Model of Consumer Choice

August 23, 2022

Underlying the demand curve is that individuals are able to make choices about which goods and services they will purchase. This model of consumer choice is built on preferences. We assume that these preferences are stable (unchanging), or at least that they are stable during the snapshot of time at which we are observing choices being made. An individual's preferences will likely change over the course of a lifetime – at the very least it seems obvious that the size of one's clothing will change over the course of a lifetime – but in this model we are just focusing on a snapshot of a particular point in time.

In this chapter we will build a model of consumer choice and discuss the conditions that need to be met for a consumer to be making optimal decisions. We will begin with an overview of the restrictions that we place on consumer preferences. Next we will discuss how these preferences are related to consumer utility. We will then develop the concept of a budget constraint. Finally, we will show how to develop the conditions that must be met for a consumer to be behaving optimally.

1 Consumer Preferences

A main presumption is that consumers get a certain benefit or satisfaction (called utility in economics) from consuming goods and services. The goal in this section is to determine the level of utility that each bundle of goods and services gives a consumer. Although the analysis extends to more than 2 goods, we will work with 2 goods for simplicity.

1.1 Properties of Consumer Preferences

There are 3 primary properties that we will deem necessary in order for our consumer preferences to be rational. The properties are defined below.

1. **Completeness** – this property says that consumers can rank their bundles such that, given 2 bundles A and B, either (1) A is at least as good as B, (2) B is at least as good as A, or (3) the consumer is indifferent between A and B. Note that if a consumer is indifferent between bundles it means he receives the same level of utility for each bundle of goods.

2. **Transitivity** – given at least 3 bundles, A, B, and C, if (1) A is at least as good as B, and (2) B is at least as good as C, then (3) it must be that A is at least as good as C.

3. **Nonsatiation** (or, as it is more commonly called, more is better) – Suppose that bundles A and B consist of two goods, good 1 and good 2. If bundle A has more of both goods than bundle B, then the consumer will prefer bundle A to bundle B. If a bundle has a larger quantity of ALL goods than another bundle, then the bundle with the larger quantity is preferred to the bundle with the smaller quantity. If the case were that bundle A had a larger quantity of good 1 than bundle B but the exact same amount of good 2, then we would say that bundle A is at least as good as bundle B. Thus, if one bundle has more of one good but the exact same of the other goods then we say that the bundle with more of the one good can be no worse than the bundle with the lesser amounts of goods.

We will assume that all consumers will have preferences that satisfy these 3 properties. The first two properties are what economists assume when they state that consumers are rational – that they can rank bundles of goods and that they have no intransitivities. Those are still big leaps of faith, but "rational" to an economist just means that people have these basic preferences; there are no assumptions on what goods individuals should prefer.

You should note that our analysis still holds if we do NOT have the more is better property. The more is better property is used for two reasons. First, it seems a reasonable assumption to make that if you have more of all goods[1] that you will be better off in the sense of having a higher utility level. Second, it makes the analysis a little more tractable.

1.2 Graphing consumer preferences

In this section we will use a graph to aid in our analysis of consumer preferences. We will focus on the positive quadrant of the Cartesian plane, as we will assume that you cannot consume negative quantities of goods. The axes of the graph will be labelled Good 1 and Good 2. Thus, each point on the graph will represent a bundle of goods consisting of an amount of Good 1 and Good 2 corresponding to that point. Figure 1.2 shows 6 bundles distinctly labelled A–F.

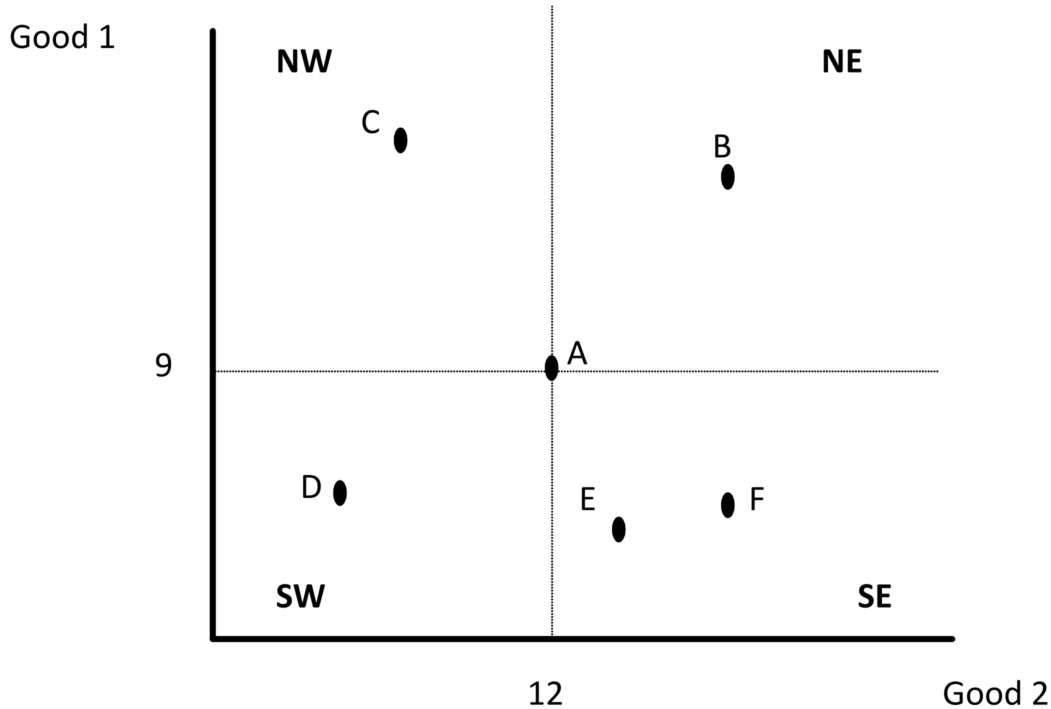

Figure 1.2: Some bundles of goods for a consumer, using Bundle A as the comparison point for other bundles.

Note that bundle A is given by the intersection of the 2 dotted lines, and it corresponds to a quantity of 9 of Good 1 and 12 of Good 2, or the ordered pair $(12, 9)$. You will also note that each section of the graph has been labelled as northeast (NE), northwest (NW), southwest (SW), or southeast (SE). These labels are in relation to point A in the graph.[2]

NE corner Now, suppose that we want to compare bundle B and bundle A based on our properties of consumer preferences. Notice that bundle B has more of both goods than bundle A. By the more is better (nonsatiation) property, it must be the case that B is preferred to A. In fact, any bundle in the NE corner of the graph is preferred to bundle A, as all of those bundles have more of both goods than bundle A.

[1] The word "good" is chosen for a reason – a good is something consumers like. If we wanted to discuss items that consumers did not like, such as pollution, we would call them "bads." Alternatively, we could always reframe a "bad" as a "good" – while most (perhaps all) consumers view pollution as a bad, we could create a good by calling it "absence of pollution."

[2] Another way to think about it is to create a new Cartesian plane with point A is the new origin. Then the NE corner is quadrant I, the NW corner is quadrant II, the SW corner is quadrant III, and the SE corner is quadrant IV.

SW corner Now, let's compare bundle D and bundle A. Because bundle A has more of both goods than bundle D, again by the more is better property we know A is preferred to D. Notice that bundle A has more of both goods than any bundle in the SW corner, which means that bundle A is preferred to any bundle in the SW corner.

NW and SE corners Notice that bundles in the NE corner (like bundle C) have more of Good 1 than bundle A, but less of Good 2. Also, bundles in the SE corner (like bundles E and F) have more of Good 2 than bundle A, but less than Good 1. This means we cannot use the more is better principle to determine which bundles in these corners are preferred to bundle A. Thus, the preference relation between bundles in the SE and NW corners and bundle A are determined by how much a particular consumer likes Good 1 and Good 2. We will use the concept of an indifference curve to determine the preference ordering of these bundles.

1.2.1 Indifference curves

An indifference curve is a plot of all the bundles that give the consumer the same level of utility (hence the name *indifference curve*, meaning that the consumer is indifferent between the bundles along the curve). Consumers have an infinite amount of indifference curves – if we were to plot all of the consumer's indifference curve we would get their indifference map. Figure 1.2.1 shows 3 indifference curves for this consumer. The curve through bundle B is labelled I_3. The curve through bundles C, A, and F is labelled I_2. The curve through bundles D and E is labelled I_1. Because C, A, and F are all on the same indifference curve, the consumer receives the same amount of utility from each bundle. Following Figure 1.2.1 are some rules for indifference curves.

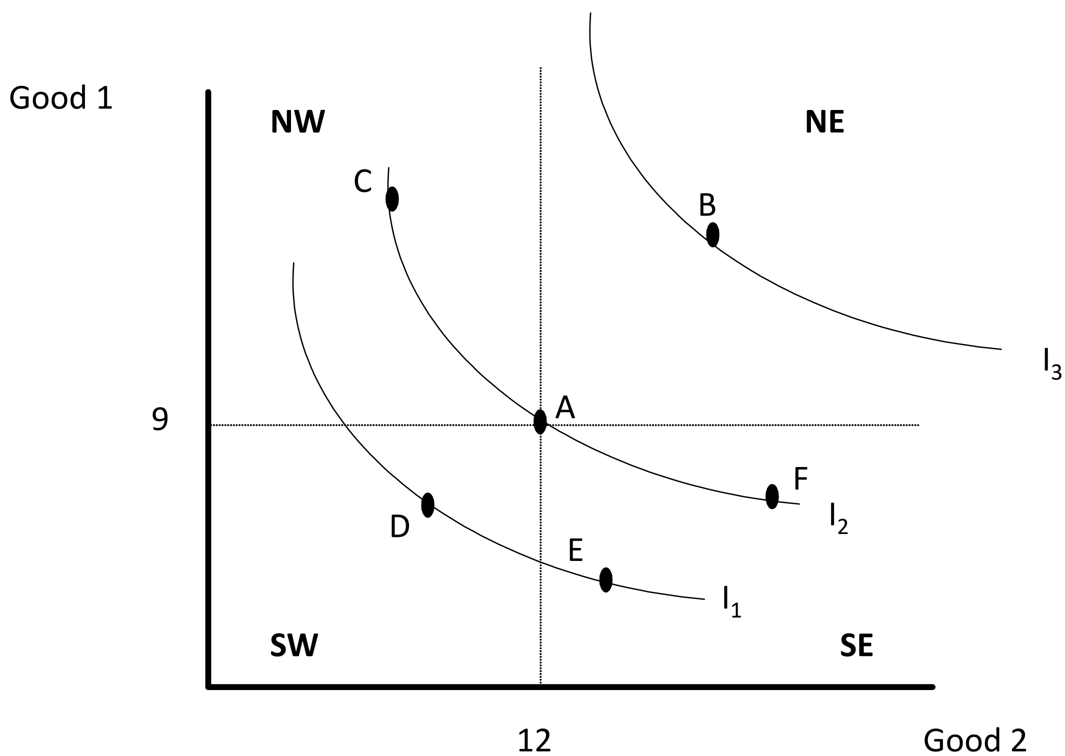

Figure 1.2.1: Example indifference curves for a consumer.

Rules for indifference curves:

1. Bundles on indifference curves farther from the origin are preferred to those closer.

   This means that the consumer would prefer to be on I_3 rather than on I_2, and would prefer to be on I_2 rather than I_1. This rule is consistent with the more is better principle. The consumer prefers bundle B to bundle A, so she must have a higher utility at bundle B than he does at bundle A. Thus, all points along the indifference curve that pass through bundle B must give a higher utility level than those that are on the indifference curve that pass through bundle A. So I_3 is preferred to I_2. A similar argument can be constructed for the relationship between I_2 and I_1.

2. There is one and only one indifference curve that passes through each bundle.

   If there was more than one indifference curve that passes through any bundle, then the consumer would be saying that a bundle gives her a utility level of 12 (from the first indifference curve passing through the bundle) as well as a utility level of 10 (from the second indifference curve passing through the same bundle). You can show that transitivity would be violated if the consumer had preferences with this feature.

3. Indifference curves may not cross.

   For starters, if they crossed then rule 2 above would be violated. You can also show that transitivity is violated by indifference curves that cross.

4. ***Indifference curves are downward sloping.***

   I have marked this rule because there are classes of indifference curves that are not exactly downward sloping. If the consumer receives zero utility from consuming one of the goods, or if the two goods are perfect complements, then the indifference curves will consist of perfectly vertical lines, perfectly horizontal lines, or a combination of the two (meaning that they are L-shaped in the case of perfect complements). We will not focus our analysis on these cases.

1.2.2 Slope of an indifference curve

The Marginal Rate of Substitution (MRS) is defined as the maximum amount of one good a consumer will give up to obtain one more unit of another good. Thus we want to find the amount of Good 1 that a person will give up in order to get one more unit of Good 2. Writing this mathematically (assuming Good 2 is on the x-axis and Good 1 is on the y-axis), we have the $MRS = \frac{-\Delta Q_1}{\Delta Q_2}$. Notice that this formula is just a slope as we simply have a change in the quantity of Good 1 divided by a change in the quantity of Good 2. Also notice that the MRS is negative because we must give up some of Good 1 in order to get more of Good 2. On a technical note, because the indifference curve is a curve and not a straight line, the slope of the indifference curve will change depending on the point at which we evaluate the slope. We will return to this concept later in the notes.

2 Utility

We have discussed indifference curves as running through bundles of goods that give the same level of utility. We will now make the concept of utility more formal. We suppose that every consumer has a "utility function" which allows her to take different bundles of goods and assign them levels of utility in such a manner that does not violate the properties of consumer preferences described above. For instance, let $U(Q_1, Q_2)$ be the consumer's utility function that determines the level of utility a consumer receives from consuming different quantities of Goods 1 and 2. A particular utility function might be:

$$U(Q_1, Q_2) = \sqrt{Q_1 * Q_2}$$

Now, for any bundle of goods A and B, we can calculate the utility level of the bundles. The table below has a few different calculations.

Q_1	Q_2	$U(Q_1, Q_2)$
9	16	12
13	13	13
12	12	12
8	18	12

Assume that the quantities in the bundles are given – then to find the utility level just plug in the quantities and calculate. You should notice that the bundles $(9, 16)$, $(12, 12)$, and $(8, 18)$ would all lie on the same indifference curve because they all have a utility level of 12. However, the bundle $(13, 13)$ would lie on a higher indifference curve because it has a utility level of 13. Note that this conforms to the more is better property because the bundle $(13, 13)$ has more of both goods than the bundle $(12, 12)$ so the consumer must prefer the bundle $(13, 13)$.

2.1 Where indifference curves come from

Indifference curves can be derived directly from utility functions. In order to do this, however, we need to use three-dimensions. Stand in the corner of a room, facing outward diagonally. Let the floor along one of the walls be the axis for the quantity of Good 1 and let the floor along the other wall be the axis for the quantity of Good 2. The crease where the walls meet is the level of utility. We can now plot the utility function since we have three dimensions. It would essentially look like a cave that starts from the origin and keeps expanding outward. Alternatively, you could think about cutting a cone into two symmetric halves. If you lay one half of the cone down it (almost) looks like what we would call a utility shell. Figure 2.1 below actually graphs the function $U(Q_1, Q_2) = \sqrt{Q_1 * Q_2}$, although it is a little difficult to see because it is supposed to be 3-D.

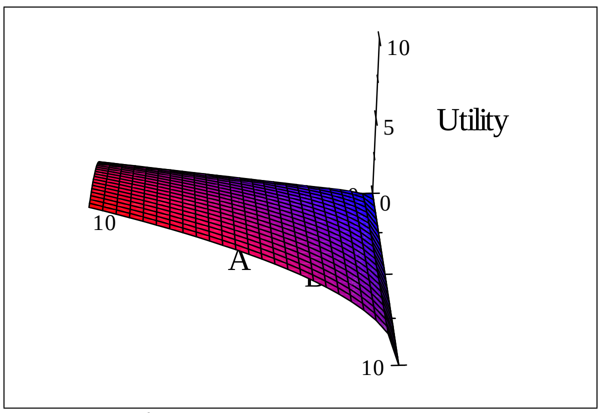

Figure 2.1: A three-dimensional graph of a utility hull with utility on the vertical axis and quantities of the goods on the other axes.

Now, suppose we pick a utility level, say 2.5, and make a nice even cut through the utility shell at 2.5. If we lay the new (now smaller) utility shell directly on the ground and trace around the bottom of the shell we will have our indifference curve for utility level 2.5. If we were to do the same at every utility level, then we would have the consumer's indifference map.

2.2 Marginal Utility

An important concept in consumer theory is marginal utility. Recall that marginal means additional – as in how much additional utility a person would get if he consumed one more unit of the good. We can define the marginal utility of Good 1 as:

$$MU_1 = \frac{\Delta U}{\Delta Q_1}$$

We can also define the marginal utility of Good 2 as:

$$MU_2 = \frac{\Delta U}{\Delta Q_2}$$

An interesting relationship then results if we find the ratio of marginal utilities:

$$\frac{MU_2}{MU_1} = \frac{\frac{\Delta U}{\Delta Q_2}}{\frac{\Delta U}{\Delta Q_1}}$$

Or:

$$\frac{MU_2}{MU_1} = \frac{\Delta Q_1}{\Delta Q_2}$$

Note that both of these changes in quantities are in the positive direction. Recall that:

$$MRS = \frac{-\Delta Q_1}{\Delta Q_2}$$

Now, if we multiply $\frac{MU_2}{MU_1}$ by (-1), we will get:

$$MRS = \frac{-MU_2}{MU_1}$$

Thus the Marginal Rate of Substitution is the negative of the ratio of marginal utilities of the goods. This result will prove useful when showing some results later.

3 Budget Constraints

We had a few goals when developing our consumer choice problem, one of which was to discuss how consumers choose the optimal bundle in a world where they have limited income. We will now discuss this concept of limited income. First, we will make a few assumptions about consumer behavior/attitude towards this limited income.

1. We begin with a fixed budget or endowment, denoted Y.
2. There is no borrowing allowed (thus, no loans or credit cards).
3. There is no saving allowed. Again, a saving-spending decision could be represented with indifference curves. We will, however, assume that all of your income must be spent now or it is lost forever. Essentially, there is no tomorrow, so no reason to save.
4. We will only consider decisions regarding 2 goods, although the analysis extends to n goods, where $n > 2$.
5. Assume that fractional amounts can be purchased. While this may not be true, it is an assumption that makes the math much, much easier.

3.1 Deriving a budget constraint

Whenever one derives a budget constraint it must be the case that we set expenditures equal to income (technically we need expenditures to be less than or equal to income). So we would have (assuming equality – which we will show will hold for the consumer who is behaving optimally):

$$Expenditures = Income$$

We know that our consumer's income is fixed at a level of Y. Suppose we have two Goods, 1 and 2. What are our expenditures on Goods 1 and 2? They are simply the price that we pay for the goods, P_1 and P_2 respectively, times the amount that we consume of those goods, Q_1 and Q_2 respectively (in this analysis it is implicitly assumed that the same price is paid for all units of the good).

So we can rewrite our budget constraint as:

$$P_1 * Q_1 + P_2 * Q_2 = Y$$

At this point you should note that the prices, P_1 and P_2, as well as the income are variables whose values are known to the consumer. What the budget constraint maps out is the different quantities of Goods 1 and 2 that the consumer can afford. Let's rewrite the budget constraint by solving for Q_1. We get:

$$Q_1 = \frac{Y}{P_1} - \frac{P_2}{P_1} Q_2$$

Notice that the budget constraint is in the form of an equation of a line, or $y = mx + b$ form (technically it's written as $y = b + mx$ above). Note that the y-intercept of the line is $\frac{Y}{P_1}$ and the slope of the line is $\left(-\frac{P_2}{P_A}\right)$. If we were given values for Y, P_1, and P_2 we could graph this line by labelling the y-axis as the quantity of Good 1 and the x-axis as the quantity of Good 2. Suppose that $Y = 50$, $P_2 = 1$, and $P_1 = 2$. Plugging in the numbers we get:

$$Q_1 = 25 - \frac{1}{2} Q_2$$

Figure 3.1 shows the plot of a budget constraint:

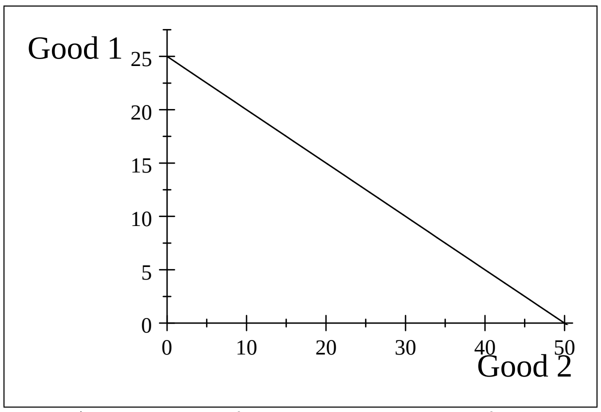

Figure 3.1: An example of a budget constraint for a consumer.

At this point it should be noted that the consumer can purchase any bundle on the budget constraint OR inside the budget constraint. Hopefully this is intuitive. If I can afford the bundle 26 units of Good 2 and 12 units of Good 1 (this bundle is on the budget constraint), then I can afford 13 units of Good 2 and 6 units of Good 1 (this bundle is inside the budget constraint). We can then define the consumer's opportunity set as the set of all the bundles that he can purchase given his income and the prices of the goods. This set is the entire triangle made by the x-axis, y-axis, and budget constraint.

3.2 Income changes and the budget constraint

Suppose that the consumer's income doubled – it is now \$100. It is assumed that prices remain the same. What will happens to the budget constraint?

We know that the generic formula for a budget constraint is:

$$Q_1 = \frac{Y}{P_1} - \frac{P_2}{P_1} Q_2$$

If only income changes, then only the y-intercept of the budget constraint is affected. The slope of the budget constraint remains the same because income does not enter the formula for the slope. If we plug the new income into the budget constraint formula, the new budget constraint is:

$$Q_A = 50 - \frac{1}{2} Q_B$$

Figure 3.2 shows the new budget constraint on the same graph as the old budget constraint:

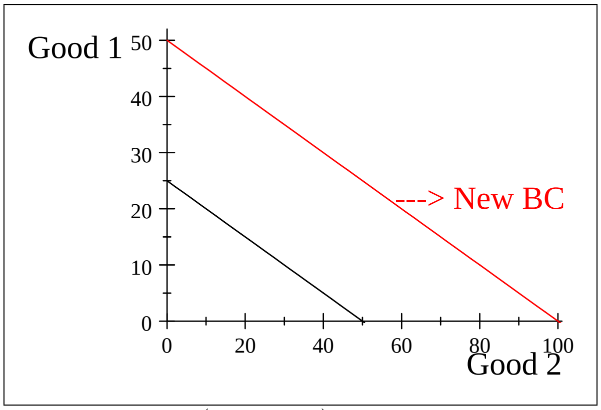

Figure 3.2: An original budget constraint (black line) and a new budget constraint (red line) after an increase in the consumer's income.

Because we had an increase in income the new budget constraint has made a parallel shift outward. This shift is reflected in the change in intercepts, the y-intercept increasing from 25 to 50 and the x-intercept increasing from 50 to 100. Notice that the consumer's opportunity set has increased as well.

3.3 Price changes and the budget constraint

Now, suppose that one of the prices change. Assume that income and the price of the other good remain constant. How does our budget constraint change?

3.3.1 Change in the price of Good 2

Suppose that we had a change in the price of Good 2. Looking at our generic formula for the budget constraint we see:

$$Q_1 = \frac{Y}{P_1} - \frac{P_2}{P_1} Q_2$$

The price of Good 2 only enters into the slope of the equation, so the y-intercept will remain the same. This should make sense, as the y-intercept tells us how much of Good 1 we can buy if we buy 0 of Good 2. Because neither income nor the price of Good 1 change we will still be able to buy exactly the same amount of Good 1 if we buy 0 of Good 2. Letting the price of Good 2 fall to 50 cents we have:

$$Q_1 = 25 - \frac{1}{4} Q_2$$

Figure 3.3.1 shows the original and new budget constraints:

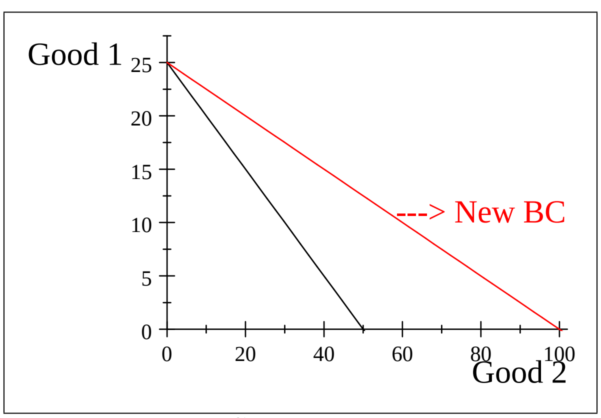

Figure 3.3.1: The effect of a price decrease in Good 2 on the consumer's budget constraint. The original budget constraint is in black; the new budget constraint is in red.

Because the price of Good 2 fell, we get a pivot effect on the budget constraint, as it swings out to the right (the red line is the new budget constraint). If the price of Good 2 rose, we would still get a pivot effect, although the budget constraint would swing in to the left.

3.3.2 Change in the price of Good 1

Because the price of Good 1 enters both the slope and y-intercept of our budget constraint we will see both of them change. However, the x-intercept will remain the same. What we will find is still a pivot effect on the budget constraint, only now the budget constraint pivots on the x-intercept.

Suppose the price of Good 1 increases to \$5. Our budget constraint is now (with the price of Good 2 being returned to it original \$1 level):

$$Q_1 = 10 - \frac{1}{5}Q_2$$

Figure 3.3.2 shows the original and new budget constraints:

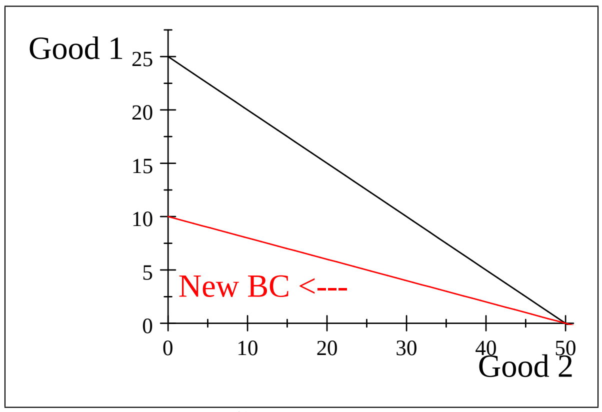

Figure 3.3.2: The effect of a price decrease in Good 2 on the consumer's budget constraint. The original budget constraint is in black; the new budget constraint is in red.

Notice that the increase in the price of Good 1 caused the budget constraint to swing inward (denoted by the red line). A decrease in the price of Good 1 would have caused the budget constraint to shift outward.

The key to both changes in the price of Good 2 and changes in the price of Good 1 is that the slope of the budget constraint changes when either changes. As we have already seen, slopes have been important in economic analysis.

3.4 Slope of the budget constraint

The slope of the budget constraint is given a specific name in economics. We call it the Marginal Rate of Transformation (MRT). The MRT tells us the rate at which the market will allow consumers to exchange goods. If the price of Good 1 is \$2 and the price of Good 2 is \$1, then the market says that if I give up purchasing one unit of Good 1 I can now purchase 2 additional units of Good 2. Mathematically then, the MRT is the $\frac{-\Delta Q_1}{\Delta Q_2}$, or how much of Good 1 I must give up in order to get more of Good 2. Note that the ΔQ_1 is negative, as we must give up some units of Good 1 to receive more units of Good 2.

You should also notice that $\frac{-\Delta Q_1}{\Delta Q_2}$ is a formula for a slope. Specifically, the MRT is the slope of the budget constraint, which is always the same at any point along the budget constraint because the budget constraint is a line. From our generic formula for the budget constraint we know that the slope is $\frac{-P_2}{P_1}$. So we now know that:

$$MRT = \frac{-P_2}{P_1}$$

This result is another important result that we will use in the next section on optimal consumer choice.

4 Optimal consumer choice

The easiest way to do this would be to set up the consumer's problem as a constrained optimization problem. We will solve the consumer's problem graphically first and then show some calculus based results.

There are two types of solutions we might find, an interior solution and a corner solution. It is easier to define a corner solution first. A corner solution occurs when a consumer buys either ONLY Good 1 or ONLY Good 2. Thus the optimal bundle (if it is a corner solution) will look like either $(0, Q_1)$ OR $(Q_2, 0)$, where Q_1 and Q_2 are both assumed to be greater than zero. At an interior solution the consumer will purchase positive quantities of both goods. We will first consider the interior solution and then the corner solution.

One important point before beginning. If the consumer is acting optimally, will she purchase a bundle inside, but not on, the budget constraint? The answer is no. The easy explanation is that if the consumer chooses to purchase a bundle inside the budget constraint then she is not spending all of his money. Essentially, she is throwing money into a lake (and we are assuming she gets no utility from throwing money into a lake or a wishing well), and why would anyone throw away money when they could get goods for it? Another explanation is that for any bundle inside the budget constraint that is being considered as the optimal bundle, a different bundle ON the budget constraint can be found that has more of BOTH goods. Thus the consumer can gain utility by moving to this bundle on the budget constraint because the more is better property of consumer preferences tells us that bundles with more of both goods are preferred to bundles without as much of both goods. So if the consumer is behaving optimally she will NOT choose a point inside the budget constraint.

4.1 Interior solution

As mentioned above, an interior solution to the consumer's problem is an optimal bundle at which the consumer purchases positive quantities of both goods. Look at Figure 4.1:

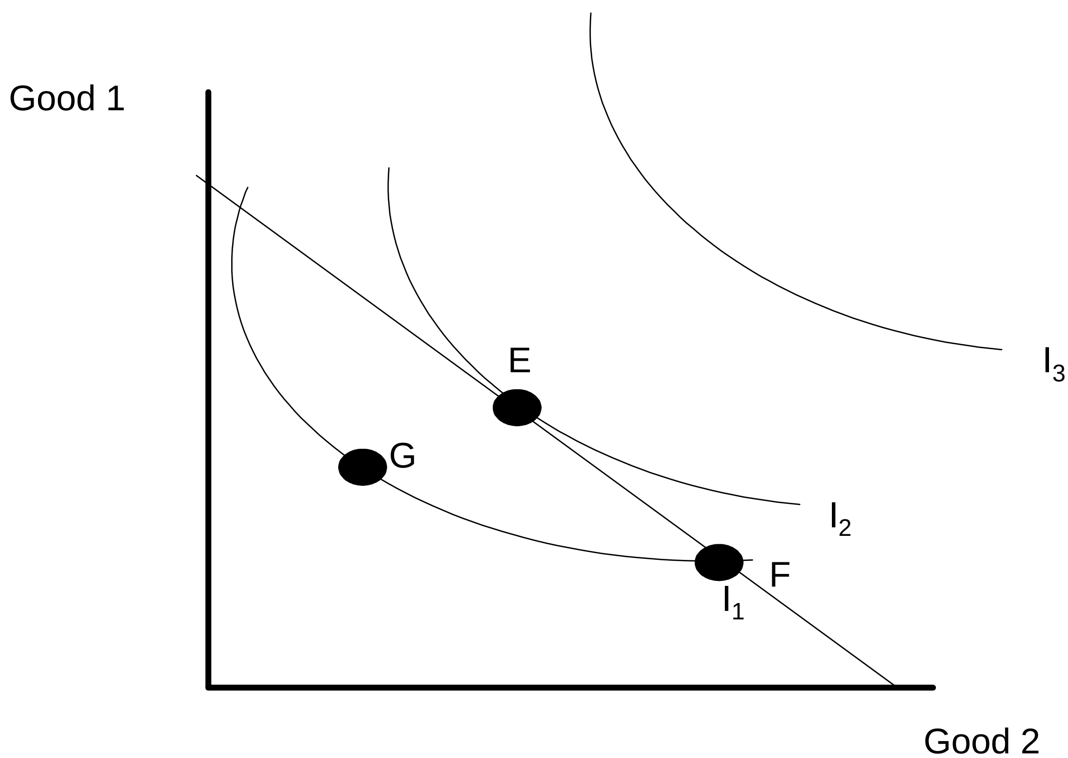

Figure 4.1: A consumer's optimal choice problem − interior solution.

We know that a consumer who is optimizing will pick a point along the budget constraint, which is the downward sloping straight line in the picture. There are 2 points labelled, E and F. Suppose the consumer chooses point F. Is he behaving in an optimal manner? That is, does he maximize his utility? A consumer maximizes his utility if he chooses a bundle such that there is no other bundle that he could have chosen, given his limited income, that would place him at a higher utility level (or on a higher indifference curve). Looking at bundle F, we notice that this consumer is indifferent between bundle F and bundle G. However, bundle G lies inside the budget constraint, so there must be an affordable bundle (call it bundle X) that he prefers to bundle G. If he prefers bundle X to bundle G, then he must prefer bundle X to bundle F. Thus, F cannot be the optimal bundle.

Now, look at bundle E. The indifference curve I_2 only touches the budget constraint once (it is tangent to the budget constraint). Note that there is no other bundle that the consumer can afford that would put him on a higher indifference curve. Thus, the optimal bundle is found by finding the indifference curve that is tangent to the budget constraint.

4.1.1 A key result for interior solutions

Recall that the slope of the budget constraint is the MRT. Also recall that the slope of the indifference curve is the MRS. A result from math class (I don't remember which one) is that if a line is tangent to a curve, then the slope of the line and the slope of the curve AT THE POINT OF TANGENCY are equal.[3] Thus, at the consumer's optimal bundle we have:

$$MRT = MRS$$

We know a few other things. We know that:

$$MRT = -\tfrac{P_2}{P_1}$$
$$MRS = \tfrac{-MU_2}{MU_1}$$

Substituting, we get:

$$-\frac{P_2}{P_1} = \frac{-MU_2}{MU_1}$$

Doing some rearranging gives us:

$$\frac{MU_1}{P_1} = \frac{MU_2}{P_2}$$

Notice what this equation tells us. At the optimal bundle, the marginal utility per dollar of each good must be the same. If it is not, the consumer can do better by shifting some dollars from the good with the lower $MU/\$$ to the good with the higher $MU/\$$. As an example, suppose that the consumer has \$7 and that he goes to Rio Bravo when they sell \$1 drinks and 10-cent wings. For simplicity, assume he must buy 10 wings at a time, so that he gets 10 wings for \$1. Now suppose that he purchases 60 wings and 1 drink. He gets through 10 wings and his 1 drink and thinks, "I would really like another drink to go with the other 50 wings that I have". Clearly he has not equated the marginal utilities per dollar for the two goods, otherwise he would not have thought this thought. In this case, if he could go back in time and reallocate his \$10 by making a different purchase, he would take some of the money spent on wings (which have a low MU because he has so many of them) and he would shift those funds to drinks (which have a high MU at the bundle (60 wings, 1 drink) because he only has one drink).

4.2 Deriving these results with calculus

Our consumer had utility function $U(Q_1, Q_2) = \sqrt{Q_1 * Q_2}$ and budget constraint $P_1 * Q_1 + P_2 * Q_2 = Y$. We can use calculus to derive the consumer's optimal choice as the consumer's goal is to maximize utility, and calculus is very useful in finding maximum points. However we cannot just take derivatives of the utility function because that function is always increasing (a little more Q_1 or Q_2 will increase utility) so the constraint plays an important role as it limits what can be purchased. Keep in mind that Y, P_1, and P_2 are constants and that the choice variables are Q_1 and Q_2 which are amounts of goods 1 and 2.

[3] As mentioned in the notes on supply and demand, economists typically care about either intersection points (as with the supply and demand model) or tangency points (as with the consumer choice model).

4.2.1 Direct substitution

How do we incorporate the budget constraint? One method would be to solve the budget constraint for either Q_1 or Q_2 and then substitute that result into the utility function. However, we can only use this method because we "know" that the budget constraint holds with equality because the consumer will spend all income. We would then have an equation with one variable and would be able to calculate the optimal choice of the remaining variable, conditional on prices and income, and then substitute that result back into the budget constraint to calculate the optimal choice of the variable for which we substituted. Earlier we solved the budget constraint for Q_1:

$$Q_1 = \frac{Y}{P_1} - \frac{P_2}{P_1}Q_2$$

Substituting into the utility function we now have:

$$\begin{aligned}
U\left(Q_1\left(Q_2\right), Q_2\right) &= \sqrt{\left(\frac{Y}{P_1} - \frac{P_2}{P_1}Q_2\right) * Q_2} \\
U\left(Q_1\left(Q_2\right), Q_2\right) &= \sqrt{\left(\frac{Y}{P_1}Q_2 - \frac{P_2}{P_1}Q_2^2\right)} \\
U\left(Q_1\left(Q_2\right), Q_2\right) &= \left(\frac{Y}{P_1}Q_2 - \frac{P_2}{P_1}Q_2^2\right)^{1/2}
\end{aligned}$$

Taking the derivative is possible but a bit of a mess and the results are not very intuitive. The result is (I took the easy route and used software):

$$\frac{dU}{dQ_2} = \frac{1}{2P_1}\frac{Y - 2P_2Q_2}{\sqrt{\frac{Q_2}{P_1}\left(Y - P_2Q_2\right)}}$$

Setting that equal to zero and solving for Q_2:

$$\begin{aligned}
\frac{dU}{dQ_2} &= \frac{1}{2P_1}\frac{Y - 2P_2Q_2}{\sqrt{\frac{Q_2}{P_1}\left(Y - P_2Q_2\right)}} \\
0 &= \frac{1}{2P_1}\frac{Y - 2P_2Q_2}{\sqrt{\frac{Q_2}{P_1}\left(Y - P_2Q_2\right)}} \\
0 &= Y - 2P_2Q_2 \\
2P_2Q_2 &= Y \\
Q_2 &= \frac{Y}{2P_2}
\end{aligned}$$

We can then substitute Q_2 back into the budget constraint and find Q_1:

$$\begin{aligned}
Q_1 &= \frac{Y}{P_1} - \frac{P_2}{P_1}Q_2 \\
Q_1 &= \frac{Y}{P_1} - \frac{P_2}{P_1}\left(\frac{Y}{2P_2}\right) \\
Q_1 &= \frac{Y}{P_1} - \frac{Y}{2P_1} \\
Q_1 &= \frac{2Y}{2P_1} - \frac{Y}{2P_1} \\
Q_1 &= \frac{Y}{2P_1}
\end{aligned}$$

For this problem, the consumer's optimal choice of (Q_1, Q_2) given Y, P_1, and P_2 is $\left(\frac{Y}{2P_1}, \frac{Y}{2P_2}\right)$. As long as P_1 and P_2 are greater than zero, given any Y, P_1, and P_2 we know exactly what this consumer would purchase.

4.2.2 Lagrangian method

Direct substitution solves the problem at hand but does not provide a lot of intuition because while the solution satisfies the main result, $\frac{MU_1}{P_1} = \frac{MU_2}{P_2}$, it is not obvious how that result is satisfied. An alternative method involves the Lagrangian approach. We set up a function called the Lagrangian which is just the utility function plus the constraint weighted by some additional variable which we will call λ.

$$\pounds(Q_1, Q_2, \lambda) = U(Q_1, Q_2) + \lambda(Y - Q_1P_1 - Q_2P_2)$$

Notice that there are now three choice variables: Q_1, Q_2, and λ. The variable λ is essentially the value of relaxing the constraint – what if the consumer did not spend all income or what is the marginal value of income? In order to ensure that $\pounds(Q_1, Q_2, \lambda) = U(Q_1, Q_2)$, so that we are maximizing the "same" function, we need $\lambda(Y - Q_1P_1 - Q_2P_2) = 0$, which means either $\lambda = 0$ or $Y - Q_1P_1 - Q_2P_2 = 0$.[4] We know the latter is true because the consumer spends all income and, if λ represents the marginal value of income, then λ should be greater than zero because a small amount of extra income will always provide additional utility. With the three choice variables, that means that three partial derivatives (one for Q_1, one for Q_2, and one for λ) will need to be taken. While that sounds worse than taking one derivative like we did in the direct substitution method, these derivatives will be less involved. Those three derivatives will then need to be set equal to zero and the system of three equations and three unknowns will need to be solved.

$$\begin{aligned}
\pounds(Q_1, Q_2, \lambda) &= U(Q_1, Q_2) + \lambda(Y - Q_1P_1 - Q_2P_2) \\
\pounds(Q_1, Q_2, \lambda) &= \sqrt{Q_1Q_2} + \lambda(Y - Q_1P_1 - Q_2P_2) \\
\pounds(Q_1, Q_2, \lambda) &= Q_1^{1/2}Q_2^{1/2} + \lambda(Y - Q_1P_1 - Q_2P_2) \\
\frac{\partial \pounds}{\partial Q_1} &= \frac{1}{2}Q_1^{-1/2}Q_2^{1/2} - \lambda P_1 \\
\frac{\partial \pounds}{\partial Q_2} &= \frac{1}{2}Q_2^{-1/2}Q_1^{1/2} - \lambda P_2 \\
\frac{\partial \pounds}{\partial \lambda} &= Y - Q_1P_1 - Q_2P_2
\end{aligned}$$

Note that the partial derivative with respect to λ is just the budget constraint. There is an additional important result from here. We have discussed the marginal utility with respect to a particular good – that is just the derivative of the utility function (the utility function, not the Lagrangian) with respect to that good. This consumer's utility function is $U(Q_1, Q_2) = Q_1^{1/2}Q_2^{1/2}$. The first part of $\frac{\partial \pounds}{\partial Q_1}$ is $\frac{1}{2}Q_1^{-1/2}Q_2^{1/2}$, which is the partial derivative of the consumer's utility function with respect to Q_1. So $\frac{1}{2}Q_1^{-1/2}Q_2^{1/2} = MU_1$. Likewise, $\frac{1}{2}Q_2^{-1/2}Q_1^{1/2} = MU_2$.

We need to select a variable to remove – we do not really care about the value of λ, so we can use $\frac{\partial \pounds}{\partial Q_1}$ and $\frac{\partial \pounds}{\partial Q_2}$ to remove λ.

$$\begin{aligned}
\frac{\partial \pounds}{\partial Q_1} &= 0 \\
\frac{1}{2}Q_1^{-1/2}Q_2^{1/2} - \lambda P_1 &= 0 \\
\frac{1}{2}Q_1^{-1/2}Q_2^{1/2} &= \lambda P_1 \\
\frac{\frac{1}{2}Q_1^{-1/2}Q_2^{1/2}}{P_1} &= \lambda \\
& \quad \mathit{or} \\
\frac{MU_1}{P_1} &= \lambda
\end{aligned}$$

[4] It is possible for both $\lambda = 0$ and $Y - Q_1P_1 - Q_2P_2 = 0$ but that is a rare case in general and not plausible given how the consumer's optimization problem is structured.

We can use the same process for $\frac{\partial \pounds}{\partial Q_2}$.

$$\begin{aligned}
\frac{\partial \pounds}{\partial Q_2} &= 0 \\
\frac{1}{2}Q_2^{-1/2}Q_1^{1/2} - \lambda P_2 &= 0 \\
\frac{1}{2}Q_2^{-1/2}Q_1^{1/2} &= \lambda P_2 \\
\frac{\frac{1}{2}Q_2^{-1/2}Q_1^{1/2}}{P_2} &= \lambda \\
& \mathit{or} \\
\frac{MU_2}{P_2} &= \lambda
\end{aligned}$$

Because $\frac{MU_1}{P_1} = \lambda$ and $\frac{MU_2}{P_2} = \lambda$, we know that $\frac{MU_1}{P_1} = \frac{MU_2}{P_2}$, which is a key result we found using the graphical approach. Setting those two equal we have:

$$\frac{MU_1}{P_1} = \frac{MU_2}{P_2}$$

$$\frac{\frac{1}{2}Q_1^{-1/2}Q_2^{1/2}}{P_1} = \frac{\frac{1}{2}Q_2^{-1/2}Q_1^{1/2}}{P_2}$$

Multiplying both sides by 2 to remove the $\frac{1}{2}$

$$\frac{Q_1^{-1/2}Q_2^{1/2}}{P_1} = \frac{Q_2^{-1/2}Q_1^{1/2}}{P_2}$$

Using the rule of exponents to put terms with negative exponents in the denominator

$$\frac{Q_2^{1/2}}{Q_1^{1/2}P_1} = \frac{Q_1^{1/2}}{Q_2^{1/2}P_2}$$

Multiplying through by $Q_1^{1/2}Q_2^{1/2}$

$$\left(Q_1^{1/2}Q_2^{1/2}\right)\frac{Q_2^{1/2}}{Q_1^{1/2}P_1} = \left(Q_1^{1/2}Q_2^{1/2}\right)\frac{Q_1^{1/2}}{Q_2^{1/2}P_2}$$

Canceling terms appropriately

$$\left(Q_2^{1/2}\right)\frac{Q_2^{1/2}}{P_1} = \left(Q_1^{1/2}\right)\frac{Q_1^{1/2}}{P_2}$$

Combining exponents of like terms

$$\frac{Q_2}{P_1} = \frac{Q_1}{P_2}$$

Solving for Q_2

$$Q_2 = \frac{P_1 Q_1}{P_2}$$

Now substituting $Q_2 = \frac{P_1 Q_1}{P_2}$ into the budget constraint:

$$\begin{aligned}
Y &= Q_1 P_1 + Q_2 P_2 \\
Y &= Q_1 P_1 + \left(\frac{P_1 Q_1}{P_2}\right) P_2 \\
Y &= Q_1 P_1 + P_1 Q_1 \\
Y &= 2 P_1 Q_1 \\
\frac{Y}{2P_1} &= Q_1
\end{aligned}$$

Substituting $\frac{Y}{2P_1} = Q_1$ into $Q_2 = \frac{P_1 Q_1}{P_2}$ gives:

$$\begin{aligned}
Q_2 &= \frac{P_1 Q_1}{P_2} \\
Q_2 &= \frac{P_1 \left(\frac{Y}{2P_1}\right)}{P_2} \\
Q_2 &= \frac{\left(\frac{Y}{2}\right)}{P_2} \\
Q_2 &= \frac{Y}{2} \div P_2 \\
Q_2 &= \frac{Y}{2} * \frac{1}{P_2} \\
Q_2 &= \frac{Y}{2P_2}
\end{aligned}$$

So again, we have that the optimal consumer solution to this problem is $(Q_1, Q_2) = \left(\frac{Y}{2P_1}, \frac{Y}{2P_2}\right)$. The answer is the same as with the direct substitution method, and while there are many more steps to the Lagrangian method, it provides some additional intuition and the derivatives are much easier to take.

5 Tying this model to markets and policy

5.1 Demand curves

We should be able to connect an individual's optimal bundle purchase to her demand curve. Recall that when creating a demand curve we hold everything in the world fixed except the price and quantity demanded of the good. In our two-good "world" with a fixed budget, "everything else" consists of the price of the other good and the individual's income, so we need to hold those two factors constant when deriving an individual's demand curve for a product.

Fix the individual's income and the price of Good 1 and vary the price of Good 2. From our earlier discussion we know that by varying the price of Good 2 will cause the budget constraint to pivot inward or outward. As the price changes, the consumer purchases a new optimal bundle that has a new quantity of Good 2. Using the prices and the new quantities we can map out a demand curve for Good 2 because all that is required to create a demand curve is the price of the good and quantity.

Another relationship between the consumer choice problem and demand curves can be seen in the Law of Diminishing Marginal Utility, which essentially states that the more an individual has of a good, the less additional benefit that individual receives from another unit of the same good. If an individual has 12 apples and receives a 13^{th} apple, the additional benefit from the 13^{th} apple is less than the additional benefit the individual received from the 12^{th} apple. The individual still receives positive utility from the 13^{th} apple (it is a good), just not as much as from the 12^{th} apple. Earlier I mentioned that when creating supply and demand graphs it is best to focus on a specific good and not a broad class of goods. If one creates a supply and demand graph of "shoes" that market will contain many different items; not just specific brands and styles, but even more general categories like running shoes, work shoes, etc. It is very likely that many people own different pairs of shoes for different occasions, and the law of diminishing marginal utility need not hold for such a broad category of shoes. It is possible to have a pair of work shoes and then purchase a pair of running shoes and receive more additional utility from the running shoes (which were a later purchase), which would seem to violate the Law of Diminishing Marginal Utility when considering the broad category of "shoes" but not when considering the more specific categories of "running shoes" or "work shoes."

5.2 Policy decisions

You may ask what this model has to do with policy decisions. Some policy decisions require examining tradeoffs individuals might make, not necessarily in purchasing two goods, but in other decisions that provide utility to the consumer. For instance, we can use this model of consumer choice to examine labor-leisure decisions, which we will do when we discuss wages and labor policy. As a policy scholar or maker, one should be thinking about the tradeoffs individuals will make when a policy is enacted and this model provides a way to think through those decisions.

5.3 Criticisms

One criticism is that utility functions are unknown to the observer, and, at times, it might be the case that the individual does not know the particular utility of an item/activity if the individual has not previously experienced the item/activity. Another is that utility functions are ordinal relationships, meaning that the utility numbers associated with different bundles only provide the rank ordering of the bundles, not a specific number. That feature of utility functions being ordinal relationships makes comparisons between individuals difficult. Utility theory is also very difficult to disprove; if one finds a violation, it is usually possible to reframe the problem to include an additional element (perhaps time or some other good) such that there is no longer a violation.

Finally, there is the belief that, because individuals in the consumer choice model are maximizing their own utility, that they are selfish. For the examples that we are using it would appear to be that way – how much of Goods 1 and 2 can the individual consume? However, an individual's utility function is not limited to containing just goods; it can also contain utility functions of other individuals (partners, children, parents, siblings, friends, etc.) However, the underlying mathematics when multiple individuals have utility functions that incorporate the utility functions of other individuals gets messy. Ultimately, the same basic relationship holds at the optimal bundle, where $\frac{MU_1}{P_1} = \frac{MU_2}{P_2}$, only the marginal utility is a lot more complicated.