CHAPTER 6

DEMAND

In the last chapter we presented the basic model of consumer choice: how maximizing utility subject to a budget constraint yields optimal choices. We saw that the optimal choices of the consumer depend on the consumer's income and the prices of the goods, and we worked a few examples to see what the optimal choices are for some simple kinds of preferences.

The consumer's **demand functions** give the optimal amounts of each of the goods as a function of the prices and income faced by the consumer. We write the demand functions as

$$x_1 = x_1(p_1, p_2, m)$$
$$x_2 = x_2(p_1, p_2, m).$$

The left-hand side of each equation stands for the quantity demanded. The right-hand side of each equation is the function that relates the prices and income to that quantity.

In this chapter we will examine how the demand for a good changes as prices and income change. Studying how a choice responds to changes in the economic environment is known as **comparative statics**, which we first described in Chapter 1. "Comparative" means that we want to compare two situations: before and after the change in the economic environment. "Statics" means that we are not concerned with any adjustment process that may be involved in moving from one choice to another; rather we will only examine the equilibrium choice.

In the case of the consumer, there are only two things in our model that affect the optimal choice: prices and income. The comparative statics questions in consumer theory therefore involve investigating how demand changes when prices and income change.

6.1 Normal and Inferior Goods

We start by considering how a consumer's demand for a good changes as his income changes. We want to know how the optimal choice at one income compares to the optimal choice at another level of income. During this exercise, we will hold the prices fixed and examine only the change in demand due to the income change.

We know how an increase in money income affects the budget line when prices are fixed—it shifts it outward in a parallel fashion. So how does this affect demand?

We would normally think that the demand for each good would increase when income increases, as shown in Figure 6.1. Economists, with a singular lack of imagination, call such goods **normal** goods. If good 1 is a normal good, then the demand for it increases when income increases, and decreases when income decreases. For a normal good the quantity demanded always changes in the same way as income changes:

$$\frac{\Delta x_1}{\Delta m} > 0.$$

If something is called normal, you can be sure that there must be a *possibility* of being abnormal. And indeed there is. Figure 6.2 presents an example of nice, well-behaved indifference curves where an increase of income results in a *reduction* in the consumption of one of the goods. Such a good is called an **inferior** good. This may be "abnormal," but when you think about it, inferior goods aren't all that unusual. There are many goods for which demand decreases as income increases; examples might include gruel, bologna, shacks, or nearly any kind of low-quality good.

Whether a good is inferior or not depends on the income level that we are examining. It might very well be that very poor people consume more bologna as their income increases. But after a point, the consumption of bologna would probably decline as income continued to increase. Since in real life the consumption of goods can increase or decrease when income increases, it is comforting to know that economic theory allows for both possibilities.

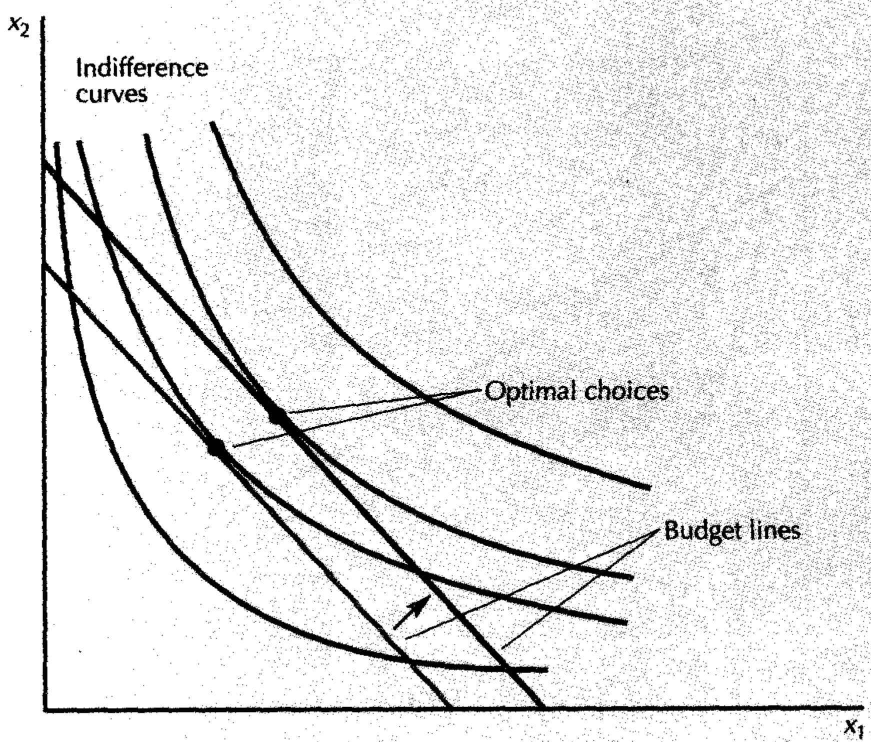

Normal goods. The demand for both goods increases when income increases, so both goods are normal goods.

6.2 Income Offer Curves and Engel Curves

We have seen that an increase in income corresponds to shifting the budget line outward in a parallel manner. We can connect together the demanded bundles that we get as we shift the budget line outward to construct the **income offer curve**. This curve illustrates the bundles of goods that are demanded at the different levels of income, as depicted in Figure 6.3A. The income offer curve is also known as the **income expansion path**. If both goods are normal goods, then the income expansion path will have a positive slope, as depicted in Figure 6.3A.

For each level of income, m, there will be some optimal choice for each of the goods. Let us focus on good 1 and consider the optimal choice at each set of prices and income, $x_1(p_1, p_2, m)$. This is simply the demand function for good 1. If we hold the prices of goods 1 and 2 fixed and look at how demand changes as we change income, we generate a curve known as the **Engel curve**. The Engel curve is a graph of the demand for one of the goods as a function of income, with all prices being held constant. For an example of an Engel curve, see Figure 6.3B.

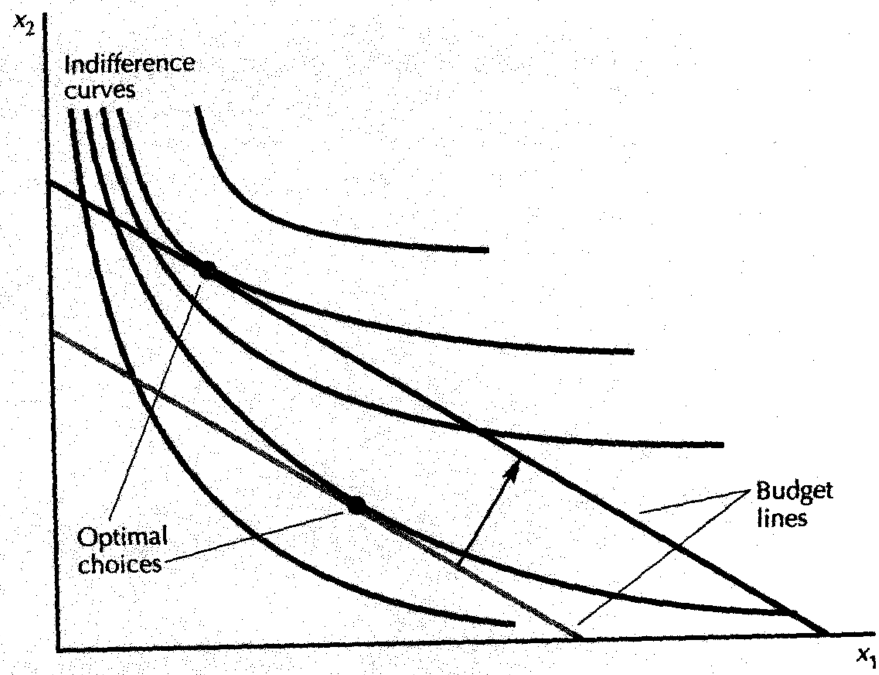

An inferior good. Good 1 is an inferior good, which means that the demand for it decreases when income increases.

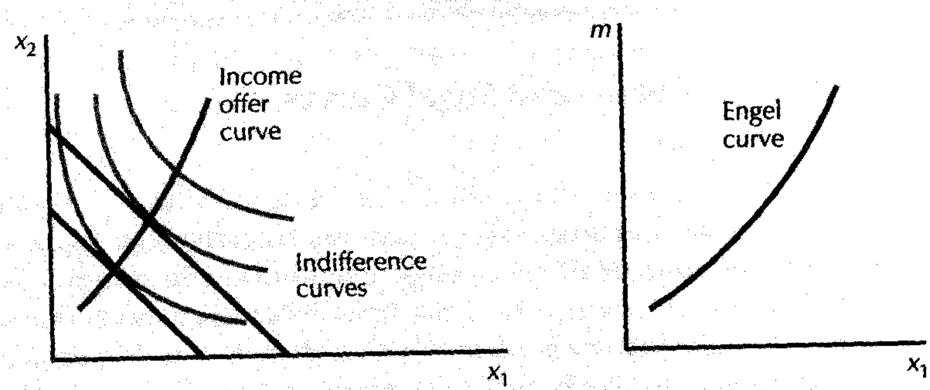

A Income offer curve

B Engel curve

How demand changes as income changes. The income offer curve (or income expansion path) shown in panel A depicts the optimal choice at different levels of income and constant prices. When we plot the optimal choice of good 1 against income, m, we get the Engel curve, depicted in panel B.

6.3 Some Examples

Let's consider some of the preferences that we examined in Chapter 5 and see what their income offer curves and Engel curves look like.

Perfect Substitutes

The case of perfect substitutes is depicted in Figure 6.4. If $p_1 < p_2$, so that the consumer is specializing in consuming good 1, then if his income increases he will increase his consumption of good 1. Thus the income offer curve is the horizontal axis, as shown in Figure 6.4A.

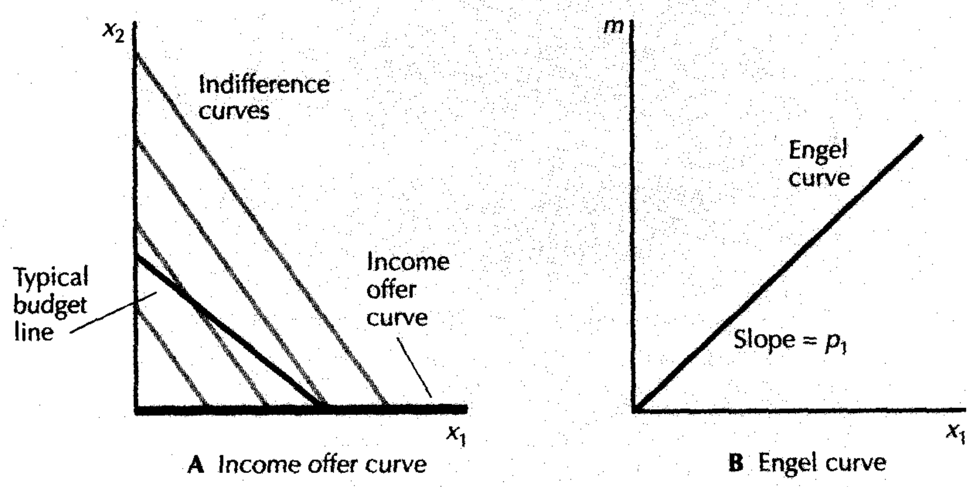

Perfect substitutes. The income offer curve (A) and an Engel curve (B) in the case of perfect substitutes.

Since the demand for good 1 is $x_1 = m/p_1$ in this case, the Engel curve will be a straight line with a slope of p_1, as depicted in Figure 6.4B. (Since m is on the vertical axis, and x_1 on the horizontal axis, we can write $m = p_1 x_1$, which makes it clear that the slope is p_1.)

Perfect Complements

The demand behavior for perfect complements is shown in Figure 6.5. Since the consumer will always consume the same amount of each good, no matter what, the income offer curve is the diagonal line through the origin as depicted in Figure 6.5A. We have seen that the demand for good 1 is $x_1 = m/(p_1 + p_2)$, so the Engel curve is a straight line with a slope of $p_1 + p_2$ as shown in Figure 6.5B.

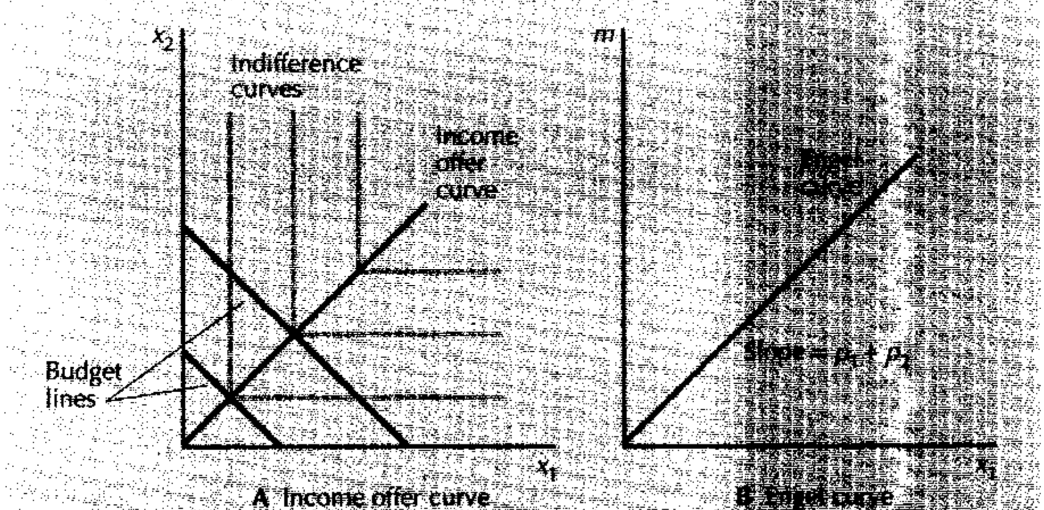

Figure 6.5 **Perfect complements.** The income offer curve (A) and an Engel curve (B) in the case of perfect complements.

Cobb-Douglas Preferences

For the case of Cobb-Douglas preferences it is easier to look at the algebraic form of the demand functions to see what the graphs will look like. If $u(x_1, x_2) = x_1^a x_2^{1-a}$, the Cobb-Douglas demand for good 1 has the form $x_1 = am/p_1$. For a fixed value of p_1, this is a *linear* function of m. Thus doubling m will double demand, tripling m will triple demand, and so on. In fact, multiplying m by any positive number t will just multiply demand by the same amount.

The demand for good 2 is $x_2 = (1-a)m/p_2$, and this is also clearly linear. The fact that the demand functions for both goods are linear functions of income means that the income expansion paths will be straight lines through the origin, as depicted in Figure 6.6A. The Engel curve for good 1 will be a straight line with a slope of p_1/a, as depicted in Figure 6.6B.

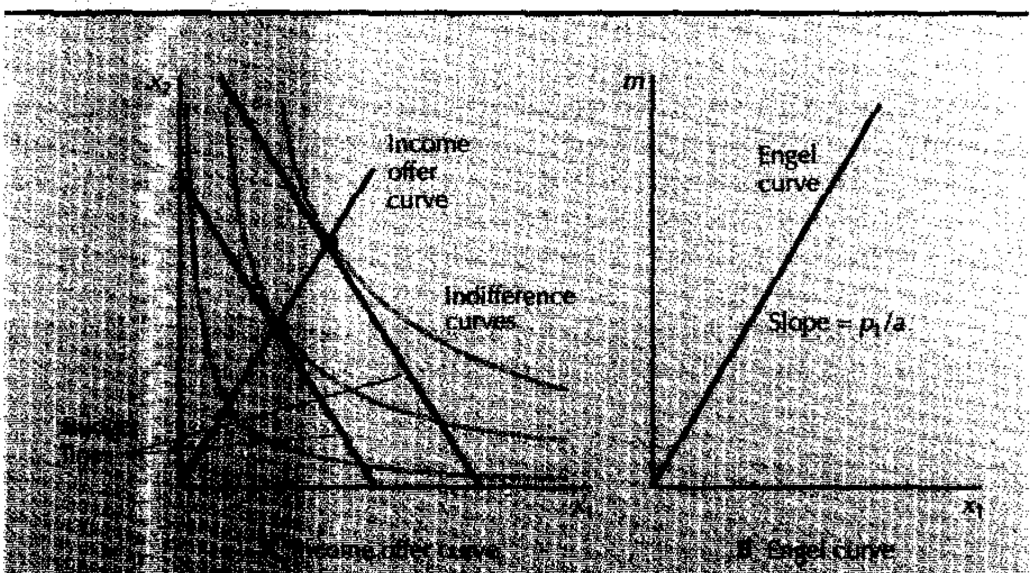

Figure 6.6 **Cobb-Douglas.** An income offer curve (A) and an Engel curve (B) for Cobb-Douglas utility.

Homothetic Preferences

All of the income offer curves and Engel curves that we have seen up to now have been straightforward—in fact they've been straight lines! This has happened because our examples have been so simple. Real Engel curves do not have to be straight lines. In general, when income goes up, the demand for a good could increase more or less rapidly than income increases. If the demand for a good goes up by a greater proportion than income, we say that it is a **luxury good**, and if it goes up by a lesser proportion than income we say that it is a **necessary good.**

The dividing line is the case where the demand for a good goes up by the same proportion as income. This is what happened in the three cases we examined above. What aspect of the consumer's preferences leads to this behavior?

Suppose that the consumer's preferences only depend on the *ratio* of good 1 to good 2. This means that if the consumer prefers (x_1, x_2) to (y_1, y_2), then she automatically prefers $(2x_1, 2x_2)$ to $(2y_1, 2y_2)$, $(3x_1, 3x_2)$ to $(3y_1, 3y_2)$, and so on, since the ratio of good 1 to good 2 is the same for all of these bundles. In fact, the consumer prefers (tx_1, tx_2) to (ty_1, ty_2) for any positive value of t. Preferences that have this property are known as **homothetic preferences**. It is not hard to show that the three examples of preferences given above—perfect substitutes, perfect complements, and Cobb-Douglas—are all homothetic preferences.

If the consumer has homothetic preferences, then the income offer curves are all straight lines through the origin, as shown in Figure 6.7. More specifically, if preferences are homothetic, it means that when income is scaled up or down by any amount $t > 0$, the demanded bundle scales up or down by the same amount. This can be established rigorously, but it is fairly clear from looking at the picture. If the indifference curve is tangent to the budget line at (x_1^*, x_2^*), then the indifference curve through (tx_1^*, tx_2^*) is tangent to the budget line that has t times as much income and the same prices. This implies that the Engel curves are straight lines as well. If you double income, you just double the demand for each good.

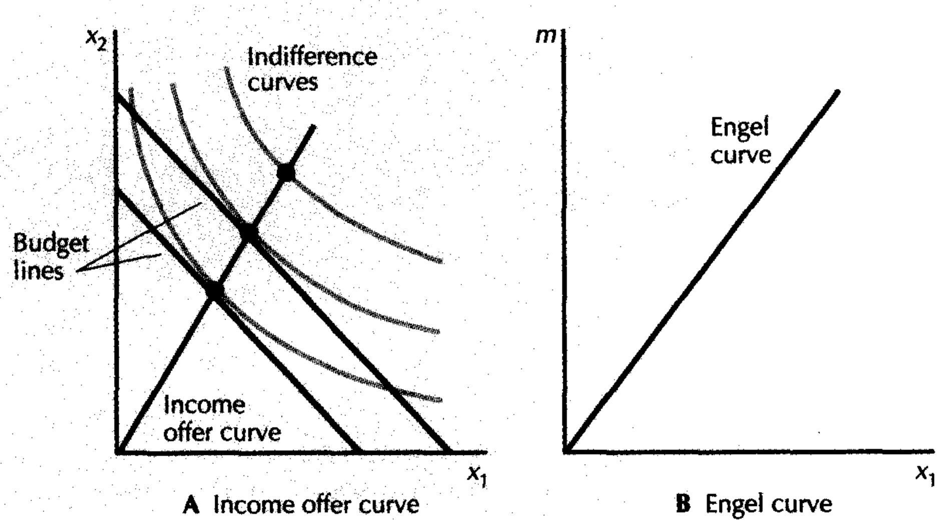

Homothetic preferences. An income offer curve (A) and an Engel curve (B) in the case of homothetic preferences.

Homothetic preferences are very convenient since the income effects are so simple. Unfortunately, homothetic preferences aren't very realistic for the same reason! But they will often be of use in our examples.

Quasilinear Preferences

Another kind of preferences that generates a special form of income offer curves and Engel curves is the case of quasilinear preferences. Recall the definition of quasilinear preferences given in Chapter 4. This is the case where all indifference curves are "shifted" versions of one indifference curve as in Figure 6.8. Equivalently, the utility function for these preferences takes the form $u(x_1, x_2) = v(x_1) + x_2$. What happens if we shift the budget line outward? In this case, if an indifference curve is tangent to the budget line at a bundle (x_1^*, x_2^*), then another indifference curve must also be tangent at $(x_1^*, x_2^* + k)$ for any constant k. Increasing income doesn't change the demand for good 1 at all, and all the extra income goes entirely to the consumption of good 2. If preferences are quasilinear, we sometimes say that there is a "zero income effect" for good 1. Thus the Engel curve for good 1 is a vertical line—as you change income, the demand for good 1 remains constant.

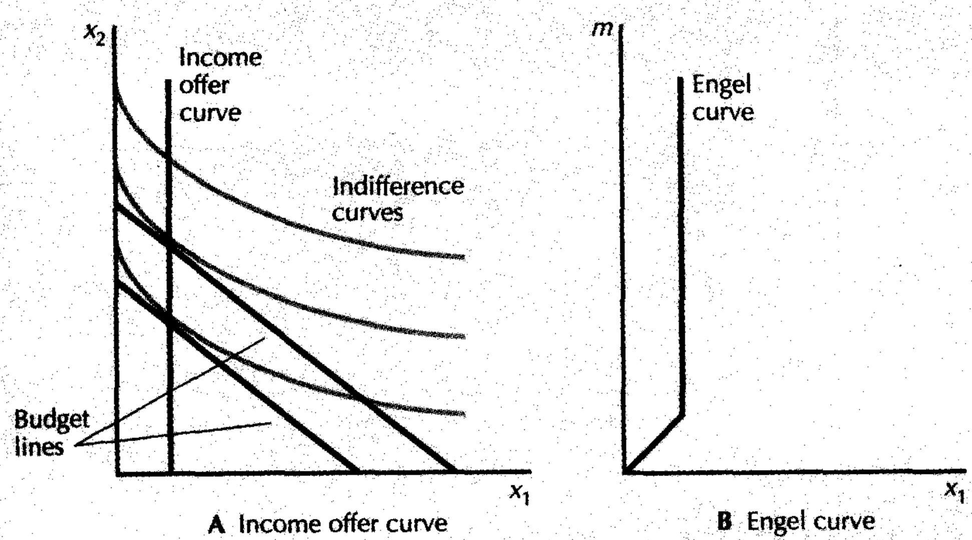

Quasilinear preferences. An income offer curve (A) and an Engel curve (B) with quasilinear preferences.

What would be a real-life situation where this kind of thing might occur? Suppose good 1 is pencils and good 2 is money to spend on other goods. Initially I may spend my income only on pencils, but when my income gets large enough, I stop buying additional pencils—all of my extra income is spent on other goods. Other examples of this sort might be salt or toothpaste. When we are examining a choice between all other goods and some single good that isn't a very large part of the consumer's budget, the quasilinear assumption may well be plausible, at least when the consumer's income is sufficiently large.

6.4 Ordinary Goods and Giffen Goods

Let us now consider price changes. Suppose that we decrease the price of good 1 and hold the price of good 2 and money income fixed. Then what can happen to the quantity demanded of good 1? Intuition tells us that the quantity demanded of good 1 should increase when its price decreases. Indeed this is the ordinary case, as depicted in Figure 6.9.

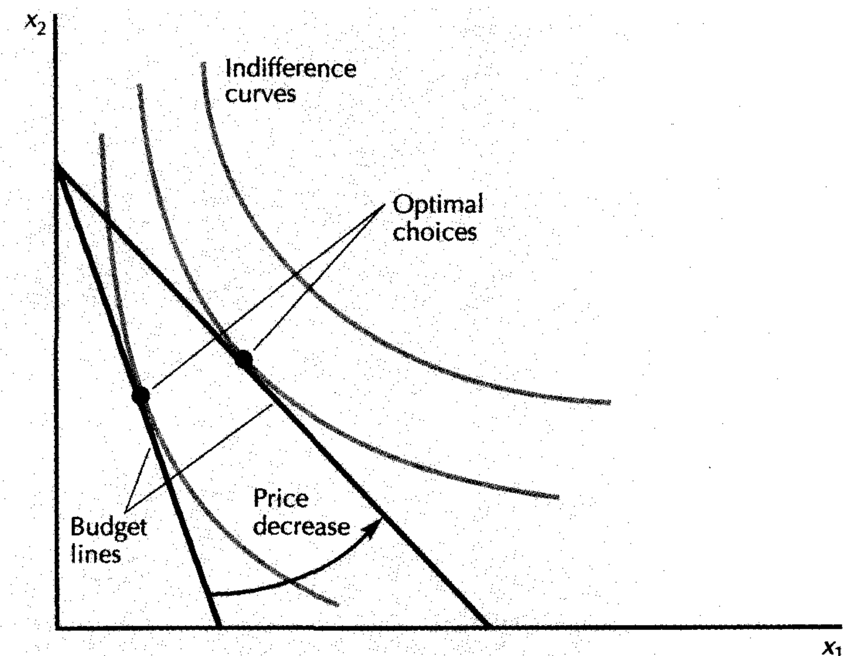

An ordinary good. Ordinarily, the demand for a good increases when its price decreases, as is the case here.

When the price of good 1 decreases, the budget line becomes flatter. Or said another way, the vertical intercept is fixed and the horizontal intercept moves to the right. In Figure 6.9, the optimal choice of good 1 moves to the right as well: the quantity demanded of good 1 has increased. But we might wonder whether this always happens this way. Is it always the case that, no matter what kind of preferences the consumer has, the demand for a good must increase when its price goes down?

As it turns out, the answer is no. It is logically possible to find well-behaved preferences for which a decrease in the price of good 1 leads to a reduction in the demand for good 1. Such a good is called a **Giffen good**,

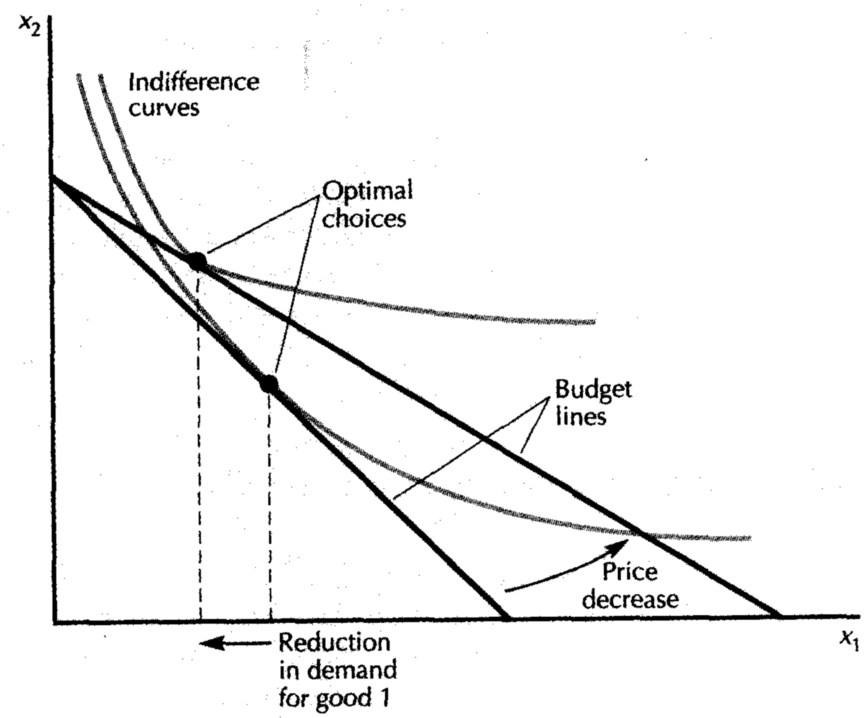

A Giffen good. Good 1 is a Giffen good, since the demand for it decreases when its price decreases.

after the nineteenth-century economist who first noted the possibility. An example is illustrated in Figure 6.10.

What is going on here in economic terms? What kind of preferences might give rise to the peculiar behavior depicted in Figure 6.10? Suppose that the two goods that you are consuming are gruel and milk and that you are currently consuming 7 bowls of gruel and 7 cups of milk a week. Now the price of gruel declines. If you consume the same 7 bowls of gruel a week, you will have money left over with which you can purchase more milk. In fact, with the extra money you have saved because of the lower price of gruel, you may decide to consume even more milk and reduce your consumption of gruel. The reduction in the price of gruel has freed up some extra money to be spent on other things—but one thing you might want to do with it is reduce your consumption of gruel! Thus the price change is to some extent *like* an income change. Even though *money* income remains constant, a change in the price of a good will change purchasing power, and thereby change demand.

So the Giffen good is not implausible purely on logical grounds, although Giffen goods are unlikely to be encountered in real-world behavior. Most goods are ordinary goods—when their price increases, the demand for them declines. We'll see why this is the ordinary situation a little later.

Incidentally, it is no accident that we used gruel as an example of both an inferior good and a Giffen good. It turns out that there is an intimate relationship between the two which we will explore in a later chapter.

But for now our exploration of consumer theory may leave you with the impression that nearly anything can happen: if income increases the demand for a good can go up or down, and if price increases the demand can go up or down. Is consumer theory compatible with *any* kind of behavior? Or are there some kinds of behavior that the economic model of consumer behavior rules out? It turns out that there *are* restrictions on behavior imposed by the maximizing model. But we'll have to wait until the next chapter to see what they are.

6.5 The Price Offer Curve and the Demand Curve

Suppose that we let the price of good 1 change while we hold p_2 and income fixed. Geometrically this involves pivoting the budget line. We can think of connecting together the optimal points to construct the **price offer curve** as illustrated in Figure 6.11A. This curve represents the bundles that would be demanded at different prices for good 1.

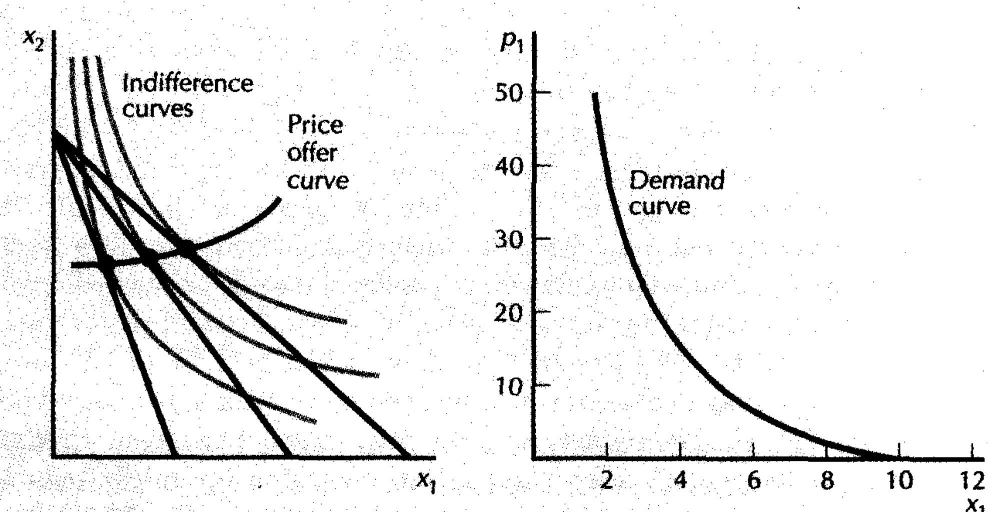

A Price offer curve

B Demand curve

The price offer curve and demand curve. Panel A contains a price offer curve, which depicts the optimal choices as the price of good 1 changes. Panel B contains the associated demand curve, which depicts a plot of the optimal choice of good 1 as a function of its price.

We can depict this same information in a different way. Again, hold the price of good 2 and money income fixed, and for each different value of p_1 plot the optimal level of consumption of good 1. The result is the **demand curve** depicted in Figure 6.11B. The demand curve is a plot of the demand function, $x_1(p_1, p_2, m)$, holding p_2 and m fixed at some predetermined values.

Ordinarily, when the price of a good increases, the demand for that good will decrease. Thus the price and quantity of a good will move in *opposite* directions, which means that the demand curve will typically have a negative slope. In terms of rates of change, we would normally have

$$\frac{\Delta x_1}{\Delta p_1} < 0,$$

which simply says that demand curves usually have a negative slope.

However, we have also seen that in the case of Giffen goods, the demand for a good may decrease when its price decreases. Thus it is possible, but not likely, to have a demand curve with a positive slope.

6.6 Some Examples

Let's look at a few examples of demand curves, using the preferences that we discussed in Chapter 3.

Perfect Substitutes

The offer curve and demand curve for perfect substitutes—the red and blue pencils example—are illustrated in Figure 6.12. As we saw in Chapter 5, the demand for good 1 is zero when $p_1 > p_2$, any amount on the budget line when $p_1 = p_2$, and m/p_1 when $p_1 < p_2$. The offer curve traces out these possibilities.

In order to find the demand curve, we fix the price of good 2 at some price p_2^* and graph the demand for good 1 versus the price of good 1 to get the shape depicted in Figure 6.12B.

Perfect Complements

The case of perfect complements—the right and left shoes example—is depicted in Figure 6.13. We know that whatever the prices are, a consumer will demand the same amount of goods 1 and 2. Thus his offer curve will be a diagonal line as depicted in Figure 6.13A.

We saw in Chapter 5 that the demand for good 1 is given by

$$x_1 = \frac{m}{p_1 + p_2}.$$

If we fix m and p_2 and plot the relationship between x_1 and p_1, we get the curve depicted in Figure 6.13B.

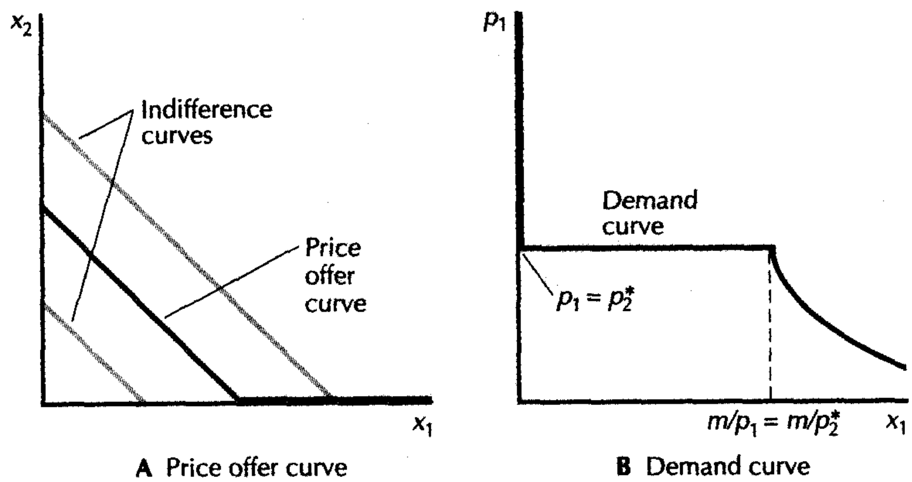

Perfect substitutes. Price offer curve (A) and demand curve (B) in the case of perfect substitutes.

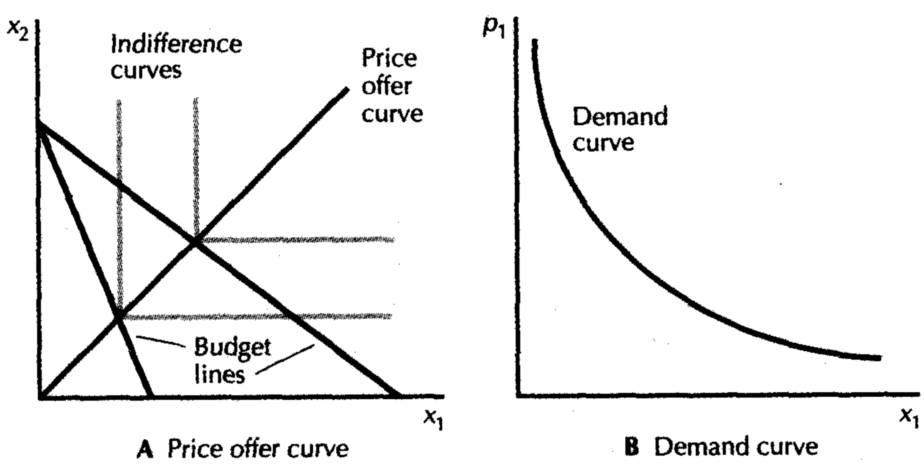

Perfect complements. Price offer curve (A) and demand curve (B) in the case of perfect complements.

A Discrete Good

Suppose that good 1 is a discrete good. If p_1 is very high then the consumer will strictly prefer to consume zero units; if p_1 is low enough the consumer will strictly prefer to consume one unit. At some price r_1, the consumer will be indifferent between consuming good 1 or not consuming it. The price at which the consumer is just indifferent to consuming or not consuming the good is called the **reservation price**.[1] The indifference curves and demand curve are depicted in Figure 6.14.

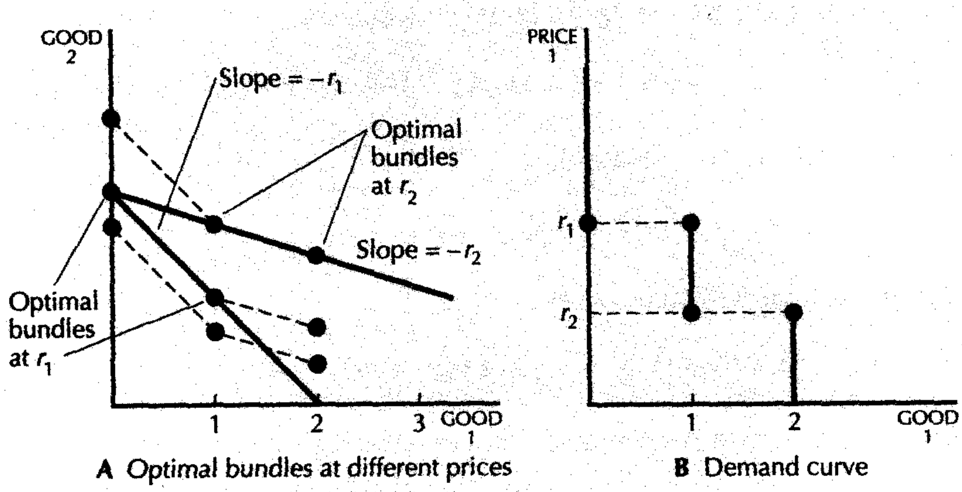

A discrete good. As the price of good 1 decreases there will be some price, the reservation price, at which the consumer is just indifferent between consuming good 1 or not consuming it. As the price decreases further, more units of the discrete good will be demanded.

It is clear from the diagram that the demand behavior can be described by a sequence of reservation prices at which the consumer is just willing to purchase another unit of the good. At a price of r_1 the consumer is willing to buy 1 unit of the good; if the price falls to r_2, he is willing to buy another unit, and so on.

These prices can be described in terms of the original utility function. For example, r_1 is the price where the consumer is just indifferent between consuming 0 or 1 unit of good 1, so it must satisfy the equation

$$u(0, m) = u(1, m - r_1). \tag{6.1}$$

Similarly r_2 satisfies the equation

$$u(1, m - r_2) = u(2, m - 2r_2). \tag{6.2}$$

[1] The term reservation price comes from auction markets. When someone wanted to sell something in an auction he would typically state a minimum price at which he was willing to sell the good. If the best price offered was below this stated price, the seller reserved the right to purchase the item himself. This price became known as the seller's reservation price and eventually came to be used to describe the price at which someone was just willing to buy or sell some item.

The left-hand side of this equation is the utility from consuming one unit of the good at a price of r_2. The right-hand side is the utility from consuming two units of the good, each of which sells for r_2.

If the utility function is quasilinear, then the formulas describing the reservation prices become somewhat simpler. If $u(x_1, x_2) = v(x_1) + x_2$, and $v(0) = 0$, then we can write equation (6.1) as

$$v(0) + m = m = v(1) + m - r_1.$$

Since $v(0) = 0$, we can solve for r_1 to find

$$r_1 = v(1). \tag{6.3}$$

Similarly, we can write equation (6.2) as

$$v(1) + m - r_2 = v(2) + m - 2r_2.$$

Canceling terms and rearranging, this expression becomes

$$r_2 = v(2) - v(1).$$

Proceeding in this manner, the reservation price for the third unit of consumption is given by

$$r_3 = v(3) - v(2)$$

and so on.

In each case, the reservation price measures the increment in utility necessary to induce the consumer to choose an additional unit of the good. Loosely speaking, the reservation prices measure the marginal utilities associated with different levels of consumption of good 1. Our assumption of convex preferences implies that the sequence of reservation prices must decrease: $r_1 > r_2 > r_3 \cdots$.

Because of the special structure of the quasilinear utility function, the reservation prices do not depend on the amount of good 2 that the consumer has. This is certainly a special case, but it makes it very easy to describe demand behavior. Given any price p, we just find where it falls in the list of reservation prices. Suppose that p falls between r_6 and r_7, for example. The fact that $r_6 > p$ means that the consumer is willing to give up p dollars per unit bought to get 6 units of good 1, and the fact that $p > r_7$ means that the consumer is not willing to give up p dollars per unit to get the seventh unit of good 1.

This argument is quite intuitive, but let's look at the math just to make sure that it is clear. Suppose that the consumer demands 6 units of good 1. We want to show that we must have $r_6 \geq p \geq r_7$.

If the consumer is maximizing utility, then we must have

$$v(6) + m - 6p \geq v(x_1) + m - px_1$$

for all possible choices of x_1. In particular, we must have that

$$v(6) + m - 6p \geq v(5) + m - 5p.$$

Rearranging this equation we have

$$r_6 = v(6) - v(5) \geq p,$$

which is half of what we wanted to show.

By the same logic,

$$v(6) + m - 6p \geq v(7) + m - 7p.$$

Rearranging this gives us

$$p \geq v(7) - v(6) = r_7,$$

which is the other half of the inequality we wanted to establish.

6.7 Substitutes and Complements

We have already used the terms substitutes and complements, but it is now appropriate to give a formal definition. Since we have seen *perfect* substitutes and *perfect* complements several times already, it seems reasonable to look at the imperfect case.

Let's think about substitutes first. We said that red pencils and blue pencils might be thought of as perfect substitutes, at least for someone who didn't care about color. But what about pencils and pens? This is a case of "imperfect" substitutes. That is, pens and pencils are, to some degree, a substitute for each other, although they aren't as perfect a substitute for each other as red pencils and blue pencils.

Similarly, we said that right shoes and left shoes were perfect complements. But what about a pair of shoes and a pair of socks? Right shoes and left shoes are nearly always consumed together, and shoes and socks are *usually* consumed together. Complementary goods are those like shoes and socks that tend to be consumed together, albeit not always.

Now that we've discussed the basic idea of complements and substitutes, we can give a precise economic definition. Recall that the demand function for good 1, say, will typically be a function of the price of both good 1 and good 2, so we write $x_1(p_1, p_2, m)$. We can ask how the demand for good 1 changes as the price of good 2 changes: does it go up or down?

If the demand for good 1 goes up when the price of good 2 goes up, then we say that good 1 is a **substitute** for good 2. In terms of rates of change, good 1 is a substitute for good 2 if

$$\frac{\Delta x_1}{\Delta p_2} > 0.$$

The idea is that when good 2 gets more expensive the consumer switches to consuming good 1: the consumer *substitutes* away from the more expensive good to the less expensive good.

On the other hand, if the demand for good 1 goes down when the price of good 2 goes up, we say that good 1 is a **complement** to good 2. This means that

$$\frac{\Delta x_1}{\Delta p_2} < 0.$$

Complements are goods that are consumed together, like coffee and sugar, so when the price of one good rises, the consumption of both goods will tend to decrease.

The cases of perfect substitutes and perfect complements illustrate these points nicely. Note that $\Delta x_1/\Delta p_2$ is positive (or zero) in the case of perfect substitutes, and that $\Delta x_1/\Delta p_2$ is negative in the case of perfect complements.

A couple of warnings are in order about these concepts. First, the two-good case is rather special when it comes to complements and substitutes. Since income is being held fixed, if you spend more money on good 1, you'll have to spend less on good 2. This puts some restrictions on the kinds of behavior that are possible. When there are more than two goods, these restrictions are not so much of a problem.

Second, although the definition of substitutes and complements in terms of consumer demand behavior seems sensible, there are some difficulties with the definitions in more general environments. For example, if we use the above definitions in a situation involving more than two goods, it is perfectly possible that good 1 may be a substitute for good 3, but good 3 may be a complement for good 1. Because of this peculiar feature, more advanced treatments typically use a somewhat different definition of substitutes and complements. The definitions given above describe concepts known as **gross substitutes** and **gross complements**; they will be sufficient for our needs.

6.8 The Inverse Demand Function

If we hold p_2 and m fixed and plot p_1 against x_1 we get the **demand curve**. As suggested above, we typically think that the demand curve slopes downwards, so that higher prices lead to less demand, although the Giffen example shows that it could be otherwise.

As long as we do have a downward-sloping demand curve, as is usual, it is meaningful to speak of the **inverse demand function**. The inverse demand function is the demand function viewing price as a function of quantity. That is, for each level of demand for good 1, the inverse demand function measures what the price of good 1 would have to be in order for the consumer to choose that level of consumption. So the inverse demand function measures the same relationship as the direct demand function, but just from another point of view. Figure 6.15 depicts the inverse demand function—or the direct demand function, depending on your point of view.

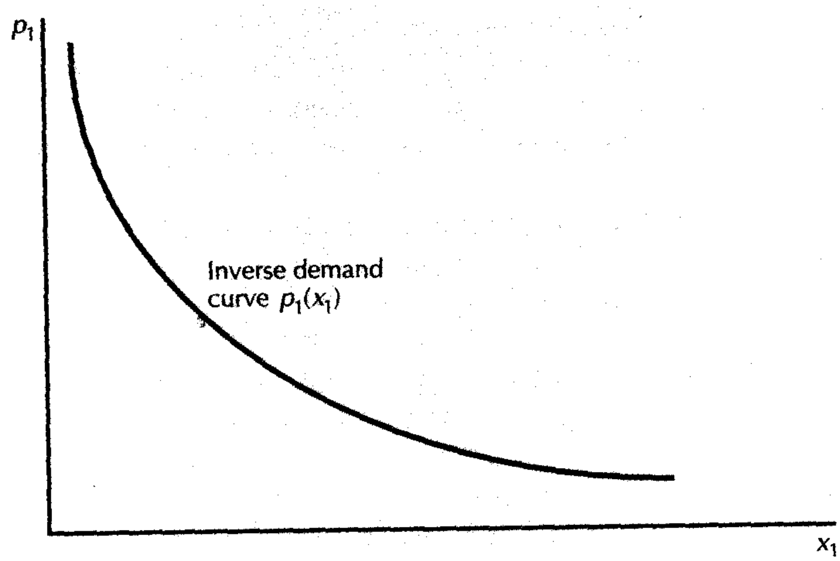

Inverse demand curve. If you view the demand curve as measuring price as a function of quantity, you have an inverse demand function.

Figure 6.15

Recall, for example, the Cobb-Douglas demand for good 1, $x_1 = am/p_1$. We could just as well write the relationship between price and quantity as $p_1 = am/x_1$. The first representation is the direct demand function; the second is the inverse demand function.

The inverse demand function has a useful economic interpretation. Recall that as long as both goods are being consumed in positive amounts, the optimal choice must satisfy the condition that the absolute value of the MRS equals the price ratio:

$$|\text{MRS}| = \frac{p_1}{p_2}.$$

This says that at the optimal level of demand for good 1, for example, we must have

$$p_1 = p_2|\text{MRS}|. \tag{6.4}$$

Thus, at the optimal level of demand for good 1, the price of good 1 is proportional to the absolute value of the MRS between good 1 and good 2.

Suppose for simplicity that the price of good 2 is one. Then equation (6.4) tells us that at the optimal level of demand, the price of good 1 measures how much the consumer is willing to give up of good 2 in order to get a little more of good 1. In this case the inverse demand function is simply measuring the absolute value of the MRS. For any optimal level of x_1 the inverse demand function tells how much of good 2 the consumer would want to have to compensate him for a small reduction in the amount of good 1. Or, turning this around, the inverse demand function measures how much the consumer would be willing to sacrifice of good 2 to make him just indifferent to having a little more of good 1.

If we think of good 2 as being money to spend on other goods, then we can think of the MRS as being how many dollars the individual would be willing to give up to have a little more of good 1. We suggested earlier that in this case, we can think of the MRS as measuring the marginal willingness to pay. Since the price of good 1 is just the MRS in this case, this means that the price of good 1 itself is measuring the marginal willingness to pay.

At each quantity x_1, the inverse demand function measures how many dollars the consumer is willing to give up for a little more of good 1; or, said another way, how many dollars the consumer was willing to give up for the last unit purchased of good 1. For a small enough amount of good 1, they come down to the same thing.

Looked at in this way, the downward-sloping demand curve has a new meaning. When x_1 is very small, the consumer is willing to give up a lot of money—that is, a lot of other goods, to acquire a little bit more of good 1. As x_1 is larger, the consumer is willing to give up less money, on the margin, to acquire a little more of good 1. Thus the marginal willingness to pay, in the sense of the marginal willingness to sacrifice good 2 for good 1, is decreasing as we increase the consumption of good 1.

Summary

1. The consumer's demand function for a good will in general depend on the prices of all goods and income.

2. A normal good is one for which the demand increases when income increases. An inferior good is one for which the demand decreases when income increases.

3. An ordinary good is one for which the demand decreases when its price increases. A Giffen good is one for which the demand increases when its price increases.

4. If the demand for good 1 increases when the price of good 2 increases, then good 1 is a substitute for good 2. If the demand for good 1 decreases in this situation, then it is a complement for good 2.

5. The inverse demand function measures the price at which a given quantity will be demanded. The height of the demand curve at a given level of consumption measures the marginal willingness to pay for an additional unit of the good at that consumption level.

REVIEW QUESTIONS

1. If the consumer is consuming exactly two goods, and she is always spending all of her money, can both of them be inferior goods?

2. Show that perfect substitutes are an example of homothetic preferences.

3. Show that Cobb-Douglas preferences are homothetic preferences.

4. The income offer curve is to the Engel curve as the price offer curve is to ...?

5. If the preferences are concave will the consumer ever consume both of the goods together?

6. Are hamburgers and buns complements or substitutes?

7. What is the form of the inverse demand function for good 1 in the case of perfect complements?

8. True or false? If the demand function is $x_1 = -p_1$, then the inverse demand function is $x = -1/p_1$.

APPENDIX

If preferences take a special form, this will mean that the demand functions that come from those preferences will take a special form. In Chapter 4 we described quasilinear preferences. These preferences involve indifference curves that are all parallel to one another and can be represented by a utility function of the form

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

The maximization problem for a utility function like this is

$$\max_{x_1, x_2} v(x_1) + x_2$$

$$\text{s.t. } p_1 x_1 + p_2 x_2 = m.$$

Solving the budget constraint for x_2 as a function of x_1 and substituting into the objective function, we have

$$\max_{x_1} v(x_1) + m/p_2 - p_1 x_1/p_2.$$

Differentiating gives us the first-order condition

$$v'(x_1^*) = \frac{p_1}{p_2}.$$

This demand function has the interesting feature that the demand for good 1 must be independent of income—just as we saw by using indifference curves. The inverse demand curve is given by

$$p_1(x_1) = v'(x_1)p_2.$$

That is, the inverse demand function for good 1 is the derivative of the utility function times p_2. Once we have the demand function for good 1, the demand function for good 2 comes from the budget constraint.

For example, let us calculate the demand functions for the utility function

$$u(x_1, x_2) = \ln x_1 + x_2.$$

Applying the first-order condition gives

$$\frac{1}{x_1} = \frac{p_1}{p_2},$$

so the direct demand function for good 1 is

$$x_1 = \frac{p_2}{p_1},$$

and the inverse demand function is

$$p_1(x_1) = \frac{p_2}{x_1}.$$

The direct demand function for good 2 comes from substituting $x_1 = p_2/p_1$ into the budget constraint:

$$x_2 = \frac{m}{p_2} - 1.$$

A warning is in order concerning these demand functions. Note that the demand for good 1 is independent of income in this example. This is a general feature of a quasilinear utility function—the demand for good 1 remains constant as income changes. However, this can only be true for some values of income. A demand function can't literally be independent of income for *all* values of income; after all, when income is zero, all demands are zero. It turns out that the quasilinear demand function derived above is only relevant when a positive amount of each good is being consumed.

In this example, when $m < p_2$, the optimal consumption of good 2 will be zero. As income increases the marginal utility of consumption of good 1 decreases. When $m = p_2$, the marginal utility from spending additional income on good 1 just equals the marginal utility from spending additional income on good 2. After that point, the consumer spends all additional income on good 2.

So a better way to write the demand for good 2 is:

$$x_2 = \begin{cases} 0 & \text{when } m \leq p_2 \\ m/p_2 - 1 & \text{when } m > p_2 \end{cases} .$$

For more on the properties of quasilinear demand functions see Hal R. Varian, *Microeconomic Analysis,* 3rd ed. (New York: Norton, 1992).