Intermediate Microeconomics: Utility

Simon Naitram, PhD

October 4, 2020

Chapter 4: Utility

Read Sections 4.0 to 4.5. Understand what cardinal utility is (4.1), but no need to focus on it.

Now that we've described preferences, how about we put it into a format that's easier to manage? The aim of this chapter is to convert our loose statements about preferences to a more concrete format. This way, we're not having to describe every single preference relation individually, we're creating a coherent set of preferences that can be represented mathematically. Most importantly, we create a simple way of *ranking preferences*.

The main tool we use is a slightly abstract one—utility.[1] Once upon a time, utility was taken to be a measure of happiness. Now we talk about utility solely as a way of describing people's preferences numerically. Much like when we were talking about preferences, all that matters is if a person values one bundle more than another (*ordinal preferences*).

[1] Here's where some philosophy comes in. Utilitarianism stems from philosophy: https://www.britannica.com/topic/utilitarianism-philosophy.

Proposition 1. Utility function. A preference relation that is complete, reflexive, transitive and continuous can be represented by a continuous utility function.

Figure 1: A continuous function

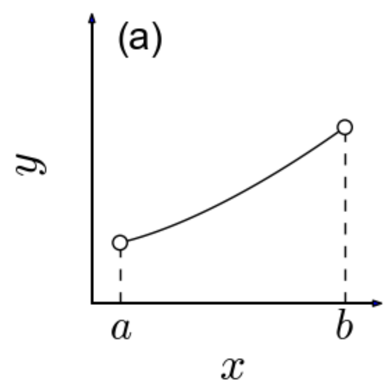

Continuity. This is another important mathematical concept that allows us to use powerful mathematical tools. Continuity is a key assumption of well-behaved functions. Take a function $y = f(x)$. If the function $f(.)$ is continuous, it means that a small change in x results in a small change in y. For a function that is discontinuous, a small change in x would result in a large change in y. A continuous function is one where you can draw it without lifting the pencil off the page. If a function is not continuous, we cannot be sure that the derivative of the function exists everywhere. For example, we cannot differentiate the function in panel (b) at $x = c$.

Figure 2: A discontinuous function

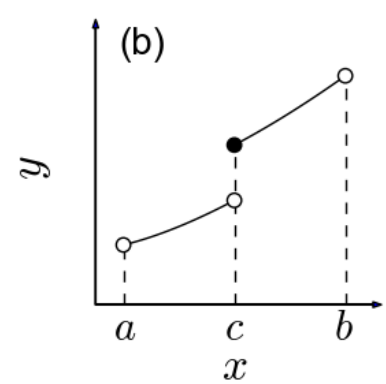

Functions. Remember what a function is. It maps some value x, using a function $U(.)$ to some new value. We will therefore write the general and (for-now) undefined utility function as $U(x)$. How does this simple function represent the preference relations we described before? By following some simple rules. A utility function represents set of preferences if when:[2]

[2] Notice we're moving away from curly lines. These are preference relations symbols. We now use simple less-than, more-than, and equal signs; because we've transformed our preference relations into numerical form.

- $x \succ y$, then $U(x) > U(y)$
- $y \succ x$, then $U(y) > U(x)$
- $x \sim y$, then $U(x) = U(y)$

Ordinal Utility We focus mainly on **ordinal utility functions**. In this case, we don't care how much bigger the number is. So if $U(x) = 6$ and $U(y) = 3$, then we know that $U(x) > U(y)$ and therefore that $x \succ y$—the consumer prefers x to y. Or we can say that they obtain greater utility from x than from y. What we can't say with ordinal utility is that the consumer prefers x twice as much as y. Now you can do this with a cardinal utility function, but we don't focus on this type because it's simply very difficult to justify what it means.

Figure 3: $u(x, y) = xy$

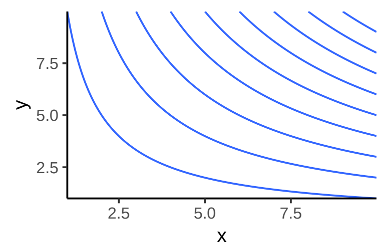

From Indifference Curve to Utility Let's relate the utility function to indifference curves. If all bundles on an indifference curve have the same level of preference, then they must have the same utility value. So if bundles $x = (x_1, x_2)$ and $y = (y_1, y_2)$ are on the same indifference curve, then

$$x \sim y \quad \text{implies} \quad U(x) = U(y). \tag{1}$$

Figure 4: $u(x, y) = ln(x) + y$

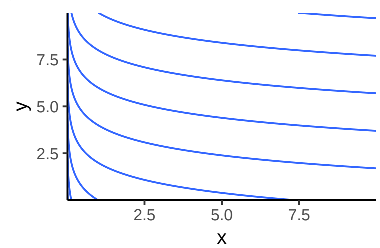

All bundles on an indifference curve give exactly the same utility. It leads to the simple implication that higher indifference curves (further to the right) have higher utility levels. This means that the idea of a utility function is simply a method of assigning numerical values to preference orderings. It means that we can also draw indifference curves using a utility function. If you're given a utility function $u(x_1, x_2)$, all you have to do is pick a constant value, and plot all the possible values of x_1 and x_2 for which the utility function is equal to that constant.

Figure 5: $u(x, y) = x^{0.5}y^{0.5}$

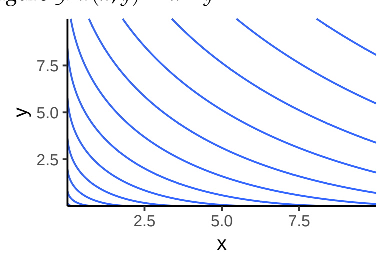

No Uniqueness Importantly, there is no single unique utility function that represents any given preference relation. Given a set of preferences represented by a utility function $U(x)$. The function can be transformed using what we call a monotonic transformation. A monotonic transformation is multiplication by a function that is always increasing. For example, mutliplying the utility function by 2 is a monotonic transformation. Crucially, a montonic transformation preserves the ordering of preferences. So if $U(x) = x^2$, then $f(U(x))$, where $f(U) = U(x) \times 2$ gives us $f(U(x)) = 2x^2$. Under this transformation, the preferences remain exactly the same. Since any positive monotonic transformation preserves the ordering of utility by preference and so represents the same preferences, then either utility function $f(x)$ or $U(x)$ represents those preferences.

Satiation We can describe a good as being a good or a bad once again. In this case, however, it gets even simpler. Our utility function of a single commodity simply traces a single line. Plotted as $u = U(x)$, the utility function can either be increasing in x (a good), or decreasing in x (a bad). A single good is frequently both a good and a bad, divided by the point of satiation.

Figure 6: Satiation in utility

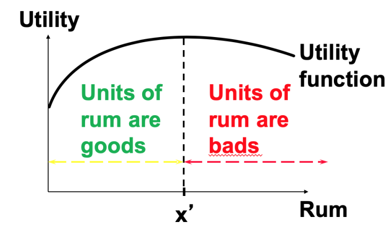

Marginal Utility If we give a consumer a little bit more of good 1, how much does his utility increase by? Once again, we're talking about the infinitesimally small changes that allow us to consider a derivative. The slope of the utility function with respect to good 1 gives us the marginal utility. As you can see in the chart above, the marginal utility is positive at first, but slowly flattens to the point of becoming negative past the point of satiation. For a utility function $u(x_1, x_2)$, we can say that:

$$\text{MU}_1 = \frac{\partial u(x_1, x_2)}{\partial x_1} \tag{2}$$

$$\text{MU}_2 = \frac{\partial u(x_1, x_2)}{\partial x_2} \tag{3}$$

Note that these are partial derivatives, so they examine the effect of a change in good i on utility holding all other goods constant. Since utility itself has no meaning numerically, then the marginal utility itself has no meaning either. Remember that we can multiply the utility function by 2,000 and still maintain the same preference ordering. The marginal utility would also increase by 2,000. So how do we scale or re-interpret the slope of this line to have some meaningful purpose?

Marginal Rate of Substitution Let's go back to the interpretation of the marginal rate of substitution. It asks: if you have to give up one unit of Good 1, how much of Good 2 do I need to give you in order to leave you at the same level of preferences? Remember the marginal utility is the change in utility for a small change in a commodity. Let's say the change in Good 1 is dx_1, and the change in Good 2 is dx_2. Then we can express the total change in utility as:

$$\frac{\partial u(x_1, x_2)}{\partial x_1} dx_1 + \frac{\partial u(x_1, x_2)}{\partial x_2} dx_2 = 0. \tag{4}$$

Note the reason this is zero is because we're investigating a change in Good 1 and Good 2 that leaves us with exactly the same utility[3] Since we already know the slope of the indifference curve is the marginal rate of substitution, and the slope can be expressed as dx_2/dx_1, then we can solve:

$$\text{MRS} = \frac{dx_2}{dx_1} = -\frac{\partial u/\partial x_1}{\partial u/\partial x_2} = -\frac{\text{MU}_1}{\text{MU}_2}. \tag{5}$$

[3] The same utility,the same indifference curve, the same level of preferences—all of these effectively mean the same thing: the consumer is indifferent between the original bundle and the new one with the changes.

That is, the marginal rate of substitution between Good 1 and Good 2 is equal to the negative ratio of the marginal utilities of the two goods. Since we're dealing with ratios, the fact that the magnitudes of the marginal utilities are meaningless is okay—all we're interested in is which marginal utility is bigger! That helps us to think of it in a different way: for tiny changes in both goods, which change gives the consumer the greater increase in utility? If it's Good 1, then the indifference curve is slanted down more steeply than a 45 deg line. This says that the consumer would need more than 1 unit of Good 2 to compensate them for a 1 unit decrease in Good 1. If it's Good 2 that gives the larger increase, then the indifference curve is slanted down less steeply than a 45 deg line. This means that the consumer would need more than 1 unit of Good 1 to compensate them for a 1 unit decrease in Good 2.

Usefulness Utility is a really fascinating concept—philosphically, intuitively, and technically. It allows us to do really interesting stuff with people's preferences, and helps us to understand it in an algebraic way. However, there will always be a discussion of whether utility actually measures anything interesting. But Charles Jones and Peter Klenow[4], in 2016, used a form of utility that included all sorts of things. Their utility function included data on consumption, leisure, inequality, and mortality. They calculated these measure for 152 countries, including Barbados. They find quite interesting things. For example: Western European living standards are closer to U.S. living standards (85% of welfare) than income suggests (67% of income) due to Europe's longer life expectancy, greater leisure time, and lower inequality. For us, the message is that utility can be used to measure the benefit derived from a range of activities, not just from consuming rum or cake. Utility is a tool that helps us to measure how well-off an individual or society is.

[4] https://pubs.aeaweb.org/doi/pdfplus/10.1257/aer.20110236

The Quasi-Linear Utility Function Traditonally the Cobb-Douglas utility function is the most commonly used, and is the utility function most textbooks focus on. However, in applied work, the quasi-linear utility function has become extremely useful. It takes us back to the idea of the numeraire. The commodity which is linear in utility is the numeraire, and the commodity which is non-linear is the one we're interested in focusing on. All it says is that 'money' as the broad numeraire always affects the consumer in exactly the same way. Again, we can think about it as: "Should I spend my money on rum, or keep it for all the other things I need to spend on?"

Panel A. Welfare and income are highly correlated at 0.96

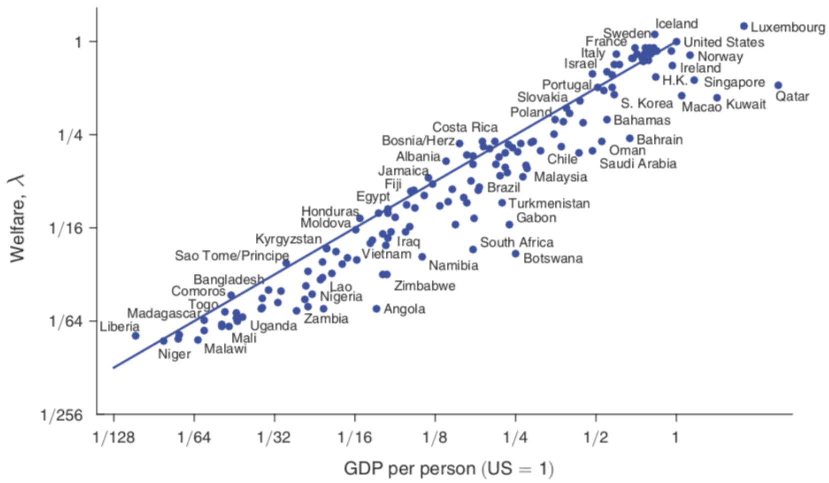

Panel B. But this masks substantial variation in the ratio of λ to GDP per capita.
The mean absolute deviation from unity is about 27%

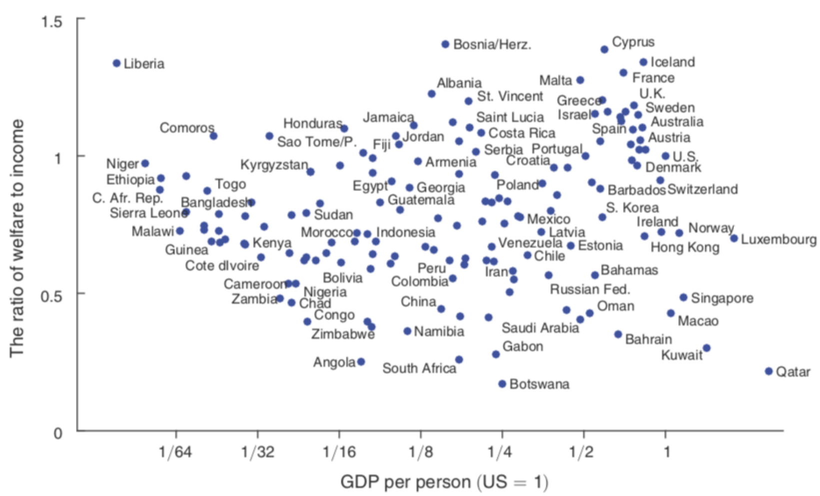

Rationality is the main assumption we make about people's preferences. We assume that people have rational preferences. Rationality says that people choose the best option available to them. You wouldn't choose \$5 over \$100 if there were no strings attached, correct? Well this is the main *behavioural* assumption we make in economics.

Figure 7: The increase in the use of the term "quasi-linear" in public economics. From Henrik Kleven

.