These notes correspond to Chapter 2 of the text.

1 Optimization

A key concept in economics is that of optimization. It's a tool that can be used for many applications, but for now we will use it for profit maximization or cost minimization. Some people may say "Well, I'm close enough, why bother with optimization?" There's a story about John D. Rockefeller, who controlled a great deal of the oil refining capacity in the late 1800s and early 1900s. He observed that workers used 40 drops of solder to seal a barrel of oil – upon learning that 39 would work, he used that number (I read that somewhere years ago – I'm not sure of the source). Now, if Standard Oil was producing 10, 20, 100, or even 1000 barrels of oil a year the cost saving would be minimal. But if they are producing 1,000 barrels a day – well, that could be significant cost savings. While you may not go through an actual calculus problem when making all of your decisions, it is important to consider all the places in which improvement could occur.

The primary focus will be on maximizing the value of the firm. There are other goals a firm may have: capturing the largest market share, reaching a sales level goal, making enough profit to not be bothered by the boss, etc. However, through the use of optimization techniques the intuition behind economic decision-making can be developed. Maximizing the lifetime value of the firm gives us:

$$Value = \sum_{t=0}^{n} \frac{Profit_t}{(1+i)^t}$$

where $Profit_t$ is the firm's profit at time t, i is the interest rate which tells how much to discount profits in the future (because \$1000 today is not the same as receiving \$1000 tomorrow – one could always invest \$1000 today at the prevailing interest rate to earn some additional money tomorrow), and n is the number of periods (days, weeks, months, years) for which the firm expects to operate.[1] Profit is simply total revenue minus total cost. Determining all of these variables (profit, n, and i) is not always easy. Our focus for now will be on the profit term.

1.1 Revenue

Total revenue is a basic concept – it is simply the product of price and quantity,

$$TR = P * q$$

Typically we assume that price is some function of quantity, so that:

$$TR = P(q) * q.$$

I am deviating slightly from the book's notation – I am assuming we are working with an individual firm, and I tend to use q to denote a firm quantity, and Q to denote an industry quantity (which is just the sum of all the q's). The simplest functional form for $P(q)$ is linear:

$$P(q) = a - bq.$$

The intercept coefficient, $a > 0$, tells us the price at which individuals would stop purchasing the good (you probably don't want to price there).[2] The slope coefficient, $-b$, tells us how much the price would need to be decreased in order to sell more units of the good. While we are working with general terms now, we will

[1] Note that there is a subtle difference between my formula and the one in the book. My formula starts at $t = 0$ while the one in the book starts at $t = 1$. Typically the first period profit is undiscounted, which is what starting at $t = 0$ would give us. It's a minor detail.

[2] There is the apocryphal story of the consultant who charged \$1 million per hour. When a friend suggested the consultant would never have any clients because the price was too high, the consultant replied, "It only takes one." In this instance the consultant doesn't really want any clients – but will take any who are willing to pay an exorbitant amount.

discuss how to estimate these in a later class meeting. Suppose $a = 24$ and $b = 1$, so that $P(q) = 24 - q$. Now we have total revenue:

$$\begin{aligned} TR &= P(q) * q \\ TR &= (24 - q) * q \\ TR &= 24q - q^2 \end{aligned}$$

We can plot this function to see its shape:

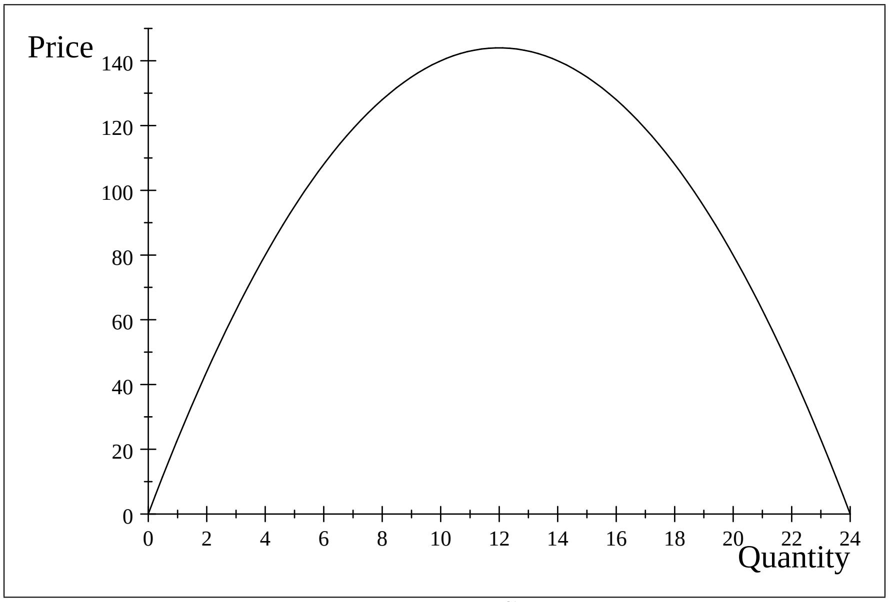

Total Revenue Curve

Another important concept is marginal revenue. The "marginal" concept will continue to be important throughout the course. Marginal simply means "additional," as in how much additional revenue occurs from an incremental change in some underlying part of the problem (in this case, quantity). Marginal revenue is simply the partial derivative of total revenue with respect to quantity:

$$MR = \frac{\partial TR}{\partial q} = 24 - 2q$$

As in solving any calculus problem, we can find the critical value (maximum or minimum) by taking the derivative of a function, setting it equal to zero, and solving. In this instance we see that revenue is maximized at $q = 12$. This result is consistent with Figure 1.1. Thus, total revenue is maximized at $q = 12$ and the maximized value of total revenue is \$144. If you were concerned with maximizing quantity sold (sales), which is sometimes a strategic goal of a firm, then you would produce 12 units in this case. Keep in mind that maximizing revenue is NOT the same as maximizing profit. We can see this pretty easily if it costs us \$500 to produce one unit of the good. If we are only generating \$144 in total revenue, and it costs \$500 to produce one unit of the good, then producing this item will not lead to a profit.

1.2 Cost

The other half of profit is cost of production. There are both long run and short run costs. Many people want to break these terms down into some time frame, such as "The short run is X time periods" or something along those lines. However, economists define these terms based upon whether or not firms can alter factors of production. Suppose a restaurant has leased a building for 2 years. For the restaurant owner, the building is very likely going to be a cost for the next two years, and it is a factor of production that cannot be altered. However, the owner can vary the number of workers as well as the amount of food ordered, and even how much electricity to use (by choosing store hours). If ANY factor of production cannot be altered, then economists say that the business is operating in the short run. If all factors of production can be altered, then economists say that a business is operating in the long run.

Another way to think about long run versus short run is to consider what types of costs make up total cost. There are two types of costs that determine total cost (TC) of production: fixed cost (FC) and variable cost (VC).

$$TC = FC + VC$$

Fixed costs are just those that do not vary with how much output is produced – using the restaurant example, the rent (or mortgage) payment for the building would be a fixed cost. It is a cost that must be paid every month, and the landlord will not care whether the restaurant has 1,000 customers per day or 10 customers per day because the contract calls for paying $\$X$ per month. Variable costs depend on how much output is produced – a restaurant serving 1,000 customers per day will likely use more staff and food than one serving 10 per day, so those costs will vary based on production. If the firm is operating in the long run, all costs will be variable; in the short run, some cost (but not all) will be fixed. For an example we will assume that $FC = 16$ and $VC = q^2$. Again, for now we are just using these numbers for purpose of an example – later we will discuss how to derive the cost function (the fixed cost part is fairly straightforward).

Like with revenue, an important concept with cost is marginal cost (MC). Again, it is simply the partial derivative of total cost with respect to quantity:

$$MC = \frac{\partial TC}{\partial q}$$

Another important concept when dealing with cost functions is average total cost (ATC).[3] Average total cost is simply total cost divided by the quantity produced,

$$ATC = \frac{TC}{q}$$

For our example, $MC = 2q$ and $ATC = q + \frac{16}{q}$. Thinking about the shape of the ATC curve, it will be U-shaped. The reason is that high average fixed costs will drive up cost for a low level of production, while high variable costs will drive up cost for high levels of production. The ATC and MC curves are plotted here:

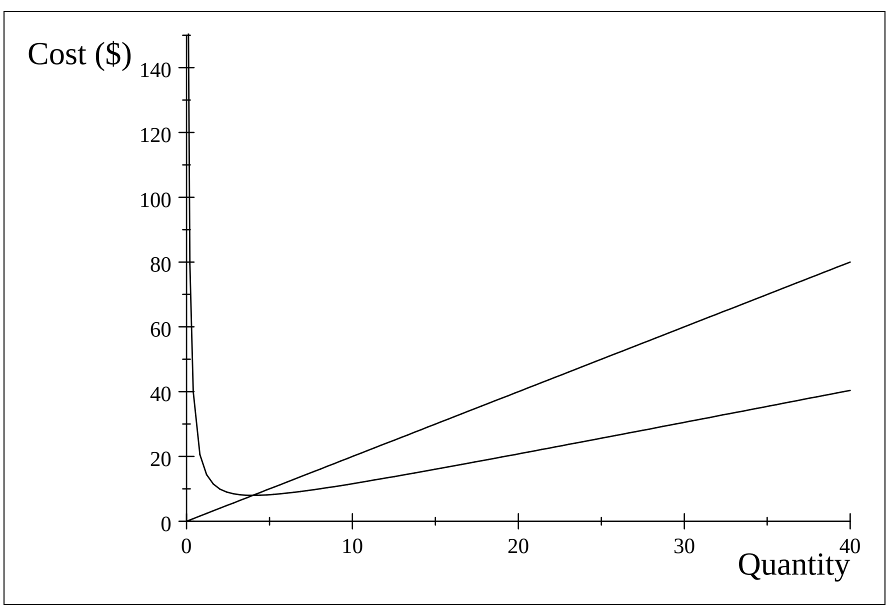

One relationship to note is that marginal cost will intersect average total cost at its minimum. The way to think about this is that if a marginal is greater than an average, then the average must be increasing; if it

[3] Chapter 2 of the text uses AC to denote "average cost." There is a point in time at which it will be important to differentiate between average total cost and average variable cost, so I will use ATC for "average total cost."

is less than the average, the average must be decreasing. As an example, consider a student who has a test score average of 95, and who scores a 90 on the next exam (so the "marginal" in this case is 90). Clearly the student's average will decrease. Sometimes I make my PhD students prove this result.

At times, a firm may wish to find the quantity at which it minimizes average total cost – the thinking is that the firm's cost per unit is the lowest at this point. This strategy could be useful if a firm does not have a very good grasp of the demand for its product, thus it may have a difficult time determining its total revenue and, as a result, its profit. We know that marginal cost intersects average total cost at its minimum, so to find the production level that minimizes average total cost simply set the two equal:

$$\begin{aligned} MC &= ATC \\ 2q &= q + \frac{16}{q} \\ 2q^2 &= q^2 + 16 \\ q^2 &= 16 \\ q &= 4 \end{aligned}$$

In this example the firm minimizes its average total cost by producing 4 units.

1.3 Profit

Combining revenue and cost gives us profit. We are focused here on profit maximization. The process is the same as maximizing revenue, only now cost is also taken into account. We will use Π to denote profit.

$$\Pi = TR - TC$$

To maximize profit, simply take the partial derivative of profit, set it equal to zero, and solve for the quantity. Alternatively, one can view profit maximization as setting marginal profit equal to zero. Generally, this result means that:

$$\begin{aligned} \frac{\partial \Pi}{\partial q} &= 0 \\ \frac{\partial \Pi}{\partial q} &= \frac{\partial TR}{\partial q} - \frac{\partial TC}{\partial q} \\ MR - MC &= 0 \\ MR &= MC \end{aligned}$$

To maximize profit, a firm should set marginal revenue equal to marginal cost. This result is the key to understanding that economic analysis is marginal analysis. For our example, we have:

$$\begin{aligned} \Pi &= TR - TC \\ \Pi &= \left(24q - q^2\right) - \left(q^2 + 16\right) \\ \Pi &= 24q - q^2 - q^2 - 16 \\ \Pi &= 24q - 2q^2 - 16 \\ \frac{\partial \Pi}{\partial q} &= 24 - 4q \\ 0 &= 24 - 4q \\ q &= 6 \end{aligned}$$

In this example the firm would maximize profit by producing 6 units. Note that the profit at 6 units is 56. Now consider the profit under our other two possibilities, revenue maximization and average total cost minimization. Under revenue maximization we had $q = 12$ and this leads to a "profit" of -16! Basically, we could have just produced stayed home, produced zero, and earned the same profit by just paying the fixed cost. Under average cost minimization we had $q = 4$ and this leads to a profit of 48. It's better than

-16, but still not as good as 56. And if the difference between 56 and 48 is not large enough to make a difference for you, imagine adding 5 zeros to the end of those numbers.

Note that at times you may not have revenue and cost functions to work with (this statement is about actually working with numbers from a business, not this class); however, you may have cost and revenue numbers associated with various levels of output. Also, maximization may tell you to produce 1,005,981$\frac{1}{8}$ units. It's unlikely that a consumer has any use for $\frac{1}{8}$ of a unit. In these instances simply make a spreadsheet and see where profit is the greatest (I know it's fairly obvious, but at times I have given problems to students and they have asked how to maximize profit without having the actual functions).

1.3.1 Incremental Concept

All of the marginal analysis that we have done concerns production. However, it fails to capture other decisions made by managers. We can extend the marginal concept by using the incremental concept to gauge other business decisions. The underlying idea is the same – consider the business decision being made, determine what the revenue from that decision is, determine what the cost is, and then determine the payoff from making that decision.

2 Constrained optimization

What we have done so far has considered unconstrained optimization. What may be useful at times is to consider constrained optimization. In this case, there is some constraint that the decision-maker must take into account. Eventually we will discuss constrained optimization for the consumer's problem (mainly to derive a fundamental result that pervades all of economics) as well as the firm's production problem. Here we will do two examples, one using a method that directly incorporates the constraint and one that does not.

Consider a firm that has the following total cost function:

$$TC = \$120X^2 + \$60Y^2 - \$20XY$$

where X and Y are quantities produced from two different production plants. The firm wants to produce 240 total units of the good, and wants to determine the minimum total cost at which it can produce these units. Because the firm wants to produce 240 units, the constraint is that $X + Y = 240$. The most direct method of determining the cost minimizing production plan is to simply rewrite the constraint, substitute it into the total cost function, and then use calculus to find the optimal solution.

$$\begin{aligned}
TC &= 120\left(240 - Y\right)^2 + 60Y^2 - 20\left(240 - Y\right)Y \\
TC &= 6912000 - 57600Y + 120Y^2 + 60Y^2 - 4800Y + 20Y^2 \\
TC &= 6912000 - 62400Y + 200Y^2 \\
\frac{\partial TC}{\partial Y} &= -62400 + 400Y \\
0 &= -62400 + 400Y \\
156 &= Y
\end{aligned}$$

Given $Y = 156$, we now know that $X = 84$. So the firm should produce 156 units at location Y and 84 at location X.

An alternative method, and one that we will use later, is the Lagrangian method. In this case, the constraint is not embedded into the objective function but added on separately. Formally, the Lagrangian is:

$$\mathcal{L}\left(X, Y, \lambda\right) = 120X^2 + 60Y^2 - 20XY + \lambda\left(240 - X - Y\right)$$

A few notes. The term λ is called the Lagrangian multiplier – this λ essentially represents the value of relaxing the constraint. The constraint, because this is a minimization problem, requires us to have zero on the left-hand side of the equation.[4] Once the Lagrangian function is set up the process is fairly straightforward – take derivatives with respect to X, Y, and λ; set the derivatives equal to zero; solve.

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial X} &= 240X - 20Y - \lambda \\
\frac{\partial \mathcal{L}}{\partial Y} &= 120Y - 20X - \lambda \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= 240 - X - Y
\end{aligned}$$

Setting all the first-order conditions equal to zero:

$$\begin{aligned}
0 &= 240X - 20Y - \lambda \\
0 &= 120Y - 20X - \lambda \\
0 &= 240 - X - Y
\end{aligned}$$

Now we have:

$$\begin{aligned}
240X - 20Y &= 120Y - 20X \\
260X &= 140Y \\
13X &= 7Y \\
X &= \frac{7}{13}Y
\end{aligned}$$

Using the constraint we have:

$$\begin{aligned}
240 - X - Y &= 0 \\
240 - \frac{7}{13}Y - Y &= 0 \\
3120 - 7Y - 13Y &= 0 \\
3120 &= 20Y \\
156 &= Y
\end{aligned}$$

To find X we have:

$$\begin{aligned}
X &= \frac{7}{13}Y \\
X &= 84
\end{aligned}$$

To find λ:

$$\begin{aligned}
120Y - 20X &= \lambda \\
120 * 156 - 20 * 84 &= \lambda \\
17040 &= \lambda
\end{aligned}$$

Thus we obtain the same answer through either method – the reason for introducing the Lagrangian method is so that we can use it later in the course to derive a few mathematical results, discuss them at an intuitive level, and then discuss how you can use that information at a practical level. The Lagrangian multiplier also gives us a rough idea of how much the objective function will change if we change the constraint by one unit (so it costs about $\$17,040$ to produce a unit of the good somewhere near 240 units).

You may wonder how minimizing the Lagrangian is the same as minimizing the TC function. Technically there is a condition that either λ or the constraint $(240 - X - Y)$ equals zero, but in our case the constraint will equal zero so minimizing the Lagrangian is the same as minimizing the TC function.

[4] There are technical details with this process that are beyond the scope of this course. I have notes for my PhD Microeconomic Theory course, BPHD 8100, that discuss these technical details if you are interested (the information is in the chapter 3 notes).