

Some Foundational Microeconomics

These notes introduce some fundamental microeconomic theory. Because macroeconomics is the aggregation of microeconomics, this material is an important foundation for the macroeconomic content that we will soon focus on. Modern macroeconomics is based on *microeconomic foundations*. This implies that macroeconomic outcomes should be directly connected to the microeconomic behavior of the firms and households which make up the economy. Consider some examples.

- i. Later, we will see that policy makers often cut taxes in order to boost aggregate economic performance. Microeconomic foundations mean that we should be able to connect an increase in households' after-tax income to how they make choices, given a budget constraint, which then may affect employment, consumption, etc.
- ii. From 2007-2016, the Federal Reserve purchased nearly \$4 trillion in assets from financial institutions. It paid for these assets with newly created money. How did this switch affect the behavior of these profit maximizing firms? Did it provide them an incentive to lend more? Did it induce them to take more or less risk?

Supply and Demand

We begin with the most important model in all of economics, supply and demand in a perfectly competitive market. A market is a collection of agents, firms and households in this case (policy makers are an important part in many markets) whose decisions cause the allocation of a certain good or service. Perfect competition means that any single household or firm is too small to affect the market price. We will use the hokey example of the market for soda.

We begin with demand. For most goods and services, households are the demanders. In the labor market, however, households are the suppliers of labor, and firms are the demanders. The demand curve is the relationship among different prices, and the amount of the product that demanders are willing to purchase at each price. We now explicitly make some assumptions about these households:

1. They are rational and make choices to maximize their utility, the total benefit to them provided by the soda. Furthermore, they understand the correct relationship between soda consumption and utility. This latter assumption is not always valid. Economists often study the impact of incorrect or incomplete information.
2. There are a large number of households and each of them is only a small part of the market. Once again, this assumption is not always true. For some goods, there may be only a few buyers. If there is only one, it is known as *monopsony*.
3. Because they are each a small part of the market, households have no ability to affect the price. This is known as *price taking*. By assumption, an individual consumer cannot effectively demand that soda companies lower their prices. Again, this assumption is not always applicable. If there are only a few demanders of the good, then they may be able to influence its price. This is known as *market power*.
4. They are able to observe all market conditions (price, utility, cost, etc.)

These assumptions offer a starting point for analyzing demand. They are not offered as a perfect representation of most markets. Issues like incomplete information, market power, etc. are important features of many real world markets. Studying these issues is, however, best done by building off of the basic model presented here.

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

Table 1 assumes a specific utility from soda:

We will now derive the demand curve. To graph it, we choose to put the price on the vertical axis and the quantity demanded on the horizontal.

Graph: Supply and Demand

As we derive the demand curve for soda, notice that it is marginal (additional) utility, not total utility, that matters at each step in the process. This is true for most examples in economics.

Suppose that the price per unit of soda is \$10. Households will choose to buy a first unit of soda because the marginal benefit is greater than the price. Likewise, households will buy a second unit because the \$90 marginal benefit (again, marginal matters, not total) is greater than the price. This remain true thorough the ninth unit. The tenth unit provides just as much marginal utility as it costs. The household is therefore indifferent to buying it. To make things easy, we will assume that households do buy the soda. Households thus demand 10 units of soda at a price of \$10. We can make this the first point on the demand curve.

Now suppose that the price is \$20. Repeating the previous exercise, households will purchase soda until its marginal benefit equals the price which occurs at 9 units. If the price is \$30, marginal utility equals the price at 8 units. As we continue this process, we see a pattern where marginal utility is demand.

In this example, the demand curve is downward sloping: households demand more soda as the price falls. While we believe that this is usually true, it results from a previously unstated assumption from the previous table. The demand curve is downward sloping because we assumed that marginal utility is decreasing. As households consume more of a good, they get more total utility, but less marginal utility. In this example, the first unit of soda is more valuable than the second which is more valuable than the third, etc. While this usually a good assumption, there are cases where it does not apply. Demand curves are thus not universally downward sloping.

Supply

We now turn our attention to the firms who supply the soda. We make the following assumptions:

1. They are rational and make choices to maximize their profits. Furthermore, they know how much of their product they will sell at each price. Once again, this assumption of complete information is not always applicable to markets.

2. There are a large number of firms and each of them is only a small part of the market. Each produces a homogenous good, that is each firm' product is indistinguishable.

These assumptions are even stronger than those made about demand. Many goods are provided by only a few suppliers. When there is just one, it is known as a *monopoly*. Likewise, there are very few markets where suppliers produce genuinely identical goods. This set of assumptions is thus presented as a baseline and microeconomics spends a great deal of time considering the implications of alternate assumptions.

3. Because they are each a small part of the market, they have no ability to affect the price and they are price takers, like households. If a firm tries to charge a price above the market price, they will lose all their business. If they charge a price below market, they will capture the entire market.

4. They are able to observe all market conditions (price, utility, costs, etc.)

Firms must, of course, pay production costs (labor, capital, etc.) to produce its product. Table 2 assumes a specific cost function.

As with utility, marginal, not total matters when solving a firm's problem.

Suppose that the market price is \$30. Firms will choose to produce a first unit because they are able to sell it for more than the \$20 it costs to produce it. The second unit costs just as much to produce, \$30, as it costs to produce. Firms are indifferent, lets assume that they produce when indifferent. So supply at \$30 is 2 units of soda.

Now suppose the price is \$50. Units 1-3 are all cheaper to produce than what they sell for and firms are indifferent as to producing the fourth unit. Supply is 4 units.

Note that supply is marginal cost.

This supply curve is upward sloping because we assumed that marginal cost is increasing. This is usually, but not always, a good assumption. Suppose that a firm has a given amount of physical

Table 2: 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

capital (machinery used in production). Perhaps it takes one worker, along with this capital, to produce one unit. Now suppose that the firm hires a second worker. The two workers must now share the capital, each has half as much to work with as when there was only one worker. They will be less productive. The firm will therefore have to hire more than one additional worker to produce a second unit of output.

There are exceptions. There are industries where marginal cost is decreasing or constant. Upward sloping supply curves are therefore not universal. They are an assumption that fits most, but not all, real world markets. We will leave discussion of downward sloping supply curves (and upward sloping demand curves) for a microeconomics class.

Equilibrium

We now introduce the concept of an economic equilibrium. This is a stable point where agents (firms and households) have no incentive to change their behavior. Here, this occurs where supply equals demand, on the graph where the two curves cross.

We have thus solved for the equilibrium quantity and price of soda. The quantity is 5 units and

the price is \$60.

Here we have solved the model using graphs. There is a set of *exogenous* factors, those that are determined outside the model. This includes the demand and supply curves. Put another way, we are not saying anything about why demand is what it is, or why costs are what they are. They are taken as given. This is in contrast to *endogenous* variables, which are those that are solved for by the model. Here, they include price and quantity. Solving the model means taking the exogenous factors and figuring out what the endogenous variables equal.

Using the Model

Having solved the model, we can now make predictions. Let's start by assuming that something happens that causes households to demand more soda at any price. Instead of demanding 4 units of soda at \$50, perhaps they now demand 6 unit. We can show this on a graph as a demand curving shifting upward.

Graph: Supply and Demand

Now, the equilibrium quantity has increased, as has the equilibrium price. An increase in demand has these effects when demand is downward sloping and supply is upward sloping. There are several factors that might cause such an increase in demand:

1. It could be the result of changing tastes. Suppose that some new health benefit of soda was discovered. People would then demand more of it.
2. Two goods are compliments if the use of one increases the utility that one obtains from the other. Suppose that the quantity of pizza increased. This would increase the marginal utility of soda because the two are compliments. Demand for soda would increase.
3. Two goods are substitutes if the use of one decreases the utility that one obtains from the other. Suppose that the quantity of beer decreased. People would wish to substitute toward soda, increasing its demand.
4. Suppose that incomes increase. For some goods, such as sports cars, we would expect demand to increase. We call these *normal goods*, and they include most luxury good and many other goods and services For other goods, such as Ramen Noodles and Natural Light Beer, demand typically decreases as incomes rise. We call these *inferior goods*. I have no idea which type soda is.

The previous examples discuss an increase in demand. What about supply? The language here confuses many students. Because the supply curve is in the same place, we do not say that supply has changed. We have, however, moved along the same demand curve. This is referred to as a “change in the quantity supplied” rather than a change in supply itself.

The model is also symmetric. If soda is instead found to have previously unknown links to serious illnesses, then demand will decrease, as will the equilibrium price and quantity.

Change in Supply

Now suppose that it becomes less expensive for a firm to supply soda. If the marginal cost of producing any amount of soda is less, then the supply curve moves down (also to the right). Here the equilibrium quantity increases, but the equilibrium price decreases.

Graph: Supply and Demand

An increase in supply comes from a reduction in marginal cost. There are multiple potential causes.

1. The price of an input may have declined. Perhaps the wages that firms pay to labor has decreased, or the rental rate of capital or land have declined, or energy prices may have declined. Or, sugar, an input to soda, could have become cheaper.
2. Something could have happened that allows firms to use fewer inputs in order to produce output. This often take the form of a technological innovation.

An Example

Consider the market for sweatpants. Suppose that, at the same time, fabric prices decline and the average household's income increases. What happens?

Graph: Supply and Demand

In this previous example, there is an increase in supply and an increase in demand if sweatpants are a normal good. If this is the case, quantity increases, but the effect on price may be positive, zero, or negative. We cannot say for sure, the answer is ambiguous, we need more information. Theoretical models yield ambiguous result all the time, it isn't a bad thing. I have found that students desperately want to say that the effect on price is zero, but there is no reason to think the effect of increased supply on price perfectly offsets that of increased demand.

If sweatpants are an inferior good, however, then demand decreases. Now price falls for sure and the effect on quantity is ambiguous.

It does not seem obvious whether sweatpants are an inferior or normal good. This is a case where we may turn to the data, in the form of empirical microeconomics, to estimate the effect of income on sweatpants demand. Empirical work is a good way of addressing theoretical ambiguity.

Are Free Markets Efficient?

There is a somewhat common misconception that economics argues that free markets always yield good outcomes. The truth is that economics spends a great deal of time discussing *market failures*, where free markets can lead to bad outcomes. How often this occurs is ultimately an empirical question.

We now introduce the notion of Pareto Efficiency to start to address these issues. Pareto efficiency means that it is not possible to make anybody better off without making at least one person worse off. Suppose, as a simple example, that \$100 may be distributed between me and you. Which of the following distributions are Pareto Efficient?

- a. \$50 for me, \$50 for you
- b. \$100 for me, \$0 for you
- c. \$25 for me, \$25 for you.

a. is Pareto Efficient. The only way that I can increase my share is to reduce yours, or vice-versa. By the same reasoning, b. is Pareto efficient as well. There can thus be multiple Pareto efficient allocations and the definition of Pareto efficiency is unrelated to equity. Because c. does not exhaust all of the \$100, I can, however, give some of the undistributed \$50 to either me or you, making at least one of us better off without doing the other any harm.

Because there are typically many Pareto efficient outcomes, an allocation being Pareto efficient does not imply that it is optimal. But an allocation not being Pareto optimal typically does suggest that it is sub-optimal, why wouldn't we want to make someone better if it does nobody else any harm?

To compare different allocations, we could employ a *social welfare function* which is a way to combine individual's welfare into a measure of societal well-being.

Question for Next Class: Compare two social welfare functions: First, $\text{Welfare} = Y_{ou}^2 + M_{e}^2$. Second, $\text{Welfare} = \sqrt{Y_{ou}} + \sqrt{M_{e}}$ Which allocation from a – c is preferred under each welfare function? Which makes more sense to you?

Our simple graph of supply and demand allows us to examine the Pareto efficiency of the market for soda. The demand curve is the marginal utility from each unit. The difference between the demand curve and the price is thus the surplus utility obtained from purchasing that unit.

Adding up all such differences, we obtain an area known as *consumer surplus*, this is the benefit to households from being able to participate in the market.

The supply curve is the cost needed to produce each additional unit. The difference between the price and cost is the surplus obtained by firms. Adding up, we obtain an area known as *producer surplus*. Producer surplus minus fixed costs are firm profits.

There is no way in this model to increase consumer surplus without decreasing producer surplus or vice-versa. The equilibrium is thus Pareto efficient. In this case, no additional governmental policies are needed to ensure optimality. The Scottish economist Adam Smith called this the “invisible hand” of the free market.

This type of result has sometimes caused some people to think that most economists have an unfettered faith in free markets and that we rarely advocate any type of policy intervention. The reality is that economists agree that free markets are efficient if the assumptions that we have made throughout developing the model are valid. There are numerous case where they are not and where, as a result, the equilibrium is not Pareto efficient. We conclude by considering a non-exhaustive list:

1. Market Power. We assumed that firms and households are too small to have any influence over the price. But what if there are barriers to entry, things that prevent additional firms or households from entering the market? In this case, either demand or supply could consist of only a small number of agents. If so, then these agent might be able to influence the price.

Agents who have market power exploit it by restricting quantity in order to affect the price in their favor. Market power can be had by either the demand side or supply side. On the supply side, examples include a monopolist who is the only supplier in a market, or a labor union which exploits market power in the labor market. Graphically, we see a higher price and lower quantity.¹

¹How exactly to best exploit market power to is left for a microeconomics course.

Graph: Supply and Demand

When there is market power on the demand side, the demander restricts quantity in order to reduce the price.

Graph: Supply and Demand

In both cases, notice that the combined areas of consumer and producer surplus are reduced compared to the case of perfect competition. This is called *deadweight loss*. This leaves room for policy to possibly improve welfare. Suppose, for example, that there is a law that can effectively eliminate monopolies. This would increase consumer surplus by more than it would decrease pro-

ducer surplus. We could then use a lump sum tax on households to compensate the former monopolist to make both parties better off than under the monopoly representing a Pareto improvement. (Although it is not obvious that such a tax is desirable). This is the basic motivation behind antitrust laws which are generally designed to prevent agents from obtaining too much market power. How much anti-trust regulation is best is partly an empirical question.


2. Lack of property rights. The model has implicitly (*i.e.* we did not list it) assumed that households can confidently use the goods and services they buy and that suppliers can confidently keep the profits that they make. Suppose instead that there was a high risk that either could be expropriated. In this case, the market might fail because agents would have to account for this risk. For this reason, most economists view protecting property rights as a fundamental purpose of government.

3. Externalities. We have assumed that the private costs and benefit (those experienced by the households purchasing the product and the firms producing it), equal the total costs and benefits. Suppose, instead, that there is an externality where the purchase of the product affects other people in the model, but that the buyer and seller do not care about this.

The classic example of an (negative) externality is pollution. Suppose that a paper mill is free to dump chemicals into public rivers. The private costs are those borne by the firm and determine its production decisions. The total costs, however, include the costs to clean up or tolerate yellow, foamy rivers that smell like shit. Firm will therefore produce too much pollution because they do not have to pay all the costs.

Graph: Supply and Demand

Figure 1: The Androscoggin River in Lewiston, 1973: Not Pareto Efficient



The solution to this type of externality is to either tax firms to align public and private costs or to restrict their production through a quota. This was the motivation behind the Clean Water Act which most agree was good policy.

Externalities may also be positive where production creates a non-private benefit. One example is production spillovers. Suppose that one firm's production produces innovation that benefits other producers in the industry. In this case, the total benefit exceeds the private benefit so that firms produce too little. Here, a common solution is a subsidy (the opposite of a tax).

Graph: Supply and Demand

4. Incomplete Information. We assumed that both suppliers and demander had good information about the market. But if some agents have incomplete or bad information, then inefficiency may result.

Consider the market for used cars. Suppose that suppliers are able to observe the true quality of cars but that potential buyers are not able to distinguish between good cars and crappy ones (lemons). Now suppose that a seller offers a buyer a car for \$1000. If they cannot observe the true value of the car, they might extract from the seller's offer that the car is worth less than \$1000. If the seller offers \$500, they might infer that it is worth less than that. In extreme cases, the market breaks down altogether. This is known as *moral hazard*.

Another example is the market for car insurance. Insurance exists to spread risk. Everyone in the pool pays premiums and these premiums are used to pay claims made by the unlucky few who are in accidents.

Suppose there are 10 drivers who each have a $\frac{1}{10}$ chance of getting into an accident that will result in \$1000 in damages. If each pays \$100 in premiums then we expect that the \$1000 will be able to pay for the damages.

But now suppose that there is one really awful driver who has a $\frac{1}{2}$ chance of getting into an accident. To cover the potential damages, the premiums will have to be higher and it is possible that the other drivers will choose not to pool with the bad driver. The market for insurance might fail. This simplified example motivates why most states mandate automotive insurance.

5. Public Goods. How much would you pay for national defense if you could choose? Most people would answer 0 because anything you contribute is unlikely to make a discernable difference. The problem is that without intervention, everyone would make a similar choice, and there would be too little national defense. This is the basic public goods problem. In a free market, people choose to free ride.

A public good is one that exhibits non-rivalry of consumption: my consumption of it does not hinder your ability to consume it, and which is non-excludable, we cannot easily prevent someone

from using it. Many goods meet these criteria imperfectly and there is often a grey area. But most economists consider other examples of public goods to include roads, infrastructure, clean air, etc.

The most common solution to the public goods problem is for a government to provide a public good. Of course, there is often disagreement over the right amount of a public good that should be provided.

6. Distortionary taxes. These are taxes which change people's behavior. Go back to the example of soda. Recently, some municipalities have enacted a tax paid by the purchasers of each unit of soda. This tax creates a gap between the amount paid by households and the amount received by suppliers from selling soda.

Graph: Supply and Demand

Instead of taxing households for buying soda, the government could have taxed the firms which are selling it. This would create a gap between the marginal cost paid by suppliers and the amount needed to actually produce the product.

Graph: Supply and Demand

Notice that it doesn't matter whether the tax is paid by the demander or the supplier. The price paid by the household is the same in both cases and the quantity is reduced by the same amount in both cases. Also note that the combined producer and consumer surplus region are reduced. The best known distortionary taxes are income taxes that collect a fraction of a worker's earnings. This tax changes their after tax wage rate and thus affects their behavior.

The theory discussed above is not controversial. In practice, honest disagreements among economists often center on whether market failures are small or large, or if large, whether government intervention is likely to significantly improve the situation.