. In fact, we agree that we have greatly simplified the task of managing a firm, but for good reason. When engineers design bridges, they first build model bridges to make sure it will work and to understand the things that can go wrong. Before a company mass produces a new item they conduct extensive market research on small groups of consumers. Before NASA built the Space Shuttle you can bet they played around with much smaller versions and simulated the outer space environment first. In a similar vein, this chapter deals with how to maximize profits for a simple price-taking firm where costs of production are known and never change. If you cannot maximize profits in this simple model of the firm, chances are, you will not be able to maximize profits for a real complex firm either. Learn the basics of a golf swing before you compete. Learn the basics of agribusiness management in simple settings before you do the real thing.

The Three Stages of Production

Strip away all the complexities of a firm, and you basically have people turning inputs into outputs. Hop growers take machinery, nitrogen, and hop seed to grow and harvest hops for use in beer. Turkey growers take buildings, corn, water, and medicine to produce turkey meat. Production, the process of transforming inputs into outputs, is an important concept in the study of the firm. A *production function* is a mathematical formula that indicates the output attained from a given level of input.

Imagine you have a plant that has not been watered for weeks. Its leaves are withered and its stem barely able to hold the plant erect. Now you go to water the plant. At first, you give it just a little bit of water. This tiny amount of water is just enough to keep the plant alive, but doesn't make it look any healthier. Now you apply a little more, and the plant begins to respond by holding up its leaves again and the stem becoming more erect. This second watering brought the plant from just surviving to recovering. Put differently, you got a bigger "bang for your buck" from the second watering than the first. Then you apply a little more, and the plant responds tremendously. Again, the "bang for your buck" was bigger from the third than the second watering. Its leaves stand out and strong, and the plant itself stands straight up. From this point on when you apply more water the plant grows, but its health doesn't make the giant leap it did before. And if you keep adding water to the plant you will eventually drown it. This story completely describes a firm's production function. The story describes the relationship between any input and its associated output, whether it be fertilizer to corn or rubber to whoopee-cushions.

Consider a slightly more complex example of a tomato canning facility. Many types of inputs are used in this facility. There is the facility itself, consisting of buildings and land. Forklifts are used to transport crates and tractor-trailers haul tomatoes to the facility. Wage workers operate the machinery. Salaried employees manage the business, perform accounting activities, and coordinate sales. Electricity, fuel, oil, packaging, and machinery parts are just a few of the additional inputs. The firm takes all these inputs, produces canned tomatoes as an output, and sells the canned tomatoes. Assume the canned tomato industry is perfectly competitive, such that the firm is a price-taker.

Some inputs are fixed in the short-run. These include the facilities, machinery, and salaried employees. Other inputs like electricity and wage labor are variable, meaning they can be easily and quickly increased or decreased. Let us explore the relationship between one of these variable inputs, labor, and output. This will serve as a metaphor between any variable input and output. The units of labor will be the number of workers employed each day. A production function tells us the number of canned tomatoes (*total product*) produced for each worker employed each day.

Suppose we begin with no workers, and increase it to one worker. The increase in total product from increasing the input level by one is referred to as the *marginal product*. Suppose the addition of one worker over zero workers increases total product from 0 to 5 cans. The marginal product of the first worker is then 5. Suppose a second worker is now added. Before, one worker had to do everything, but with two workers each one can specialize in specific activities. One may concentrate on loading raw tomatoes into the machine while the other concentrates on unloading canned tomatoes. Specialization allows each worker to be more productive and total product surges from 5 to 20 units. The marginal product of the second worker is 15, and is greater than the marginal product of the first worker. From hiring the second worker, the firm gets "a bigger bang for their buck" than the first worker. <Insert Margin Item 13.A>
<Insert Figure 13.A>

Initially, when more laborers are added they specialize and productivity for all workers rises. As a result, the marginal product of each additional worker is greater than the previous worker. More importantly, the marginal product of the additional worker is greater than the average product of all workers (see Figure 13.A). In Figure 13.A, the marginal product of the first worker is 5, but is 15 for the second worker. So long as the marginal product is rising, or if the marginal product is greater than the average product, we say the firm is in the *first stage of production*. The first stage is important because, if the marginal product is greater than the average product, this indicates that each additional worker is more productive than the average existing worker, so adding that worker will always increase the average productivity in the canning facility. An analogy can be made to Beckham joining a soccer team. Beckham is more productive than the average soccer player, so adding Beckham to a soccer team will increase the average productivity of the team players (because the average now includes Beckham, and Beckham is better than all team members).

Eventually as more workers are added, there will come a point where adding another worker increases output, but not as much as the previous worker. Production facilities are generally designed to match workers with machines, and a certain number of workers is designed to operate each machine. Once the number of workers exceed this number, those new workers are performing tasks that we could say are "less important." One tomato canning machine may be operated efficiently using two workers, and one forklift is operated by only one worker. Another worker is useful, but does not contribute as much as previous workers. The marginal product is positive, but is a smaller value. For example, going from two to three workers increases production from 20 to 30 in the figure above. The marginal product of the third worker is ten, which is less than the marginal product of 15 for the second worker. The marginal product can be declining, but so long as the marginal product is greater than the average product, the firm is still in the first stage of production. Even though the marginal product is declining, the marginal is greater than the average product, so the average product is rising.

For example, after Beckham is added to the soccer team, suppose we add another player named Rodrigo. Rodrigo is good, not as good as Beckham, but better than the average player on the team. While Rodrigo's marginal product might be less than Bechham's, because Rodrigo is better than the average team player, adding him to the team will still increase the average performance of the soccer team players. The marginal product is declining, but it is still greater than the average product.

In the latter part of the first stage, the marginal product will continue to decline and the average product will continue to rise. Eventually, the marginal and average products will equal, and that point begins the s*econd stage of production.* Adding workers now still contributes to production, but both the marginal and the average product are declining. The canning facility is getting crowded. More people can still be hired, and there are jobs for them to do, but these jobs are not as important as the previously filled jobs. While they contribute towards the facility's production of canned tomatoes, their contribution is not as large as previous hires.

Finally, there will come a point where adding a worker actually detracts from production. Imagine we kept adding workers in a factory to the point where it was too crowded to even move. If you cannot move, you cannot work, and those new workers cause production to fall. This is the *third stage of production,* where the marginal product of an input is negative. <Insert Margin Item 13.B>

The three stages of production are illustrated in Figure 13.B. It is useful to consider another example to illustrate the production function. Consider the production of wheat using nitrogen as an input. All the other inputs of production (machinery, labor, pesticides) are assumed fixed, and we want to evaluate how wheat yield responds to changes in nitrogen use. Without applications of chemical nitrogen to wheat, very little yield will be realized. But wheat responds to the first couple pounds of nitrogen by greatly increasing the plant growth. Another couple pounds and wheat responds even more by even greater growth. The marginal product of nitrogen is increasing at low levels of use; meaning the marginal product of the 4th pound is greater than the marginal product of the 3rd pound. This is the first stage of production. As you add more nitrogen, wheat will still respond by growing taller, but the response is not as pronounced. The marginal product is still positive, and is falling and eventually becomes less than the average product, at which point you are in the second stage of production. Finally, you will hit a stage where adding nitrogen does not increase wheat yields, and may even decrease it. This is the third stage. If you are not familiar with fertilizer and crops, just reconsider our first example of watering a plant. The plant responds to the watering at first, but there comes a point when more water does not increase plant growth, and too much water will drown the plant and detract from growth. <Insert Figure 13.B>

The presence of the three stages of production give rise to the S-shaped production function in 13.B. In the first stage, marginal product is rising as input use is increased. The marginal product is positive, but declines to a values less than the average product in the second stage. At the third stage, the marginal product is negative and greater input use detracts from total product. The marginal product is really just the slope of the production function at a particular point. To illustrate, the slope between points on the production function is illustrated by the red dotted lines in 13.B. The slope tells us how fast production changes with input use. In the first stage, the slope is positive and becomes steeper as more input is used. The slope is still positive, but becomes less steep in the second stage. Notice that the slope would be highest at the exact same input level that the marginal product curve is at its peak. This is no coincidence, marginal product is the slope of the production function. Finally, at the third stage, the slope of the production function becomes negative.

In addition to marginal products, the average product is an important concept. As before the relationship between marginal and average product will be illustrated with an analogy. Suppose we have a baseball team and output is measured by total number of homeruns hit by all players during a game. At first we have no players, but then we start adding players. Every new player is better than the previous player, meaning player 2 is better than player 1, player 3 is better than player 2, etc. Each new player hits more homeruns than the previous player. The marginal product of each player is rising, and we are in the first stage of production. Then, after the fifth player is added, we start adding new players such that each new player hits worse than the previous player. Player 6 is worse than player 5, player 7 is worse than player 6, etc. Each of these new players hit homeruns, so they increase output, but they increase output by less than the previous player. Eventually, as more players of increasing inferiority are added, one player will hit less homeruns than the average player on the team. This is stage two of production.

Average product is measured as the total output divided by total inputs used. In our baseball analogy, it equals total homeruns divided by the total number of players. In stage one, each new player is better than the average player, and average product is rising. Adding better players increases the average performance of the team. Once we hit stage two, each new player is worse than the previous player. The new players may be contributing to the team by hitting homeruns, but not as many homeruns as the average team player, and the new players drag the average performance down.

Optimal Input Use

This is a good point to stop and look at actual production data. Figure 13.C shows lbs of nitrogen applied to wheat and the corresponding wheat yields for Oklahoma. The yields are taken from actual Oklahoma yield measurements and nitrogen use.[1] We cannot directly measure the marginal product of nitrogen at each level of nitrogen, because nitrogen is not increased in increments of one. It is increased in increments of 20 lbs. However, we can obtain an estimate of the marginal product in a range of input use through the formula $MP = \Delta q / \Delta x$ where Δ means "change", q is total product, and x is input level. Thus, $\Delta q / \Delta x$ is the change in total product divided by the change in lbs of nitrogen. To see how the formula operates, suppose that you work five more hours and your output increases by 100. The change in output is 100 and the change in input use (hours worked) is 5. The marginal product formula yields $MP = \Delta q / \Delta x = 100/5 = 20$. What this means is that for each additional hour worked your output increases by approximately 20 units. That is the very definition of marginal product.

This formula simply takes the slope between two points on the production function (remember, slope equals rise over run, $\Delta q / \Delta x$). For example, when nitrogen is increased from 0 to 20, the marginal product within this range is $MP = \Delta q / \Delta x = (30.5 - 23.0) / (20 - 0) = 0.375$. That is, when using between zero and 20 lbs of nitrogen, an extra pound of nitrogen increases wheat yields by about 0.375 bushels per acre.

Using this formula we can see how the marginal product of nitrogen changes with nitrogen use. In Figure 13.C the marginal product is declining as more nitrogen is used for all nitrogen levels. The figure also shows the average product of nitrogen. Moreover, the marginal product is everywhere less than the average product. That is, it appears production is in the second stage of production throughout. What happened to the first stage? Production must have shifted from the first stage to the second stage somewhere between 0 and 20 lbs of nitrogen, so it cannot be detected in the data. What happened to the third stage? If more than 100 lbs of nitrogen were applied, and one kept adding more and more, there would come a point when wheat yields did not change, and would eventually start falling. While we cannot observe the first and third production stages in these data, it clearly demonstrates the second stage, which turns out to be the most important stage anyway. The impact of nitrogen on wheat yields declines the more nitrogen that is used. Therefore, there will come a point where

[1] The data are based off experimental plots, but are adjusted to reflect typical nitrogen use rates.

increasing nitrogen use is not profitable. At this point, the increase in yield does not justify the additional nitrogen purchase. Let us employ the concept of marginal product to articulate the profit maximizing level of nitrogen use. An additional pound of nitrogen should be used if the value it provides is greater than its cost, just like you do not purchase an item unless you value it more than its price. The value of one more pound of nitrogen is the additional yield it provides (the marginal product) times the value of that yield (the market price of wheat). Not surprisingly, we refer to this as the *marginal value* of input use. The market price for wheat is usually around $3.25 / bu, so the marginal value of nitrogen can be calculated as the marginal product times $3.25, as shown in Figure 13.D. The marginal cost of nitrogen is roughly $0.15. <Insert Margin Item 13.C> <Insert Figure 13.D>

Using the above table, suppose the farmer initially plans on applying 20 lbs of nitrogen. Should more than 20 lbs be used? At 20 lbs, the marginal value of one more pound is $1.22. That is, by increasing nitrogen use from 20 to 21 pounds per acre, revenue increases by $1.22 per acre. The marginal cost of increasing nitrogen use from 20 to 21 lbs is only $0.15, so profits increase by $1.22 - $0.15 = $1.07 / acre. The verdict is clear: more than 20 lbs should be applied. This reveals a very important production concept: an additional input should be used whenever the marginal value is greater than the input price. Consulting Figure 13.D, the marginal value of nitrogen is greater than the input price up to 60 lbs per acre. If nitrogen could only be purchased and applied in increments of 20 lbs per acre, the farmer should apply 60 lbs, and the optimal yield (i.e. profit maximizing yield) is 37.8 bushels per acre. <Insert Margin Item 13.D> <Insert Figure 13.D>

This is an important concept. The profit maximizing level of output is *less* than the production maximizing level of output. Put differently, firms to not want to produce at their maximum production level. At 60 lbs of nitrogen per acre, one could increase yield by using more nitrogen, but this would lower the firm's profits.

When Should Cattle Be Sold?

Feedlots are in the business of purchasing adolescent cattle, feeding them a high protein diet and selling them when they are big enough for slaughter. Cattle within a feedlot are referred to as live-cattle (also fed-cattle). Feedlot profits depend critically on selling cattle at the optimal weight. Each day cattle are on feed, they gain weight, which increases the revenue from each head of cattle sold. If the market price for live-cattle is $0.75 per pound, then each extra pound increases revenues by $0.75. However, for each day cattle are kept in the feedlot they must be fed, and that feed cost money. It costs roughly $1.41 for each additional day a cow is fed. To maximize profits, the manager keeps cattle in the feedlot so long as the value of the extra pounds produced outweigh the costs of feeding the cattle. Our input is "*days on feed*", and the output is pounds per cow.

One of the authors collected data on cattle weights and the number of days cattle were on feed to determine the marginal product of days on feed. Using the regression analysis technique discussed in Chapter 6, a formula for marginal product was calculated as

Marginal Product (MP) = 4.36 – 0.0157(*DOF*)

where *DOF* stands for *days on feed*. The first day cattle are brought into the feedlot *DOF* = *1*, and the 50th day of being in the feedlot *DOF* = 50. The marginal product equation is interpreted as follows. Suppose a lot of cattle has been held in the feedlot for 25 days, and the manager is considering keeping them one day longer. She can expect each head of cattle to gain MP = 4.36 – 0.0157(25) = 3.9675 lbs from that extra day. If the cattle have been on feed for 200 days, one day of extra feed only results in MP = 4.36 – 0.0157(200) = 1.22 lbs in extra weight. Eventually, there will come a point where the value of an extra pound is less than the cost of feed. At that point, the feedlot manager should sell the cattle. As before, the marginal value of *days on feed* is the marginal product times the market price of live-cattle, denoted P_{LC}. In 2004, the average live-cattle price was around $0.75 / lb. The marginal cost or input price of *days on feed* is denoted p_{DOF} and is about $1.41. That is, keeping cattle on feed one additional day costs $1.41 per head. This includes all the feed, labor, veterinary, and other costs associated with feeding an animal one additional day. Profits are maximized by keeping cattle on feed so long as the marginal value of *days on feed* is greater than the price. The greater the *days on feed*, the lower will be the marginal value because cattle put on less weight as they age. Eventually, the marginal value will just equal the price, and the cattle will be sold.

The input decision is demonstrated in Figure 13.E. For low input levels, the marginal product is high, higher than the input cost. So long as marginal value is greater than the input price, more of that input is used and profits rise. For example, when cattle have been on feed 50 days, the marginal product is 3.9675 lbs. Holding cattle one more day increases cattle weights by 3.9675 lbs per cow. At a price of $0.75 per lb, this translates into $2.98 per head of greater revenues. The cost of holding cattle one more day is only $1.41, so increasing days on feed from 50 to 51 days increases profits by $2.98 - $1.41 = $1.57 per head. Using the same logic, the feedlot manager should keep increasing *days on feed* until the marginal value equals the input price. Using the calculations below, profits are maximized by keeping cattle on feed approximately 158 days. <Insert Figure 13.E>

Notice that the optimal *days on feed* is given by the intersection of the marginal value curve and the input price. Thus, the marginal value curve is the input demand curve for the firm. This relationship can be used to study how input use changes. First, consider the obvious case where the input price rises. To equate marginal value with the input price the firm must use less of the input. As less of the input is used, marginal product rises, aligning marginal value with the new input price. Using the same logic, if the price of an input falls the quantity demanded by the firm rises.

Now consider the case where there is technological progress, and more output can be produced using the same amount of input. An example is the use of growth hormones in cattle, which increases the rate at which cattle convert feed to muscle. This is another way of saying growth hormones increase the marginal product at all input levels. Given that the marginal value equals marginal product times cost, if marginal product rises the marginal value curve shifts upward. The intersection of input demand and input price is now at a higher input level. The firm buys more of the input.

Consider again the business of producing live-cattle, where *days on feed* is an input. Some cattle are administered regular, low doses of antibiotics in their feed and water to promote growth. Because this poses a health hazard, some lawmakers have considered banning this practice. What would happen to the optimal *days on feed* if these antibiotics were eliminated? Without these antibiotics, cattle would grow at a slower rate, putting on fewer pounds each day. This is just another was of saying the marginal product would fall. If the marginal product falls, then the marginal value curve shifts downward. The firm's demand for *days on feed* falls, and live-cattle are sold and slaughtered at a younger age. This suggests that a ban on antibiotics not only hurts antibiotic makers and cattle producers, but corn producers as well. As the marginal product of *days on feed* falls, feedlots feed their cattle fewer days, reducing their demand for feed and lowering their corn purchases. Finally, suppose that the price of live-cattle rises. Each additional pound gained by cattle now translates into greater revenues. Since the demand curve is marginal product times the output price, the firms' demand for *days on feed* rises and cattle are slaughtered at an older age. <Insert Figure 13.F>

Costs of Production

Production is the process of turning inputs into outputs. Inputs come at a price, so the ability of a firm to produce at a low cost depends on how well they use their inputs. Increasing production requires the use of more inputs, yet a firm does not always have control over the level of all inputs. A milk processing facility uses expensive buildings, machinery, milk, labor, and raw materials as its inputs. If a milk processing facility wishes to process greater volumes of milk, it can increase the use of all those inputs. However, it takes time for a new building to be built. While the firm can immediately vary its use of milk, labor, and most raw materials, it cannot immediately increase or decrease the number of its processing facilities (like milk tanks, building, piping). In a period of time, some of its inputs are held fixed and cannot be varied. We refer to this as the *short-run*, when at least one input is fixed. Given time, the firm can build more facilities. Given time, the firm can increase or decrease the level of *all* inputs used. The length of time it takes to vary all inputs is referred to as the *long-run*. As an example, suppose the milk processing facility can freely vary all inputs at any time except its processing facility (buildings and milk processing machinery). It can easily increase or decrease labor, volume of milk, and raw materials. However, it takes two years for a new processing facility to be built. The short-run is then less than two years and the long-run is greater than two years. <Insert Margin Item 13.E>

At any point in time, some inputs can be varied and some inputs may be fixed. Increasing production requires greater use of those variable inputs, and the costs of those variable inputs are referred to as *variable costs*. Variable costs are costs that vary with production. When production rises, so do variable costs. Conversely, *fixed costs* refer to the cost of inputs that do not change with the level of output. Using our milk processing facility, recall that in the short-run the number of processing facilities cannot be changed. The firm likely took a loan to pay for the building, and each month it must make a payment to the bank. That payment is the same regardless of how much milk is processed, making it a fixed cost.

In this section we wish to build a simple economic model illustrating how input use translates into production costs. Let production or total product produced by the firm be denoted q. Total costs equal variable costs plus fixed costs. Denote fixed costs as FC. Suppose that the only variable cost (VC) is labor. The amount of labor used (# of hours) is given by L and the price of labor is the wage rate w ($ per hour). Variable costs in this example equal labor use time the wage rate: variable costs = wL. Fixed costs equal FC regardless of the output level or number of hours worked.

$$\text{Total Costs (TC)} = VC + \text{FC} = wL + \text{FC}$$

In previous chapters we extensively relied on the concept of *marginal cost*, which is the additional cost incurred from increasing production by one unit. The formula for measuring marginal cost is the change in total cost divided by the change in output: $\Delta TC / \Delta q$ where Δ means "change." To illustrate the formula, suppose increasing production by one increases costs by ten. The marginal cost is then $\Delta TC / \Delta q = 10 / 1 = 10$. If increasing production by two increases costs by 16, a measure of marginal cost is $\Delta TC / \Delta q = 16 / 2 = 8$. On average, it costs $8 for each additional unit produced. In our example, the only way to increase production is to increase labor, as all other inputs are held fixed. The change in total costs equals the change in labor times the wage rate. If the wage rate is $10 per hour and 50 more hours are employed, costs rise by $w\Delta L$ = ($10)(50) = 500.

$$\text{Change in total costs} = \Delta\text{TC} = w\Delta L$$

Then, noting that marginal cost is the change in total cost divided by the change in quantity produced, we obtain marginal cost by dividing both sides of the previous equation by Δq.

$$\text{Marginal Cost (MC)} = \frac{\Delta\text{TC}}{\Delta q} = w\frac{\Delta L}{\Delta q} = \frac{w}{\left(\Delta q / \Delta L\right)} = \frac{w}{\text{MP}}$$

Notice that the term $\left(\Delta q / \Delta L\right)$ is simply the marginal product. Marginal cost can simply be stated as the input price divided by the marginal product of that input. This makes sense. The higher the price of inputs, the greater the cost of purchasing those inputs and producing a product. The higher the price of labor, the greater the cost of using laborers to take raw vegetables and package them into ready-to-eat salads. The formula also tells us that the greater the marginal product the lower the marginal cost. The more productive your laborers are the less your cost of using those workers to produce a product. At this point we should revisit our previous discussing of marginal product and the three stages of production. The marginal product curve is recreated in Figure 13.G below. Marginal cost is inversely related to marginal product. In the earlier portion of the first stage of production marginal product is rising, which means marginal cost must be falling. Marginal product starts falling in the latter portion of the first stage, and so marginal cost starts rising. At some point in marginal products' descent, it will equal the average product. This is where the second stage begins, and as you can see from Figure 13.G, this is also the point where marginal cost equals average cost.

The concept of average cost is also important. First consider average variable costs (AVC), which is simply variable costs divided by output. Using some algebra, we obtain the relation

$$\text{AVC} = \frac{\text{VC}}{q} = \frac{wL}{q} = \frac{w}{\left(q / L\right)} = \frac{w}{AP}.$$

Average variable costs are simply the input price divided by the average product. Before, we showed that at low levels of input use average product is rising. Inputs are becoming more productive, so average variable costs naturally fall as a result. However, at some point the average product will begin falling, which increases average variable costs. Average variable costs are then inversely related to average product, as shown in the below figure. <Insert Figure 13.G>

Finally there is average total costs (ATC), which are total costs divided by total product. The only difference between average variable costs and average costs are average fixed costs, as shown below.

$$\mathrm{ATC} = \frac{\mathrm{VC}}{q} + \frac{\mathrm{FC}}{q} = \frac{wL}{q} + \frac{\mathrm{FC}}{q} = \frac{w}{\left(q/L\right)} + \frac{\mathrm{FC}}{q} = \frac{w}{AP} + \frac{\mathrm{FC}}{q}$$

Fixed costs (FC) are constant, so as output rises (q becomes large) the average fixed costs (FC / q) will become small. For large levels of output, the difference between average variable costs and average costs are small. As illustrated in Figure 13.H, at zero output there is a significant difference between the average variable cost and average cost curve. This difference is fixed costs. However as total product become large the two curves converge. <Insert Figure 13.H>

A Numerical Example

To help illustrate, consider a numerical example. Refer back to the hypothetical production function in Figure 13.A, which is recreated below. The input is labor, which we will assume costs $5 per worker. Fixed costs are $15. Variable costs are calculated simply as the number of workers times the per worker cost of $5, and total costs are variable costs plus the $15 fixed cost. Marginal cost, average variable cost, and average cost are calculated as described in the formulas above, and recreated below.

Marginal Cost = (change in total product) / (change in total costs)
Average Variable Cost = (variable cost) / (total product)
Average Cost = (total cost = variable cost + fixed cost) / (total product)

Marginal cost is decreasing a first, then starts to increase. Average variable costs follow a similar pattern, although its minimum is at a higher quantity. The average cost curve lies above the average variable cost curve. For the remainder of this chapter we will focus exclusively on these curves, but note that their shape is derived from the three stages of the production function and input prices. <Insert Figure 13.I>

Maximizing Profit

In the previous sections we saw that firms will continue to use more inputs until the marginal value of the input equals its price. While we concentrated on one input at a time, in reality managers seek to equate marginal value and price for *all* inputs. This maximizes the firm's profits. In this section, we want to study how a firm maximizes profits when it is a price-taker. The firm cannot control prices, neither the price it receives for output nor the price it pays for input. We assume that however much the firm is producing, it is producing that amount at the least cost possible. The only variable left to discuss is how much the firm should produce. That is, given a marginal cost, average variable cost, and average cost curve, what is the profit maximizing quantity? The answer turns out to be surprisingly simple, following the two rules below.

Determining the Profit-Maximizing Output Level

Step 1: If price is greater than the minimum average variable cost, go to Step 2. Otherwise, produce nothing.

Step 2: Find the quantity where price equals marginal cost and produce that quantity.

Consider the equation for profit: Profit = *Revenues - VC - FC*, where *VC* is variable and *FC* is fixed costs. We can multiply and divide *VC* by the output level q. This leaves the equation unchanged, as $(VC/q)*q = VC$. However, noting that VC/q is simply average variable costs, it allows a convenient representation for profit.

$$\text{Profit } = \text{Revenues} - \left(\frac{VC}{q}\right)(q) - FC = \text{Revenues} - (AVC)(q) + FC$$

If we then note that revenue equal price times quantity, $(P)(q)$, profits can be rewritten as

$$\text{Profit } = (P)(q) - (AVC)(q) + FC = [(P) - (AVC)](q) - FC\,.$$

If the firm does not produce anything, making $q = 0$, profits simply equal $-FC$. However, so long as it can produce some quantity where average variable costs is less than profits, it can make profits greater than $-FC$ and should do so. In the equation above, if the firm can produce at some quantity where $[(P) - (AVC)]$ is positive it make more profits than just $-FC$. It might not be making money, but at least it will not be losing as much as it could. If price is less than average variable costs and a firm produces anyway, this would be like buying ground beef for $3.00, making it into a burger and selling the burger for $2.00. Not only do you have to pay fixed costs like rent on the building but you loss money on the individual burger. But if the variable costs of a burger are $4.00 and you sell the burger for $5.00, you can use that $1.00 profit to help pay fixed costs like building rent. You might not cover all your fixed costs, but at least you made money to help pay your fixed cost. If revenues are greater than variable costs the firm is better off producing. It might cover all its fixed costs to and make a profit or it might not. But if revenues are less than variable costs the firm should not produce because it will lose its fixed costs and more. A smart firm, at worse, will never lose more than its fixed costs. <Insert Figure 13.J>

Thus, if price is P_1 in Figure 13.J the firm will not produce anything. Price is below the average variable cost line, meaning no matter how much the firm produces it will never cover its variable costs. It is better off just not producing anything and paying the fixed cost. However, if price is P_2 the firm can at least cover its variable cost. It should produce some quantity. Specifically, it should produce q_2 units. As a rule the firm should produce another unit whenever price is greater than marginal cost. If it only costs you $100 to produce another unit and you receive a price of $150, of course you should produce another unit. Following this unit, a firm keeps increasing production until price equals marginal cost. At this point producing another unit increases costs more than price, and detracts from profits. If price if P_2, the firm produces where price equals marginal cost, at q_2. Using the algebra described above, we can write profits as

$$\text{Profit} = [(P)-(AC)](q).$$

At a price of P_2 and quantity of q_2, the average cost curve is higher than price. The $[(P)-(AC)](q)$ is negative and the firm is losing money. However, since it can still cover its variable costs it is best off cutting its loses by producing q_2 units. Next observe what happens if price is P_3. The firm again produces where price equals marginal cost (given that price is greater than the minimum average variable cost) which is at q_3. Now, price is greater than average cost and the firm enjoys a profit.

Towards the Long-Run

In our discussion above we considered how costs changed if firms increased or decreased output. Throughout, we assumed that at least one input was being held fixed. In one example, labor could be varied but the number of tomato canning processing facilities could not. In another example, the nitrogen applied to a crop could be varied, but the number of tractors, combines, and grain bins could not. As the firm moves from the short-run to the long-run, inputs that were formerly fixed are now under the control of the firm. The tomato canning facility can now increase the number of processing facilities in addition to labor, and the wheat farmer can purchase more equipment in addition to nitrogen. A useful, general view of the firm is one in which raw materials and labor are variable in the short-run, but capital (e.g. buildings, machinery, offices) cannot. However, in the long-run capital can be changed as well. In the long-run, all inputs are variable. There are no fixed costs, so the average variable cost and average cost curve become one in the same.

For simplicity, let us assume that there are only two inputs: labor and capital. Capital is fixed in the short-run, but labor is not. In the long-run, both labor and capital can be varied. As you might expect, labor and capital generally complement each other. They are complements in production, meaning an increase in the use of one enhances the productivity of the other. Try painting a house with a paintbrush and then try it with a spray gun and you will understand this complementarily. Your natural painting abilities have not changed a bit, but with a spray gun in hand, your productivity soars.

The average cost curves drawn previously assume that capital is fixed. It was like assuming we varied the number of painters while holding the number of spray guns constant. But if we can increase both the number of painters and number of spray guns we greatly increase our painting productivity. Now able to paint two houses per day instead of one, our cost of painting each house declines. This is referred to as economies of scale (or increasing returns to scale); where the firm expands its production ability by increasing the use of all inputs and finds that its average costs fall.
<Insert Figure 13.L>

Economies of scale typically disappear at some production level. In Chapter 4 we discussed how beef processing firms have learned they can greatly reduce their average costs by "getting big." Building bigger processing facilities that rely more heavily on machinery and automated processes, they can transform live-cattle into beef at a lower per pound cost. Again, this is called increasing returns to scale. When increasing returns to scale exist, firms will get bigger. Their lower costs allow them to sell their output at a lower price, driving their competitors who did not "get big" out of business. However, if a single beef processing firm keeps growing and growing, they need greater and greater supplies of cattle. To induce cattle producers to increase their production, they must receive a higher price for their cattle. As the price of cattle rises, so too does the beef processing firm's costs, and average costs begin to rise. We refer to this as *diseconomies of scale* or *decreasing returns to scale*, where greater output leads to an increase in average production costs. In general economists say that at lower levels of output there are increasing returns to scale, but as output continually expands the firm will eventually realize decreasing returns to scale. Think of football. If you start off small, gaining weight will usually improve your playing ability. But once you hit a certain weight you are no longer big but obese, and additional weight hurts your playing ability. As you might expect there is an optimal weight that maximizes your playing ability. Similarly, there is one firm size that minimizes firm costs. This point is called the *minimum efficient scale*.

Between increasing and decreasing returns to scale is the *minimum efficient scale*, the point where long-run average costs are at its lowest. In Figure 13.L there are two types of average cost curves. The long-run average costs illustrates how costs change over the long-run when all inputs can be varied. The other average costs curves are short-run curves, dictating how costs change with production when at least one input is fixed. Not surprisingly, the long-run average costs are always lower. This is because in the long-run the firm has control over more inputs, and therefore has more options at its disposal to keep prices low. <Insert Figure 13.L.>

This figure should be interpreted in the following way. Take Point A. If the firm wishes to increase production in the short-run, at least one input is held fixed. This hinders the ability of the firm to expand efficiently. In a tomato canning facility there are two ways to expand production. The first is to make everyone work overtime, exerting more effort in the same factory. This is not very efficient because workers become exhausted and machinery runs more than it is designed to run. Another is to open a new factory and hire new workers. This second method is usually more efficient (so long as the output increase is to be permanent), meaning it allows you to produce at a lower cost per unit. The first method refers to a short-run output expansion, like moving from Point A to Point B in Figure 13.L. Since capital (the processing facility) is being held fixed, the increase output moves along the average cost curve AC. The long-run expansion increases both hours worked by laborers and capital, and so average costs are even lower. The firm moves along the long-run average cost curve; from Point A to C. The same output expansion is achieved at a lower average cost.

The difference between the short- and long-run boils down to the number of options available to the firm. A firm wants to increase output, and we presume it increases output at the lowest cost. The more inputs the firm can vary, the more options it had to find a less expensive means for expanding output. In the short-run the firm's options are limited because some inputs are held fixed. The long-run gives the firm more ways to increase output at a lower cost, and if a less expensive way can be found the firm will pursue it. Just like a consumer can be made happier by being given more choices, a firm can decrease its cost by being given more freedom to alter input usage.

These average cost curves are simply a snapshot of the firm. In reality, the curves themselves change as input prices change and with technological advancements. Consider the brewing industry between 1950 and today. Beer production is usually measured in millions of barrels, where one barrel is 31 gallons. In 1950 the minimum efficient scale for brewers was 0.1 million barrels. This was the lowest point on the

long-run average cost curve for the typical firm. Since then, technological innovations have allowed brewers to realize economies of scale past 0.1 million barrels. After these technological changes, the minimum efficient size for brewing increased to 8 million barrels in 1970 and 18 million barrels in 2000. While the demand for beer has grown since 1950, mainly since incomes and populations have risen, demand has not grown as fast as the minimum efficient scale. As a result, it takes fewer brewers to meet the market demand for beer. Some brewers had to leave the market, and it would undoubtedly be those with higher costs. As the figure below shows, the increase in the minimum efficient scale had a pronounced effect on the number of mass-producing brewers (number of brewers excluding your local micro-brewery). This number has fallen from 350 in 1950 to 24 today. Today, most all the beer consumed in the U.S. is produced by one of three brewers: Anheuser-Busch, Miller, and Coors.[2] This is in large part due to a growing minimum efficient scale. It pays for firms to get bigger, and they do. Yet consumer demand does not grow as fast, so some firms go out-of-business.
<Insert Figure 13.M>

<u>Summary</u>

In the introduction to this chapter we discussed how engineers build models of bridges before the real thing, and NASA simulates how space shuttles will fly in space before actually flying there. If you cannot build a model bridge, you cannot build the real thing. And if you cannot manage model or simulated businesses profitably your chances at the real thing are small. Agribusiness managers face numerous, complex problems; such as, how many inputs like fertilizer or labor should be used each week, how much output should be produced each month, and how much capital should be procured to ensure the long-run profitability of the firm?

To prepare students for decision-making in these complex environments, this chapter built a "model firm." We assumed both the output price and input prices are fixed, and that the firm produces a single identical product. Based on these assumptions the "model firm" became a collect of marginal product curves and cost curves. The chapter then discussed how to determine the optimal input use, production levels, and how firms' costs differ in the short- and long-run.

[2] Tremblay and Tremblay.

Margin Item 13.A

Marginal Product: the additional output realized from increasing the use of an input by one unit.

Margin Item 13.B

The marginal product is positive and greater than the average product in the first stage, positive and less than the average product in the second stage, and negative in the third stage of production.

Margin Item 13.C

The marginal value of an input equals the marginal product of the input times the price of the good being produced. It is the increase in the value of output from increasing input use by one.

Margin Item 13.D

An additional unit of input should be used whenever the marginal value is greater than the input cost.

Margin Item 13.E

Short-Run: a period of time over which at least one input is being held fixed
Long-Run: a period of time after which no inputs are held fixed

Margin Item 13.F

Variable Cost: a cost that changes with production
Fixed Cost: a cost that remains the same regardless of the production level