

CHAPTER 20

COST MINIMIZATION

Our goal is to study the behavior of profit-maximizing firms in both competitive and noncompetitive market environments. In the last chapter we began our investigation of profit-maximizing behavior in a competitive environment by examining the profit-maximization problem directly.

However, some important insights can be gained through a more indirect approach. Our strategy will be to break up the profit-maximization problem into two pieces. First, we will look at the problem of how to minimize the costs of producing any given level of output, and then we will look at how to choose the most profitable level of output. In this chapter we'll look at the first step—minimizing the costs of producing a given level of output.

20.1 Cost Minimization

Suppose that we have two factors of production that have prices w_1 and w_2 , and that we want to figure out the cheapest way to produce a given level of output, y . If we let x_1 and x_2 measure the amounts used of the

two factors and let $f(x_1, x_2)$ be the production function for the firm, we can write this problem as

$$\min_{x_1, x_2} w_1 x_1 + w_2 x_2$$

$$\text{such that } f(x_1, x_2) = y.$$

The same warnings apply as in the preceding chapter concerning this sort of analysis: make sure that you have included *all* costs of production in the calculation of costs, and make sure that everything is being measured on a compatible time scale.

The solution to this cost-minimization problem—the minimum costs necessary to achieve the desired level of output—will depend on w_1 , w_2 , and y , so we write it as $c(w_1, w_2, y)$. This function is known as the **cost function** and will be of considerable interest to us. The cost function $c(w_1, w_2, y)$ measures the minimal costs of producing y units of output when factor prices are (w_1, w_2) .

In order to understand the solution to this problem, let us depict the costs and the technological constraints facing the firm on the same diagram. The isoquants give us the technological constraints—all the combinations of x_1 and x_2 that can produce y .

Suppose that we want to plot all the combinations of inputs that have some given level of cost, C . We can write this as

$$w_1 x_1 + w_2 x_2 = C,$$

which can be rearranged to give


$$x_2 = \frac{C}{w_2} - \frac{w_1}{w_2} x_1.$$

It is easy to see that this is a straight line with a slope of $-w_1/w_2$ and a vertical intercept of C/w_2 . As we let the number C vary we get a whole family of **isocost lines**. Every point on an isocost curve has the same cost, C , and higher isocost lines are associated with higher costs.

Thus our cost-minimization problem can be rephrased as: find the point on the isoquant that has the lowest possible isocost line associated with it. Such a point is illustrated in Figure 20.1.

Note that if the optimal solution involves using some of each factor, and if the isoquant is a nice smooth curve, then the cost-minimizing point will be characterized by a tangency condition: the slope of the isoquant must be equal to the slope of the isocost curve. Or, using the terminology of Chapter 18, the *technical rate of substitution must equal the factor price ratio*:

$$-\frac{MP_1(x_1^*, x_2^*)}{MP_2(x_1^*, x_2^*)} = \text{TRS}(x_1^*, x_2^*) = -\frac{w_1}{w_2}. \quad (20.1)$$



Cost minimization. The choice of factors that minimize production costs can be determined by finding the point on the isoquant that has the lowest associated isocost curve.

(If we have a boundary solution where one of the two factors isn't used, this tangency condition need not be met. Similarly, if the production function has "kinks," the tangency condition has no meaning. These exceptions are just like the situation with the consumer, so we won't emphasize these cases in this chapter.)

The algebra that lies behind equation (20.1) is not difficult. Consider any change in the pattern of production $(\Delta x_1, \Delta x_2)$ that keeps output constant. Such a change must satisfy

$$MP_1(x_1^*, x_2^*)\Delta x_1 + MP_2(x_1^*, x_2^*)\Delta x_2 = 0. \quad (20.2)$$

Note that Δx_1 and Δx_2 must be of opposite signs; if you increase the amount used of factor 1 you must decrease the amount used of factor 2 in order to keep output constant.

If we are at the cost minimum, then this change cannot lower costs, so we have

$$w_1\Delta x_1 + w_2\Delta x_2 \geq 0. \quad (20.3)$$

Now consider the change $(-\Delta x_1, -\Delta x_2)$. This also produces a constant level of output, and it too cannot lower costs. This implies that

$$-w_1\Delta x_1 - w_2\Delta x_2 \geq 0. \quad (20.4)$$

Putting expressions (20.3) and (20.4) together gives us

$$w_1\Delta x_1 + w_2\Delta x_2 = 0. \quad (20.5)$$

Solving equations (20.2) and (20.5) for $\Delta x_2/\Delta x_1$ gives

$$\frac{\Delta x_2}{\Delta x_1} = -\frac{w_1}{w_2} = -\frac{MP_1(x_1^*, x_2^*)}{MP_2(x_1^*, x_2^*)},$$

which is just the condition for cost minimization derived above by a geometric argument.

Note that Figure 20.1 bears a certain resemblance to the solution to the consumer-choice problem depicted earlier. Although the solutions look the same, they really aren't the same kind of problem. In the consumer problem, the straight line was the budget constraint, and the consumer moved along the budget constraint to find the most-preferred position. In the producer problem, the isoquant is the technological constraint and the producer moves along the isoquant to find the optimal position.

The choices of inputs that yield minimal costs for the firm will in general depend on the input prices and the level of output that the firm wants to produce, so we write these choices as $x_1(w_1, w_2, y)$ and $x_2(w_1, w_2, y)$. These are called the **conditional factor demand functions**, or **derived factor demands**. They measure the relationship between the prices and output and the optimal factor choice of the firm, *conditional* on the firm producing a given level of output, y .

Note carefully the difference between the *conditional* factor demands and the profit-maximizing factor demands discussed in the last chapter. The conditional factor demands give the cost-minimizing choices for a given *level* of output; the profit-maximizing factor demands give the profit-maximizing choices for a given *price* of output.

Conditional factor demands are usually not directly observed; they are a hypothetical construct. They answer the question of how much of each factor *would* the firm use if it wanted to produce a given level of output in the cheapest way. However, the conditional factor demands are useful as a way of separating the problem of determining the optimal level of output from the problem of determining the most cost-effective method of production.

EXAMPLE: Minimizing Costs for Specific Technologies

Suppose that we consider a technology where the factors are perfect complements, so that $f(x_1, x_2) = \min\{x_1, x_2\}$. Then if we want to produce y units of output, we clearly need y units of x_1 and y units of x_2 . Thus the minimal costs of production will be

$$c(w_1, w_2, y) = w_1y + w_2y = (w_1 + w_2)y.$$

What about the perfect substitutes technology, $f(x_1, x_2) = x_1 + x_2$? Since goods 1 and 2 are perfect substitutes in production it is clear that the firm will use whichever is cheaper. Thus the minimum cost of producing y units of output will be $w_1 y$ or $w_2 y$, whichever is less. In other words:

$$c(w_1, w_2, y) = \min\{w_1 y, w_2 y\} = \min\{w_1, w_2\} y.$$

Finally, we consider the Cobb-Douglas technology, which is described by the formula $f(x_1, x_2) = x_1^a x_2^b$. In this case we can use calculus techniques to show that the cost function will have the form

$$c(w_1, w_2, y) = K w_1^{\frac{a}{a+b}} w_2^{\frac{b}{a+b}} y^{\frac{1}{a+b}},$$

where K is a constant that depends on a and b . The details of the calculation are presented in the Appendix.

20.2 Revealed Cost Minimization

The assumption that the firm chooses factors to minimize the cost of producing output will have implications for how the observed choices change as factor prices change.

Suppose that we observe two sets of prices, (w_1^t, w_2^t) and (w_1^s, w_2^s) , and the associated choices of the firm, (x_1^t, x_2^t) and (x_1^s, x_2^s) . Suppose that each of these choices produces the same output level y . Then if each choice is a cost-minimizing choice at its associated prices, we must have

$$w_1^t x_1^t + w_2^t x_2^t \leq w_1^t x_1^s + w_2^t x_2^s$$

and

$$w_1^s x_1^s + w_2^s x_2^s \leq w_1^s x_1^t + w_2^s x_2^t.$$

If the firm is always choosing the cost-minimizing way to produce y units of output, then its choices at times t and s must satisfy these inequalities. We will refer to these inequalities as the **Weak Axiom of Cost Minimization (WACM)**.

Write the second equation as

$$-w_1^s x_1^t - w_2^s x_2^t \leq -w_1^s x_1^s - w_2^s x_2^s$$

and add it to the first equation to get

$$(w_1^t - w_1^s)x_1^t + (w_2^t - w_2^s)x_2^t \leq (w_1^t - w_1^s)x_1^s + (w_2^t - w_2^s)x_2^s,$$

which can be rearranged to give us

$$(w_1^t - w_1^s)(x_1^t - x_1^s) + (w_2^t - w_2^s)(x_2^t - x_2^s) \leq 0.$$

Using the delta notation to depict the *changes* in the factor demands and factor prices, we have

$$\Delta w_1 \Delta x_1 + \Delta w_2 \Delta x_2 \leq 0.$$

This equation follows solely from the assumption of cost-minimizing behavior. It implies restrictions on how the firm's behavior can change when input prices change and output remains constant.

For example, if the price of the first factor increases and the price of the second factor stays constant, then $\Delta w_2 = 0$, so the inequality becomes

$$\Delta w_1 \Delta x_1 \leq 0.$$

If the price of factor 1 increases, then this inequality implies that the demand for factor 1 must decrease; thus the conditional factor demand functions must slope down.

What can we say about how the minimal costs change as we change the parameters of the problem? It is easy to see that costs must increase if either factor price increases: if one good becomes more expensive and the other stays the same, the minimal costs cannot go down and in general will increase. Similarly, if the firm chooses to produce more output and factor prices remain constant, the firm's costs will have to increase.

20.3 Returns to Scale and the Cost Function

In Chapter 18 we discussed the idea of returns to scale for the production function. Recall that a technology is said to have increasing, decreasing, or constant returns to scale as $f(tx_1, tx_2)$ is greater, less than, or equal to $tf(x_1, x_2)$ for all $t > 1$. It turns out that there is a nice relation between the kind of returns to scale exhibited by the production function and the behavior of the cost function.

Suppose first that we have the natural case of constant returns to scale. Imagine that we have solved the cost-minimization problem to produce 1 unit of output, so that we know the **unit cost function**, $c(w_1, w_2, 1)$. Now what is the cheapest way to produce y units of output? Simple: we just use y times as much of every input as we were using to produce 1 unit of output. This would mean that the minimal cost to produce y units of output would just be $c(w_1, w_2, 1)y$. In the case of constant returns to scale, the cost function is linear in output.

What if we have increasing returns to scale? In this case it turns out that costs increase less than linearly in output. If the firm decides to produce twice as much output, it can do so at *less* than twice the cost, as long as the factor prices remain fixed. This is a natural implication of the idea of increasing returns to scale: if the firm doubles its inputs, it will more than

double its output. Thus if it wants to produce double the output, it will be able to do so by using less than twice as much of every input.

But using twice as much of every input will exactly double costs. So using less than twice as much of every input will make costs go up by less than twice as much: this is just saying that the cost function will increase less than linearly with respect to output.

Similarly, if the technology exhibits decreasing returns to scale, the cost function will increase more than linearly with respect to output. If output doubles, costs will more than double.

These facts can be expressed in terms of the behavior of the **average cost function**. The average cost function is simply the cost *per unit* to produce y units of output:

$$AC(y) = \frac{c(w_1, w_2, y)}{y}.$$

If the technology exhibits constant returns to scale, then we saw above that the cost function had the form $c(w_1, w_2, y) = c(w_1, w_2, 1)y$. This means that the average cost function will be

$$AC(w_1, w_2, y) = \frac{c(w_1, w_2, 1)y}{y} = c(w_1, w_2, 1).$$

That is, the cost per unit of output will be constant no matter what level of output the firm wants to produce.

If the technology exhibits increasing returns to scale, then the costs will increase less than linearly with respect to output, so the average costs will be declining in output: as output increases, the average costs of production will tend to fall.

Similarly, if the technology exhibits decreasing returns to scale, then average costs will rise as output increases.

As we saw earlier, a given technology can have *regions* of increasing, constant, or decreasing returns to scale—output can increase more rapidly, equally rapidly, or less rapidly than the scale of operation of the firm at different levels of production. Similarly, the cost function can increase less rapidly, equally rapidly, or more rapidly than output at different levels of production. This implies that the average cost function may decrease, remain constant, or increase over different levels of output. In the next chapter we will explore these possibilities in more detail.

From now on we will be most concerned with the behavior of the cost function with respect to the output variable. For the most part we will regard the factor prices as being fixed at some predetermined levels and only think of costs as depending on the output choice of the firm. Thus for the remainder of the book we will write the cost function as a function of output alone: $c(y)$.

20.4 Long-Run and Short-Run Costs

The cost function is defined as the minimum cost of achieving a given level of output. Often it is important to distinguish the minimum costs if the firm is allowed to adjust all of its factors of production from the minimum costs if the firm is only allowed to adjust some of its factors.

We have defined the short run to be a time period where some of the factors of production must be used in a fixed amount. In the long run, all factors are free to vary. The **short-run cost function** is defined as the minimum cost to produce a given level of output, only adjusting the variable factors of production. The **long-run cost function** gives the minimum cost of producing a given level of output, adjusting *all* of the factors of production.

Suppose that in the short run factor 2 is fixed at some predetermined level $\bar{x}_2$, but in the long run it is free to vary. Then the short-run cost function is defined by

$$c_s(y, \bar{x}_2) = \min_{x_1} w_1 x_1 + w_2 \bar{x}_2$$

$$\text{such that } f(x_1, \bar{x}_2) = y.$$

Note that in general the minimum cost to produce y units of output in the short run will depend on the amount and cost of the fixed factor that is available.

In the case of two factors, this minimization problem is easy to solve: we just find the smallest amount of x_1 such that $f(x_1, \bar{x}_2) = y$. However, if there are many factors of production that are variable in the short run the cost-minimization problem will involve more elaborate calculation.

The short-run factor demand function for factor 1 is the amount of factor 1 that minimizes costs. In general it will depend on the factor prices and on the levels of the fixed factors as well, so we write the short-run factor demands as

$$x_1 = x_1^s(w_1, w_2, \bar{x}_2, y)$$

$$x_2 = \bar{x}_2.$$

These equations just say, for example, that if the building size is fixed in the short run, then the number of workers that a firm wants to hire at any given set of prices and output choice will typically depend on the size of the building.

Note that by definition of the short-run cost function

$$c_s(y, \bar{x}_2) = w_1 x_1^s(w_1, w_2, \bar{x}_2, y) + w_2 \bar{x}_2.$$

This just says that the minimum cost of producing output y is the cost associated with using the cost-minimizing choice of inputs. This is true by definition but turns out to be useful nevertheless.

The long-run cost function in this example is defined by

$$c(y) = \min_{x_1, x_2} w_1 x_1 + w_2 x_2$$

$$\text{such that } f(x_1, x_2) = y.$$

Here both factors are free to vary. Long-run costs depend only on the level of output that the firm wants to produce along with factor prices. We write the long-run cost function as $c(y)$, and write the long-run factor demands as

$$\begin{aligned} x_1 &= x_1(w_1, w_2, y) \\ x_2 &= x_2(w_1, w_2, y). \end{aligned}$$

We can also write the long-run cost function as

$$c(y) = w_1 x_1(w_1, w_2, y) + w_2 x_2(w_1, w_2, y).$$

Just as before, this simply says that the minimum costs are the costs that the firm gets by using the cost-minimizing choice of factors.

There is an interesting relation between the short-run and the long-run cost functions that we will use in the next chapter. For simplicity, let us suppose that factor prices are fixed at some predetermined levels and write the long-run factor demands as

$$\begin{aligned} x_1 &= x_1(y) \\ x_2 &= x_2(y). \end{aligned}$$

Then the long-run cost function can also be written as

$$c(y) = c_s(y, x_2(y)).$$

To see why this is true, just think about what it means. The equation says that the minimum costs when all factors are variable is just the minimum cost when factor 2 is fixed *at the level that minimizes long-run costs*. It follows that the long-run demand for the variable factor—the cost-minimizing choice—is given by

$$x_1(w_1, w_2, y) = x_1^s(w_1, w_2, x_2(y), y).$$

This equation says that the cost-minimizing amount of the variable factor in the long run is that amount that the firm would choose in the short run—if it happened to have the long-run cost-minimizing amount of the fixed factor.

20.5 Fixed and Quasi-Fixed Costs

In Chapter 19 we made the distinction between fixed factors and quasi-fixed factors. Fixed factors are factors that must receive payment whether or not any output is produced. Quasi-fixed factors must be paid only if the firm decides to produce a positive amount of output.

It is natural to define fixed costs and quasi-fixed costs in a similar manner. **Fixed costs** are costs associated with the fixed factors: they are independent of the level of output, and, in particular, they must be paid whether or not the firm produces output. **Quasi-fixed costs** are costs that are also independent of the level of output, but only need to be paid if the firm produces a positive amount of output.

There are no fixed costs in the long run, by definition. However, there may easily be quasi-fixed costs in the long run. If it is necessary to spend a fixed amount of money before any output at all can be produced, then quasi-fixed costs will be present.

20.6 Sunk Costs

Sunk costs are another kind of fixed costs. The concept is best explained by example. Suppose that you have decided to lease an office for a year. The monthly rent that you have committed to pay is a fixed cost, since you are obligated to pay it regardless of the amount of output you produce. Now suppose that you decide to refurbish the office by painting it and buying furniture. The cost for paint is a fixed cost, but it is also a **sunk cost** since it is a payment that is made and cannot be recovered. The cost of buying the furniture, on the other hand, is not entirely sunk, since you can resell the furniture when you are done with it. It's only the *difference* between the cost of new and used furniture that is sunk.

To spell this out in more detail, suppose that you borrow \$20,000 at the beginning of the year at, say, 10 percent interest. You sign a lease to rent an office and pay \$12,000 in advance rent for next year. You spend \$6,000 on office furniture and \$2,000 to paint the office. At the end of the year you pay back the \$20,000 loan plus the \$2,000 interest payment and sell the used office furniture for \$5,000.

Your total sunk costs consist of the \$12,000 rent, the \$2,000 of interest, the \$2,000 of paint, but only \$1,000 for the furniture, since \$5,000 of the original furniture expenditure is recoverable.

The difference between sunk costs and recoverable costs can be quite significant. A \$100,000 expenditure to purchase five light trucks sounds like a lot of money, but if they can later be sold on the used truck market for \$80,000, the actual sunk cost is only \$20,000. A \$100,000 expenditure

on a custom-made press for stamping out gizmos that has a zero resale value is quite different; in this case the entire expenditure is sunk.

The best way to keep these issues straight is to make sure to treat all expenditures on a flow basis: how much does it cost to do business for a year? That way, one is less likely to forget the resale value of capital equipment and more likely to keep the distinction between sunk costs and recoverable costs clear.

Summary

1. The cost function, $c(w_1, w_2, y)$, measures the minimum costs of producing a given level of output at given factor prices.
2. Cost-minimizing behavior imposes observable restrictions on choices that firms make. In particular, conditional factor demand functions will be negatively sloped.
3. There is an intimate relationship between the returns to scale exhibited by the technology and the behavior of the cost function. *Increasing* returns to scale implies *decreasing* average cost, *decreasing* returns to scale implies *increasing* average cost, and *constant* returns to scale implies *constant* average cost.
4. Sunk costs are costs that are not recoverable.

REVIEW QUESTIONS

1. Prove that a profit-maximizing firm will always minimize costs.
2. If a firm is producing where $MP_1/w_1 > MP_2/w_2$, what can it do to reduce costs but maintain the same output?
3. Suppose that a cost-minimizing firm uses two inputs that are perfect substitutes. If the two inputs are priced the same, what do the conditional factor demands look like for the inputs?
4. The price of paper used by a cost-minimizing firm increases. The firm responds to this price change by changing its demand for certain inputs, but it keeps its output constant. What happens to the firm's use of paper?
5. If a firm uses n inputs ($n > 2$), what inequality does the theory of revealed cost minimization imply about changes in factor prices (Δw_i) and the changes in factor demands (Δx_i) for a given level of output?

APPENDIX

Let us study the cost-minimization problem posed in the text using the optimization techniques introduced in Chapter 5. The problem is a constrained-minimization problem of the form

$$\min_{x_1, x_2} w_1 x_1 + w_2 x_2$$

$$\text{such that } f(x_1, x_2) = y.$$

Recall that we had several techniques to solve this kind of problem. One way was to substitute the constraint into the objective function. This can still be used when we have a specific functional form for $f(x_1, x_2)$, but isn't much use in the general case.

The second method was the method of Lagrange multipliers and that works fine. To apply this method we set up the Lagrangian

$$L = w_1 x_1 + w_2 x_2 - \lambda(f(x_1, x_2) - y)$$

and differentiate with respect to x_1 , x_2 and λ . This gives us the first-order conditions:

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

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

$$f(x_1, x_2) - y = 0.$$

The last condition is simply the constraint. We can rearrange the first two equations and divide the first equation by the second equation to get

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

Note that this is the same first-order condition that we derived in the text: the technical rate of substitution must equal the factor price ratio.

Let's apply this method to the Cobb-Douglas production function:

$$f(x_1, x_2) = x_1^a x_2^b.$$

The cost-minimization problem is then

$$\min_{x_1, x_2} w_1 x_1 + w_2 x_2$$

$$\text{such that } x_1^a x_2^b = y.$$

Here we have a specific functional form, and we can solve it using either the substitution method or the Lagrangian method. The substitution method would involve first solving the constraint for x_2 as a function of x_1 :

$$x_2 = (y x_1^{-a})^{1/b}$$

and then substituting this into the objective function to get the unconstrained minimization problem

$$\min_{x_1} w_1 x_1 + w_2 (y x_1^{-a})^{1/b}.$$

We could now differentiate with respect to x_1 and set resulting derivative equal to zero, as usual. The resulting equation can be solved to get x_1 as a function of w_1 , w_2 , and y , to get the conditional factor demand for x_1 . This isn't hard to do, but the algebra is messy, so we won't write down the details.

We will, however, solve the Lagrangian problem. The three first-order conditions are

$$\begin{aligned} w_1 &= \lambda a x_1^{a-1} x_2^b \\ w_2 &= \lambda b x_1^a x_2^{b-1} \\ y &= x_1^a x_2^b. \end{aligned}$$

Multiply the first equation by x_1 and the second equation by x_2 to get

$$\begin{aligned} w_1 x_1 &= \lambda a x_1^a x_2^b = \lambda a y \\ w_2 x_2 &= \lambda b x_1^a x_2^b = \lambda b y, \end{aligned}$$

so that

$$x_1 = \lambda \frac{a y}{w_1} \quad (20.6)$$

$$x_2 = \lambda \frac{b y}{w_2}. \quad (20.7)$$

Now we use the third equation to solve for λ . Substituting the solutions for x_1 and x_2 into the third first-order condition, we have

$$\left(\frac{\lambda a y}{w_1} \right)^a \left(\frac{\lambda b y}{w_2} \right)^b = y.$$

We can solve this equation for λ to get the rather formidable expression

$$\lambda = (a^{-a} b^{-b} w_1^a w_2^b y^{1-a-b})^{\frac{1}{a+b}},$$

which, along with equations (20.6) and (20.7), gives us our final solutions for x_1 and x_2 . These factor demand functions will take the form

$$\begin{aligned} x_1(w_1, w_2, y) &= \left(\frac{a}{b} \right)^{\frac{b}{a+b}} w_1^{\frac{-b}{a+b}} w_2^{\frac{b}{a+b}} y^{\frac{1}{a+b}} \\ x_2(w_1, w_2, y) &= \left(\frac{a}{b} \right)^{-\frac{a}{a+b}} w_1^{\frac{a}{a+b}} w_2^{\frac{-a}{a+b}} y^{\frac{1}{a+b}}. \end{aligned}$$

The cost function can be found by writing down the costs when the firm makes the cost-minimizing choices. That is,

$$c(w_1, w_2, y) = w_1 x_1(w_1, w_2, y) + w_2 x_2(w_1, w_2, y).$$

Some tedious algebra shows that

$$c(w_1, w_2, y) = \left[\left(\frac{a}{b} \right)^{\frac{b}{a+b}} + \left(\frac{a}{b} \right)^{\frac{-a}{a+b}} \right] w_1^{\frac{a}{a+b}} w_2^{\frac{b}{a+b}} y^{\frac{1}{a+b}}.$$

(Don't worry, this formula won't be on the final exam. It is presented only to demonstrate how to get an explicit solution to the cost-minimization problem by applying the method of Lagrange multipliers.)

Note that costs will increase more than, equal to, or less than linearly with output as $a + b$ is less than, equal to, or greater than 1. This makes sense since the Cobb-Douglas technology exhibits decreasing, constant, or increasing returns to scale depending on the value of $a + b$.