Lecture Notes 1

Tom Flesher

January 24, 2015

Chapter 1

Supply and Demand

No two concepts are as fundamental to economics as supply and demand. Arguably, all of economics comes down to the choices exemplified by these two simple functions. As a result, even project managers and engineers taking a required course will benefit from an overview of these important concepts.

Economics is the study of choices made by individuals. In the economic language, we call these individuals "agents" and treat them as though they are solving a simple maximization problem. The two types of agents are "households" and "firms." Households and firms interact in two markets, called *the market for goods and services* and *the factors market*. This course focuses on choices made by firms in the market for goods and services.

1.1 Demand

In the factors market, households sell their factors of production - specifically, labor - to firms. The firms pay a price, called a wage, to the worker in exchange for this labor and use the labor as an input in the production of goods and services. The wage is determined by the workers' productivity levels.

In the market for goods and services, households are buyers of the same goods that their neighbors worked to produce. Every worker's income results from a household spending money to purchase goods from a firm, which then uses that revenue to pay workers' wages. In order to stay profitable, a firm must sell goods.

In economics, every agent faces an optimization problem, generally bounded by certain constraints.

1.1.1 The Household's Maximization Problem

Households seek to get the most bang for their buck - in other words, to gain the most happiness for each dollar spent. Consequently, a household seeks to buy goods that are useful or valuable to that household. When faced with a choice of where to obtain those goods, it seeks to buy them as cheaply as possible. Given an unlimited income, a household could sustain itself quite easily. Economists represent happiness using a fanciful concept called *utility.* Utility is designed to represent happiness in a way that can be quickly converted into dollars, so you may think of a one-unit increase in utility as equivalent to the happiness increase from gaining \$1.

The household's objective function is the function which they seek to maximize. This is called the **utility function** and can take a variety of forms. The simplest, called *linear utility*, simply assigns a utility value to each good and adds up the utility of all goods purchased. An example of a linear utility function is:

$$u(\text{apples, bananas} = \alpha \times \text{apples} + \beta \times \text{bananas}$$

This implies that consuming each apple is worth \$$\alpha$ to the household and consuming each banana is worth \$$\beta$.

Suppose that $\alpha = 2$ and $\beta = 1$. If this were the case, the household would almost certainly spend its entire income on apples, because each apple is worth twice as much utility as each banana. However, that is not a realistic set of assumptions; though some of us have wished we could spend an entire week's income on beer,[1] we typically have preferences that guide us toward making more reasonable and varied purchases. A more common type of utility function in this case is called *concave utility*, which allows for *diminishing returns* to consumption of a good. A function that has a positive first derivative and a negative second derivative is concave, by definition. An example of a concave utility function is:

$$u(\text{apples, bananas} = \alpha \times \sqrt{\text{apples}} + \beta \times \sqrt{\text{bananas}}$$

This means that even if $\alpha = 2$ and $\beta = 1$, there will come a point at which the gain in utility per dollar spent will be higher for purchasing a banana than for an apple. Suppose that apples and bananas each cost \$1. The gain in utility for purchasing an additional good is called *marginal utility* of a good and can be determined as $\frac{\Delta u(\cdot)}{\Delta Q}$, where $u(\cdot)$ is the utility function and Q is the quantity of the good. Taking a limit, we can also represent it as $\frac{\partial u(\cdot)}{\partial Q}$, or the partial derivative of the utility function with respect to the good in question. For now, however, let's confine ourselves to the discrete case, since looking at unit changes in quantity more closely models human behavior. If we assume $\alpha = 2$ and $\beta = 1$, then buying one apple generates \$2 worth of utility. The consumer is then faced with a choice: buy another apple for total utility of $2 \times sqrt(2) = 2.82$, or buy a banana for total utility of 3 (2 from the apple and 1 from the banana)? Because the marginal utility of the first banana is greater than the marginal utility for the second apple, it makes sense to purchase the banana before purchasing a second apple.

Households do not, of course, have an unlimited income, and so the choice to buy a banana before a second apple may become the choice to purchase a banana

[1] Or, in some cases, heavily-peated Islay scotch

instead of a second apple. Households are constrained by their income, usually represented as a wage. (More complicated models might involve investment or rental income, but wages will suffice for now.) A household's primary constraint is that income must be greater than or equal to spending:

$$\text{Wage} \times \text{Hours Worked} \geq \text{Household Spending}$$

Household spending can be broken down by good, allowing for each good to have its own price. However, doing so serves no purpose at the moment. Though households can overcome this fundamental constraint through borrowing, such an optimization is a two-period problem. Multi-period problems are discussed in Part 2, The Time Value Of Money.

Thanks to the way we defined utility - a unit of utility is equal to $1 worth of happiness - we now have enough information to describe demand.

1.1.2 The Demand Correspondence

Demand represents a buyer's marginal utility for purchasing another unit of the good, subject to the buyer's ability. In other words, demand is the buyer's marginal willingness to pay. Because demand depends on a buyer's willingness to pay, it will take into account both the buyer's marginal utility and his ability to pay.

A **demand correspondence** is a formal mathematical relation between price and quantity demanded. For a single good, it usually (in the linear case) takes the form

$$Q^D = a - b \times P$$

where a represents the quantity that a buyer would consume if the price were $0 and b representing how readily buyers are dissuaded from purchasing by price. The a parameter can thus be thought of as representing the marginal utility to be gained by consumption of a good and the b parameter represents the income constraint. Graphing the demand correspondence creates a **demand curve**.

Shifts in the demand curve can then, roughly, be sorted into shifts to the a parameter (which represent changes in the baseline demand for the good) and shifts to the b parameter (which represent changes in the consumer's ability to afford a good).

Complementary goods are goods which people prefer to use together, such as baseballs and baseball bats, coffee and cream, and

One common error made by students new to economics is treating "a shift in demand" as equivalent to "a change in the quantity that will be purchased." Changing the price can change the *quantity demanded*, abbreviated Q^D. This doesn't require any change to a consumer's preferences. However, the consumer's preferences or the consumer's income can change. When they do, the consumer's marginal utility or ability to pay - the a or b parameter - will change with them. When these notes say "demand," they refer to the demand correspondence and not the quantity demanded.

razors and shaving cream. When the price of a complementary good falls, it becomes easier to consume the complement, and the user then wishes to consume more of the good in question. If the price of cream decreases, I may choose to drink more coffee. If, on the other hand, the price of cream increases, I have a choice - I may choose to continue drinking my present amount of coffee, but the price of cream will erode my income, leaving me less able to drink all the coffee I desire. Alternatively, I may choose to drink coffee with less cream or no cream at all, making the experience less enjoyable and leaving my desire for coffee lower than before. Either way, I choose to consume less, showing a decrease in my demand for coffee.

Substitute goods are goods which people prefer to use one or the other of, such as coffee and tea, pens and pencils, or staples and paperclips. When the price of a substitute good falls, it becomes more difficult to consume the good in question. If the price of coffee increases, I may choose to replace some of my beverage drinking with tea in order to save money. As a result, when my demand for coffee decreases, my demand for tea will increase.

Changes in income directly affect the ability to pay. When the income earned by a household decreases, demand for most goods, called *normal goods*, will decrease. A household having income problems might, for example, delay the purchase of new clothing until clothes are unwearable, while a worker who gets a raise may go on a shopping spree. Not all goods follow this pattern. Demand for some goods, called *inferior goods*, rises when income decreases and falls when income increases; ramen noodles, a common food for college students, are a classic example of inferior goods. When income rises, students often substitute away from ramen noodles toward other, tastier foods.[2]

The **Law of Demand** is a statement of the relationship between prices and quantity demanded. It states that the relationship between prices and quantity demanded is negative - that the demand correspondence slopes downward.

1.2 Supply

Now that we understand households - who sell in the factors market and buy in the goods and services market - we can discuss their counterparts and trading partners: the firms.

[2] Or, in some cases, heavily-peated Islay scotch

1.2.1 The Firm's Maximization Problem

The firm is a profit-seeking enterprise. We'll delve deeper into profit later, but for now, define profit as total revenue minus total cost:[3]

$$\Pi(Q, P) = TR - TC$$

The firm uses profit to determine its quantity supplied, Q^S, according to a **supply correspondence** that usually takes the form $Q^S = c + dP$, where c is the quantity that the firm would supply even if the price were 0 and d measures the sensitivity of supply to price. (Usually, c is 0, or even negative.) When the firm chooses which quantity to product, it may need to take price into account. A perfectly-competitive firm - for example, a firm selling commodities - will take the spot price of a commodity as a given, but most markets involve some degree of product differentiation and as a result firms cater to specific sections of the market. They can entice new buyers by lowering the price or extract more revenue from current buyers by raising the price of the good. In chapter 3, we will discuss price-elasticity of demand in order to analyze these choices.

Several changes can affect supply, but they all come back to one main question: What does it cost to product this good?

Complements in production are similar to the complementary goods discussed under the Demand section above. However, rather than being goods that people like to consume together, complements in production are goods that are easy to produce together. An example can be found in the production of beef: when a cow is slaughtered for its meat, the cow's hide can then be tanned and used as leather for the production of clothing, bags, or other consumer goods. Thus, when the price consumers are willing to pay for leather increases, it becomes more attractive for farmers to raise cattle and the supply of beef will increase. When the price of complements in production rises, the supply of a good will increase, because it becomes less costly to produce that good.

Substitutes in production are goods that can be produced instead of each other. Complements in production generally use similar input goods and similar processes to produce, so the end product may vary without significant change to the process. Consider the cattle farmer we discussed in the previous example. Since cows and sheep are similar animals - they both graze in pastures, produce milk, and can be eaten for meat - it would likely be relatively simple to convert a cattle farm to a sheep farm. Consequently, if the price consumers are willing to pay for lamb or mutton[4] increases, farmers may substitute away from producing cattle and toward producing sheep. When the price of substitutes in production

[3] In the profit, revenue, and cost functions, the quantity variable Q doesn't have a superscript denoting it as quantity demanded or quantity supplied. That's because profit doesn't depend on what the seller wants to sell or what the buyer wants to buy - it depends on what's actually sold. $Q = Q^S = Q^D$, so before calculating any of these, make sure you know what quantity you're using. If there is no exchange but you are given a price, assume that $Q = Q^D$.

[4] Lamb is the meat of immature sheep, whereas mutton is the meat of grown sheep.

rises, the supply of a good will decrease, because it becomes more costly to produce that good.[5]

Finally, ongoing input prices have the strongest immediate effect on supply. Air travel is a clear recent example. The main variable cost of providing air travel is the fuel. During the early 2000s, the price of petroleum rose sharply due to unstable geopolitical conditions in the Middle East. As a result, airlines were forced to raise the prices of flights to avoid losing money. Since each flight became more costly to provide, the quantity provided at any price - Q^S - decreased. When the price of input goods increases, supply decreases.

The **Law of Supply** is a statement of the relationship between prices and quantity supplied. It states that the relationship between prices and quantity supplied is negative - that the supply correspondence slopes upward. There are several reasons for this. The most relevant is that the input goods used to produce final goods are typically purchased ahead of time. When prices rise, then, there is increased profit for each item sold. Since firms attempt to optimize profit, the firms will then choose to sell more units at the higher price to increase their share of the profit.

1.3 Equilibrium

An equilibrium is a price at which the firm's supplied quantity equals the households' (total) demanded quantity. At the equilibrium price, the firm is happy to sell the equilibrium quantity. Though it might not be the firm's profit-maximizing quantity, varying the quantity would affect either the number of buyers or the buyers' marginal willingness to pay.

To find an equilibrium, we need a way to determine the firm's supply correspondence and the households' demand correspondence. The simplest case involves being given both correspondences, but even if we are given a utility function and a cost function, it is possible to derive supply and demand functions from them.

First, note that given a price, we can determine quantity demanded and quantity supplied. That means we can treat P as a parameter, meaning that it is not a variable but a fixed value that we just happen not to know yet, and we will call that equilibrium price P^*. Second, note that in order for an equilibrium to exist, the quantity supplied must be the same as the quantity demanded, as functions of that price. Since $Q^S(P^*) = Q^D(P^*)$, we can simply relabel that quantity Q^*. Next, set up and solve the following system of equations:

$$Q^* = Q^D = a - bP^*$$

$$Q^* = Q^S = c + dP^*$$

$$a - bP^* = c + dP^* \Rightarrow \frac{a-c}{b+d} = P^*$$

[5]No, not directly, but the cost of the best alternative drops. If this seems counterintuitive, see the section on opportunity cost, below.

Now that we know P^*, we can then determine Q^* by plugging P^* into either the demand or supply correspondence (or both, to check your math).

1.4 Opportunity Cost

In Engineering Economics especially, the concept of opportunity cost is important to making decisions. Considering the value of other options allows any economic agent to evaluate not only the value of the benefit gained but also the cost expended. However, taking one action often necessitates foregoing another. For example, you are currently reading these notes.[6] Think about all of the other options you have for spending the time you are currently exhausting on studying: you could be napping, drinking at the Changing Times Pub, binge-watching *Doctor Who*, studying for a more important class, tutoring your neighbor in calculus, or engaging in intellectually stimulating conversation with your husband, wife, boyfriend, girlfriend, significant other, or polyamorous family unit. However, since you are currently studying Engineering Economics, it must be the case that you consider that activity more important, or less costly, or both, than the alternatives.

Delve deeper, though. Realistically, you could only do one of those other activities. While it's true you could try to tutor at the Changing Times, or sleep while watching Netflix, you can't give each activity your full attention, and thus the benefits would be limited. (For example, it's difficult to enjoy a glass of scotch while simultaneously trying to explain the product rule to a younger student.) Thus, the opportunity you're foregoing is the best available option, not all of the others.

Opportunity cost is the value of the best foregone alternative. What would you be doing if you weren't studying? Let's assume you'd choose to tutor your neighbor at a rate of $25 per hour. Since that would be your choice, you don't need to consider the value of watching *Doctor Who* or studying differential equations - since you can only engage in one activity at a time, the fact that you'd choose to tutor means they no longer matter to our calculation. Since you are giving up the opportunity to make $25 in order to gain an hour studying Engineering Economics, it must be true that you value Engineering Economics more than the $25.

"But wait!" you protest. "I would rather have $25!" Would you really? Consider: do you have a tutoring client? If not, then tutoring isn't really a foregone alternative, since it was never an option. If so, then why are you not tutoring them? Common answers might be that you need to get a good grade in this class because it is a requirement to graduate, or that your scholarship depends on your GPA, or that your parents will kill you if you fail a class. Even though you probably did not explicitly consider all of those reasons, your decision to study rather than to do something else means that you value the benefits that come from studying, if not the studying itself, more highly than the benefits of doing something else.

[6] Presumably because you are required to.

A more formal definition of opportunity cost is a ratio of what is given up to what is gained. One example often involves the limited time that students have to study for multiple classes. If you have allocated a limited amount of time to study, then you may be able to adjust your scores in one class, but not both. Imagine that you have two exams tomorrow, one in English and one in History, and that you only have time to study for one.[7] If you go in cold to each exam, your final average will be a 70% – enough to pass with a C, but just barely. Studying for English will net you a final average of 80%, while studying for History will net you a final average of 90%. Since studying for History is the best alternative to studying for English, the opportunity cost of studying for English is (90-70) = 20 points (given up) on your History grade. Accordingly, the opportunity cost of studying for History is (80-70) = 10 points on your English grade. Since, however, it is not necessarily true that you must study for only one exam, it is possible to calculate the opportunity cost in a more granular way. In studying for English, 10 points in English are gained, but 20 points in History are given up. The opportunity cost is thus the ratio:

$$\text{Opportunity Cost} = \frac{\text{given up}}{\text{gained}} = \frac{20}{10} = 2$$

Ergo, the opportunity cost of raising your English grade 1 point is the chance to raise your History grade 2 points; in other words, the potential to gain 2 points in History is given up to gain 1 point in English. The opportunity cost of raising your History grade 1 point is then $\frac{1}{2}$ point on your English grade.[8]

1.4.1 The Do-Nothing Alternative

Consider the following situation: You are at a casino. You have a crisp new $100 bill in your pocket and an hour before your friend arrives. There are several options available: blackjack, poker, and slot machines. Each has its advantages and disadvantages. Blackjack offers a 45% probability that you will double your money over the next hour, but a 55% probability you will lose it all. Based on your understanding of statistics, you know this means you should expect to have about $90 at the end of the hour. Poker is a better proposition - since it is a game of skill, you have a 60% chance of earning an extra $50 (for a total of $150), but a 40% chance of losing all of your money. That means you can expect to have about $90 in your pocket at the end of the hour. Slot machines, to go to the other extreme, are a highly negative expected-value proposition. You stand a 1% chance of winning $1000, but a 99% chance of losing all of your money. As a result, you could expect to have about $1 in your pocket at the end of the hour.

Thinking like an economist, you quickly winnow your options down to blackjack or poker, since you cannot abide such a risky proposition. Then, however,

[7]This may seem unrealistic, but assume a cursory overview of broad topics is factored into your final score already.

[8]Demonstrate this as an exercise.

you're stuck - the expected values are the same. Which game is it rational to play?

Similarly, consider this problem raised in a freshman course on ethics: You are on your way out of a coffee shop carrying a double shot of espresso and a $1 bill you received as change. Two homeless people, one man and one woman, each step toward you and simultaneously ask you for the dollar. Since you don't have any coins, you cannot split the value between the two people. Who should you give your dollar to?[9]

What do these two situations have in common? In each of them, you are attempting to choose between two options that result in negative consequences for you. In the gambling scenario, you have two options, each with the expectation of losing $10. In the coffee shop scenario, you have two people each asking for $1. In neither case is there a compelling reason to choose one option over the other. The underlying assumption, though, is that we must choose an option at all.

The **do-nothing alternative** is often (but not always) a hidden option when making choices. For example, in the gambling scenario, you have the option to literally do nothing for an hour until your friend arrives. This leaves you $100 with certainty. In the coffee shop scenario, you have the option to politely refuse each person's request, leaving you free to keep your dollar. Not every situation allows a do-nothing option; for example, a baseball manager faced with the option of starting his worst starting pitcher or a pitcher who is usually used only in long relief cannot opt to simply start no pitcher. However, a voter who is disgusted with all available candidates may bemoan his "forced" vote for the lesser of two evils without acknowledging that he has the option simply not to vote at all. The do-nothing option is often low-cost but has low returns as well, making it a great way to avoid choosing the best of a bad lot, but a lousy choice for a firm seeking growth.

[9]This was met with considerable debate about the probability that the homeless woman had children.