Elasticities and Costs

These notes continue with our treatment of introductory microeconomics. Much of economics considers how economic behavior responds to changes in prices, conditions, polices, etc. Elasticities are useful tools to quantify these effects. These notes also begin to consider why markets respond the way they do to changes in costs or utilities, and how they use inputs such as capital, labor etc.

Elasticity

We begin with "demand elasticity," formally the *price elasticity of demand.* Suppose that the price of a good or service changes. If this leads to a large change in the quantity demanded, the good is said to be elastic. This corresponds to a relatively flat demand curve. If a price change leads to small changes in the quantity demanded, however, then demand is said to be inelastic. This corresponds to a steep demand curve.

Recall the example of soda:

Table 1: Utility of Soda

Units	Marginal Utility	Total Utility
1	$100	$100
2	$90	$190
3	$80	$270
4	$70	$340
5	$60	$400
6	$50	$450
7	$40	$490
8	$30	$520
9	$20	$540
10	$10	$550

Formally, demand elasticity is the percentage change in the quantity demand divided by the percentage change in price. Lets consider the demand elasticity of soda when the price rises from \$50 to \$60. We can start by calculating the price change measured as a percentage:

$$\Delta P = \frac{new - old}{old} = \frac{60 - 50}{50} = 0.20 = 20\% \tag{1}$$

We then calculate percentage change in the quantity demanded:

$$\Delta Q = \frac{new - old}{old} = \frac{5 - 6}{6} = -.1666 = -16.7\% \tag{2}$$

The demand elasticity is then:

$$e^d = \frac{\Delta Q}{\Delta P} = \frac{-.1666}{.2} = -.833 \tag{3}$$

We can then say that a 1% increase in the price causes consumers to demand -0.83% *less* soda. The negative sign captures this inverse relationship where higher prices causes consumers to demand less of the good. This is related to demand being downward sloping and is common for most goods. Sometimes this negative sign is omitted when the elasticity is reported.

One question is why do we use percentage changes when calculating elasticities? We do so to make the measure unit-independent. Suppose that we instead said that a \$1 increase in price led to a one-unit decease in soda demand. This result then depends on whether a unit is a two-liter bottle, a ton, etc. Using percentages avoids this problem.

Now let us calculate the demand elasticity for a price change from \$90 to \$100:

$$\Delta P = \frac{new - old}{old} = \frac{100 - 90}{90} = 0.111 = 11.1\% \tag{4}$$

We then calculate percentage change in the quantity demanded:

$$\Delta Q = \frac{new - old}{old} = \frac{1 - 2}{2} = -.5 = -50\% \tag{5}$$

The demand elasticity is then:

$$e^d = \frac{\Delta Q}{\Delta P} = \frac{-.5}{.111} = -4.5 \tag{6}$$

Notice that different points on the demand curve can yield different demand elasticities. This is normal. Reported elasticites are usually calculated at equilibrium, where supply meets demand.

Here are some estimated elasticities for common goods:[1]

Table 2: Examples of Estimated Demand Elasticities

Good	e^d
Salt	-0.1
Coffee	-0.25
Gasoline (short run)	-0.2
Gasoline (long run)	-.0.7
Movies	-0.9
Private Education	-1.1
Restaurant Meals	-2.3
Tomatoes	-4.6

A number of factors can impact how a good's demand elasticity:

i. Is it a necessity or luxury? If a good is a necessity, then it is likely to be quite inelastic. One extreme example is a life-saving medicine. If its price changes, households are unlikely to change their demand. Another example is coffee (-0.25) which is so delicious that a world without it would make the survivors envy the dead.[2] Because it is easier to forgo luxuries, however, they tend to be more elastic.

ii. Does it have good substitutes. If good substitutes, exist, then a good tends to be more elastic. Consider meals at restaurants (-2.3) If prices rise, households can fairly easily substitute towards eating at home, hence it is an elastic good. There are, however, no good substitutes for coffee; tea is fucking gross. Hence, it is an inelastic good.

[1]Taken from: Anderson, P., McLellan, R., Overton, J. and G Wolfram. 1997. "Price Elasticity of Demand."

[2]New Bates students must try the Italian Bakery.

iii. Goods that cannot be postponed tend to be more inelastic. In the short-run, gasoline (-0.2) is an example and the immediate effect of price changes on demand tends to be small. Elasticity can, however, depend on the time horizon. With time, households can find alternate transportation methods or substitute towards a more fuel efficient car. Gasoline is thus less inelastic (-0.7) in the long-run.

Demand elasticity has important implications for markets. Suppose that a good becomes more expensive to produce (*i.e.* its marginal cost increases). Consider the impact on two different goods, one very elastic (the demand curve is flat) and the other being very inelastic (the demand curve is vertical):

Graph: Cost shock in two different markets:

Notice that when demand is very inelastic, the entire price change is passed on to consumers. Consider a life-saving medicine. because consumers will pay almost any price for it, if it becomes more expensive to produce, firms can simply pass the added cost on to consumers. If demand is very elastic, however, then they do not have the ability to do so.

Implications for tax policy:

This result has important implications for tax policy. One way that costs could increase is if the government requires that producers pay a tax for each unit sold. A government might do this

Table 3: Cost of Soda

Units	Marginal Cost	Total Cost
1	$20	$20
2	$30	$50
3	$40	$90
4	$50	$140
5	$60	$200
6	$70	$270
7	$80	$350
8	$90	$440
9	$100	$540
10	$110	$650

to raise revenue or to impact consumer behavior (*e.g.* a "sin tax" on tobacco or alcohol). If this taxed is imposed on inelastic goods, however, then the main impact is likely to be that consumers pay higher prices.

Going back to the example of soda, recall the cost schedule:

Suppose that the government decides to tax soda, either to raise revenue or to try and discourage soda consumption. One way that it could do so is to charge a $10 per unit tax on producers. This raises their marginal cost by $10, shifting supply up (or equivalently to the left). We can calculate the impact on price and quantity.

Graph: tax on soda producers:

An alternate approach is that the government could instead require that consumers pay the tax (like a sales tax). In this case, the tax reduces marginal utility by $10 because consumers must pay this amount for each unit that they buy. They still, for example, obtain the same utility from each delicious glass of Diet Dr. Pepper, which tastes more like Dr. Pepper. But that is then reduced by the amount of the tax.

Graph: tax on soda consumers:

Notice that the result is the same. This is generally true in our model of supply and demand. Who the tax is imposed on makes no difference to the market allocation. The elasticity of the good, however, does. It does not matter if we tax cigarette sellers or smokers.

An important application of this result is Social Security in the United States. This is a program where almost all workers pay a fraction of their income (known as payroll taxes) to the Federal government, In exchange, they receive a public pension upon retirement along with some other benefits. When this program was established in 1935, the government wanted the tax burden to be shared by workers and firms. It thus imposed a tax (currently 6.2% up until some income threshold) on *both* workers and firms. But our model suggests that this distinction should not effect the labor market and it would make no difference were the entire 12.4% paid directly by workers or firms. For this reason, economists view the Social Security tax as closer to 12.4%.

Other Elasticities

Most students find calculating elasticities to be scintillating and are thus ecstatic to find out that there are several others worth discussing:

1. The *income elasticity of demand* is the percentage change in the quantity demanded for a 1% increase in income. If this is negative, then the good is "inferior." Examples include goods that may have higher quality substitutes-Ramen Noodles, Natural Light, etc. A "normal" good has a positive elasticity- higher incomes lead to more demand. Sometimes, an income elasticity of demand above one is used as a definition of a luxury good.

An extremely inferior good can theoretically have such a negative income elasticity that its demand curve can become upward sloping. This is known as a *Giffen Good*. Suppose that a household is highly dependent on a single, inexpensive food for its sustenance. Now suppose that the price of that good rises. There are two effects. First, that good has become more expensive. The *substitution effect* suggests that households will demand less of the good because they substitute towards other foods. The *income effect*, however, suggests that households are poorer because of the price rise. If the good is inferior, the income effect will reduce increase their demand. If it is very inferior, then the overall effect may be that higher prices lead to higher demand, an upward sloping demand curves.

Giffen goods are a curiosity and hard to identify. One current candidate is rice in some Asian countries which is an inferior good and makes up a large part of peoples' diets, making the income effect important.[3]

2. The *cross elasticity of demand* refers to how a 1% change in the price of good Y affects the demand for good X. If this is positive, then the goods are substitutes. An increase in the price of Scotch Whisky likely leads to an increase in the demand for wine as households substitute their

[3] Lekhe1, F., Islam, L., Islam, S., and A. Akter. 2020. "Giffen Behavior for Rice Consumption in Rural Bangladesh." *International Journal of Applied Economics*, Vol. 11(1): 48-59.

choice of libations. A negative elasticity, however, implies that the goods are compliments. A 1% increase in the price of guns likely leads to reduced demand for ammunition.

3. The *price elasticity of supply* describes how firms' supply responds to changes in price. We go back to the soda example, again with the damn soda example. Suppose that the price of soda rises from $30 to $40.

$$\Delta P = \frac{new - old}{old} = \frac{40 - 30}{30} = \frac{1}{3} = 33.3\% \tag{7}$$

We then calculate percentage change in the quantity demanded:

$$\Delta Q = \frac{new - old}{old} = \frac{3 - 2}{2} = \frac{1}{2} = 50\% \tag{8}$$

The demand elasticity is then:

$$e^{a} = \frac{\Delta Q}{\Delta P} = \frac{.5}{.333} = 1.5 \tag{9}$$

Notice that the supply elasticity is positive, this reflects the upward sloping supply curve.

More generally, an elasticity just refers to how the demand or supply for a good or service changes (in percentage terms) to a 1% change to some other economic factor. Different elasticities show up all the time in economics.

Production Functions and Costs

We now dig deeper into the firm side of the market. We begin with the concept of a *production function*. This is the idea that there is a relationship between how much a firm produces and the inputs that it uses to produce its output. There are many inputs to the productive process:

1. Labor (L). Labor is the work used to produce output. It may be measured in hours or the number of workers. Firms and workers meet in labor markets where households supply labor and firms demand it. The price of labor is known as the wage.

Labor is often thought of as an input that firms are able to adjust in the short-run. We saw, for example, in the spring 2020 that firms were able to quickly adjust their workforce in response to restrictions and changes in demand caused by covid-19. This flexibility came with social costs, the U.S. employment rate rose to 15% in April 2020.

2. Capital (K). Capital refers to goods that are used to produce final goods rather than those which are consumed themselves. Common categories include business equipment and machinery, software, intellectual property, and commercial structures.

Capital itself is owned by households. This includes a private business where a household owns the associated capital. It also refers to stock ownership where households purchase ownership stakes in corporations. This is ownership of those firms' capital and entitles the shareholders to a share of the firms' profits. Income from capital is more concentrated in relatively high income households compared to income (wage) from labor.

We often view capital as less flexible in the short-run than labor. Like all of the assumptions made in this class, this not intended to be a perfect description of the world, rather it is offered as a reasonable approximation.

3. Land. Like capital, land is seen as being fairly fixed in the short-run.

4. Energy. Higher energy prices make production more costly. Like labor, energy is fairly flexible in the short-run.

5. Technology. Technological innovation drives costs lower in the long-run. From the perspective of most firms, however, it is out of their control.

6. Policies. Government polices including those related to taxation, regulation, trade, etc. can also impact the production process. These are mostly taken as given by firms.

7. Intermediate Goods. Many firms production processes depend on buying intermediate goods and using these to create final goods. A simple example is a homebuilder buying lumber to build a home. If intermediate goods are unavailable, this can impact the firm's production function.

For simplicity, we will consider a production function that depends on only labor and capital. This is a simplification. Other inputs are interesting and could be added to the model. For an intro class, however, we value simplicity:

$$output = Q(F, L) \tag{10}$$

We will assume that labor is flexible in the short-run. Firms can decide how much labor they want to employ at the market wage. This isn't always true, some countries restrict firms' ability to reduce their labor force and union contracts can also impose restrictions. We will also assume that capital is flexible, but only in the long-run. Once again, this is a simplification.

A firm's *revenue* is what it receives by selling its product. Revenue is just price multiplied by quantity:

$$revenue = PxQ(F, L) \tag{11}$$

A firm's accounting costs are the payments it must make to the owners of labor and capital. In this example, let us assume that the firm rents capital from households at a rental rate denoted r. Its costs are then wages: wxL and capital: rxK. The firm's accounting profit is then simply its revenue minus costs.

$$profit = PxQ(F, L) - wxL - rxK \tag{12}$$

Accounting profits (or losses) are what most people think of as profits and are what is reported when, for example, a corporation announces its profits (sometimes called "earnings"). But economists like to think of one other cost, the "opportunity cost" facing the firm. Opportunity cost refers what the cost of forgoing the next best option. Suppose, for example, that if this firm owner did not run their firm, they would choose to be employed as professional hand model earning $100,000 per period. The firm's *economic profit* simply includes this as a cost:

$$econprofit = PxQ(F, L) - wxL - rxK - oppcost \tag{13}$$

If the firm's economic profit is positive, then the firm is maxing out its income. If it is negative, then they could be making more income doing something else. It is possible that they may still decide to produce. The owner of the firm may get utility from operating their business as opposed to living in the high-flying and cut-throat world of hand modelling.

Cost curves will be important when we talk about market structure. Consider an example where a firm has a fixed capital stock equal to 10 units which each costs $10 to operate. The firm's capital bill is thus $100. The firm can produce different levels of output depending on how much labor it hires. Let us assume that each unit of labor costs $15/hour.

Table 4: Sample Cost Structure: $K = 10$

Units	Labor	Total Costs	Ave. Cost
1	1	115	115
2	3	130	65
3	6	175	58.3
4	10	235	58.8
5	15	310	62
6	21	415	69.2
7	27	510	72.9

It is useful to graph the average total cost curve. It has an inverse-U shape.

Graph: ATC

There are a few features that explain this shape. First, the firm's capital bill is inflexible in the short-run. As a result, producing low levels of out put yields a very high average cost because its total costs are dominated by capital. Producing additional output thus initially lowers average total costs. Second is the decreasing *marginal product of labor*. Notice that the first worker produces one unit of output. This is their marginal product. The first worker is very productive because they have the entire capital stock to work with. Now consider the next two workers, for a total of three. Each now has only one-third of the firm's capital stock to work with and they are less productive. It takes two *additional* workers to produce the output. Each worker's marginal product (the extra output they create) is just one-half of a unit.

As workers become very unproductive, it takes more of them to produce additional output. The average total cost curve thus eventually starts to increase once again.

Suppose instead that the had a lower capital stock so that $K = 5$ and the capital bill is thus 50. If capital and labor are substitutes, then the firm will have to hire more labor to produce any level of output:

Table 5: Sample Cost Structure: $K = 5$

Units	Labor	Total Costs	Ave. Cost
1	2	80	40
2	5	125	62.5
3	10	180	60
4	20	350	82.5
5	40	650	130
6	80	1250	208.3
7	200	3050	435.7

Note that actual cost functions are likely to be dependent on the details of the specific industry. Lower capital causes average total cost to bottom out for a lower level of output. In this case, that minimum is lower than the first example, although this is not true in general.

Now consider the long run. When capital is flexible too, the firm can choose any short-term average cost curve to be on. For any quantity, it will always choose to be on the short-run average total cost that has bottomed out for that quantity. The long run average total cost curve simply passes through all the short run average total cost curves.

Graph: Short and Long Run Average Total Costs

The long run ATC curve is also an inverse-U, but it is flatter than any of the short run ATC curves. We can divide the long run ATC curve into three regions:

1. The downward sloping part of the curve reflects *economies of scale*. If an industry is on this part of the curve, then a smaller number of firms, each producing more output, will result in lower average costs. We might expect the industry to consolidate into a smaller number of firms. In this example, there may be advantages to size. A firm like Walmart, that already has an efficient distribution system (a ind of capital), may be able to sell books more cheaply than a specialized bookstore.

2. A flat part in the middle (which may not exist) is known as *constant returns to scale*. Here, bigger firms do not produce at a lower or higher cost.

3 An upward sloping part reflects *diseconomies of scale*. Here, bigger firms are less efficient at producing. We might expect large firms to break up into smaller ones.

Different industries may have very different cost structures. These difference help to explain why some industries are dominated by a few large firms while others are made up of many smaller ones.