Consumer Choice

I. Definitions

The purpose of this section is to carefully define the specific terms used in the consumer choice model.

A. Types of Commodities

1. Goods versus Bads

During the course of the semester it will be common to refer to "goods" when discussing such things as production and consumption, demand and supply, and so forth. From the point of view of consumers "goods" are those items that when consumed make a consumer better off. Thus, increasing the consumption of a good increases a consumer's welfare.

On the other hand, "bads" are just the opposite – those items that when consumed make a consumer worse off. Thus, decreasing the consumption of a bad increases a consumer's welfare.

What are some examples of goods and bads? In fact, individual consumers get to decide for themselves whether a particular commodity is a good or a bad. For example, some consumers may view peanut butter as a good while others view it as a bad. However, there do exist a number of commodities that people generally view as bads.

For example, crime is generally considered as a bad – to everyone except the criminal. Thus, if you're mugged or burglarized most consumers view this as decreasing their welfare. Other examples of generalized bads would include fire, accidents, wars and invasions, and pollution.

The general way in which bads will be treated in this class will be by looking at the "good" created by diminishing the consumption of the bad. Thus, police protection in an attempt to diminish crime is a good. Likewise, fire protection, insurance, national defense, and pollution controls would be goods because they all diminish the consumption of a bad.

2. Goods versus Services

In the comparison between goods and services, both goods and services are similar in that they are considered to increase a consumer's welfare. Hence, they are both "goods" in the sense discussed above. However, in this comparison goods are normally considered to be commodities. Thus, goods have a physical reality, existing independent of the producer. There are, of course, numerous examples of goods in this sense, including automobiles, oranges, books, etc.

On the other hand, services, whose consumption also increases an individual's welfare, do not physically exist. Rather, they consist of actions taken by one or more individuals to benefit the consumer. For example, a dentist and his or her assistant take actions on your behalf when you visit them for a dental checkup. Other examples of such services would include legal services, physician services, accounting services, etc.

In general, unless explicitly noted otherwise, the term "goods" will be used to identify either commodities or services that fit the first definition – anything that increases a consumer's welfare with consumption.

B. Utility

Recall that earlier, when discussing actors in the market, we identified households or consumers as acting to maximize their "satisfaction." At that point, the goal of maximizing satisfaction and the concept of satisfaction itself was only vaguely defined. Now, however, that concept is crucial in our discussion of how consumers behave in making decisions.

Understanding this concept more fully requires us to define what is meant by "satisfaction." The step in the process is to rename the concept to the more conventional one of "utility" and to examine more closely what is meant by the concept of utility or satisfaction.

Utility is an attempt to quantify the satisfaction a consumer gains from consuming a good or goods. But how can satisfaction be accurately measured? There exist two types of utility that reflect two different notions of how utility or satisfaction could be measured.

1. Ordinal Utility

Under ordinal utility, we assume that individual consumers have preferences about goods that reflect the underlying satisfaction or utility they gain from consuming these goods. We also assume that consumers can order or rank their preferences. Thus, a consumer may prefer two oranges to an orange or an apple – or the reverse may be true. In comparing any two goods, or groups/bundles of goods, a consumer is able to decide which of the groups is preferred to the other. Ordinal utility also allows for the possibility that the consumer is indifferent between the consumption of any two bundles of goods.

That consumers have the ability to rank or order one's preferences, as required under ordinal utility, seems like a reasonable assumption to make about consumers to most students. In our daily lives, we face these kinds of decisions all the time and, given a world of scarcity, must decide between different goods and different bundles of goods. That these decisions are, at least partially, based upon one's preferences for the different goods seems not only reasonable but also inescapable.

Although the concept of ordinal utility is undoubtedly realistic, unfortunately it is not a particularly useful concept for our current purposes – predicting how consumers will behave in different situations. As a result, we now turn to a more useful but perhaps less realistic type of utility.

2. Cardinal Utility

Cardinal utility assumes that individuals can not only rank their preferences but that explicit numbers measuring utility can be assigned to the consumption of different goods and bundles of goods. A higher number represents higher utility and satisfaction.

There are several advantages, theoretically, to cardinal utility. For example, cardinal utility theoretically allows interpersonal comparisons via the utility numbers. Thus, if individual A's utility for consuming an orange is 10 while equaling 15 for individual B,

then B values the orange more than A. Of course, realistically there exists the problem of both whether the underlying assumption that consumers actually have such utility numbers as well as the problem of, even if they do, objectively determining what they are for different consumers. We will discuss the latter a bit later in this chapter.

The most important advantage to cardinal utility is not that it can be objectively measured, but that it allows us to predict how consumers behave as they make their decisions. Pragmatically, the issue of whether or not consumers actually have utility, which determines their consumption, is much less important than whether or not models based upon the assumption that they do can accurately predict consumer choices. Whether or not the assumption is correct, consumers behave as if they make decisions based upon cardinal utility.

C. Total utility and marginal utility

The main purpose of our discussion of utility is to develop a model of consumer choice. In our model of consumer choice, we will assume that consumers' goal is to maximize their cardinal utility. Within the framework of this model of consumer choice, we must first define two concepts:

1. Total Utility

Given the discussion above, the concept of Total Utility (TU) is quite straightforward. A given consumer's total utility simply measures in numbers the utility that an individual consumer gains from the consumption of a good or goods, with higher numbers indicating that the consumer's welfare has increased. Table 1 illustrates this concept for the consumption of oranges for consumer A. Notice that as this individual increases consumption of oranges that her total utility also increases, indicating that increased consumption makes the consumer better off.

Table 1

Total Utility from Consuming Oranges for Consumer A

Quantity of Oranges Consumed	Total Utility
0	0
1	100
2	250
3	350
4	400

2. Marginal Utility

Marginal Utility (MU) is a different, but related concept. The word "marginal" might lead you to think that marginal utility is utility that is of lower quality. However, in the context of utility, "marginal" essentially means the "extra" or "additional" utility, Thus, marginal utility measures how much extra utility the consumer gains by consuming one more unit of a good.

For example, we can see that the consumer in Table 1 increases her total utility by 100 when she consumes the first orange and by 150 when she consumes the second orange. The marginal utility derived from the Total Utility presented in Table 1 is given in Table 2. Notice that for each orange consumed, while total utility includes the utility gained from previous consumption, the marginal utility only measures the utility gained from consuming the last orange.

Table 2

Total and Marginal Utility from Consuming Oranges for Consumer A

Quantity of Oranges Consumed	Total Utility	Marginal Utility
0	0	--
1	100	100
2	250	150
3	350	100
4	400	50

Table 2 illustrates the concept of marginal utility when consumption rises by one orange. In this case, the marginal utility simply equals the change in total utility from the previous orange consumed. However, suppose consumption rises by multiple numbers of oranges at once.

Consider, for example, what marginal utility would equal when consumption rises by two oranges at each step, with the numbers for total utility remaining the same as in Tables 1 and 2. Thus, total utility would still equal zero when consumption is zero and would rise to 100 initially, but this would be for two oranges not one. However, the definition of marginal utility requires one to measure the extra utility gained by consuming one, not two, additional units of the good. We can take into account the fact that consumption rose by two and not one by simply dividing the change in total utility by the change in consumption. This yields the desired result – the extra utility gained by consuming only one of the additional oranges. As a result, mathematically marginal utility is found by (recall that the symbol Δ means "change in"):

(1) $$\text{Marginal Utility} = \frac{\Delta TU}{\Delta Q}$$

As a result, the marginal utility when quantity consumed increases by two rather than by one each time is given in Table 3.

Table 3

Total and Marginal Utility from Consuming Oranges for Consumer A

Quantity of Oranges Consumed	Total Utility	Marginal Utility (DTU/DQ)
0	0	--
2	100	50
4	250	75
6	350	50
8	400	25

The concept of "marginal" will be your constant companion throughout the rest of microeconomics. Later, we will discuss similar concepts such as marginal product, marginal cost, marginal revenue, etc. All of these concepts have a definition that is similar, although not exactly the same, to marginal utility. An actor in the market, typically a household or a firm, is engaging in an activity. The activity yields or causes some other variable to change. For marginal utility, households are engaging in consumption, which causes their total utility to change. The marginal concept simply measures the extra amount of the variable, such as utility, caused by a one-unit change in the activity, such as consumption of a good.

Thus, the only thing that will change with these various types of marginal concepts will be: (1) which actor in the market is involved in (2) which activity and (3) which variable is affected. However, the definition will always remain essentially the same as will the equation used to calculate the actual numbers. Given that the marginal concept is so numerous in microeconomics and so important to individual decision-making, it is extremely important that you both understand the concept generally and in specific for each type of marginal.

D. The law of diminishing marginal utility

Next, consider what both total and marginal utility might look like, graphically. As consumption increases, what happens to total utility, for example – does it increase, decrease, or remain constant? Given our definition of utility above, it seems reasonable that total utility at least must increase as the quantity consumed increases as is shown in Graph 1. But does it really?

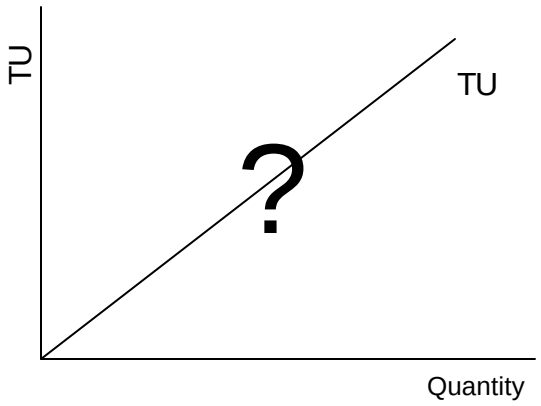

Graph 1

Actually, the best way to understand what happens to total utility as consumption increases is to carefully consider what happens to marginal utility as consumption increases. To understand this concept think carefully about how you would behave – you're a consumer and your intuitive understanding of consuming and markets can help you understand this and other concepts.

Suppose you consume one pizza. Obviously you get some utility or satisfaction from this consumption, but then consider what happens to the utility you get from each individual pizza (that's marginal utility) as you eat more and more pizzas. For this exercise to make sense, you must also assume that the length of time you have to eat all the pizzas remains constant. Thus, in one day you eat one pizza and then a second in the same day. Then you eat a third pizza in the same day, and a fourth, and a fifth, and… What happens eventually to the amount of utility or satisfaction you gain from each successive pizza you eat? While the marginal utility might rise initially, in this limited time period eventually you just start getting sick and tired of pizzas. When that happens your marginal utility falls. This process is known as the law of diminishing marginal utility and is stated as follows:

> Eventually, as the consumption of a good increases, ceteris paribus, the marginal utility must decrease.

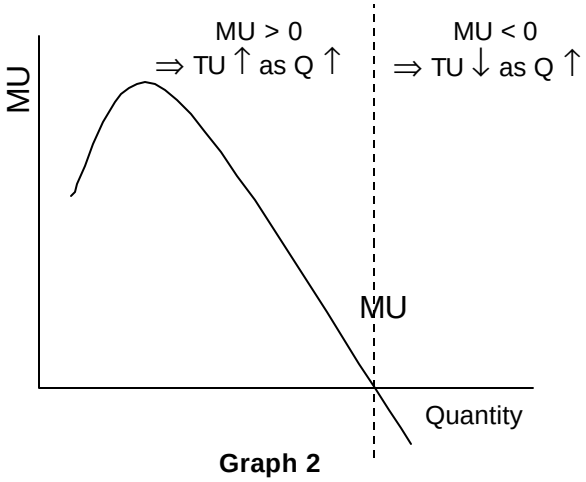

Graph 2

Graph 2 illustrates the law of diminishing marginal utility. While marginal utility initially rises with increased consumption, eventually, marginal utility falls.

E. The relationship between Marginal and Total Utility

Graph 2 also helps us to understand what Total Utility looks like. Recall that we find marginal utility by dividing the change in total utility by the change in consumption (Q). However, this is simply the slope of the total utility curve – the rise (ΔTU) divided by the run (ΔQ). Clearly when marginal utility is positive then total utility must increase as consumption increases and, likewise, when marginal utility is negative then total utility must decrease as consumption increases.

Further, given that marginal utility is the slope of the total utility curve, then when marginal utility is increasing, the total utility curve is getting steeper. The reverse is also true, when marginal utility is decreasing, then total utility must be getting flatter. Graph 3 shows the relationship between marginal and total utility. Notice that at each of the dotted lines, both curves change:

1. Line 1 – Marginal Utility reaches a maximum – Total Utility begins to become flatter.
2. Line 2 – Marginal Utility equals zero – Total Utility reaches a maximum.

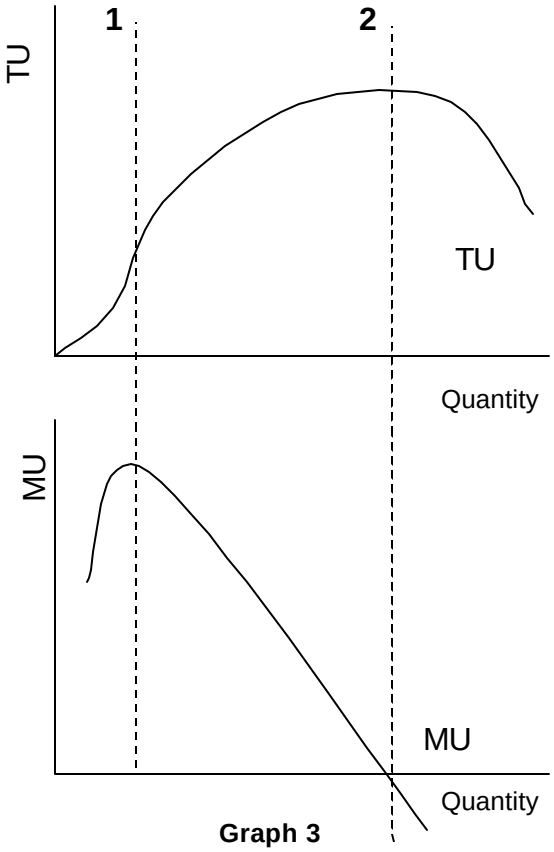

Graph 3

II. Maximizing Utility

As discussed above, households or consumers are assumed to act to maximize "satisfaction" or utility.

A. How do consumers act to maximize total utility?

The simple utility maximization model focuses first of all on the variables that matter to consumers in their consumption decisions. These variables include:

- The consumer's preferences (measured by their utility).
- The prices of goods the consumer buys.
- The consumer's income.

The model includes the first two relevant variables, preferences and prices, by dividing each good's marginal utility by its price: MU/P. Before using this in the model, it is important to carefully define it's meaning. Suppose that at the current level of consumption, a consumer's marginal utility for a good equals 200 while the price equals $20. Dividing 200 by 20 obviously yields 10, but what does this 10 tell us? Essentially what we've done is taken the 200 units of utility that consumption of the good gives the consumer and divide it into 20 equal parts – one for each dollar spent on the good. Thus, the MU/P equals the utility the consumer will gain by spending one more dollar on the good.

How does income enter into the model? Clearly, given that the goal is to maximize utility, the consumer with a dollar to spend will spend it on the good with the highest MU/P. Thus, the consumer will spend her income first on those goods with the highest value, those goods with the highest MU/P. If the MU/P were equal for two goods then the consumer would be indifferent between them and would randomly choose one over the other.

Where does the consumer end up by following this general rule? A consumer will maximize her/her utility by consuming such that:

(2) $MU_1/P_1 = MU_2/P_2 = MU_3/P_3 = ….. MU_K/P_K$

In other words, a consumer maximizes her total utility by consuming such that the MU/P is equal for all the goods that she consumes (good 1, good 2, good 3, …., and good K). It is relatively easy to see that this would be the case simply by considering what would happen if equation 2 were not true. For simplicity, let's assume that the consumer is currently only consuming two goods and that at those current consumption levels: $MU_1/P_1 = 20$ while $MU_2/P_2 = 10$.

Clearly, if the consumer took one dollar from consumption of the second good, losing 10 units of utility, and spent it instead on the first good, gaining 20 units of utility, he would increase his total utility by a net of 10 units of utility.

This simple demonstration illustrates two important principles:

1. Whenever the MU/P are unequal for goods then the consumer can increase their total utility by buying less of the good whose MU/P is smaller and more of the good whose MU/P is larger. For example, suppose that $MU_1/P_1 > MU_2/P_2$. In this case, the consumer will change her consumption patterns by increasing consumption of good 1 and decreasing consumption of good 2.

2. According to the law of diminishing marginal utility, when a consumer increases their consumption of a good then his marginal utility decreases and the reverse. Consider again the example given in principle 1 above, where $MU_1/P_1 > MU_2/P_2$ and consumers responded by increasing their consumption for good 1 and decreasing their consumption of good 2. Changing consumption in this manner will cause MU_1 and, hence, MU_1/P_1 to decrease while MU_2 and MU_2/P_2 to increase. Both of these changes will move the consumer back to an equality between MU_1/P_1 and MU_2/P_2.

   Thus, when MU_1/P_1 and MU_2/P_2 are not equal, the consumer responds by changing his consumption in a way that will eventually make them equal.

B. Marginal Utility and Demand

Earlier, in our discussion of cardinal utility, we noted that it is difficult to actually measure a person's utility mostly because preferences are so subjective. Furthermore, individuals sometimes have an incentive to be less than truthful about their preferences. Suppose, for example, a government agency is distributing welfare payments on the basis of potential recipients' preferences or utility. That type of welfare program would give individuals an incentive to inflate the utility they would gain from the program and illustrates the problems inherent in measuring utility.

Regardless of the subjectivity of preferences and utility, there does exist an objective method to measure utility although doing so is only possible when consumers are making decisions within a market. Consider the demand curve to the right.

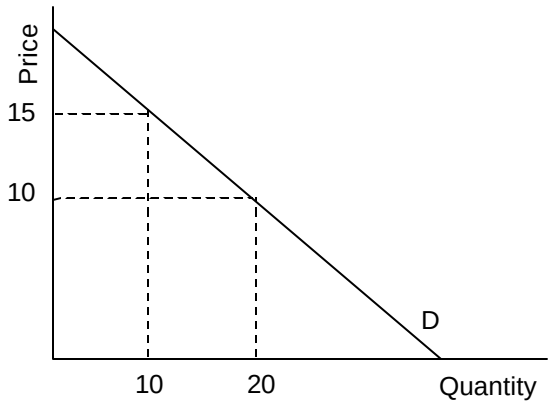

Graph 4

The normal way of looking at this demand curve is to say that at a price of ten dollars the consumer would buy twenty units of the good. However, if the price rises above ten, the consumer will not buy the twentieth unit. Thus, the demand curve also tells us the maximum price the consumer would be willing to pay to buy any given unit of the good, such as the twentieth. For example, the maximum price the consumer in Graph 4 would be willing to pay for the tenth unit would be fifteen dollars, again given by the demand curve at the given quantity.

One final step remains before illustrating how we can measure utility objectively, which is to ask what is the value to the consumer in dollars of their consumption. Again, in Graph 4, the maximum price the consumer would pay for the tenth unit equals fifteen dollars and is given by the demand at a quantity of ten. Does the consumer value this tenth unit more, less, or the same as the maximum amount he'd be willing to pay for it? Clearly, the consumer's value for this tenth unit must equal fifteen dollars. If the consumer valued the tenth unit less than fifteen dollars then she would not be willing to pay fifteen dollars for the good. Likewise, if the consumer valued the tenth unit more than fifteen dollars then she would also be willing to pay more for the good.

Thus, the demand curve measures the extra value, measured in dollars, that a consumer gains from consuming an additional unit of the good. Marginal Utility is defined to equal the extra utility a consumer gains from consuming an additional unit of the good. The only difference between these two definitions is that the former, the value, is measured in dollars while the unit that measures utility has never been defined. However, if utility is also measured in dollars then the two definitions become the same.

The above discussion illustrates an important point; demand is an objective measure of marginal utility. That is, simply asking individuals about their preferences and utility suffers from the serious flaw that the measure is subjective and individuals may not behave consistent with that measure in the real world. However, if we are trying to measure behavior, why not simply look at how they behave in a market when making actual consumption decisions? That is what economists do when they estimate demand curves based on the market behavior of individuals.

C. Consumer Surplus

Above, we discussed the maximum price that consumers were willing to pay for a good. However, as all consumers understand, consumers do not always pay the maximum price for a good. Consumer surplus conceptually recognizes this fact. In fact, consumer surplus is the equivalent of a firm's profit for consumers. Profit is the difference between a firm's revenue and it's costs. Consumer Surplus is the difference between a consumer's value and costs. Since value is measured by price in the market, as discussed in the previous section, consumer surplus exists whenever the consumer actually pays less for a good (his cost) than the maximum price he would pay for a good (his value.)

Consider, for example, the situation in the graph 4 above. If the market price equals ten dollars then this consumer would buy 20 units of the good. However, the consumer values the tenth unit of the good at fifteen dollars. The difference between the price the consumer actually paid ($10) and the price she would have been willing to pay ($15) for the tenth unit equals consumer surplus for that unit. Thus, the consumer in Graph 4 has a consumer surplus of five dollars for the tenth unit consumed.

This example illustrates two important points about consumer surplus:

1. For any given unit of consumption, the consumer surplus equals the difference between the demand curve (the unit's value) and the actual price of the good. The downward sloping demand curve reflects the consumer's diminishing marginal utility as consumption increases. Diminishing marginal utility clearly implies that the first units consumed will have higher levels of consumer surplus than the last units consumed.

2. Total consumer surplus simply measures the difference between demand (value) and actual price for all units of the good consumed. In Graph 4, we can find this value simply by adding up the consumer surplus for each of the 20 units consumed when price equals ten dollars. In fact, the total consumer surplus is given by the area of the triangle that lies between the demand curve and a horizontal line at the market price, as shown in Graph 5.

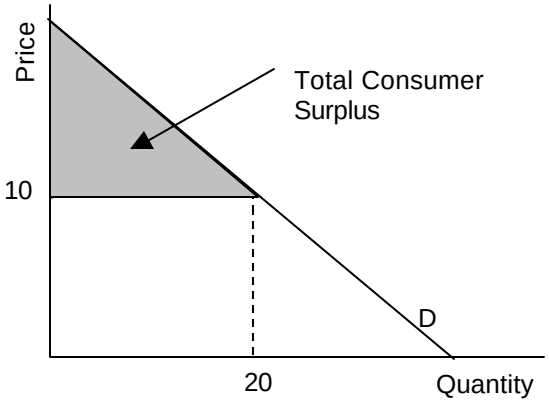

Graph 5

D. The Model in Action

The true test of an economic model is how well it predicts human behavior. The utility consumer choice model, while having some shortcomings, does predict human behavior well. For example, consider what we all as consumers know about how consumers behave. For simplicity, assume that a consumer only consumes two goods. Further, assume that the consumer is initially consuming each good so as to maximize his total utility – that is, $MU_1/P_1 = MU_2/P_2$. Finally, suppose that the price of good 1 (P_1) increases.

In this situation, we want to carefully consider how a consumer would actually respond to the increase in the price of good 1 and (2) and compare that to how the model predicts the consumer would respond. If the two are the same, this is an indication that the model is useful in predicting actual behavior in the real world.

- How does the consumer actually respond to the price increase?

  In the real world, we would observe two impacts caused by an increase in P_1. The first impact is just the straightforward impact that we all observe in our daily lives that economists call the law of demand. As the price of the good increases, consumers respond by buying less of the good.

  Second, the change in relative prices may also change consumption of the second good depending upon the relationship between the two goods. As we discussed in our section on Demand, if the two goods are substitutes then the consumer responds to an increase in P_1 by also increasing her consumption of the now relatively cheaper second good. However, if the two goods are complements, then just the opposite happens – as P_1 increases the consumer responds by buying less of both goods because they are consumed together. Of course, if the two goods are not related in consumption, are neither complements nor substitutes, then an increase in P_1 has no impact on the consumption of the second good.

- How does the consumer respond according to the model?

  Recall that initially the consumer is assumed to be maximizing her total utility so that $MU_1/P_1 = MU_2/P_2$. This changes when P_1 increases. The ratio on the right of the equation doesn't change while the ratio on the left decreases (because it is being divided by a larger number) so that now $MU_1/P_1 < MU_2/P_2$. As discussed in Section IIA. above, when the two MU/P are unequal to maximize utility the consumer must buy more of the good whose utility per dollar is larger and less of the good whose utility per dollar is smaller.

  Thus, in this case the model predicts that the consumer will respond to the increase in P_1 by buying less of good 1 and more of good 2. Notice that this prediction is consistent with what is observed in the real world, assuming that the two goods are substitutes in consumption.