CHAPTER 4

UTILITY

In Victorian days, philosophers and economists talked blithely of "utility" as an indicator of a person's overall well-being. Utility was thought of as a numeric measure of a person's happiness. Given this idea, it was natural to think of consumers making choices so as to maximize their utility, that is, to make themselves as happy as possible.

The trouble is that these classical economists never really described how we were to measure utility. How are we supposed to quantify the "amount" of utility associated with different choices? Is one person's utility the same as another's? What would it mean to say that an extra candy bar would give me twice as much utility as an extra carrot? Does the concept of utility have any independent meaning other than its being what people maximize?

Because of these conceptual problems, economists have abandoned the old-fashioned view of utility as being a measure of happiness. Instead, the theory of consumer behavior has been reformulated entirely in terms of **consumer preferences**, and utility is seen only as a *way to describe preferences.*

Economists gradually came to recognize that all that mattered about utility as far as choice behavior was concerned was whether one bundle had a higher utility than another—how much higher didn't really matter.

Originally, preferences were defined in terms of utility: to say a bundle (x_1, x_2) was preferred to a bundle (y_1, y_2) meant that the x-bundle had a higher utility than the y-bundle. But now we tend to think of things the other way around. The *preferences* of the consumer are the fundamental description useful for analyzing choice, and utility is simply a way of describing preferences.

A **utility function** is a way of assigning a number to every possible consumption bundle such that more-preferred bundles get assigned larger numbers than less-preferred bundles. That is, a bundle (x_1, x_2) is preferred to a bundle (y_1, y_2) if and only if the utility of (x_1, x_2) is larger than the utility of (y_1, y_2): in symbols, $(x_1, x_2) \succ (y_1, y_2)$ if and only if $u(x_1, x_2) > u(y_1, y_2)$.

The only property of a utility assignment that is important is how it *orders* the bundles of goods. The magnitude of the utility function is only important insofar as it *ranks* the different consumption bundles; the size of the utility difference between any two consumption bundles doesn't matter. Because of this emphasis on ordering bundles of goods, this kind of utility is referred to as **ordinal utility**.

Consider for example Table 4.1, where we have illustrated several different ways of assigning utilities to three bundles of goods, all of which order the bundles in the same way. In this example, the consumer prefers A to B and B to C. All of the ways indicated are valid utility functions that describe the same preferences because they all have the property that A is assigned a higher number than B, which in turn is assigned a higher number than C.

Different ways to assign utilities.

Bundle	U_1	U_2	U_3
A	3	17	-1
B	2	10	-2
C	1	.002	-3

Since only the ranking of the bundles matters, there can be no unique way to assign utilities to bundles of goods. If we can find one way to assign utility numbers to bundles of goods, we can find an infinite number of ways to do it. If $u(x_1, x_2)$ represents a way to assign utility numbers to the bundles (x_1, x_2), then multiplying $u(x_1, x_2)$ by 2 (or any other positive number) is just as good a way to assign utilities.

Multiplication by 2 is an example of a **monotonic transformation**. A monotonic transformation is a way of transforming one set of numbers into another set of numbers in a way that preserves the order of the numbers.

We typically represent a monotonic transformation by a function $f(u)$ that transforms each number u into some other number $f(u)$, in a way that preserves the order of the numbers in the sense that $u_1 > u_2$ implies $f(u_1) > f(u_2)$. A monotonic transformation and a monotonic function are essentially the same thing.

Examples of monotonic transformations are multiplication by a positive number (e.g., $f(u) = 3u$), adding any number (e.g., $f(u) = u + 17$), raising u to an odd power (e.g., $f(u) = u^3$), and so on.[1]

The rate of change of $f(u)$ as u changes can be measured by looking at the change in f between two values of u, divided by the change in u:

$$\frac{\Delta f}{\Delta u} = \frac{f(u_2) - f(u_1)}{u_2 - u_1}.$$

For a monotonic transformation, $f(u_2) - f(u_1)$ always has the same sign as $u_2 - u_1$. Thus a monotonic function always has a positive rate of change. This means that the graph of a monotonic function will always have a positive slope, as depicted in Figure 4.1A.

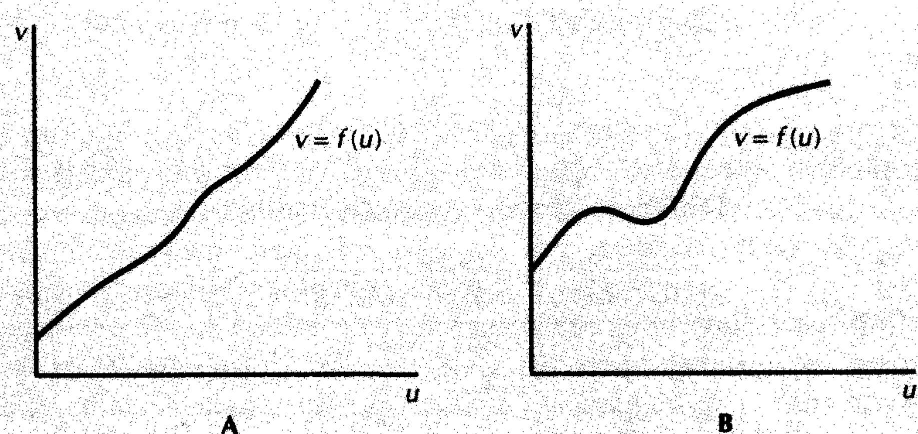

A positive monotonic transformation. Panel A illustrates a monotonic function—one that is always increasing. Panel B illustrates a function that is *not* monotonic, since it sometimes increases and sometimes decreases.

[1] What we are calling a "monotonic transformation" is, strictly speaking, called a "positive monotonic transformation," in order to distinguish it from a "negative monotonic transformation," which is one that *reverses* the order of the numbers. Monotonic transformations are sometimes called "monotonous transformations," which seems unfair, since they can actually be quite interesting.

If $f(u)$ is *any* monotonic transformation of a utility function that represents some particular preferences, then $f(u(x_1, x_2))$ is also a utility function that represents those same preferences.

Why? The argument is given in the following three statements:

1. To say that $u(x_1, x_2)$ represents some particular preferences means that $u(x_1, x_2) > u(y_1, y_2)$ if and only if $(x_1, x_2) \succ (y_1, y_2)$.
2. But if $f(u)$ is a monotonic transformation, then $u(x_1, x_2) > u(y_1, y_2)$ if and only if $f(u(x_1, x_2)) > f(u(y_1, y_2))$.
3. Therefore, $f(u(x_1, x_2)) > f(u(y_1, y_2))$ if and only if $(x_1, x_2) \succ (y_1, y_2)$, so the function $f(u)$ represents the preferences in the same way as the original utility function $u(x_1, x_2)$.

We summarize this discussion by stating the following principle: *a monotonic transformation of a utility function is a utility function that represents the same preferences as the original utility function.*

Geometrically, a utility function is a way to label indifference curves. Since every bundle on an indifference curve must have the same utility, a utility function is a way of assigning numbers to the different indifference curves in a way that higher indifference curves get assigned larger numbers. Seen from this point of view a monotonic transformation is just a relabeling of indifference curves. As long as indifference curves containing more-preferred bundles get a larger label than indifference curves containing less-preferred bundles, the labeling will represent the same preferences.

4.1 Cardinal Utility

There are some theories of utility that attach a significance to the magnitude of utility. These are known as **cardinal utility** theories. In a theory of cardinal utility, the size of the utility difference between two bundles of goods is supposed to have some sort of significance.

We know how to tell whether a given person prefers one bundle of goods to another: we simply offer him or her a choice between the two bundles and see which one is chosen. Thus we know how to assign an ordinal utility to the two bundles of goods: we just assign a higher utility to the chosen bundle than to the rejected bundle. Any assignment that does this will be a utility function. Thus we have an operational criterion for determining whether one bundle has a higher utility than another bundle for some individual.

But how do we tell if a person likes one bundle twice as much as another? How could you even tell if *you* like one bundle twice as much as another?

One could propose various definitions for this kind of assignment: I like one bundle twice as much as another if I am willing to pay twice as much for it. Or, I like one bundle twice as much as another if I am willing to run twice as far to get it, or to wait twice as long, or to gamble for it at twice the odds.

There is nothing wrong with any of these definitions; each one would give rise to a way of assigning utility levels in which the magnitude of the numbers assigned had some operational significance. But there isn't much right about them either. Although each of them is a possible interpretation of what it means to want one thing twice as much as another, none of them appears to be an especially compelling interpretation of that statement.

Even if we did find a way of assigning utility magnitudes that seemed to be especially compelling, what good would it do us in describing choice behavior? To tell whether one bundle or another will be chosen, we only have to know which is preferred—which has the larger utility. Knowing how much larger doesn't add anything to our description of choice. Since cardinal utility isn't needed to describe choice behavior and there is no compelling way to assign cardinal utilities anyway, we will stick with a purely ordinal utility framework.

4.2 Constructing a Utility Function

But are we assured that there is any way to assign ordinal utilities? Given a preference ordering can we always find a utility function that will order bundles of goods in the same way as those preferences? Is there a utility function that describes any reasonable preference ordering?

Not all kinds of preferences can be represented by a utility function. For example, suppose that someone had intransitive preferences so that $A \succ B \succ C \succ A$. Then a utility function for these preferences would have to consist of numbers $u(A)$, $u(B)$, and $u(C)$ such that $u(A) > u(B) > u(C) > u(A)$. But this is impossible.

However, if we rule out perverse cases like intransitive preferences, it turns out that we will typically be able to find a utility function to represent preferences. We will illustrate one construction here, and another one in Chapter 14.

Suppose that we are given an indifference map as in Figure 4.2. We know that a utility function is a way to label the indifference curves such that higher indifference curves get larger numbers. How can we do this?

One easy way is to draw the diagonal line illustrated and label each indifference curve with its distance from the origin measured along the line.

How do we know that this is a utility function? It is not hard to see that if preferences are monotonic then the line through the origin must intersect every indifference curve exactly once. Thus every bundle is getting a label, and those bundles on higher indifference curves are getting larger labels—and that's all it takes to be a utility function.

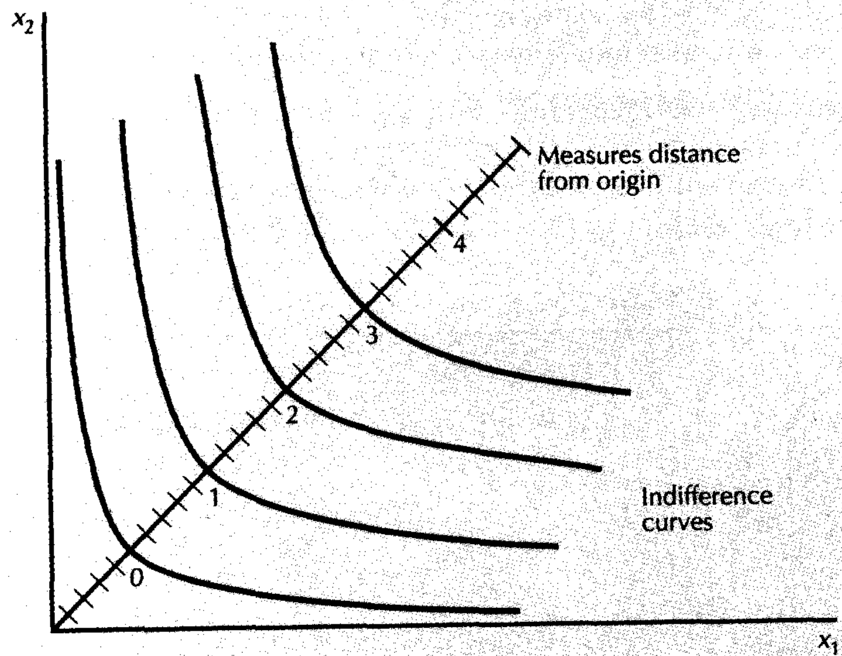

Constructing a utility function from indifference curves. Draw a diagonal line and label each indifference curve with how far it is from the origin measured along the line.

This gives us one way to find a labeling of indifference curves, at least as long as preferences are monotonic. This won't always be the most natural way in any given case, but at least it shows that the idea of an ordinal utility function is pretty general: nearly any kind of "reasonable" preferences can be represented by a utility function.

4.3 Some Examples of Utility Functions

In Chapter 3 we described some examples of preferences and the indifference curves that represented them. We can also represent these preferences by utility functions. If you are given a utility function, $u(x_1, x_2)$, it is relatively easy to draw the indifference curves: you just plot all the points (x_1, x_2) such that $u(x_1, x_2)$ equals a constant. In mathematics, the set of all (x_1, x_2) such that $u(x_1, x_2)$ equals a constant is called a **level set**. For each different value of the constant, you get a different indifference curve.

EXAMPLE: Indifference Curves from Utility

Suppose that the utility function is given by: $u(x_1, x_2) = x_1 x_2$. What do the indifference curves look like?

We know that a typical indifference curve is just the set of all x_1 and x_2 such that $k = x_1 x_2$ for some constant k. Solving for x_2 as a function of x_1, we see that a typical indifference curve has the formula:

$$x_2 = \frac{k}{x_1}.$$

This curve is depicted in Figure 4.3 for $k = 1, 2, 3 \cdots$.

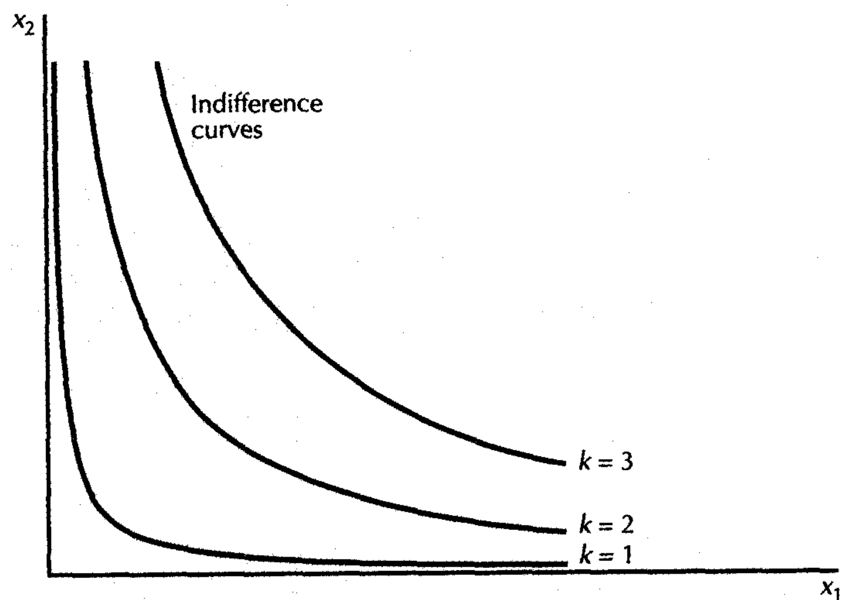

Indifference curves. The indifference curves $k = x_1 x_2$ for different values of k.

Let's consider another example. Suppose that we were given a utility function $v(x_1, x_2) = x_1^2 x_2^2$. What do its indifference curves look like? By the standard rules of algebra we know that:

$$v(x_1, x_2) = x_1^2 x_2^2 = (x_1 x_2)^2 = u(x_1, x_2)^2.$$

Thus the utility function $v(x_1, x_2)$ is just the square of the utility function $u(x_1, x_2)$. Since $u(x_1, x_2)$ cannot be negative, it follows that $v(x_1, x_2)$ is a monotonic transformation of the previous utility function, $u(x_1, x_2)$. This means that the utility function $v(x_1, x_2) = x_1^2 x_2^2$ has to have exactly the same shaped indifference curves as those depicted in Figure 4.3. The labeling of the indifference curves will be different—the labels that were $1, 2, 3, \cdots$ will now be $1, 4, 9, \cdots$—but the set of bundles that has $v(x_1, x_2) =$

9 is exactly the same as the set of bundles that has $u(x_1, x_2) = 3$. Thus $v(x_1, x_2)$ describes exactly the same preferences as $u(x_1, x_2)$ since it *orders* all of the bundles in the same way.

Going the other direction—finding a utility function that represents some indifference curves—is somewhat more difficult. There are two ways to proceed. The first way is mathematical. Given the indifference curves, we want to find a function that is constant along each indifference curve and that assigns higher values to higher indifference curves.

The second way is a bit more intuitive. Given a description of the preferences, we try to think about what the consumer is trying to maximize—what combination of the goods describes the choice behavior of the consumer. This may seem a little vague at the moment, but it will be more meaningful after we discuss a few examples.

Perfect Substitutes

Remember the red pencil and blue pencil example? All that mattered to the consumer was the total number of pencils. Thus it is natural to measure utility by the total number of pencils. Therefore we provisionally pick the utility function $u(x_1, x_2) = x_1 + x_2$. Does this work? Just ask two things: is this utility function constant along the indifference curves? Does it assign a higher label to more-preferred bundles? The answer to both questions is yes, so we have a utility function.

Of course, this isn't the only utility function that we could use. We could also use the *square* of the number of pencils. Thus the utility function $v(x_1, x_2) = (x_1 + x_2)^2 = x_1^2 + 2x_1x_2 + x_2^2$ will also represent the perfect-substitutes preferences, as would any other monotonic transformation of $u(x_1, x_2)$.

What if the consumer is willing to substitute good 1 for good 2 at a rate that is different from one-to-one? Suppose, for example, that the consumer would require *two* units of good 2 to compensate him for giving up one unit of good 1. This means that good 1 is *twice* as valuable to the consumer as good 2. The utility function therefore takes the form $u(x_1, x_2) = 2x_1 + x_2$. Note that this utility yields indifference curves with a slope of -2.

In general, preferences for perfect substitutes can be represented by a utility function of the form

$$u(x_1, x_2) = ax_1 + bx_2.$$

Here a and b are some positive numbers that measure the "value" of goods 1 and 2 to the consumer. Note that the slope of a typical indifference curve is given by $-a/b$.

Perfect Complements

This is the left shoe–right shoe case. In these preferences the consumer only cares about the number of *pairs* of shoes he has, so it is natural to choose the number of pairs of shoes as the utility function. The number of complete pairs of shoes that you have is the *minimum* of the number of right shoes you have, x_1, and the number of left shoes you have, x_2. Thus the utility function for perfect complements takes the form $u(x_1, x_2) = \min\{x_1, x_2\}$.

To verify that this utility function actually works, pick a bundle of goods such as $(10, 10)$. If we add one more unit of good 1 we get $(11, 10)$, which should leave us on the same indifference curve. Does it? Yes, since $\min\{10, 10\} = \min\{11, 10\} = 10$.

So $u(x_1, x_2) = \min\{x_1, x_2\}$ is a possible utility function to describe perfect complements. As usual, any monotonic transformation would be suitable as well.

What about the case where the consumer wants to consume the goods in some proportion other than one-to-one? For example, what about the consumer who always uses 2 teaspoons of sugar with each cup of tea? If x_1 is the number of cups of tea available and x_2 is the number of teaspoons of sugar available, then the number of correctly sweetened cups of tea will be $\min\{x_1, \frac{1}{2}x_2\}$.

This is a little tricky so we should stop to think about it. If the number of cups of tea is greater than half the number of teaspoons of sugar, then we know that we won't be able to put 2 teaspoons of sugar in each cup. In this case, we will only end up with $\frac{1}{2}x_2$ correctly sweetened cups of tea. (Substitute some numbers in for x_1 and x_2 to convince yourself.)

Of course, any monotonic transformation of this utility function will describe the same preferences. For example, we might want to multiply by 2 to get rid of the fraction. This gives us the utility function $u(x_1, x_2) = \min\{2x_1, x_2\}$.

In general, a utility function that describes perfect-complement preferences is given by

$$u(x_1, x_2) = \min\{ax_1, bx_2\},$$

where a and b are positive numbers that indicate the proportions in which the goods are consumed.

Quasilinear Preferences

Here's a shape of indifference curves that we haven't seen before. Suppose that a consumer has indifference curves that are vertical translates of one another, as in Figure 4.4. This means that all of the indifference curves are just vertically "shifted" versions of one indifference curve. It follows that the equation for an indifference curve takes the form $x_2 = k - v(x_1)$, where k is a different constant for each indifference curve. This equation says that the height of each indifference curve is some function of x_1, $-v(x_1)$, plus a constant k. Higher values of k give higher indifference curves. (The minus sign is only a convention; we'll see why it is convenient below.)

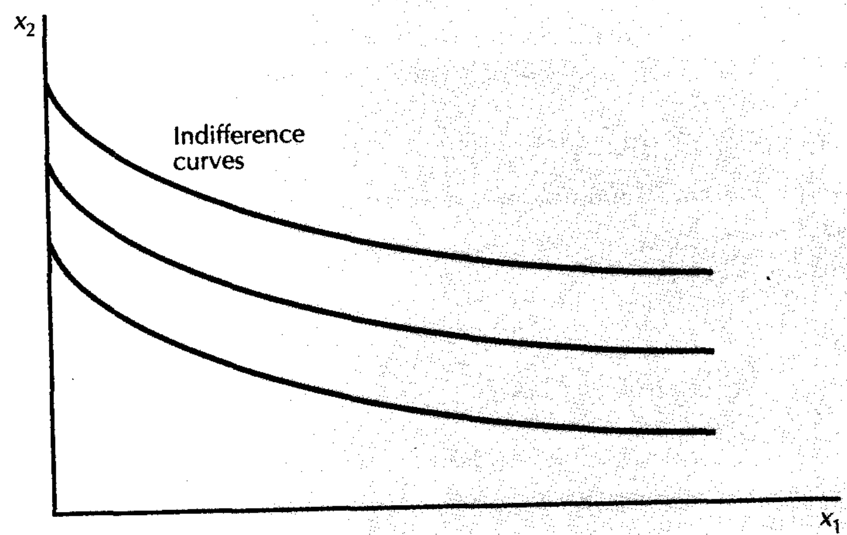

Quasilinear preferences. Each indifference curve is a vertically shifted version of a single indifference curve.

The natural way to label indifference curves here is with k—roughly speaking, the height of the indifference curve along the vertical axis. Solving for k and setting it equal to utility, we have

$$u(x_1, x_2) = k = v(x_1) + x_2.$$

In this case the utility function is linear in good 2, but (possibly) nonlinear in good 1; hence the name **quasilinear utility**, meaning "partly linear" utility. Specific examples of quasilinear utility would be $u(x_1, x_2) = \sqrt{x_1} + x_2$, or $u(x_1, x_2) = \ln x_1 + x_2$. Quasilinear utility functions are not particularly realistic, but they are very easy to work with, as we'll see in several examples later on in the book.

Cobb-Douglas Preferences

Another commonly used utility function is the **Cobb-Douglas** utility function

$$u(x_1, x_2) = x_1^c x_2^d,$$

where c and d are positive numbers that describe the preferences of the consumer.[2]

The Cobb-Douglas utility function will be useful in several examples. The preferences represented by the Cobb-Douglas utility function have the general shape depicted in Figure 4.5. In Figure 4.5A, we have illustrated the indifference curves for $c = 1/2$, $d = 1/2$. In Figure 4.5B, we have illustrated the indifference curves for $c = 1/5$, $d = 4/5$. Note how different values of the parameters c and d lead to different shapes of the indifference curves.

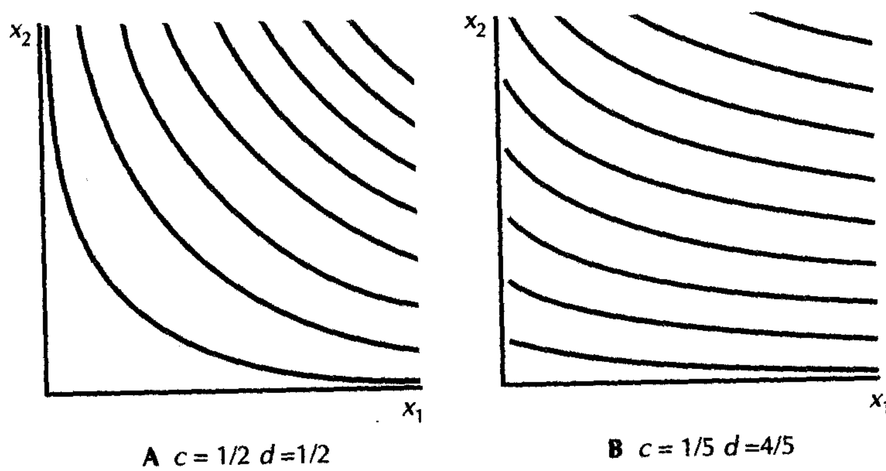

Cobb-Douglas indifference curves. Panel A shows the case where $c = 1/2$, $d = 1/2$ and panel B shows the case where $c = 1/5$, $d = 4/5$.

Cobb-Douglas indifference curves look just like the nice convex monotonic indifference curves that we referred to as "well-behaved indifference curves" in Chapter 3. Cobb-Douglas preferences are the standard example of indifference curves that look well-behaved, and in fact the formula describing them is about the simplest algebraic expression that generates well-behaved preferences. We'll find Cobb-Douglas preferences quite useful to present algebraic examples of the economic ideas we'll study later.

Of course a monotonic transformation of the Cobb-Douglas utility function will represent exactly the same preferences, and it is useful to see a couple of examples of these transformations.

2 Paul Douglas was a twentieth-century economist at the University of Chicago who later became a U.S. senator. Charles Cobb was a mathematician at Amherst College. The Cobb-Douglas functional form was originally used to study production behavior.

First, if we take the natural log of utility, the product of the terms will become a sum so that we have

$$v(x_1, x_2) = \ln(x_1^c x_2^d) = c \ln x_1 + d \ln x_2.$$

The indifference curves for this utility function will look just like the ones for the first Cobb-Douglas function, since the logarithm is a monotonic transformation. (For a brief review of natural logarithms, see the Mathematical Appendix at the end of the book.)

For the second example, suppose that we start with the Cobb-Douglas form

$$v(x_1, x_2) = x_1^c x_2^d.$$

Then raising utility to the $1/(c + d)$ power, we have

$$x_1^{\frac{c}{c+d}} x_2^{\frac{d}{c+d}}.$$

Now define a new number

$$a = \frac{c}{c+d}.$$

We can now write our utility function as

$$v(x_1, x_2) = x_1^a x_2^{1-a}.$$

This means that we can always take a monotonic transformation of the Cobb-Douglas utility function that make the exponents sum to 1. This will turn out to have a useful interpretation later on.

The Cobb-Douglas utility function can be expressed in a variety of ways; you should learn to recognize them, as this family of preferences is very useful for examples.

4.4 Marginal Utility

Consider a consumer who is consuming some bundle of goods, (x_1, x_2). How does this consumer's utility change as we give him or her a little more of good 1? This rate of change is called the **marginal utility** with respect to good 1. We write it as MU_1 and think of it as being a ratio,

$$MU_1 = \frac{\Delta U}{\Delta x_1} = \frac{u(x_1 + \Delta x_1, x_2) - u(x_1, x_2)}{\Delta x_1},$$

that measures the rate of change in utility (ΔU) associated with a small change in the amount of good 1 (Δx_1). Note that the amount of good 2 is held fixed in this calculation.[3]

[3] See the appendix to this chapter for a calculus treatment of marginal utility.

This definition implies that to calculate the change in utility associated with a small change in consumption of good 1, we can just multiply the change in consumption by the marginal utility of the good:

$$\Delta U = MU_1 \Delta x_1.$$

The marginal utility with respect to good 2 is defined in a similar manner:

$$MU_2 = \frac{\Delta U}{\Delta x_2} = \frac{u(x_1, x_2 + \Delta x_2) - u(x_1, x_2)}{\Delta x_2}.$$

Note that when we compute the marginal utility with respect to good 2 we keep the amount of good 1 constant. We can calculate the change in utility associated with a change in the consumption of good 2 by the formula

$$\Delta U = MU_2 \Delta x_2.$$

It is important to realize that the magnitude of marginal utility depends on the magnitude of utility. Thus it depends on the particular way that we choose to measure utility. If we multiplied utility by 2, then marginal utility would also be multiplied by 2. We would still have a perfectly valid utility function in that it would represent the same preferences, but it would just be scaled differently.

This means that marginal utility itself has no behavioral content. How can we calculate marginal utility from a consumer's choice behavior? We can't. Choice behavior only reveals information about the way a consumer *ranks* different bundles of goods. Marginal utility depends on the particular utility function that we use to reflect the preference ordering and its magnitude has no particular significance. However, it turns out that marginal utility can be used to calculate something that does have behavioral content, as we will see in the next section.

4.5 Marginal Utility and MRS

A utility function $u(x_1, x_2)$ can be used to measure the marginal rate of substitution (MRS) defined in Chapter 3. Recall that the MRS measures the slope of the indifference curve at a given bundle of goods; it can be interpreted as the rate at which a consumer is just willing to substitute a small amount of good 2 for good 1.

This interpretation gives us a simple way to calculate the MRS. Consider a change in the consumption of each good, $(\Delta x_1, \Delta x_2)$, that keeps utility constant—that is, a change in consumption that moves us along the indifference curve. Then we must have

$$MU_1 \Delta x_1 + MU_2 \Delta x_2 = \Delta U = 0.$$

Solving for the slope of the indifference curve we have

$$\text{MRS} = \frac{\Delta x_2}{\Delta x_1} = -\frac{MU_1}{MU_2}. \tag{4.1}$$

(Note that we have 2 over 1 on the left-hand side of the equation and 1 over 2 on the right-hand side. Don't get confused!)

The algebraic sign of the MRS is negative: if you get more of good 1 you have to get *less* of good 2 in order to keep the same level of utility. However, it gets very tedious to keep track of that pesky minus sign, so economists often refer to the MRS by its absolute value—that is, as a positive number. We'll follow this convention as long as no confusion will result.

Now here is the interesting thing about the MRS calculation: the MRS can be measured by observing a person's actual behavior—we find that rate of exchange where he or she is just willing to stay put, as described in Chapter 3.

The utility function, and therefore the marginal utility function, is not uniquely determined. Any monotonic transformation of a utility function leaves you with another equally valid utility function. Thus, if we multiply utility by 2, for example, the marginal utility is multiplied by 2. Thus the magnitude of the marginal utility function depends on the choice of utility function, which is arbitrary. It doesn't depend on behavior alone; instead it depends on the utility function that we use to describe behavior.

But the *ratio* of marginal utilities gives us an observable magnitude—namely the marginal rate of substitution. The ratio of marginal utilities is independent of the particular transformation of the utility function you choose to use. Look at what happens if you multiply utility by 2. The MRS becomes

$$\text{MRS} = -\frac{2MU_1}{2MU_2}.$$

The 2s just cancel out, so the MRS remains the same.

The same sort of thing occurs when we take any monotonic transformation of a utility function. Taking a monotonic transformation is just relabeling the indifference curves, and the calculation for the MRS described above is concerned with moving along a given indifference curve. Even though the marginal utilities are changed by monotonic transformations, the *ratio* of marginal utilities is independent of the particular way chosen to represent the preferences.

4.6 Utility for Commuting

Utility functions are basically ways of describing choice behavior: if a bundle of goods X is chosen when a bundle of goods Y is available, then X must have a higher utility than Y. By examining choices consumers make we can estimate a utility function to describe their behavior.

This idea has been widely applied in the field of transportation economics to study consumers' commuting behavior. In most large cities commuters have a choice between taking public transit or driving to work. Each of these alternatives can be thought of as representing a bundle of different characteristics: travel time, waiting time, out-of-pocket costs, comfort, convenience, and so on. We could let x_1 be the amount of travel time involved in each kind of transportation, x_2 the amount of waiting time for each kind, and so on.

If $(x_1, x_2, \ldots, x_n)$ represents the values of n different characteristics of driving, say, and $(y_1, y_2, \ldots, y_n)$ represents the values of taking the bus, we can consider a model where the consumer decides to drive or take the bus depending on whether he prefers one bundle of characteristics to the other.

More specifically, let us suppose that the average consumer's preferences for characteristics can be represented by a utility function of the form

$$U(x_1, x_2, \ldots, x_n) = \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_n x_n,$$

where the coefficients β_1, β_2, and so on are unknown parameters. Any monotonic transformation of this utility function would describe the choice behavior equally well, of course, but the linear form is especially easy to work with from a statistical point of view.

Suppose now that we observe a number of similar consumers making choices between driving and taking the bus based on the particular pattern of commute times, costs, and so on that they face. There are statistical techniques that can be used to find the values of the coefficients β_i for $i = 1, \ldots, n$ that best fit the observed pattern of choices by a set of consumers. These statistical techniques give a way to estimate the utility function for different transportation modes.

One study reports a utility function that had the form[4]

$$U(TW, TT, C) = -0.147TW - 0.0411TT - 2.24C, \tag{4.2}$$

where

TW = total walking time to and from bus or car
TT = total time of trip in minutes
C = total cost of trip in dollars

The estimated utility function in the Domenich-McFadden book correctly described the choice between auto and bus transport for 93 percent of the households in their sample.

[4] See Thomas Domenich and Daniel McFadden, *Urban Travel Demand* (North-Holland Publishing Company, 1975). The estimation procedure in this book also incorporated various demographic characteristics of the households in addition to the purely economic variables described here. Daniel McFadden was awarded the Nobel Prize in economics in 2000 for his work in developing techniques to estimate models of this sort.

The coefficients on the variables in Equation (4.2) describe the weight that an average household places on the various characteristics of their commuting trips; that is, the marginal utility of each characteristic. The *ratio* of one coefficient to another measures the marginal rate of substitution between one characteristic and another. For example, the ratio of the marginal utility of walking time to the marginal utility of total time indicates that walking time is viewed as being roughly 3 times as onerous as travel time by the average consumer. In other words, the consumer would be willing to substitute 3 minutes of additional travel time to save 1 minute of walking time.

Similarly, the ratio of cost to travel time indicates the average consumer's tradeoff between these two variables. In this study, the average commuter valued a minute of commute time at $0.0411/2.24 = 0.0183$ dollars per minute, which is \$1.10 per hour. For comparison, the hourly wage for the average commuter in 1967, the year of the study, was about \$2.85 an hour.

Such estimated utility functions can be very valuable for determining whether or not it is worthwhile to make some change in the public transportation system. For example, in the above utility function one of the significant factors explaining mode choice is the time involved in taking the trip. The city transit authority can, at some cost, add more buses to reduce this travel time. But will the number of extra riders warrant the increased expense?

Given a utility function and a sample of consumers we can forecast which consumers will drive and which consumers will choose to take the bus. This will give us some idea as to whether the revenue will be sufficient to cover the extra cost.

Furthermore, we can use the marginal rate of substitution to estimate the *value* that each consumer places on the reduced travel time. We saw above that in the Domenich-McFadden study the average commuter in 1967 valued commute time at a rate of \$1.10 per hour. Thus the commuter should be willing to pay about \$0.37 to cut 20 minutes from his or her trip. This number gives us a measure of the dollar benefit of providing more timely bus service. This benefit must be compared to the cost of providing more timely bus service in order to determine if such provision is worthwhile. Having a quantitative measure of benefit will certainly be helpful in making a rational decision about transport policy.

Summary

1. A utility function is simply a way to represent or summarize a preference ordering. The numerical magnitudes of utility levels have no intrinsic meaning.

2. Thus, given any one utility function, any monotonic transformation of it will represent the same preferences.

3. The marginal rate of substitution, MRS, can be calculated from the utility function via the formula MRS $= \Delta x_2/\Delta x_1 = -MU_1/MU_2$.

REVIEW QUESTIONS

1. The text said that raising a number to an odd power was a monotonic transformation. What about raising a number to an even power? Is this a monotonic transformation? (Hint: consider the case $f(u) = u^2$.)

2. Which of the following are monotonic transformations? (1) $u = 2v - 13$; (2) $u = -1/v^2$; (3) $u = 1/v^2$; (4) $u = \ln v$; (5) $u = -e^{-v}$; (6) $u = v^2$; (7) $u = v^2$ for $v > 0$; (8) $u = v^2$ for $v < 0$.

3. We claimed in the text that if preferences were monotonic, then a diagonal line through the origin would intersect each indifference curve exactly once. Can you prove this rigorously? (Hint: what would happen if it intersected some indifference curve twice?)

4. What kind of preferences are represented by a utility function of the form $u(x_1, x_2) = \sqrt{x_1 + x_2}$? What about the utility function $v(x_1, x_2) = 13x_1 + 13x_2$?

5. What kind of preferences are represented by a utility function of the form $u(x_1, x_2) = x_1 + \sqrt{x_2}$? Is the utility function $v(x_1, x_2) = x_1^2 + 2x_1\sqrt{x_2} + x_2$ a monotonic transformation of $u(x_1, x_2)$?

6. Consider the utility function $u(x_1, x_2) = \sqrt{x_1 x_2}$. What kind of preferences does it represent? Is the function $v(x_1, x_2) = x_1^2 x_2$ a monotonic transformation of $u(x_1, x_2)$? Is the function $w(x_1, x_2) = x_1^2 x_2^2$ a monotonic transformation of $u(x_1, x_2)$?

7. Can you explain why taking a monotonic transformation of a utility function doesn't change the marginal rate of substitution?

APPENDIX

First, let us clarify what is meant by "marginal utility." As elsewhere in economics, "marginal" just means a derivative. So the marginal utility of good 1 is just

$$MU_1 = \lim_{\Delta x_1 \to 0} \frac{u(x_1 + \Delta x_1, x_2) - u(x_1, x_2)}{\Delta x_1} = \frac{\partial u(x_1, x_2)}{\partial x_1}.$$

Note that we have used the *partial* derivative here, since the marginal utility of good 1 is computed holding good 2 fixed.

Now we can rephrase the derivation of the MRS given in the text using calculus. We'll do it two ways: first by using differentials, and second by using implicit functions.

For the first method, we consider making a change (dx_1, dx_2) that keeps utility constant. So we want

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

The first term measures the increase in utility from the small change dx_1, and the second term measures the increase in utility from the small change dx_2. We want to pick these changes so that the total change in utility, du, is zero. Solving for dx_2/dx_1 gives us

$$\frac{dx_2}{dx_1} = -\frac{\partial u(x_1, x_2)/\partial x_1}{\partial u(x_1, x_2)/\partial x_2},$$

which is just the calculus analog of equation (4.1) in the text.

As for the second method, we now think of the indifference curve as being described by a function $x_2(x_1)$. That is, for each value of x_1, the function $x_2(x_1)$ tells us how much x_2 we need to get on that specific indifference curve. Thus the function $x_2(x_1)$ has to satisfy the identity

$$u(x_1, x_2(x_1)) \equiv k,$$

where k is the utility label of the indifference curve in question.

We can differentiate both sides of this identity with respect to x_1 to get

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

Notice that x_1 occurs in two places in this identity, so changing x_1 will change the function in two ways, and we have to take the derivative at each place that x_1 appears.

We then solve this equation for $\partial x_2(x_1)/\partial x_1$ to find

$$\frac{\partial x_2(x_1)}{\partial x_1} = -\frac{\partial u(x_1, x_2)/\partial x_1}{\partial u(x_1, x_2)/\partial x_2},$$

just as we had before.

The implicit function method is a little more rigorous, but the differential method is more direct, as long as you don't do something silly.

Suppose that we take a monotonic transformation of a utility function, say, $v(x_1, x_2) = f(u(x_1, x_2))$. Let's calculate the MRS for this utility function. Using the chain rule

$$\begin{aligned} \text{MRS} = -\frac{\partial v/\partial x_1}{\partial v/\partial x_2} &= -\frac{\partial f/\partial u}{\partial f/\partial u} \frac{\partial u/\partial x_1}{\partial u/\partial x_2} \\ &= -\frac{\partial u/\partial x_1}{\partial u/\partial x_2} \end{aligned}$$

since the $\partial f/\partial u$ term cancels out from both the numerator and denominator. This shows that the MRS is independent of the utility representation.

This gives a useful way to recognize preferences that are represented by different utility functions: given two utility functions, just compute the marginal rates of substitution and see if they are the same. If they are, then the two utility functions have the same indifference curves. If the direction of increasing preference is the same for each utility function, then the underlying preferences must be the same.

EXAMPLE: Cobb-Douglas Preferences

The MRS for Cobb-Douglas preferences is easy to calculate by using the formula derived above.

If we choose the log representation where

$$u(x_1, x_2) = c \ln x_1 + d \ln x_2,$$

then we have

$$\begin{aligned} \text{MRS} &= -\frac{\partial u(x_1, x_2)/\partial x_1}{\partial u(x_1, x_2)/\partial x_2} \\ &= -\frac{c/x_1}{d/x_2} \\ &= -\frac{c}{d}\frac{x_2}{x_1}. \end{aligned}$$

Note that the MRS only depends on the ratio of the two parameters and the quantity of the two goods in this case.

What if we choose the exponent representation where

$$u(x_1, x_2) = x_1^c x_2^d?$$

Then we have

$$\begin{aligned} \text{MRS} &= -\frac{\partial u(x_1, x_2)/\partial x_1}{\partial u(x_1, x_2)/\partial x_2} \\ &= -\frac{c x_1^{c-1} x_2^d}{d x_1^c x_2^{d-1}} \\ &= -\frac{c x_2}{d x_1}, \end{aligned}$$

which is the same as we had before. Of course you knew all along that a monotonic transformation couldn't change the marginal rate of substitution!