11
Demand

Consider the source of our economic behavior. What exactly motivates us to engage in economic activity? A large part of the answer to this question lies in one word: ***consumption.*** We need to have goods and services to survive and also want to enjoy a standard of living commensurate with both our expectations and our dreams. For some people, a small amount of consumption will do, but most of us want to have more of the good things that our economic system can provide.

Utility

From this basic assumption about human behavior, we can infer that any level of consumption brings with it a certain amount of satisfaction—or, as economists say, ***utility.*** Utility can be a difficult thing to measure precisely. It is also difficult to make comparisons of utility between people. Your utility from consuming an automobile or a banana is probably going to be different from mine or from anyone else's. Indeed, we might discover some surprising results if we could actually measure utility between people. Depriving you of your fourth Cadillac, for example, might result in a greater loss of satisfaction to you than I would

feel if I were to lose my one Ford. Even though you have four Cadillacs, the loss of one might depress you because all of your neighbors have five or six.

Nevertheless, we can say some important things about utility, as long as we confine our discussion to *one individual*. For example, Mary Smith might tell us that she "seems" to be getting about 5 units of utility from eating a hamburger and around 2 units from consuming a banana. To determine if she is accurate about her utility estimates, we will give her $0.70 and offer her the opportunity to buy fractions of hamburgers or bananas. If we price a full hamburger at $0.50 and a banana at $0.20 and Mary spends her $0.70 for one hamburger and one banana, we can assume that her rough utility ratio of 5 : 2 is correct.

Now let's make another utility observation, but this time we will only consider the consumption of one product, such as the Cadillacs you own. Even though we can't determine how much utility your fourth Cadillac gives you compared to the utility my one Ford gives me, we can say that you probably have received more utility from your *first* Cadillac than you will from your fourth. Or, to consider a more realistic example, Mary's third hamburger (in the short run) will probably give her less satisfaction than her second, and her second hamburger will probably give her less satisfaction than her first.

But since I know more about my own utility levels than yours, or Mary Smith's, I will turn to a more detailed illustration of this principle, using my own utility preferences. Thus, after considerable thought and experimentation, I come up with the following **marginal utility chart**:

Hamburger	Marginal Utility (MU_h)
1	7
2	3
3	1
4	0

From these figures, you can see that my first hamburger gives me 7 units of utility. The second provides 3 additional units, but by the time I eat my third, my *additional utility* is only 1 unit. Economists call this additional utility (which is assigned to the consumption of a specific unit) ***marginal utility,*** or simply MU. Obviously, I'm starting to get pretty full after two hamburgers; after my third hamburger, I am completely full, so the fourth will give me no marginal utility whatsoever. If somebody *gave* me that fourth hamburger, I'd leave it on my plate. What would your utility chart for hamburgers look like? You might want to construct one just for fun.

The Law of Diminishing Marginal Utility

Can we make any generalization about this utility pattern? Apparently, as a person consumes more of a given product, that

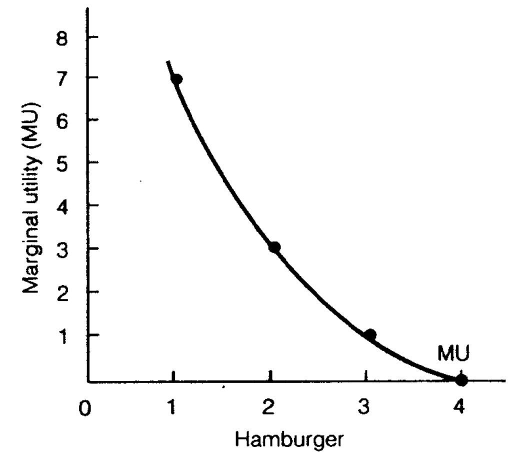

Figure 11-1 The *law of diminishing marginal utility:* as a person consumes more and more hamburgers, he or she experiences less and less marginal utility from each additional hamburger.

product has less marginal utility for that individual. This principle will probably be true, no matter what person or product we examine. Economists call this universal principle the **law of diminishing marginal utility**.

We can apply this principle to many types of human experience. Diminishing marginal utility is worth remembering if you happen to be on a diet, since that tenth spoonful of ice cream (or fourth cookie) will undoubtedly give you less satisfaction than the first. It may also help to explain why good marriages sometimes go bad or why that once "exciting" job eventually becomes boring.

Sometimes it's helpful to visualize economic relationships such as this one. Using the MU data just given, I can graph my *diminishing marginal utility curve* for hamburgers as shown in Figure 11-1.

Keep in mind that this graph is *my* MU curve for hamburgers; your curve, or someone else's, would probably look different. Spend a moment sketching different possible MU curves. Try one for a "Wimpy," who can lovingly eat a dozen hamburgers before he gets full. What about a vegetarian? What about your MU curve?

Now we can begin to make some interesting observations about how a consumer might behave when given a choice between two products. How, for example, would the law of diminishing marginal utility help me to maximize my satisfactions, given a limited income?

To see how this is done, let's add a second item to my menu—milkshakes. Let's assume that my marginal utility chart for milkshakes looks something like this:

Milkshake	Marginal Utility (MU_m)
1	12
2	3
3	1/2

We will make this consumer example more realistic by adding an income constraint. To keep our example simple, we will assume that I am given, say, $4 per day and that the price of a milkshake is $1 and the price of a hamburger is $1. With this information, how do I go about maximizing my total utility? How do I spend my limited income in a way that I can enjoy the highest possible level of total satisfaction?

The best method of maximizing my utility would be to use what economists call ***marginal decision making***—to make a separate decision for each dollar at my disposal. Thus, I take my first dollar and ask the question, "Where will this dollar give me the greatest marginal utility?"

If you compare the utility chart for milkshakes with the utility chart for hamburgers, you will see that I ought to spend my first dollar on a milkshake, because that first milkshake will give me 12 units of utility. (If I spend that same dollar on a hamburger, it will only give me 7 units of utility.) My marginal decision making now leads me to ask, "How can I best spend my second dollar?" "My third dollar?" "My fourth dollar?" A summary of my decisions follows.

My second dollar will be spent on a hamburger. The third dollar is a "toss-up," since both products will give me the same marginal utility (3 units) per dollar. (In the example, I will choose a milkshake for my third dollar.) My fourth dollar, in a sense, "balances things out," so that *once my total income of $4 is spent,* the marginal utility per dollar's worth of each product

	Milkshake		Hamburger		
dollar 1 →	1	12 units	1	7 units	← dollar 2
dollar 3 →	2	3 units	2	3 units	← dollar 4
	3	1/2 unit	3	1 unit	
			4	0 units	

(3 units per $1) will be equal. Obviously, if I could obtain more utility from spending my last dollar on another product (french fries, for example), I would want to do so.

Let's assume that we have precise information about all the products we wish to consume (and therefore know the marginal utility of each product) and that we can spend our money exactly as we wish (even for fractions of hamburgers or milkshakes). Then we can attain the highest possible total satisfaction *if* the marginal utility per dollar's worth (marginal utility divided by product price, or MU/*P*) of product A is equal to the MU/*P* of product B, which is equal to the MU/*P* of product C, and so on. When these ratios are equal, so that

$$\frac{MU_A}{P_A} = \frac{MU_B}{P_B} = \frac{MU_C}{P_C} = \frac{MU_D}{P_D}$$

A = hamburgers
B = milkshakes
C = french fries
D = other things

then we, as consumers, have, in a sense, "solved" our maximization problem. We have spent our limited income in a way that gives us the maximum amount of utility.

So far, so good. But there are still people who are bothered by our inability to measure utility precisely, as we just tried to do in this example. Fortunately, we can use another method, called the ***indifference-curve*** *approach,* to determine consumer efficiency without giving actual utility values. Let's take a look.

Indifference Curves

Economists have a little of "the psychologist" in them, as well as a little of "the newspaper reporter." An economic researcher could conceivably run around asking people about their income levels, consumption habits, work preferences, personal values, and so on. The answers to these questions, in turn, would give the researcher insight into how consumers behave under various economic conditions. At some point, our economist-psychologist-reporter might even be able to discover some generalized principles (such as the law of diminishing marginal utility) that could become the basis of an important economic theory. The *indifference-curve approach* is such a technique. It allows us to ask some simple questions and derive some interesting generalizations and conclusions from the answers. Let's look at an example.

You are the economic researcher, and you ask me the following question: "If I gave you 1 milkshake and 3 hamburgers, you would derive a certain amount of utility from that combination, right?" I answer "Yes."

Then you go on: "Let's call that amount of utility your **total utility level** *Y*. Now if I reduce the number of hamburgers to 2, how many *additional* milkshakes would you need to keep yourself at total utility level *Y*?" Suppose I answer, "I will need an extra 1/4 milkshake to make up for the lost hamburger." This means that I am *totally indifferent* about whether I consume a combination of 3 hamburgers and 1 milkshake or a combination of 2 hamburgers and 1-1/4 milkshakes.

As a final question, you might ask me how many milkshakes I would need if I consumed only 1 hamburger but wanted to stay at total utility level *Y*? Let's say that I would need 2 full milkshakes to be indifferent to the other combinations. The following chart shows the results of your research (h = hamburger; m = milkshake):

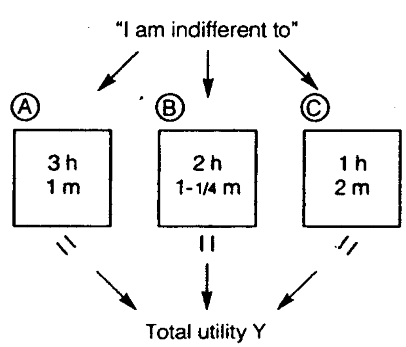

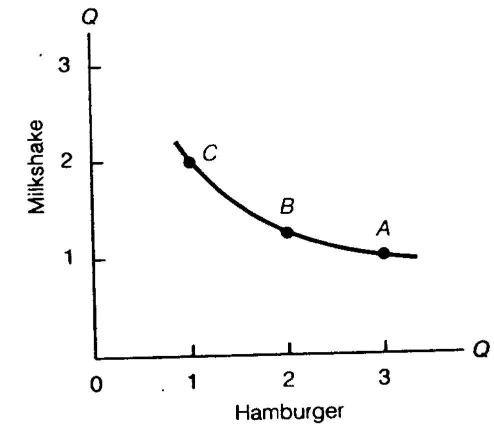

FIGURE 11-2 At points along the *indifference curve*, any combination of milkshakes and hamburgers offers the consumer an equivalent amount of total utility. The consumer is therefore indifferent to consuming at points *A*, *B*, or *C* (or at any other point on the indifference curve).

Thus, I would be equally well off with combination A, B, or C. I would be indifferent about consuming at any of these combinations because my total utility stays the same in each case. Now let's graph these "points" of indifference, measuring hamburgers on the horizontal axis and milkshakes on the vertical axis (see Figure 11-2). Connecting the points that represent combinations A, B, and C gives us a smooth **indifference curve** showing all the points of indifference. Suppose, for example, that you ask me, "Would you prefer to consume at point *A*, *B*, or *C* on the graph, or somewhere in between?" I would have to answer by saying, "I'm indifferent; all points give me equal satisfaction."

My indifference curve in Figure 11-2 might be described as having a "bow-like" or lazy C shape. Economists say that such a curve is "convex to the origin" (the origin is always in the lower left-hand corner of the graph). You may wonder why my indifference curve (or anybody else's, for that matter) has this general shape. Why isn't it a straight line or a "dome"? We will be able to answer this question once we understand the law of diminishing marginal utility. Let's see how it works.

Look closely again at point *A* in Figure 11-2, which represents the combination of 3 hamburgers and 1 milkshake. If I happen to be consuming at point *A*, I am obviously "full" of hamburgers; remember that this third hamburger added very little to my overall satisfaction. Since this third hamburger is

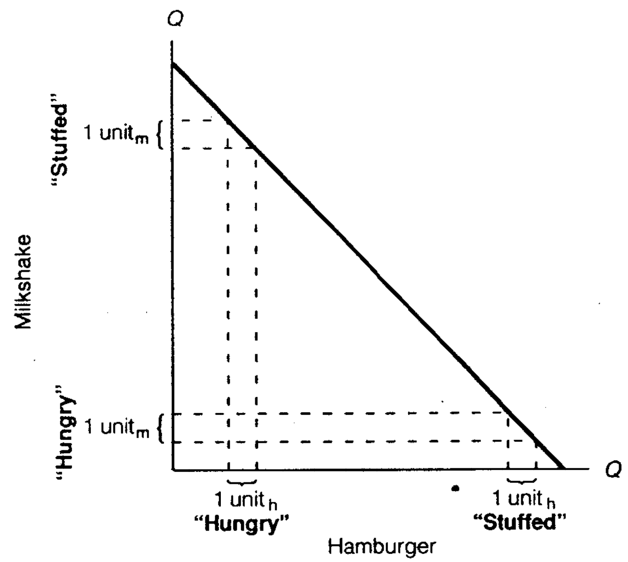

FIGURE 11-3 A straight-line indifference curve violates the law of diminishing marginal utility. It implies that consumers will trade off equal amounts of milkshakes, whether they are "stuffed" with hamburgers or "hungry" for them.

not that important to me, I am willing to "trade it off," or substitute it, for just one-fourth of a milkshake if I move to consumption point *B*.

Now let's look at point *C*, which represents 1 hamburger and 2 milkshakes. Here, hamburgers suddenly become much "dearer" to me, while milkshakes are less important due to the law of diminishing marginal utility. Thus, at point *C*, I am willing to give up a greater quantity of milkshakes (3/4 of a milkshake) to get that second hamburger.

If the law of diminishing marginal utility is working (and we are assuming that it is), then the indifference curve will have to be convex to the origin. An interesting test of this reasoning is to intentionally convert an indifference curve to a straight line (like the one in Figure 11-3) and then prove that this *cannot* be a valid shape.

Can you see why the line in Figure 11-3 violates the law of diminishing marginal utility? A straight line implies *equal* trade-offs of hamburgers for milkshakes, whether I am "hungry" for burgers or "stuffed" with burgers. But such a *continuous* one-for-one trade-off just doesn't conform to reality. The only configuration consistent with the law of diminishing marginal utility (and the actual behavior of consumers) is our original convex curve.

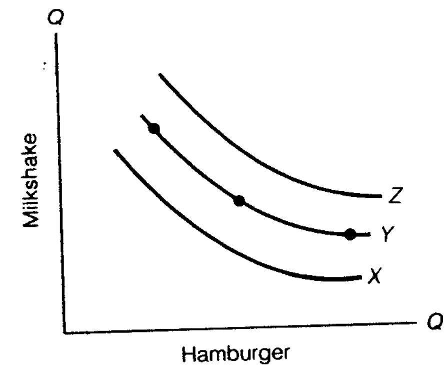

FIGURE 11-4 An *indifference map* is a unique "fingerprint," showing an individual's preferences—in this case, for milkshakes and hamburgers. In general, as indifference increases, total utility increases.

Indifference Map

Needless to say, we could draw thousands of indifference curves, each one reflecting the unique consumption preferences of a different individual. We could even show many *different levels* of utility for the same person. Up to this point, we have only discussed a single indifference curve that represents total utility at the *Y* level. There is also a curve somewhere below the *Y* level that would represent a lower total level of satisfaction (let's call it total utility at the *X* level) and another curve showing even higher satisfaction than *Y* (we'll call it total utility at the *Z* level). Of course, we could draw many other indifference curves above, below, and between these three curves. These different levels of utility, represented by a series of indifference curves, make up an **indifference map**—a kind of "consumer preference fingerprint"—on which each person displays a unique set of indifference curves. An indifference map for milkshakes and hamburgers is shown in Figure 11-4.

One important advantage of using an indifference map instead of our earlier marginal decision-making approach is that we now no longer need to assign actual utility values to the consumption of different products. For example, we really do not

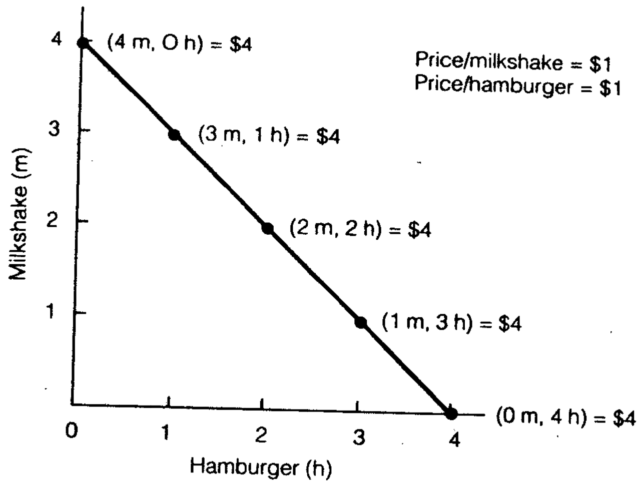

Figure 11-5 The *budget line:* here, a $4 budget allows the consumer to select from a variety of milkshake-hamburger combinations. If both milkshakes and hamburgers are priced at $1 per unit, then the budget line indicates all the consumption possibilities that a consumer can purchase with $4.

know how many *units of utility* the *X* curve represents, but we can still say with some certainty that indifference curve *X* is lower than indifference curve *Y*. Thus, all other things being equal, *an individual prefers to consume on the highest indifference curve.* The highest level of satisfaction in this illustration is indifference curve *Z*.

The Budget Line

Let's return to the example in which I have an income of $4 per day and hamburgers and milkshakes cost $1 apiece. If I spend all of my daily income on hamburgers and buy no milkshakes, I can buy 4 hamburgers. This means that my $4 budget allows me to operate at a consumption level of 4 hamburgers and 0 milkshakes. Of course my $4 could also buy 2 hamburgers and 2 milkshakes, or 3 hamburgers and 1 milkshake, or 3 milkshakes and 1 hamburger. These combinations all represent possible consumption levels, given my $4 income. Plotting these various combinations results in a **budget line,** which shows every combination of hamburgers and milkshakes that I can purchase for $4. Figure 11-5 summarizes these data.

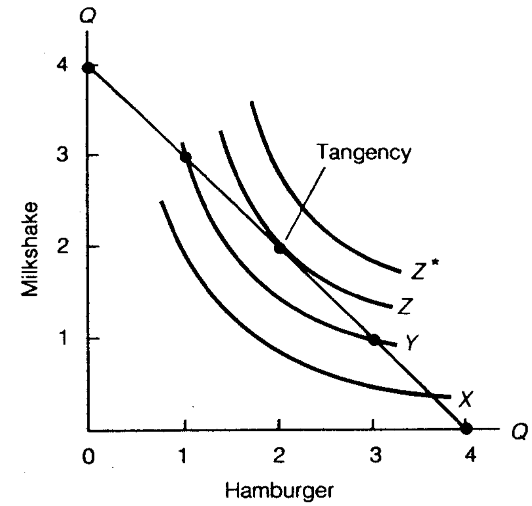

Figure 11-6 A consumer obtains the greatest total utility from a $4 income at the point at which the highest indifference curve *Z* just touches (is tangent to) the $4 budget line.

Sometimes it can be helpful to look at a budget line as if it were a kind of "economic straightjacket"—a visual representation of the "cruel world" of economic reality. Of course we would all like our budgets to be larger, but they are not. We have just so much money to spend, and we must limit our consumption to the possible product combinations that lie somewhere on the budget line.

So even though we are limited by budget restrictions, we still have a certain amount of *choice* in terms of selecting the right "bundle" of goods that will give us the greatest satisfaction. This is simply another way of looking at the fundamental economic problem of **utility maximization**. How do we solve the problem this time? How can I be sure I have chosen the best product combination with my $4 income?

To answer these questions, all we need to do is combine the budget line with the indifference map. The combined system is shown in Figure 11-6.

The solution can easily be seen on the graph. First note that it *is possible* to consume on indifference curve *Y*. It crosses the $4 budget line in two places. But why should the consumer stop on indifference curve *Y* when it is also possible to climb up to indifference curve *Z*? Notice that there is only one point where

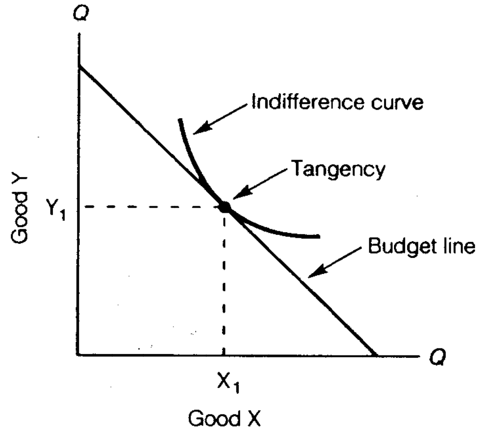

Figure 11-7 The total utility obtained from consuming good Y and good X is maximized at the point at which the indifference curve is tangent to the budget line.

the budget line is just tangent to indifference curve *Z*; this point of tangency represents the approximate consumption of 2 hamburgers and 2 milkshakes.

Yet someone might logically ask, "If higher indifference curves represent higher satisfaction, why don't you just move to the highest indifference curve of all (*Z*)?" The answer, of course, is that *Z** is not consistent with the $4 budget constraint. As you can see, the highest indifference curve does not coincide with the budget at any point on the graph. The very highest possible level of utility that I can attain with my $4 budget is the *Z* level shown by the point of tangency in Figure 11-6. Thus, we can say that *individuals maximize their utility by consuming at the point where the indifference curve is tangent to the budget line*.

Total utility will always be maximized at this point of tangency. It does not matter which indifference map we are looking at, what the dollar income constraint is, or what particular goods we are using in our example. In fact, it might be helpful to graph a more "generic" representation, using the more generalized X and Y goods (see Figure 11-7).

Now that we have established the point of utility maximization, or ***consumer efficiency,*** we can demonstrate a variety of economic relationships. You'll be surprised to see how versatile the indifference-curve/budget-line format can be in describing consumer behavior. For example, how can we analyze a change in a consumer's income? Let's take a look.

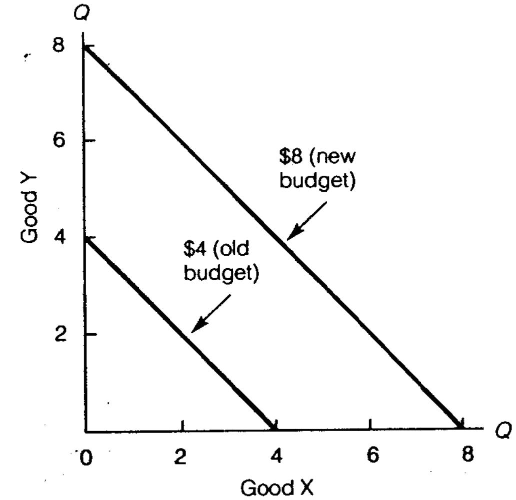

Figure 11-8 An expanding budget is represented by a parallel, rightward shift in the budget line.

Income Change

We will assume that our typical consumer (Chester Olson), who had an income of $4 per day, is given a raise to $8 per day. How do we show this change within the format we developed earlier?

Such an income increase can be shown by a *rightward shift* in Chester's budget line. Keeping the price of good X at $1 and the price of good Y at $1, we can see that the new point of reference on the *x* axis (for the $8 income) is now 8 units. This means that if Chester spent all his income on good X, he would be able to buy 8 units. The same is true of good Y. In Figure 11-8, we can easily see the difference between the old and the new budget lines.

Now let's take a variety of income levels and examine the points of tangency with their respective indifference curves. The points of tangency in Figure 11-9(a) show us exactly where Chester will maximize satisfaction at different levels of income.

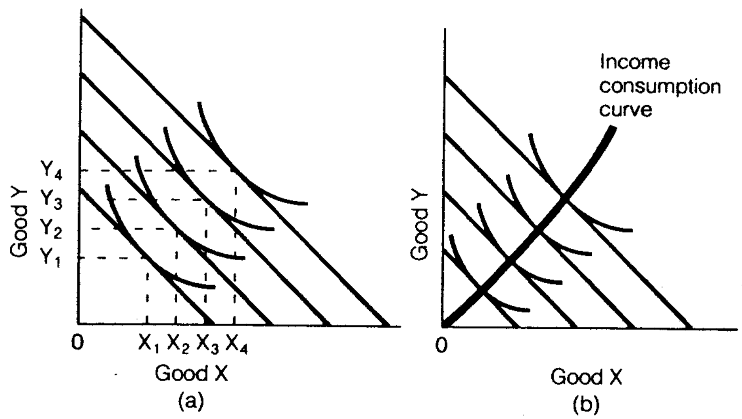

FIGURE 11-9 (a) Chester maximizes total utility as his income increases. (b) Connecting the points at which the indifference curves are tangent to the budget lines yields the *income-consumption curve*.

In Figure 11-9(b), note that we have connected all the points of tangency with a continuous line. A glance at this ***income-consumption curve*** shows the relative preference between the two goods as Chester's income increases.

In Figure 11-9(b), it looks as if Chester's preferences are fairly "balanced" between the two products. In other words, both good X and good Y in this example are what economists call ***normal goods***. When income levels increase, consumers tend to buy *more* of a normal good.

Economists also recognize the possibility that when income levels increase, consumers may purchase *less* of an ***inferior good***. Macaroni, powdered-milk, used cars over ten years old, retread tires, and used books are all inferior goods. Note that an inferior good does not always have to be inferior in terms of quality; powdered milk, for example, is highly nutritious and is often recommended for low-fat diets. Generally speaking, though, an inferior good tends to be a "poor person's product"; families tolerate these goods at low-income levels, but as their incomes rise, they tend to discard inferior goods in favor of normal goods. In Figure 11-10, we have expanded Chester's income from $2 to $8 per day. We can see that Chester purchases the inferior good (macaroni) less and less as his income expands. The income-consumption curve rises in a *leftward* direction, indicating that Chester is maximizing his satisfaction with fewer and fewer inferior goods as his budget increases.

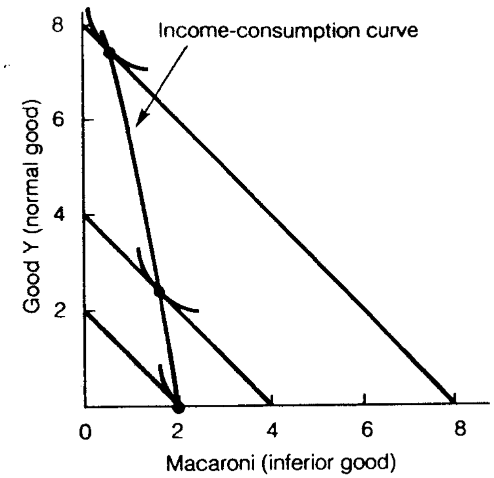

FIGURE 11-10 If income increases and consumers begin to buy less of a certain good, economists say that the product is an *inferior good*. Here, the consumption of macaroni, an inferior good, decreases as income increases.

Demand

Income changes are interesting microeconomic concepts but do not help us get to the heart of the demand curve, which is the major goal of this chapter. What is the key variable to understanding demand?

To answer this question, it might be helpful to review the fundamental nature of a demand curve. Recall from Chapter 3 that **demand** is the relationship between the *price P* of a good and the *quantity Q* of the good that is purchased. Can we derive a demand curve for hamburgers, using our friend Chester Olson as an illustration? Fortunately, this is not too difficult.

Perhaps the easiest method is to play "economist-reporter" and simply ask Chester how many hamburgers he will buy at different prices. Let's assume he tells us that he will buy 1 hamburger if the unit price is $2, 2 hamburgers if the unit price goes down to $1, and 4 hamburgers if the unit price drops to $0.50:

QUANTITY (*Q*)	PRICE (*P*)
1	$2
2	1
4	0.50

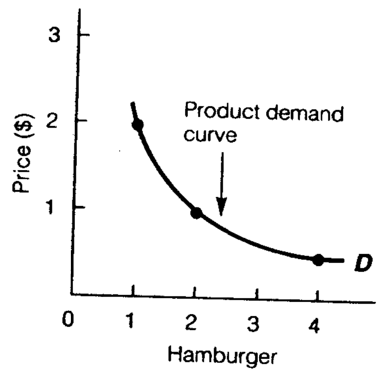

Figure 11-11 Information about consumer demand can be obtained by simply asking a person how much of a good he or she will buy at different prices and then plotting the price-quantity points on a demand curve.

This "direct research" method of determining Chester's hamburger demand gives us the curve shown in Figure 11-11.

We can also use the indifference-curve/budget-line format to determine Chester's hamburger demand. First, we find out what Chester's indifference map looks like; then all we have to do is *vary the unit price of hamburger* and observe how these price changes affect Chester's consumption.

To see how to show a price change on our indifference-curve/budget-line graph, let's return to our example of a $4 budget and an original price of $1 for a hamburger. If we lower the price of a hamburger to $0.50, Chester can buy a maximum of 8 hamburgers with his $4 income. We can also easily determine where the budget line intersects the *x* axis if we divide Chester's income by the price of a hamburger. Thus, if the price of a hamburger increases to $2, Chester can buy a maximum of 2 hamburgers with his $4 income ($4/$2 = 2).

Each time the unit price of hamburger changes, *the slope of the budget line changes.* At lower hamburger prices, it generally has a lower slope; at higher prices, the budget line becomes steeper. (We are assuming, of course, that the unit price of good Y does not increase or decrease.) We can see the slope changing in Figure 11-12.

Now all we have to do is trace Chester's indifference-curve map, drawing in the indifference curves that are tangent to the different budget lines. When we connect these points of tangency with a line, we have what economists commonly call a ***price-consumption curve,*** as seen in Figure 11-13(a).

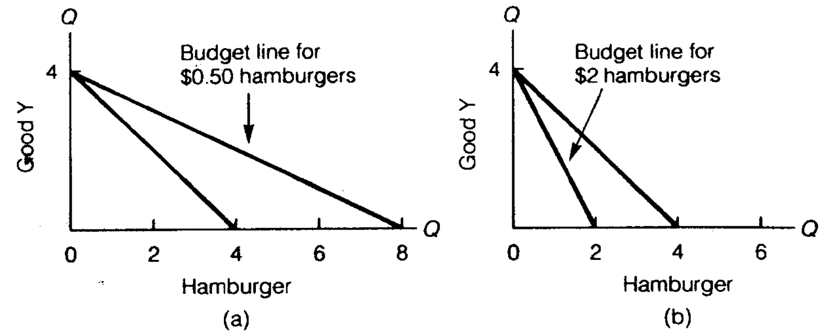

Figure 11-12 Changing the unit price of a hamburger: (a) if the price is lower ($0.50), then the $4 budget line is less sloped; (b) if the price is higher ($2), then the budget line is steeper.

Deriving the Demand Curve

After our long labors, we are now close to deriving a demand curve from an indifference-curve system. In fact, the last connecting link is really quite simple; perhaps you have already spotted it. All we need to do to find Chester's demand is to read off the number of hamburgers that he will consume at the different unit prices of hamburger, shown in Figure 11-13(a). Thus, at the $2 price (the steepest budget line), we see that Chester will demand 1 hamburger. If we draw a vertical dashed line down to the horizontal (hamburger) axis from the point of tangency of the lowest sloped line (representing $0.50 hamburgers), we see that Chester will demand 4 hamburgers.

These results, which are summarized by the product demand curve in Figure 11-13(b), are in exact agreement with our experimental method of finding demand by direct research (see Figure 11-11). Both of these approaches to finding Chester's demand are valid, but the indifference-curve method took us back to our study of consumption and was built up, in a sense, "from scratch." We began with the consumer's desire to maximize their utility within a budget limitation. We then observed the effects of the law of diminishing marginal utility, which became incorporated into the special C-shape (convex to the origin) of the indifference curves. At this point, we added budget lines (income constraints) to our indifference-curve graph and noted the point of consumer efficiency at the tangency of each budget line and indifference curve. Finally, we changed the price

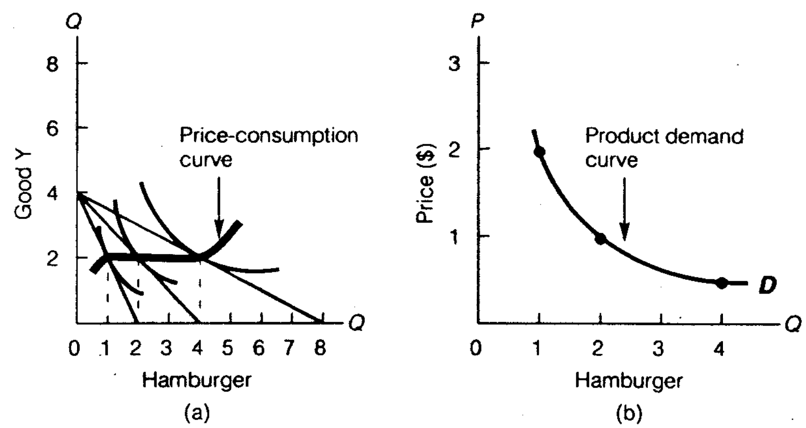

FIGURE 11-13 (a) If we change the price of a hamburger (represented by the three budget lines with different slopes) and then connect the points at which the budget lines are tangent to the indifference curves, we can derive three points on a *price-consumption curve.* (b) From the information given in the graph in (a), we can determine how many hamburgers will be purchased at the three different prices. When we graph these points, we obtain a *product demand curve* for hamburgers. Compare this method with the one used in Figure 11-11.

of good X (hamburgers) and noted the change in quantity demanded. From that information we derived the demand curve shown in Figure 11-13(b).

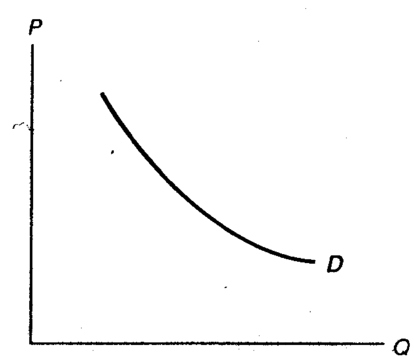

This completes our study of demand. We have discovered where a sample demand curve "comes from" and have learned some interesting variations related to maximizing satisfaction (inferior goods versus normal goods) when a person's income changes. We have also developed some valuable tools that will make our study of production theory much easier later in the book. But before we continue, let's look again at our accomplishment—an individual's product demand curve (a friendlier curve now)—and remind ourselves how far we have come.