The Behavior of Consumers

3

When you study economics, people expect you to be able to predict things. So here's an exercise to test your prediction skills: A man goes into a restaurant. What does he order?

If that seems impossible to answer, it's because you don't have enough information. In fact, you're missing *two* kinds of information: First, what does the man like to eat? Second, what's available on the menu? In other words, you need to know something about the man's *preferences* and you need to know something about his *opportunities*.

Neither sort of information is terribly useful on its own. You might know that the man prefers fish to pork and pork to pasta, but you still can't predict what he'll order if you don't know what's on the menu. Or, you might know that the menu lists lion, giraffe, buffalo, musk ox, wild boar, black bear, Malaysian frog's legs, yak, elk, ostrich, and Egyptian cobra,[1] but you still can't predict what he'll order unless you know how he feels about eating lion.

The restaurant patron has a relatively simple choice to make: Which of several items should I order? A shopper at a grocery store has to decide something a little more complicated: How *much* of each item should I buy? Instead of the 10 or 12 choices that might appear on a restaurant menu, the shopper faces in principle an infinite number of possibilities. We think of "4 pounds of cherries and 3 pounds of plums" as a single possible menu item. Other menu items might include "3 pounds of cherries and 4 pounds of plums" or "3.5 pounds of cherries and 3.5 pounds of plums." The available menu consists of the combinations that the consumer can afford to buy, given his income and given the prices of the goods.

In this chapter, we'll develop methods of keeping track of consumers' preferences across such combinations of goods. Then we'll develop

[1]These are the menu items at the Panache Restaurant in Killington, Vermont.

methods for keeping track of the available menu, and finally we'll combine the two sorts of information to make predictions about consumers' behavior.

3.1 Tastes

The Latin proverb "*De gustibus non est disputandum*" can be translated as "There's no accounting for tastes." Some people like antique wooden furniture, whereas others prefer brass. You are likely to get a variety of answers if you ask different friends whether they would prefer to live in a world without Bach or in a world without clean sheets.

Economists accept the wisdom of the proverb and make no attempt to account for tastes. Why people prefer the things that they do is an interesting topic, but it is not one that we will explore. We take people's tastes as given and see what can be said about them.

Indifference Curves

Imagine a consumer named Beth who lives in a world with only two goods: eggs and root beer. You might imagine asking Beth which of these two goods she likes better. But although this question sounds sensible at first, it really isn't. There are many reasons why not. First, the answer is likely to depend on what quantities of eggs and root beer are being compared. Second, the answer is likely to depend on how much of each Beth happens to own already. The question is open to several interpretations: Are we asking which good Beth would least like to do without altogether, or are we asking which she would rather receive for her birthday?

Here is a better question: We can ask Beth whether she'd rather own a basket of 3 eggs and 5 root beers or a basket with 4 eggs and 2 root beers. In principle, we could discover the answer by taking away all of Beth's possessions and then offering her a choice between the two baskets. A question makes sense when some (possibly imaginary) experiment is capable of revealing the answer.

Of course, there are many possible baskets besides the ones we've described. We can display all of them simultaneously on a graph, as in Exhibit 3.1. Each point on that graph represents a basket containing a certain number of eggs and a certain number of root beers. For example, point A represents a basket with 3 eggs and 5 root beers—the first of the 2 baskets we offered to Beth.

EXERCISE 3.1

Describe the baskets represented by points B, C, and D. Which represents the second basket of our imaginary experiment?

What can we say about Beth's preferences among these baskets? Compare basket A to basket B, for example. Which would she prefer to own? Basket B contains more eggs than basket A (4 units instead of 3) and also more root beers (7 instead of 5). *If we assume that eggs and root beer are both* **goods**—items that Beth would prefer to have more of whenever she can—then the choice is unambiguous. Basket B is better than basket A.

Goods Items of which the consumer would prefer to have more rather than less.

EXHIBIT 3.1 Basket of Goods

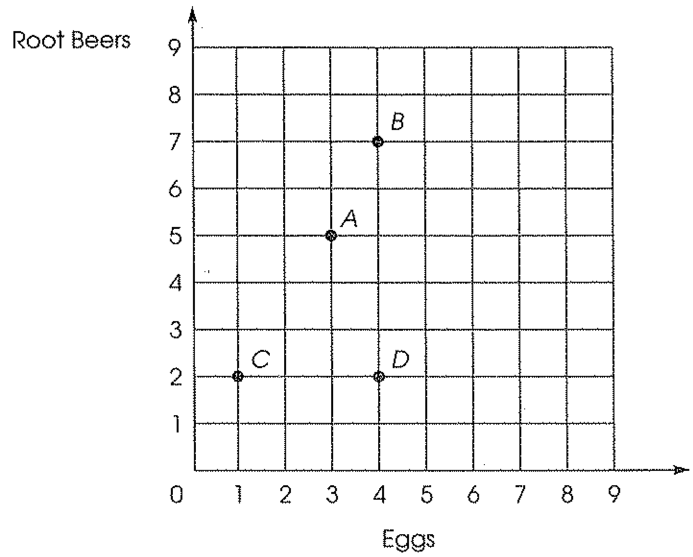

Each point on the graph represents a basket containing a certain number of eggs and a certain number of root beers. For example, point *A* corresponds to 3 eggs and 5 root beers.

EXERCISE 3.2

Which is preferable, basket *A* or basket *C*? How do you know?

When it comes to comparing basket *A* with basket *D*, the choice is less clear. Basket *D* has more eggs (4 units versus 3) but less root beer (2 versus 5). Which will Beth prefer? At this point we cannot possibly say. She might like *A* better than *D*, or *D* better than *A*. It is also possible (though not necessary or even likely) that she would happen to like them both equally.

Now consider this question: Where should we look to find the baskets that Beth likes exactly as much as *A*? They can't be to the "northeast" of *A* (like *B*) because the baskets there are all preferred to *A*. They can't be to the "southwest" of *A* (like *C*) because *A* is preferred to all of those baskets. They must all be to either the "northwest" or the "southeast" of *A* (like *D*). This doesn't mean that *D* is necessarily one of them, just that they lie in the same general direction from *A* that *D* does.

If we draw in a few of the baskets that Beth likes just as well as she likes *A*, they might look like the points shown in panel A of Exhibit 3.2. Because each of these baskets is *exactly as good* as *A*, they must all be *exactly as good* as each other. This means that each one must lie either to the northwest or to the southeast of each other one, which accounts for the downward slope that is apparent in the picture.

The baskets shown in panel A of Exhibit 3.2 are only a few of those that Beth likes just as well as *A*. There are many other such baskets as well. The collection of all such baskets forms a curve, shown in black in panel B of the exhibit. From our discussion in the preceding paragraph, we know that the curve will be downward sloping. Because Beth is indifferent between any two points on this curve, it is called an **indifference curve**.

Indifference curve A collection of baskets, all of which the consumer considers equally desirable.

EXHIBIT 3.2 Comparing Baskets

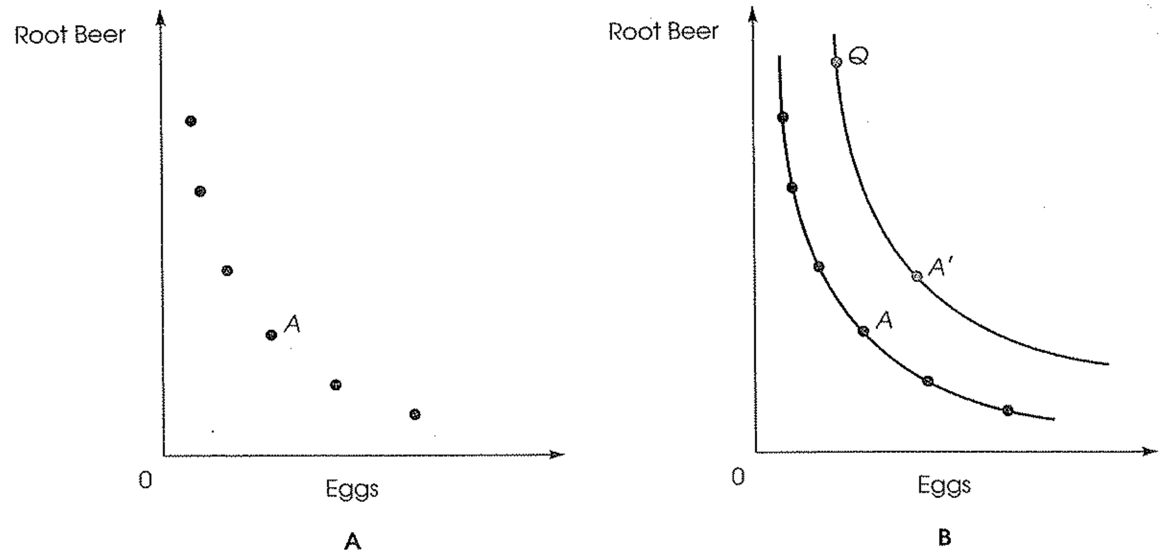

Panel A shows several baskets that Beth considers to be equally desirable. None of these can lie to the northeast or southwest of any other one, because if it did, one would be clearly preferable to the other. As a result, they all lie to the northwest and southeast of each other, accounting for the downward slope.

The black indifference curve in panel B includes the points from panel A, as well as all of the other baskets that Beth considers equally as desirable as these. The colored indifference curve shows a different set of baskets, all of which are equally as desirable as each other. Knowledge of Beth's indifference curves allows us to make inferences about her preferences that would otherwise be impossible. For example, we know that Beth likes *Q* and *A*′ equally because they are on the same indifference curve and that *A*′ is preferable to *A* because it contains more of everything. We may infer that Beth prefers *Q* to *A*.

There is nothing special about basket *A*. We could as easily have begun with a different basket, such as *A*′ in panel B of Exhibit 3.2. That panel depicts both the indifference curve through *A* (in black) and the indifference curve through *A*′ (in color).

The indifference curves do not have to have the same shape, but they do both have to slope downward.

If we know a consumer's indifference curves, we can make inferences that would not be possible otherwise. Try comparing basket *A* to basket *Q* in panel B of Exhibit 3.2. Basket *A* has more eggs than basket *Q*, but basket *Q* has more root beer than basket *A*. Without more information, we cannot say which one Beth will prefer. But the indifference curves provide that additional information. We know that Beth likes *Q* and *A*′ equally, because they are on the same indifference curve. We know that she likes *A*′ better than she likes *A*, because it is to the northeast of *A* and therefore contains more of everything than *A* does. We may infer that she likes *Q* better than she likes *A*.

In general, a basket is preferable to another precisely when it is on a higher indifference curve, where *higher* means "above and to the right."

Relationships among Indifference Curves

Of course, Beth has more than two indifference curves. Indeed, we can draw an indifference curve through any point that we choose to start with.

Because of this, *the indifference curves fill the entire plane.* (More precisely, they fill the entire quadrant of the plane in which both coordinates are positive.)

An important feature of indifference curves is that *indifference curves never cross.* To understand why this must be true, imagine a consumer with two indifference curves that cross, as in Exhibit 3.3.

From the fact that baskets *P* and *Q* are on the same (black) indifference curve, we know that the consumer likes these baskets equally well. From the fact that baskets *R* and *Q* are on the same (brown) indifference curve, we know that he also likes these equally well. Putting these facts together, we conclude that he likes *P* and *R* equally well. But this is impossible, since *R* is to the northeast of *P* and therefore contains more of both goods. In other words, if indifference curves cross, impossible things will happen. We conclude that indifference curves don't cross.

Marginal Values

We have said that indifference curves slope downward, but we haven't yet said anything about how steep the slope is. In this section, we will interpret the slope of the indifference curve. The first step is to understand how indifference curves can tell us whether certain trades are desirable.

Desirable and Undesirable Trades

Suppose you have 7 eggs and 2 root beers; this basket is represented by point *C* in Exhibit 3.4. Your friend Jeremy offers to trade you 2 root beers for an egg. If you accept his offer, you'll end up at point *P*. (That is, you'll give Jeremy an egg, leaving you with 6 eggs, and he'll give you 2 root beers, leaving you with 4 root beers. Point *P* illustrates your new basket.)

EXHIBIT 3.3 Indifference Curves Never Cross

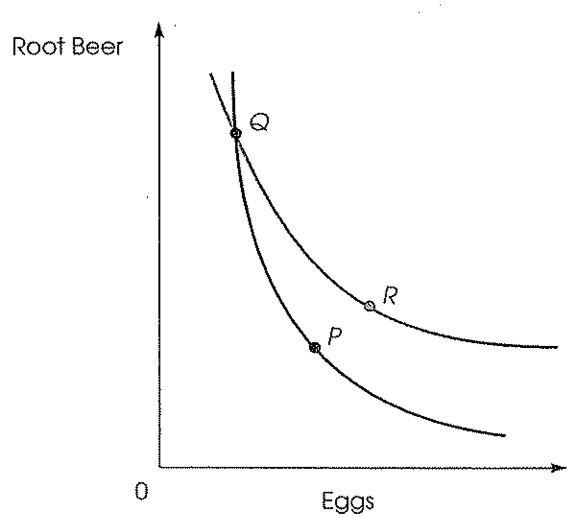

Crossing indifference curves, such as those shown in the graph, cannot occur. The consumer likes *P* and *Q* equally well because they are both on the same (black) indifference curve. He also likes *R* and *Q* equally well because they are both on the same (colored) indifference curve. We may infer that he likes *P* and *R* equally well, which we know to be false (in fact, *R* is preferred to *P*). Thus, the graph cannot be correct.

EXHIBIT 3.4 Marginal Value

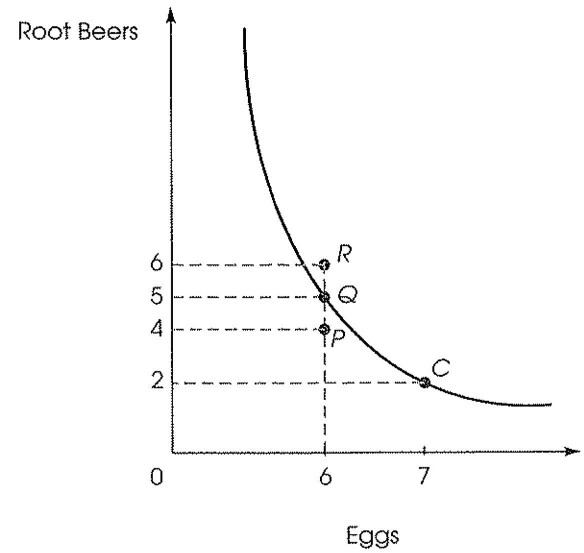

Suppose you start with basket *C*. If someone offers to trade you 2 root beers for an egg, you can move to basket *P*, which is worse; so you'll reject this trade. The minimum price you'd accept for an egg is 3 root beers, moving you to basket *Q*. Thus (to you), the marginal value of an egg is 3 root beers.

Will you accept Jeremy's offer? It depends on your preferences. Suppose, for example, that you have the indifference curve shown in Exhibit 3.4. Then you will *not* accept Jeremy's offer, because—according to your preferences—point *P* is inferior to point *C*.

In other words, when Jeremy says, "I'll give you 2 root beers for an egg," you'll say, "No thanks; I'd rather keep the egg." In ordinary language, we'd say that your seventh egg is worth more to you than 2 root beers.

Suppose Jeremy tries again, by offering you 4 root beers for an egg instead of 2 root beers. Now do you accept the trade? If you do, you'll end up at point *R*, above your original indifference curve. This trade is desirable; it makes you happier; you got more for your egg than you thought it was worth.

EXERCISE 3.3 Explain why Jeremy's new offer brings you to point *R*.

Finally, what if Jeremy had offered you exactly 3 root beers for your seventh egg? This brings you to point *Q*, which is exactly as desirable as your original point *C*. That is, trading an egg for 3 root beers makes you neither better nor worse off than you were to begin with. This makes it reasonable to say that your seventh egg is *worth* exactly 3 root beers (to you). We say that (to you) the **marginal value** of an egg is 3 root beers.

Marginal value of *X* in terms of *Y* The number of *Y*s for which the consumer would be just willing to trade one *X*.

In general, the marginal value that you place on good *X* (in terms of good *Y*) is defined to be the number of *Y*s for which you'd be just willing to trade one *X*.[2] (The adjective *marginal* refers to the fact that you are trading just *one X*.)

Given a consumer's initial basket and the indifference curve through that basket, you can always compute the marginal value of the horizontal good by traveling leftward 1 unit and then seeing how far upward you must travel to reach the indifference curve. In Exhibit 3.4, this means starting at point *C*, traveling leftward 1 egg (from 7 to 6), and then observing that you must travel upward 3 root beers (from 2 to 5); thus—as we have already said—the marginal value of an egg is 3 root beers.

EXERCISE 3.4

How can you use the indifference curve of Exhibit 3.4 to illustrate the marginal value of *root beers* in terms of *eggs?*

Marginal Value as a Slope

Exhibit 3.5 illustrates the indifference curves of two consumers, each starting with basket *C*. We can use these indifference curves to compute the marginal value of an egg to each consumer. For Jack, the marginal value of an egg is 6 root beers; for Jill, the marginal value of an egg is 1 root beer.

EXHIBIT 3.5 Marginal Value as a Slope

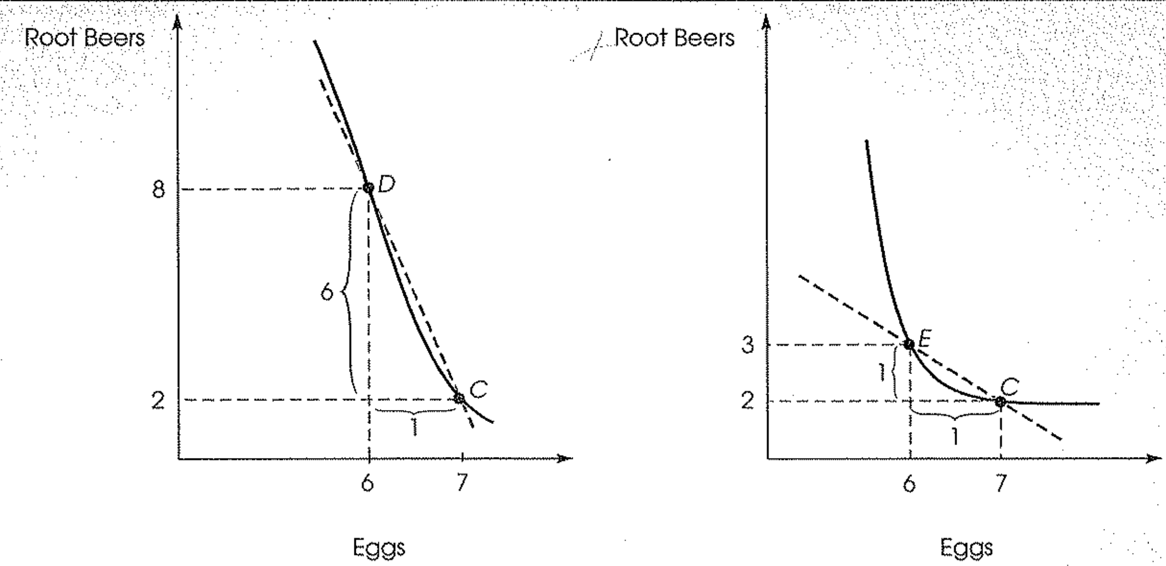

A: Jack's indifference curve

B: Jill's indifference curve

Jack and Jill each start with basket *C*. To Jack, the marginal value of an egg is 6 root beers. Thus, trading 1 egg for 6 root beers leaves him on the same indifference curve; in other words, his indifference curve goes through point *D* and so has slope −6. To Jill, the marginal value of an egg is 1 root beer, so her indifference curve goes through point *E* and has slope −1. In general, the marginal value of an egg is equal to the absolute value of the slope of the indifference curve.

---

[2]In many textbooks, the marginal value is called the *marginal rate of substitution* or *MRS*. Unfortunately, there is quite a bit of confusion associated with this term. The quantity that we've called the marginal value of *X* in terms of *Y* is sometimes called the marginal rate of substitution between *X* and *Y*, and sometimes called the marginal rate of substitution between *Y* and *X*. To avoid this confusion, we will stick with the term *marginal value*.

EXERCISE 3.5

Explain how to compute these marginal values from the graphs in Exhibit 3.5.

Now let's forget about marginal values for a moment and ask a purely geometric question: What is the *slope* of Jack's indifference curve at point *C*? By the slope of a curve we mean the slope of a line tangent to that curve. The tangent line at *C* is well approximated by the illustrated line through *C* and *D*. So we want to compute the slope of that line.

Recalling that the slope of a line is given by the *rise over the run*, we see that in this case the slope is $-6/1 = -6$. The numerator 6 is the vertical distance between points *C* and *D*, the denominator 1 is the horizontal distance, and there is a minus sign because the curve is downward sloping. The absolute value of this slope is 6 (or, more precisely, 6 root beers per egg). Recall that according to Jack, this is exactly the marginal value of an egg.

Likewise, in panel B the line through *C* and *E* has a slope with absolute value 1, which according to Jill is the marginal value of an egg.

It is no coincidence that these slopes are equal to the corresponding marginal values. In panel A, for example, we compute the marginal value of an egg as the vertical distance from *D* to *C* (that is, 6), while we compute the absolute value of the slope as that same vertical distance divided by the horizontal distance, which is 1. But dividing by 1 leaves the number 6 unchanged.

In general, then, **for a consumer with basket *C*, the marginal value of an egg is equal to the slope of the indifference curve at point C.** Consequently, **the steeper the indifference curve, the greater the marginal value of an egg.**

The Shape of Indifference Curves

A starving person with a refrigerator full of root beer is likely to value an egg more highly (in terms of root beer) than a thirsty person with a refrigerator full of eggs. Because marginal value is reflected by the slopes of indifference curves, we can translate this statement into geometry: As a general rule, we expect indifference curves to be steep near baskets containing few eggs and many root beers and to be shallow near baskets containing many eggs and few root beers.

Consider the two sets of indifference curves shown in Exhibit 3.6. Both sets slope downward. The first set slopes steeply in the area where baskets contain few eggs and many root beers (that is, in the "northwest" part of the figure) and shallowly in the area where baskets contain few root beers and many eggs. This consumer conforms to the general rule of the preceding paragraph.

Another consumer might have the indifference curves shown in panel B of Exhibit 3.6. This consumer values eggs highly when she has many eggs and few root beers, but values eggs much less when she has few eggs and many root beers. Such tastes are possible, but they seem unlikely.

Therefore, we will always assume that indifference curves are shaped like those in panel A rather than those in panel B. That is, we assume that indifference curves bow inward toward the origin. This property is

EXHIBIT 3.6 The Curvature of Indifference Curves

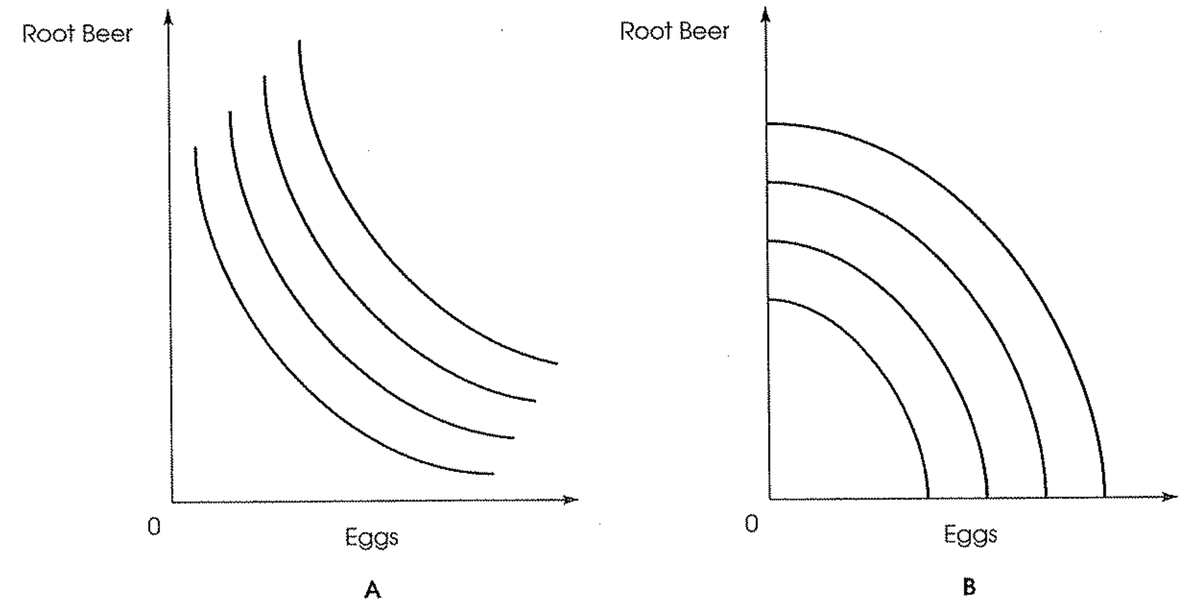

The indifference curves in panel A are convex (bowed in toward the origin), indicating that when the consumer has few eggs and many root beers (in the "northwest" part of the diagram), she places a high marginal value on eggs—that is, you'd have to offer her a lot of root beer to get her to part with an egg. We assume that indifference curves have this shape, rather than the alternative shape illustrated in panel B.

expressed by saying that indifference curves are **convex**. At the end of Section 3.2 we will give another, independent justification for assuming convexity.

Convex Bowed in toward the origin, like the curves in panel A of Exhibit 3.6.

EXERCISE 3.6

Under what circumstances do you expect the consumer to value additional root beers highly relative to additional eggs? Combine this answer with your answer to Exercise 3.4 to draw a conclusion about where the indifference curves should be steep and where they should be shallow. Does your conclusion give further support to our assumption that indifference curves are convex, or does it suggest a reason to doubt that assumption?

More on Indifference Curves

Properties of Indifference Curves: A Summary

Here are the fundamental facts about a given consumer's indifference curves:

Indifference curves slope downward, they fill the plane, they never cross, and they are convex.

A consumer's indifference curves between two goods encode everything that there is to say about the consumer's tastes regarding those goods.

A different consumer is likely to have a different family of indifference curves (also satisfying the fundamental facts). This is just another way of saying that tastes may differ across individuals.

> We have assumed that eggs and root beer are both *goods*—items you'd always prefer to have more of—and we've concluded that indifference curves are downward sloping and convex. A different assumption would lead to different conclusions. End-of-chapter problems 5 and 6 will lead you through the analysis when one or both of the goods is replaced by a *bad*—something you'd prefer to have less of. (In problems 3 and 4, you'll encounter other special circumstances in which the shapes of indifference curves can differ from what is pictured in Exhibit 3.6A).

The Composite-Good Convention

In order to draw indifference curve diagrams, we must assume that there are only two goods in the world. This might appear to be a severe limitation, yet in fact it is not. In many applications we will want to concentrate our attention on a single good—say, eggs. In that case we divide the world into two classes of goods, namely, "eggs" and "things that are not eggs," otherwise known as "all other goods." This allows us to draw indifference curves between eggs (on the horizontal axis) and all other goods (on the vertical).

There remains the problem of units. What is a single unit of *all other goods*? The simplest solution to this problem is to measure all other goods in terms of their dollar value.

Composite-good convention
The lumping together of all goods but one into a single portmanteau good.

When we lump together all things that are not eggs and measure it in a single unit like dollars, we say that we are using the **composite-good convention.**

In the presence of the composite-good convention, the slope of an indifference curve is the marginal value of an egg in terms of other goods, with the other goods measured in dollars. Thus, it is the minimum number of dollars for which the consumer would be willing to trade an egg.

3.2 The Budget Line and the Consumer's Choice

In order to predict a consumer's behavior, we need to know two things. First, we need to know the consumer's tastes, which is the same thing as saying that we need to know his indifference curves. Second, we need to know the options available to the consumer. In other words, we need to know his budget.

The Budget Line

Continue to assume a world with two goods. Instead of calling them eggs and root beers, we're going to start calling them X and Y. You may continue to think of them as eggs and root beers if you wish. In order to determine which baskets our consumer can afford, we need to know three things: the price of X, the price of Y, and the consumer's income.

Rather than make up specific numbers, let's make up names for the three things we need to know:

P_X = the price of X in dollars
P_Y = the price of Y in dollars
I = the consumer's income in dollars

Now let's suppose that the consumer is considering the purchase of a particular basket. Suppose that the basket contains x units of X and y units of Y. (Keep in mind that the capital letters X and Y are the *names* of the goods and the small letters x and y are the *quantities*.) How much will it cost the consumer to acquire this basket? The x units of X at a price of P_X dollars apiece will cost $P_X \cdot x$ dollars. The y units of Y at a price of P_Y dollars apiece will cost $P_Y \cdot y$ dollars. The total price of the basket is

$$P_X \cdot x + P_Y \cdot y \text{ dollars}$$

Under what circumstances can the consumer afford to acquire this particular basket? Clearly, he can acquire it only if the price of the basket does not exceed his income. In other words, he can afford the basket precisely if

$$P_X \cdot x + P_Y \cdot y \leq I$$

In fact, we can say a little more. Let's take seriously our assumption that X and Y are the only goods in the world. (In view of the composite-good convention, this assumption is not as outrageous as it seems.) Then the consumer will have to spend his entire income on X and Y,[3] and must choose a basket that costs exactly I dollars. The consumer can have the basket in question precisely if

$$P_X \cdot x + P_Y \cdot y = I$$

It is important to distinguish the meanings of the various symbols in this equation. P_X, P_Y, and I are particular, fixed numbers that the consumer faces. The letters x and y are variables that can represent the contents of any basket. As the consumer considers purchasing various baskets, the values of x and y change. For each basket he plugs the relevant values of x and y into the equation, and he asks if the equation is true. Asking "Does this basket make the equation true?" is exactly the same as asking "Can I afford to purchase this basket?"

The line described by the equation $P_X \cdot x + P_Y \cdot y = I$ is a picture of all the baskets that the consumer can afford. It is called the consumer's **budget line.**

Budget line The set of all baskets that the consumer can afford, given prices and his or her income.

[3]It is possible that the consumer would want to save some income, but in that case we would want to consider savings as another good. If we are using the composite-good convention, we can include savings along with "all other goods."

Another way to write the equation of the budget line (using some simple algebraic manipulations) is

$$y = -\frac{P_X}{P_Y} \cdot x + \frac{I}{P_Y}$$

If you remember that P_X, P_Y, and I are constants and that x and y are variables, you may recognize this as the equation of a line with slope $-P_X/P_Y$ and y-intercept I/P_Y. The points on that line are those that satisfy the equation and are therefore those that represent baskets that the consumer can buy. Exhibit 3.7 shows the budget line.

Here is an easy way to remember how to draw the budget line. If you were the consumer and you bought no Xs at all, how many Ys could you afford? Since your income is I and Ys sell at a price of P_Y apiece, the answer is I/P_Y. This means that the point $(0, I/P_Y)$ must be on the budget line. If you bought no Ys at all, how many Xs could you afford? The answer is I/P_X. This means that the point $(I/P_X, 0)$ must be on the budget line. The budget line must be the line connecting the points $(0, I/P_Y)$ and $(I/P_X, 0)$.

What if P_X, P_Y, and I were all to double simultaneously? This would have no effect on the ratios I/P_Y and I/P_X. It follows that a simultaneous doubling of all prices and income would have no effect on the budget line. This accords with our expectation that only relative prices matter.

The geometry of the budget line reflects everything there is to know about the opportunities facing the consumer. For example, the slope of the budget line is $-P_X/P_Y$, and the ratio P_X/P_Y is the relative price of X in terms of Y. Therefore, the budget line will be steep when X is expensive relative to Y, and it will be shallow when X is inexpensive relative to Y.

EXHIBIT 3.7 The Budget Line

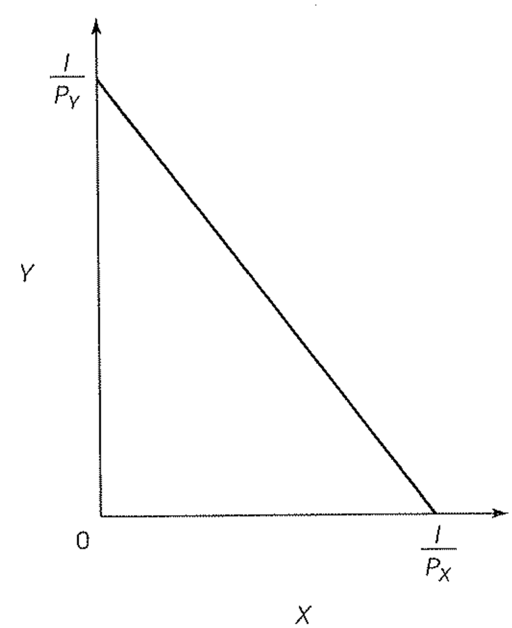

The consumer's budget line depicts the various baskets that he can afford with his income.

The Consumer's Choice

The Geometry of the Consumer's Choice

The budget line conveys an entirely different kind of information than the indifference curves do. The indifference curves reflect the consumer's preferences without regard to what he can actually afford to buy. The budget line shows which baskets he can afford to buy (that is, it shows his opportunities) without regard to his preferences. To determine how the consumer will actually behave, we must combine these two kinds of information. To this end, we have drawn the indifference curves and the budget line on the same graph, as in Exhibit 3.8.

We now have enough information to determine which basket this consumer will choose. Look at the baskets pictured. Of these, *F* is on the highest indifference curve and the one that the consumer would most like to own. (There are also many baskets not pictured that the consumer would like even more than *F*.) Unfortunately, he can't afford basket *F*—it's outside the budget line. By contrast, point *E* is inside the budget line and would fail to exhaust his income; therefore, *E* is ruled out as well. The baskets that the consumer can acquire are the ones on his budget line. In Exhibit 3.8 these include *A, B, O, C,* and *D*.

Of these, he will choose the one on the highest possible indifference curve. It is clear from the picture that this choice is *O*. In fact, *O* is not just

EXHIBIT 3.8 The Consumer's Optimum

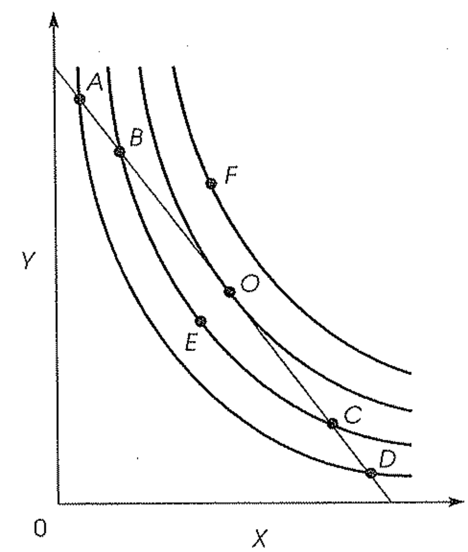

The consumer must choose one of the baskets that is on his budget line, such as *A, B, O, C,* or *D*. Of these, he will choose the one that is on the highest indifference curve, namely, *O*. Thus, the consumer is led to choose the basket at the point where his budget line is tangent to an indifference curve. This point is called the consumer's optimum.

At the consumer's optimum, the relative price of *X* in terms of *Y* (given by the slope of the budget line) and the marginal value of *X* in terms of *Y* (given by the slope of the tangent line to the indifference curve) are equal. The geometric reason for this is that the budget line is the tangent line to the indifference curve. The economic reason for it is that whenever the relative price is different from the marginal value, the consumer will continue to make exchanges until the two become equal.

the best choice among the five baskets we have considered but the best choice of any basket on the budget line. From the picture, the following is clear:

The basket the consumer chooses will always be located where his budget line is tangent to one of his indifference curves.

Optimum (plural: optima) The most preferred of the baskets on the budget line.

This basket is called the consumer's **optimum**. Because there is only one such point, the budget line and the indifference curves give sufficient information for us to predict which basket the consumer will choose.

The Economics of the Consumer's Choice

We can analyze the consumer's problem from a different perspective and still reach the same conclusion about the location of his optimum.

Referring to Exhibit 3.8, suppose that the consumer owns basket *A*. How much *Y* would this consumer be willing to trade for an additional unit of *X*? The answer is given by the marginal value of a unit of *X* (in terms of *Y*), which is measured by the absolute value of the slope of his or her indifference curve at *A*.

How much *Y* would this consumer actually have to sacrifice in order to acquire an additional unit of *X*? The answer is given by the relative price of *X* in terms of *Y*, which is the ratio P_X/P_Y, the absolute value of the slope of his budget line.

Of these two, which is greater, the marginal value or the relative price? At point *A* the indifference curve is steeper than the budget line. Consequently, the amount of *Y* that the consumer is *willing* to pay for a unit of *X* exceeds the amount of *Y* that he actually *has to* pay for a unit of *X*. In such a situation, buying a unit of *X* is an attractive proposition. The consumer will exchange *Y*s for *X*s at the going relative price, ending up with more *X* and less *Y* than he started with. This will bring him to a point like *B*.

Now the same reasoning applies again. At *B* it is still the case that the marginal value exceeds the relative price. The consumer will want to buy another unit of X, which will move him further down the budget line.

This process will continue until the consumer reaches point *O*. At that point the price that he is willing to pay for *X* and the price at which he is able to purchase *X* have become equal. There is no longer anything to be gained from additional trades.

A similar process occurs if the consumer starts out with basket *D*. Here the marginal value of *X* is less than the relative price of *X*; the consumer values his last unit of *X* at less than the number of *Y*s that can be exchanged for it in the marketplace. In this case, he will happily trade away his last unit of *X*, ending up with more *Y*s and fewer *X*s, at a point like *C*.

As long as the marginal value of *X* is less than the relative price of *X*, the consumer will trade *X*s for *Y*s. This process stops when the marginal value and the relative price become equal, at point *O*.

Whenever the marginal value of *X* exceeds the relative price of *X*, the consumer will want to buy *X*s, moving down the budget line. Whenever the marginal value is less than the relative price, the consumer will want to sell *X*s, moving up the budget line. The only point at which he can settle is *O*, where the marginal value and the relative price are exactly equal.

Thus, the economic reasoning leads to the same conclusion as the geometric reasoning: Of the points available to the consumer, the optimum occurs where his budget line is tangent to one of his indifference curves.

Corner Solutions

There is an exception to the rule that the consumer's optimum always occurs at a tangency. This exception is illustrated in Exhibit 3.9. In this case there is no tangency for the consumer to choose.

To predict the consumer's choice in this situation, we can use simple geometry. We know that the consumer must choose a basket on his budget line. Of all of these baskets, we can see from the picture that the one lying on the highest indifference curve is *P*. Therefore, the consumer chooses basket *P*.

Here is an alternative path to the same conclusion: Suppose the consumer begins with basket *S*. At this point his indifference curve is less steep than his budget line. To this consumer the marginal value of *X* in terms of *Y* is less than the relative price of *X* in terms of *Y*. The last unit of *X* is worth less to him than it will bring in the marketplace. Therefore, he trades *X* for *Y*, moving to a point like *R*. Now the same reasoning applies again, leading the consumer to move first to *Q* and then to *P*. The same reasoning would apply no matter what the original basket was.

The situation depicted in Exhibit 3.9 is called a **corner solution** because the consumer's optimum occurs in a corner of the diagram. As you can see from the picture, he consumes no *X* whatsoever and spends all of his income on *Y*.

Corner solution An optimum occurring on one of the axes when there is no tangency between the budget line and an indifference curve.

EXHIBIT 3.9 A Corner Solution

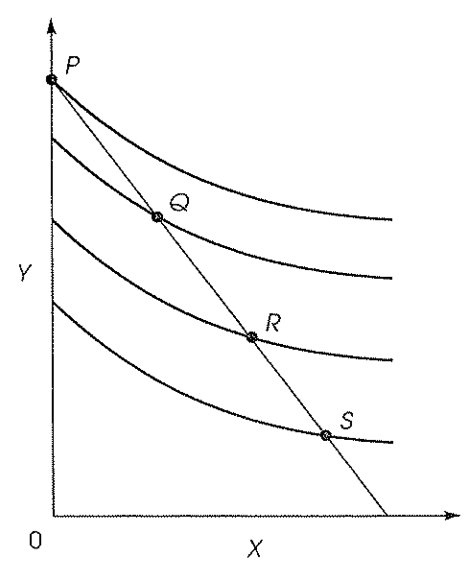

If the consumer's indifference curves look like those pictured, there is no tangency between his budget line and any of his indifference curves. Of all the points on the budget line, the consumer will choose the most desirable, namely, *P*. At any other point on the budget line the marginal value of *X* in terms of *Y* is less than the relative price, so the consumer can sell *X*s for more than they are worth to him and will continue to do so until he has sold all of his *X*s, ending up in the corner at *P*.

More on the Shape of Indifference Curves

In Section 3.1 we justified the assumption that indifference curves are convex with an appeal to the idea of marginal value. Now we can give an additional reason for making this assumption.

Suppose a consumer has the indifference curves illustrated in Exhibit 3.10. Will this consumer choose to purchase the basket at point *O*? No! He can do better. Points *C* and *D* are both available to him (they are on his budget line) and they are on a higher indifference curve than *O*. And can he do better than *C* and *D*? Yes. Every movement "outward" along the budget line, away from *O* and toward one of the axes, improves the consumer's welfare. For this reason he will always want to choose a basket on one of the axes—a corner solution. In this case he will choose basket *A*.

EXERCISE 3.7

Why does the consumer choose basket *A* rather than basket *E*? How would the budget line have to look for him to choose a point on the *X*-axis rather than the *Y*-axis?

Because this consumer always selects a corner solution, he consumes either zero units of *X* or zero units of *Y*. But goods that consumers choose to purchase none of are not very interesting from the viewpoint of economics. So now we have our additional reason for assuming that indifference curves are convex. They might not be—but in this case one of the goods in question would not be consumed at all, and we would prefer to turn our attention to goods that *are* consumed. Therefore, we usually confine our attention to convex indifference curves.

EXHIBIT 3.10 The Consumer's Choice with Nonconvex Indifference Curves

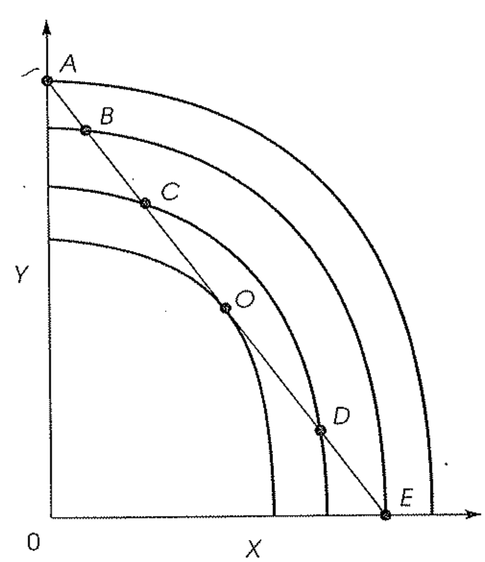

Nonconvex indifference curves always lead to a corner solution. The consumer pictured here will choose point *A*, which is on the highest possible indifference curve.

3.3 Applications of Indifference Curves

Now let's put our new tools to use. In this section we'll see several applications of indifference curve analysis.

Standards of Living

Economic conditions change all the time. Incomes go up and down, and so do prices. How do we tell which changes are good for the consumer and which are bad?

Sometimes it's easy. If your friend Harold's income goes up while prices remain unchanged, his life has certainly improved. If his income stays fixed while all prices rise, he's worse off than before. But what if some prices rise while others fall? Is that good or bad for Harold?

Sometimes there's not enough information to answer that question. Other times there is. Let's take an example: Harold consumes goods X and Y. Their prices are $P_X = \$3$ and $P_Y = \$4$. He chooses to buy 4 units of X and 2 of Y, exhausting his income of \$20. Now the price of X rises to \$4 while the price of Y falls to \$2, and his income stays fixed at \$20. Is Harold better or worse off than before?

To answer, start by drawing Harold's original budget line, marked *Original* in Exhibit 3.11. Given his income of \$20 and given $P_X = \$3$, Harold can afford up to $6\frac{2}{3}$ Xs if he buys no Ys. Because $P_Y = \$4$, he can afford up to 5 Ys if he buys no Xs. Those calculations determine the two endpoints of his *Original* budget line. We are told that Harold chooses basket $O = (4, 2)$, so there must be an indifference curve tangent to the *Original* budget line at that point, as illustrated in the exhibit.

EXHIBIT 3.11 Harold's Original and New Budget Lines

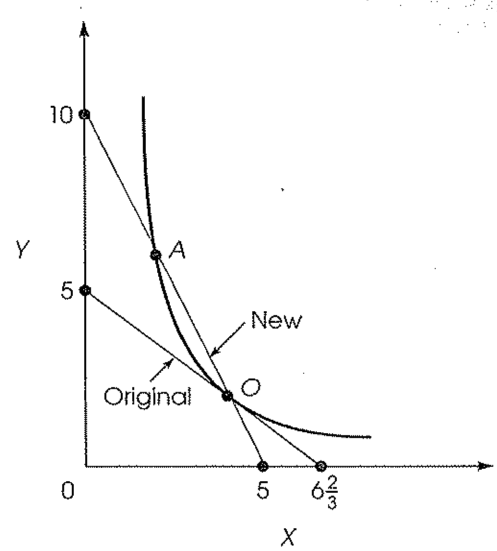

The graph shows Harold's *original* and *new* budget lines. We know that an indifference curve is tangent to the *original* line at the point O. We can calculate that the *new* budget line passes through the point O.

Now let's draw Harold's *New* budget line, after the prices change to $P_X = \$4$ and $P_Y = \$2$. The endpoints are at $X = 5$ and $Y = 10$, as shown in the exhibit. But knowing the endpoints is not enough to draw the *New* budget line accurately. We also have to think about whether it passes above, below, or through the point O.

To settle this question, ask whether Harold can afford basket O at the *New* prices. With $P_X = \$4$ and $P_Y = \$2$, basket O costs $(\$4 \times 4) + (\$2 \times 2) = \$20$, which is exactly Harold's income. So O must be on his *New* budget line, or, to put it another way, his *New* budget line must pass through point O. That's how we've drawn it in Exhibit 3.11.

Now let's find Harold's *New* optimum point—the point where his *New* budget line is tangent to an indifference curve. The first thing we can do is rule out point O. That's because *a smooth curve cannot be tangent to two different lines at the same point*—an important fact about geometry that will be useful to keep in mind.

So where is Harold's new optimum? Panel A of Exhibit 3.12 explores some possibilities. The two dashed curves are not possible, because either of them would have to cross the original indifference curve. That means Harold can't have an optimum in the region above and to the left of A or in the region below and to the right of O. Instead, his new optimum must lie between A and O, for example, at P in panel B. If you look at the panel, you'll see that P must lie on a higher indifference curve than O. Therefore, the price changes must have made Harold better off.

Price Indices

To measure how people are affected by price changes, the U.S. Department of Labor, through its Bureau of Labor Statistics, reports

EXHIBIT 3.12 Finding the New Optimum

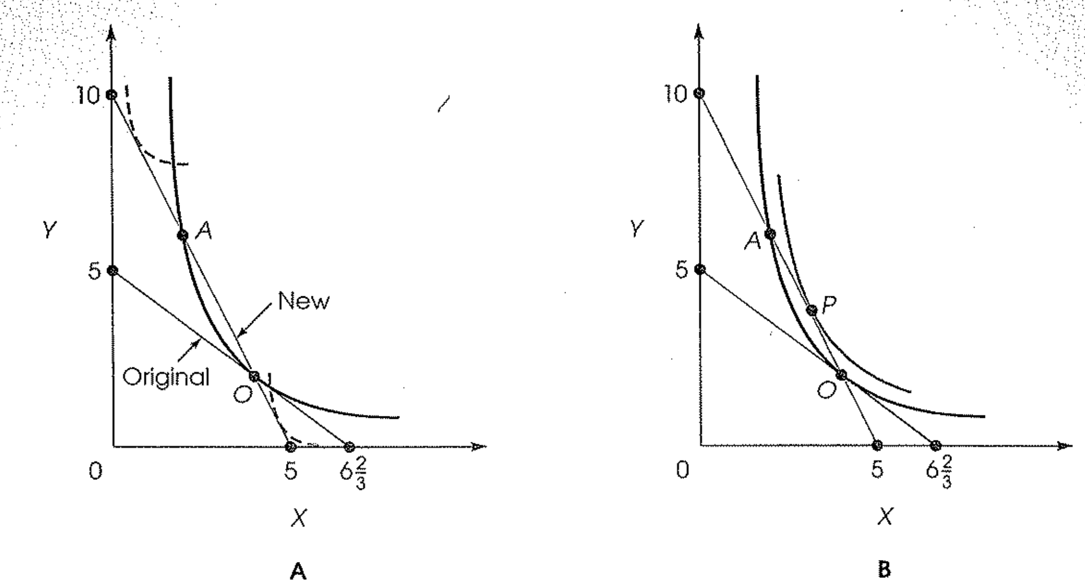

The dashed indifference curves in panel A cannot be correct, because they cross the indifference curve through O. The only correct way to draw an indifference curve tangent to the *new* budget line is with the tangency between A and O, at a point like P, as in panel B. The new indifference curve is then necessarily higher than the old one, so you are better off at the new optimum.

estimates of changes in the "cost of living," also called the "price level." Roughly, they do this by tracking the cost of a given basket over time. If the basket gets more expensive, they say that the cost of living has gone up (which suggests that people are worse off); if the basket gets cheaper, they say that the cost of living has gone down (which suggests that people are better off).

The big problem with this procedure is that the answer you get depends on which basket you choose to track. Look again at Exhibit 3.12. Basket *O* costs $20 under the original prices and $20 under the new prices. If you track basket *O*, you'll say the cost of living hasn't changed at all. That's misleading, because, as we've just seen, Harold is definitely happier with the new prices than with the old ones.

If you tracked basket *P*, you'd get a different answer. Basket *P* is outside the *Original* budget line, which means it must cost more than $20 at the original prices. But it is exactly on the *New* budget line, meaning it costs just $20 at the new prices. So basket *P* does get cheaper over time, and if you used it to measure the cost of living you'd say that the cost of living had come down.

The cost of living measurement that you get by tracking the original basket (in this case *O*) is called a **Laspeyres price index** (pronounced "Laspears"), and it tends to make things look worse than they are. The cost of living measurement that you get by tracking the new basket (in this case *P*) is called a **Paasche price index** (pronounced "Posh"), and it tends to make thing look better than they are. Unfortunately, there is no perfect way to measure changes in the cost of living.

Laspeyres price index A price index based on the basket consumed in the earlier period.

Paasche price index A price index based on the basket consumed in the later period.

Differences in Tastes

Germans eat a lot of starch. Italians eat more tomatoes. Greeks use olive oil and the French use hollandaise. Why doesn't everyone eat the same diet?

There are only two possible answers: People in different countries must have either different tastes or different opportunities (or both). Maybe Italians eat tomatoes because they like them better than Germans do—that's a difference in tastes. Or maybe Italians eat tomatoes because tomatoes are cheaper in Italy, or because Italians are too poor to afford a German diet—those are differences in opportunities.

How do we tell which theory is right? There's no question that prices and incomes differ across countries, so there's no question that there are differences in opportunities. The question is whether those differences in opportunities suffice to explain the different choices people make, or whether their tastes must also differ.

Start with a fictional example: Suppose Albert lives in Rome, where tomatoes sell for $2 a pound and potatoes sell for $1 a pound. He earns $10 a day, with which he buys 4 tomatoes and 2 potatoes. Betty lives in Berlin, where tomatoes sell for $3 a pound and potatoes sell for $6 a pound. She earns $45 a day, with which she buys 1 tomato and 7 potatoes. Using these numbers, let's figure out whether Albert and Betty could have identical tastes.

The first step is to plot Albert and Betty's budget lines, which we've done in panel A of Exhibit 3.13. Albert's optimum point (4,2) is labeled *A*, and Betty's optimum point (1,7) is labeled *B*.

EXERCISE 3.8 Make sure the budget lines are drawn correctly.

Now let's change the problem slightly. Suppose that instead of buying 1 tomato and 7 potatoes, Betty buys 5 tomatoes and 5 potatoes. Then the picture looks like panel B in Exhibit 3.13. Here we can't tell whether the indifference curves eventually cross, and we can't tell whether Albert and Betty have identical tastes.

> To conclude that Albert and Betty have identical tastes, we would have to know that they share all their indifference curves. There are several reasons why we can't draw this conclusion from Exhibit 3.13B. First, we have no idea whether the two pictured indifference curves eventually cross. Second, even if they don't cross, it doesn't follow that Albert and Betty share these indifference curves; it only follows that they *might*. Third, even if Albert and Betty share the two pictured indifference curves, it doesn't follow that they share *all* their indifference curves. So the picture does not contain nearly enough information to answer the question of whether Albert and Betty's tastes are identical.

Now that we've completed our detour into fiction, what about the real world? To seek evidence of taste differences across European countries, we can replace Albert and Betty with "the average German" and "the average Italian," and we can use realistic numbers for the prices of tomatoes and potatoes. Then we can repeat the exercise with Germans and Greeks, or Greeks and Italians, or Poles and Hungarians, and with more than just two goods. If we ever get a picture like panel A of Exhibit 3.13, we've spotted a taste difference.

EXHIBIT 3.13 Comparing Preferences

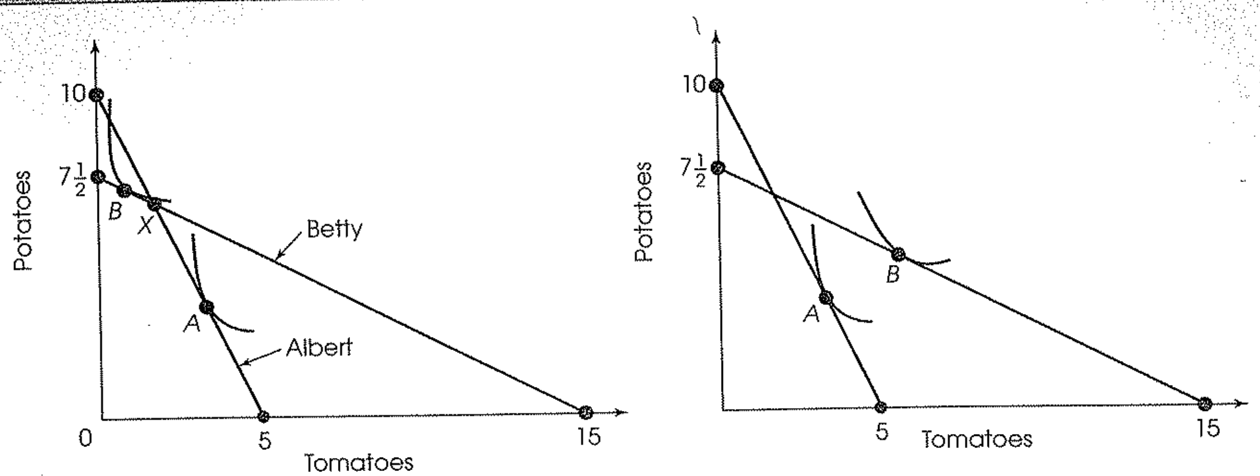

A. One possibility

B. Another possibility

In Panel A, Albert's indifference curve (tangent at *A*) must eventually cross Betty's indifference curve (tangent at *B*). Therefore, Albert and Betty cannot possibly have the same tastes. In panel B, the indifference curves might or might not cross and Albert and Betty might or might not have identical tastes.

Harvard Professor Hendrik Houthakker carried out this exercise and found no evidence of any taste differences. In other words, when Professor Houthakker drew his graphs, none looked like panel A of Exhibit 3.13. Instead, every one of his pictures leaves open the possibility that tastes could be either the same or different.

On the one hand, that doesn't prove anything. On the other hand, the more times you look for something and fail to find it, the more you're entitled to suspect it's not really there. Professor Houthakker searched repeatedly for evidence of taste differences and failed to find them. That doesn't prove tastes are remarkably similar across countries, but it is certainly evidence in that direction.

Here's a similar question: Do people's tastes change over time? For example, did the average Englishman in 1950 have different tastes than the average Englishman in 1900? We can use the same techniques: In Exhibit 3.13, replace Albert and Betty with "the average Englishman in the year 1950" and "the average Englishman in the year 1900." Look at not just tomatoes and potatoes but other pairs of goods. A picture like panel A of Exhibit 3.13 would show that tastes had changed over that half-century.

Using 127 different goods in every possible pairing, there are many hundreds of cases where the budget lines cross, raising the possibility of a configuration like panel A of Exhibit 3.13. In no case does that configuration actually occur.[4] In other words, there are a lot of opportunities to observe a taste change and no actual observations. Again, that doesn't prove anything, but it is highly suggestive.

What's the Best Way to Be Taxed?

Which would you rather pay: A percentage income tax (under which the government takes a certain percent of your earning) or a head tax (under which the government takes a certain number of dollars every day, regardless of how much you earn)?

Obviously, the answer depends at least partly on the size of the taxes. A 1% income tax is probably better than a $10,000 daily head tax, whereas a $1 daily head tax is probably better than a 90% income tax.

So let's make the comparison fair by assuming that each tax costs you the same amount of money. Here's a specific example: Suppose you earn $10 an hour. The government can impose a 50% income tax, in which case you will work 10 hours a day, earn $100, and pay $50 in taxes. Or the government can impose a $50 daily head tax. Which do you prefer?

To answer, we begin by drawing your budget line between *leisure* (measured in hours) and *income* (measured in dollars). If you don't work at all, you'll have 24 hours a day leisure and zero income, represented by the point (24, 0) in panel A of Exhibit 3.14. If you work 24 hours a day, you'll have zero leisure and $240 in income, represented by point (0, 240). The line connecting these points is your budget line, labeled *Original* in the exhibit. By choosing the number of hours you want to work, you can achieve any point on this line.

Now suppose the government imposes a 50% income tax. You can still achieve the point (24, 0) by not working at all, but you can no longer achieve (0, 240). If you work 24 hours a day, your post-tax income is now

[4] S. Landsburg, "Taste Change in the United Kingdom," *Journal of Political Economy* 89 (1981):92–104.

EXHIBIT 3.14 An Income Tax versus a Head Tax

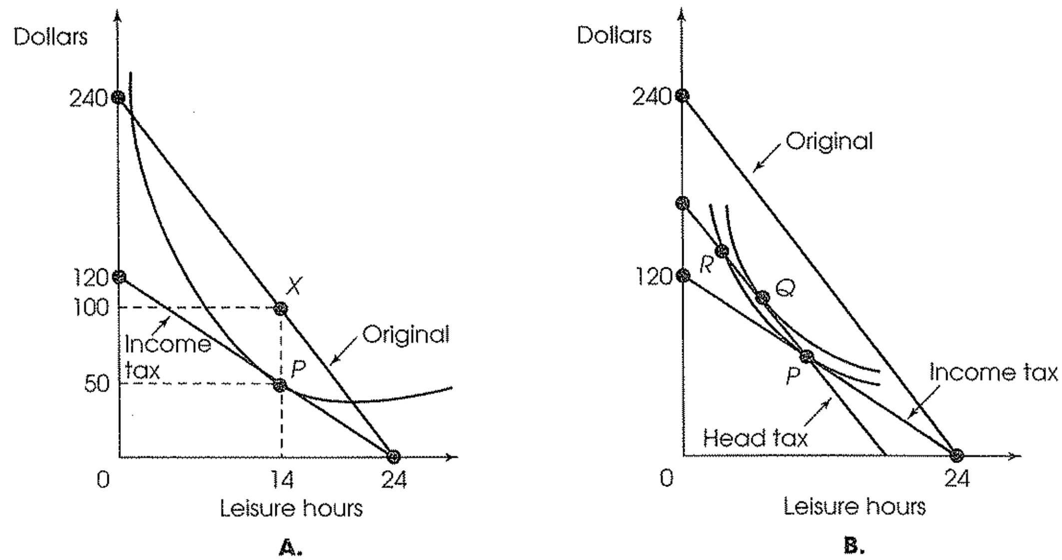

Panel A shows your *original* (untaxed) budget line and your *income tax* budget line. The optimum or the *income tax* line is at *P*, where your after-tax income is $50.

Panel B shows the *head tax* budget line, which lies a vertical distance $50 below the *original* budget line and consequently passes through point *P*. The optimum on the *head tax* line must be at a point like *Q* between *P* and *R*, and it is consequently on a higher indifference curve. The head tax is thus preferable to the income tax.

only $120. So your new budget line is the one labeled *Income tax* in panel A of exhibit 3.14.

We have assumed that under the income tax, you choose to work 10 hours a day. That means you have 14 hours of leisure and after-tax earnings of $50, represented by point *P* in the exhibit. Because you *choose* point *P*, we can conclude that there must be an indifference curve tangent to the budget line at that point, as shown.

Note that if you worked 10 hours *without* being taxed, you'd have earnings of $100. Thus, point $X = (14,100)$ must be on the *Original* budget line. (Of course, this point is no longer available to you.)

> Sometimes students think that X must be the optimum on the *Original* budget line. There is no reason to believe that is true. If the income tax were abolished, you would quite likely work some number of hours other than 10, which is the same thing as saying that the optimum on the *Original* line is somewhere other than at *X*.

Now that we've illustrated the effect of income tax, let's turn to the effect of the head tax. First we abolish the income tax, returning you to the *Original* budget line; then we impose a head tax of $50 a day, causing the budget line to drop a vertical distance $50, to the *Head tax* line in panel B of Exhibit 3.14.

To draw the *Head tax* line accurately, we have to ask whether it passes above, below, or through the point *P*. The answer: We already know that point P is exactly \$50 below point X, so when the *Original* line drops \$50, it passes exactly through point P. And that's how we've drawn it in the exhibit.

Now that you're paying a head tax, you will choose an optimum along the *Head tax* budget line. That optimum, labeled Q in the exhibit, must fall between P and R. Otherwise, the indifference curves would be forced to cross.

It's clear from the picture that Q must be on a higher indifference curve than P. It follows that you're happier paying the head tax than you are paying the income tax.

Discussion

Since your head tax bill is the same size as your income tax bill, you might be tempted to think that either tax is equally unpleasant. To see why this is false, consider your position at point P under the income tax. Here your marginal value of leisure is \$5. If you forgo that last hour of leisure by working another hour, you will take home \$5 in wages, gaining nothing. However, when the income tax is abolished and replaced by the head tax, you have the opportunity to forgo an hour of leisure in exchange for an additional \$10 in wages. This new opportunity is an attractive one. By accepting it, you move up and to the left along the *head tax* budget line, improving your situation. You continue to move in that direction until you reach point Q, where your marginal value of leisure is exactly \$10 per hour.

Summary

A consumer's behavior depends on his tastes and his opportunities. His tastes are encoded in his indifference curves and his opportunities are encoded in his budget line. By combining this information in a single graph, we can predict the consumer's behavior.

Each consumer has a family of indifference curves. Each curve in the family consists of baskets among which he is indifferent. His indifference curves slope downward, fill the plane, never cross, and are convex. A different consumer will have a different family of indifference curves, also satisfying these properties.

The slope of an indifference curve is equal (in absolute value) to the marginal value of X in terms of Y. That is, it is the number of units of Y for which the consumer is just willing to trade one unit of X.

As the consumer moves along an indifference curve in the direction of more X and less Y, we expect that the marginal value of X will decrease. This accounts for the convexity of indifference curves.

The consumer's budget line depends on his income and the prices of the goods that he buys. Its equation is

$$P_X \cdot x + P_Y \cdot y = I$$

where P_X and P_Y are the prices of X and Y and I is the consumer's income. The slope of the budget line is equal (in absolute value) to the relative price of X in terms of Y.

The consumer's optimum occurs where his budget line is tangent to one of his indifference curves. This is the point at which he attains the