# CHAPTER 3

# PREFERENCES

We saw in Chapter 2 that the economic model of consumer behavior is very simple: people choose the best things they can afford. The last chapter was devoted to clarifying the meaning of "can afford," and this chapter will be devoted to clarifying the economic concept of "best things."

We call the objects of consumer choice **consumption bundles**. This is a complete list of the goods and services that are involved in the choice problem that we are investigating. The word "complete" deserves emphasis: when you analyze a consumer's choice problem, make sure that you include all of the appropriate goods in the definition of the consumption bundle.

If we are analyzing consumer choice at the broadest level, we would want not only a complete list of the goods that a consumer might consume, but also a description of when, where, and under what circumstances they would become available. After all, people care about how much food they will have tomorrow as well as how much food they have today. A raft in the middle of the Atlantic Ocean is very different from a raft in the middle of the Sahara Desert. And an umbrella when it is raining is quite a different good from an umbrella on a sunny day. It is often useful to think of the "same" good available in different locations or circumstances as a different good, since the consumer may value the good differently in those situations.

However, when we limit our attention to a simple choice problem, the relevant goods are usually pretty obvious. We'll often adopt the idea described earlier of using just two goods and calling one of them "all other goods" so that we can focus on the tradeoff between one good and everything else. In this way we can consider consumption choices involving many goods and still use two-dimensional diagrams.

So let us take our consumption bundle to consist of two goods, and let $x_1$ denote the amount of one good and $x_2$ the amount of the other. The complete consumption bundle is therefore denoted by $(x_1, x_2)$. As noted before, we will occasionally abbreviate this consumption bundle by $X$.

## 3.1 Consumer Preferences

We will suppose that given any two consumption bundles, $(x_1, x_2)$ and $(y_1, y_2)$, the consumer can rank them as to their desirability. That is, the consumer can determine that one of the consumption bundles is strictly better than the other, or decide that she is indifferent between the two bundles.

We will use the symbol $\succ$ to mean that one bundle is **strictly preferred** to another, so that $(x_1, x_2) \succ (y_1, y_2)$ should be interpreted as saying that the consumer **strictly prefers** $(x_1, x_2)$ to $(y_1, y_2)$, in the sense that she definitely wants the x-bundle rather than the y-bundle. This preference relation is meant to be an operational notion. If the consumer prefers one bundle to another, it means that he or she would choose one over the other, given the opportunity. Thus the idea of preference is based on the consumer's *behavior*. In order to tell whether one bundle is preferred to another, we see how the consumer behaves in choice situations involving the two bundles. If she always chooses $(x_1, x_2)$ when $(y_1, y_2)$ is available, then it is natural to say that this consumer prefers $(x_1, x_2)$ to $(y_1, y_2)$.

If the consumer is **indifferent** between two bundles of goods, we use the symbol $\sim$ and write $(x_1, x_2) \sim (y_1, y_2)$. Indifference means that the consumer would be just as satisfied, according to her own preferences, consuming the bundle $(x_1, x_2)$ as she would be consuming the other bundle, $(y_1, y_2)$.

If the consumer prefers or is indifferent between the two bundles we say that she **weakly prefers** $(x_1, x_2)$ to $(y_1, y_2)$ and write $(x_1, x_2) \succeq (y_1, y_2)$.

These relations of strict preference, weak preference, and indifference are not independent concepts; the relations are themselves related! For example, if $(x_1, x_2) \succeq (y_1, y_2)$ and $(y_1, y_2) \succeq (x_1, x_2)$ we can conclude that $(x_1, x_2) \sim (y_1, y_2)$. That is, if the consumer thinks that $(x_1, x_2)$ is at least as good as $(y_1, y_2)$ *and* that $(y_1, y_2)$ is at least as good as $(x_1, x_2)$, then the consumer must be indifferent between the two bundles of goods.

Similarly, if $(x_1, x_2) \succeq (y_1, y_2)$ but we know that it is *not* the case that $(x_1, x_2) \sim (y_1, y_2)$, we can conclude that we must have $(x_1, x_2) \succ (y_1, y_2)$. This just says that if the consumer thinks that $(x_1, x_2)$ is at least as good as $(y_1, y_2)$, and she is not indifferent between the two bundles, then it must be that she thinks that $(x_1, x_2)$ is strictly better than $(y_1, y_2)$.

## 3.2 Assumptions about Preferences

Economists usually make some assumptions about the "consistency" of consumers' preferences. For example, it seems unreasonable—not to say contradictory—to have a situation where $(x_1, x_2) \succ (y_1, y_2)$ and, at the same time, $(y_1, y_2) \succ (x_1, x_2)$. For this would mean that the consumer strictly prefers the x-bundle to the y-bundle ... and vice versa.

So we usually make some assumptions about how the preference relations work. Some of the assumptions about preferences are so fundamental that we can refer to them as "axioms" of consumer theory. Here are three such axioms about consumer preference.

**Complete.** We assume that any two bundles can be compared. That is, given any x-bundle and any y-bundle, we assume that $(x_1, x_2) \succeq (y_1, y_2)$, or $(y_1, y_2) \succeq (x_1, x_2)$, or both, in which case the consumer is indifferent between the two bundles.

**Reflexive.** We assume that any bundle is at least as good as itself: $(x_1, x_2) \succeq (x_1, x_2)$.

**Transitive.** If $(x_1, x_2) \succeq (y_1, y_2)$ and $(y_1, y_2) \succeq (z_1, z_2)$, then we assume that $(x_1, x_2) \succeq (z_1, z_2)$. In other words, if the consumer thinks that $X$ is at least as good as $Y$ and that $Y$ is at least as good as $Z$, then the consumer thinks that $X$ is at least as good as $Z$.

The first axiom, completeness, is hardly objectionable, at least for the kinds of choices economists generally examine. To say that any two bundles can be compared is simply to say that the consumer is able to make a choice between any two given bundles. One might imagine extreme situations involving life or death choices where ranking the alternatives might be difficult, or even impossible, but these choices are, for the most part, outside the domain of economic analysis.

The second axiom, reflexivity, is trivial. Any bundle is certainly at least as good as an identical bundle. Parents of small children may occasionally observe behavior that violates this assumption, but it seems plausible for most adult behavior.

The third axiom, transitivity, is more problematic. It isn't clear that transitivity of preferences is *necessarily* a property that preferences would have to have. The assumption that preferences are transitive doesn't seem compelling on grounds of pure logic alone. In fact it's not. Transitivity is a hypothesis about people's choice behavior, not a statement of pure logic. Whether it is a basic fact of logic or not isn't the point: it is whether or not it is a reasonably accurate description of how people behave that matters.

What would you think about a person who said that he preferred a bundle $X$ to $Y$, and preferred $Y$ to $Z$, but then also said that he preferred $Z$ to $X$? This would certainly be taken as evidence of peculiar behavior.

More importantly, how would this consumer behave if faced with choices among the three bundles $X$, $Y$, and $Z$? If we asked him to choose his most preferred bundle, he would have quite a problem, for whatever bundle he chose, there would always be one that was preferred to it. If we are to have a theory where people are making "best" choices, preferences must satisfy the transitivity axiom or something very much like it. If preferences were not transitive there could well be a set of bundles for which there is no best choice.

## 3.3 Indifference Curves

It turns out that the whole theory of consumer choice can be formulated in terms of preferences that satisfy the three axioms described above, plus a few more technical assumptions. However, we will find it convenient to describe preferences graphically by using a construction known as **indifference curves**.

Consider Figure 3.1 where we have illustrated two axes representing a consumer's consumption of goods 1 and 2. Let us pick a certain consumption bundle $(x_1, x_2)$ and shade in all of the consumption bundles that are weakly preferred to $(x_1, x_2)$. This is called the **weakly preferred set**. The bundles on the boundary of this set—the bundles for which the consumer is just indifferent to $(x_1, x_2)$—form the **indifference curve**.

We can draw an indifference curve through any consumption bundle we want. The indifference curve through a consumption bundle consists of all bundles of goods that leave the consumer indifferent to the given bundle.

One problem with using indifference curves to describe preferences is that they only show you the bundles that the consumer perceives as being indifferent to each other—they don't show you which bundles are better and which bundles are worse. It is sometimes useful to draw small arrows on the indifference curves to indicate the direction of the preferred bundles. We won't do this in every case, but we will do it in a few of the examples where confusion might arise.

If we make no further assumptions about preferences, indifference curves can take very peculiar shapes indeed. But even at this level of generality, we can state an important principle about indifference curves: *indifference curves representing distinct levels of preference cannot cross.* That is, the situation depicted in Figure 3.2 cannot occur.

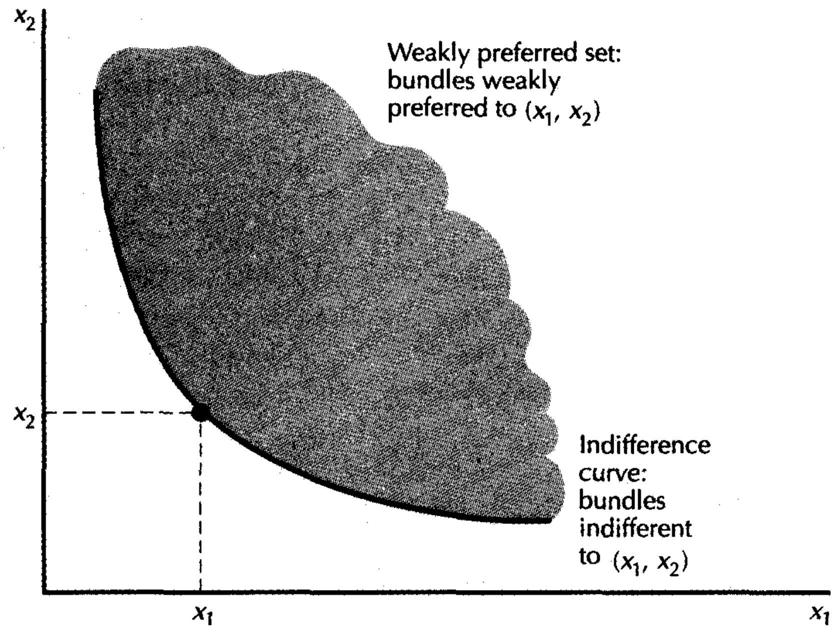


**Weakly preferred set.** The shaded area consists of all bundles that are at least as good as the bundle $(x_1, x_2)$.

In order to prove this, let us choose three bundles of goods, $X$, $Y$, and $Z$, such that $X$ lies only on one indifference curve, $Y$ lies only on the other indifference curve, and $Z$ lies at the intersection of the indifference curves. By assumption the indifference curves represent distinct levels of preference, so one of the bundles, say $X$, is strictly preferred to the other bundle, $Y$. We know that $X \sim Z$ and $Z \sim Y$, and the axiom of transitivity therefore implies that $X \sim Y$. But this contradicts the assumption that $X \succ Y$. This contradiction establishes the result—indifference curves representing distinct levels of preference cannot cross.

What other properties do indifference curves have? In the abstract, the answer is: not many. Indifference curves are a way to describe preferences. Nearly any "reasonable" preferences that you can think of can be depicted by indifference curves. The trick is to learn what kinds of preferences give rise to what shapes of indifference curves.

## 3.4 Examples of Preferences

Let us try to relate preferences to indifference curves through some examples. We'll describe some preferences and then see what the indifference curves that represent them look like.

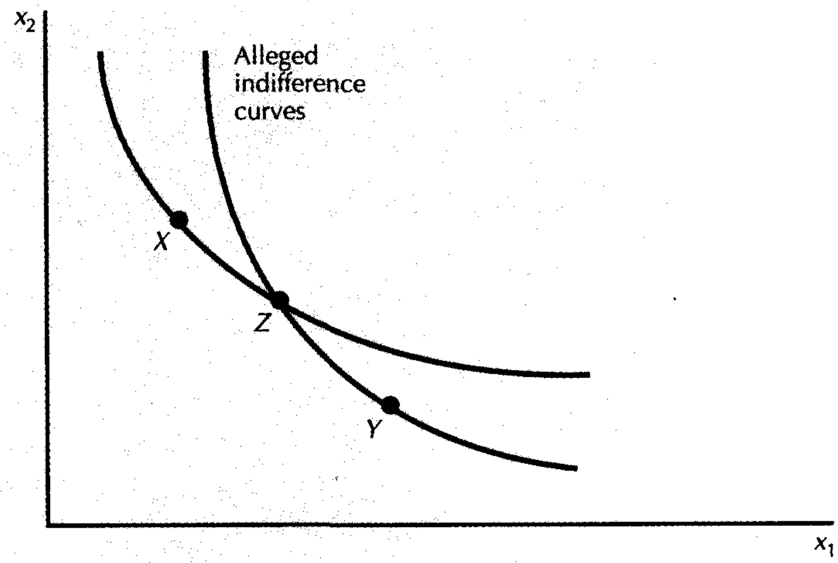


**Indifference curves cannot cross.** If they did, $X$, $Y$, and $Z$ would all have to be indifferent to each other and thus could not lie on distinct indifference curves.

There is a general procedure for constructing indifference curves given a "verbal" description of the preferences. First plop your pencil down on the graph at some consumption bundle $(x_1, x_2)$. Now think about giving a little more of good 1, $\Delta x_1$, to the consumer, moving him to $(x_1 + \Delta x_1, x_2)$. Now ask yourself how would you have to *change* the consumption of $x_2$ to make the consumer indifferent to the original consumption point? Call this change $\Delta x_2$. Ask yourself the question "For a given change in good 1, how does good 2 have to change to make the consumer just indifferent between $(x_1 + \Delta x_1, x_2 + \Delta x_2)$ and $(x_1, x_2)$?" Once you have determined this movement at one consumption bundle you have drawn a piece of the indifference curve. Now try it at another bundle, and so on, until you develop a clear picture of the overall shape of the indifference curves.

## Perfect Substitutes

Two goods are **perfect substitutes** if the consumer is willing to substitute one good for the other at a *constant* rate. The simplest case of perfect substitutes occurs when the consumer is willing to substitute the goods on a one-to-one basis.

Suppose, for example, that we are considering a choice between red pencils and blue pencils, and the consumer involved likes pencils, but doesn't care about color at all. Pick a consumption bundle, say $(10, 10)$. Then for this consumer, any other consumption bundle that has 20 pencils in it is just as good as (10, 10). Mathematically speaking, any consumption bundle $(x_1, x_2)$ such that $x_1 + x_2 = 20$ will be on this consumer's indifference curve through (10, 10). Thus the indifference curves for this consumer are all parallel straight lines with a slope of $-1$, as depicted in Figure 3.3. Bundles with more total pencils are preferred to bundles with fewer total pencils, so the direction of increasing preference is up and to the right, as illustrated in Figure 3.3.

How does this work in terms of general procedure for drawing indifference curves? If we are at (10, 10), and we increase the amount of the first good by one unit to 11, how much do we have to change the second good to get back to the original indifference curve? The answer is clearly that we have to decrease the second good by 1 unit. Thus the indifference curve through (10, 10) has a slope of $-1$. The same procedure can be carried out at any bundle of goods with the same results—in this case all the indifference curves have a constant slope of $-1$.

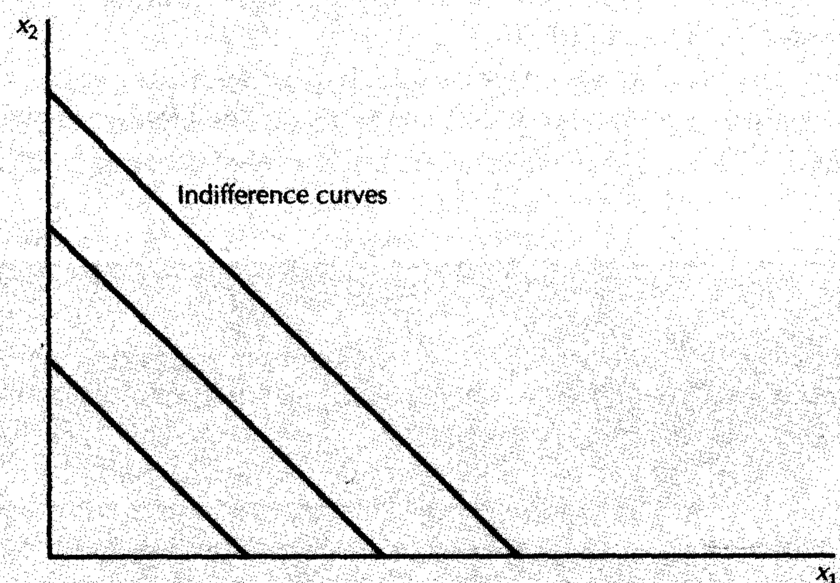


**Perfect substitutes.** The consumer only cares about the total number of pencils, not about their colors. Thus the indifference curves are straight lines with a slope of $-1$.

The important fact about perfect substitutes is that the indifference curves have a *constant* slope. Suppose, for example, that we graphed blue pencils on the vertical axis and *pairs* of red pencils on the horizontal axis. The indifference curves for these two goods would have a slope of $-2$, since the consumer would be willing to give up two blue pencils to get one more *pair* of red pencils.

In the textbook we'll primarily consider the case where goods are perfect substitutes on a one-for-one basis, and leave the treatment of the general case for the workbook.

## Perfect Complements

**Perfect complements** are goods that are always consumed together in fixed proportions. In some sense the goods "complement" each other. A nice example is that of right shoes and left shoes. The consumer likes shoes, but always wears right and left shoes together. Having only one out of a pair of shoes doesn't do the consumer a bit of good.

Let us draw the indifference curves for perfect complements. Suppose we pick the consumption bundle $(10, 10)$. Now add 1 more right shoe, so we have $(11, 10)$. By assumption this leaves the consumer indifferent to the original position: the extra shoe doesn't do him any good. The same thing happens if we add one more left shoe: the consumer is also indifferent between $(10, 11)$ and $(10, 10)$.

Thus the indifference curves are L-shaped, with the vertex of the L occurring where the number of left shoes equals the number of right shoes as in Figure 3.4.

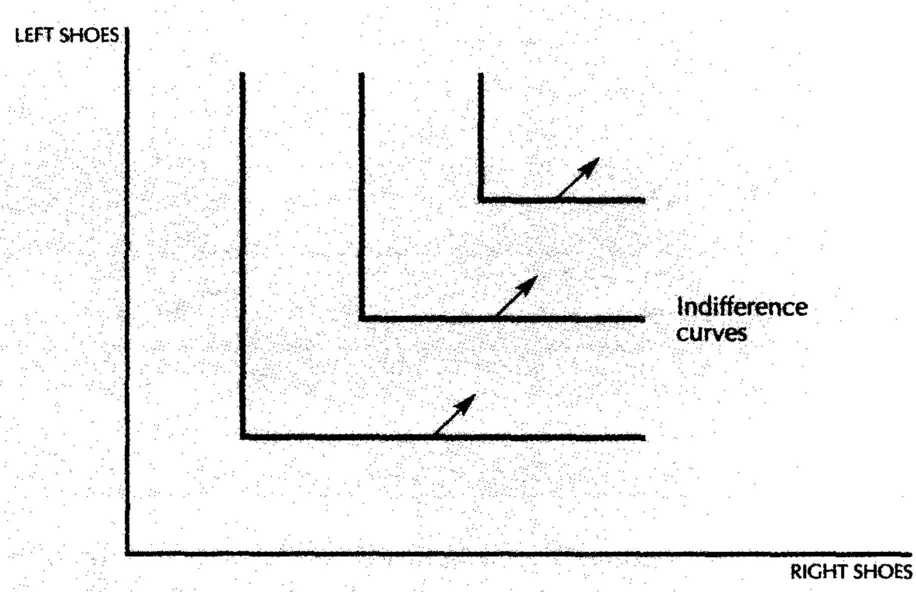


**Perfect complements.** The consumer always wants to consume the goods in fixed proportions to each other. Thus the indifference curves are L-shaped.

Increasing both the number of left shoes and the number of right shoes at the same time will move the consumer to a more preferred position, so the direction of increasing preference is again up and to the right, as illustrated in the diagram.

The important thing about perfect complements is that the consumer prefers to consume the goods in fixed proportions, not necessarily that the proportion is one-to-one. If a consumer always uses two teaspoons of sugar in her cup of tea, and doesn't use sugar for anything else, then the indifference curves will still be L-shaped. In this case the corners of the L will occur at (2 teaspoons sugar, 1 cup tea), (4 teaspoons sugar, 2 cups tea) and so on, rather than at (1 right shoe, 1 left shoe), (2 right shoes, 2 left shoes), and so on.

In the textbook we'll primarily consider the case where the goods are consumed in proportions of one-for-one and leave the treatment of the general case for the workbook.

## Bads

A **bad** is a commodity that the consumer doesn't like. For example, suppose that the commodities in question are now pepperoni and anchovies—and the consumer loves pepperoni but dislikes anchovies. But let us suppose there is some possible tradeoff between pepperoni and anchovies. That is, there would be some amount of pepperoni on a pizza that would compensate the consumer for having to consume a given amount of anchovies. How could we represent these preferences using indifference curves?

Pick a bundle $(x_1, x_2)$ consisting of some pepperoni and some anchovies. If we give the consumer more anchovies, what do we have to do with the pepperoni to keep him on the same indifference curve? Clearly, we have to give him some extra pepperoni to compensate him for having to put up with the anchovies. Thus this consumer must have indifference curves that slope up and to the right as depicted in Figure 3.5.

The direction of increasing preference is down and to the right—that is, toward the direction of decreased anchovy consumption and increased pepperoni consumption, just as the arrows in the diagram illustrate.

## Neutrals

A good is a **neutral good** if the consumer doesn't care about it one way or the other. What if a consumer is just neutral about anchovies?[1] In this case his indifference curves will be vertical lines as depicted in Figure 3.6.

[1] Is anybody neutral about anchovies?

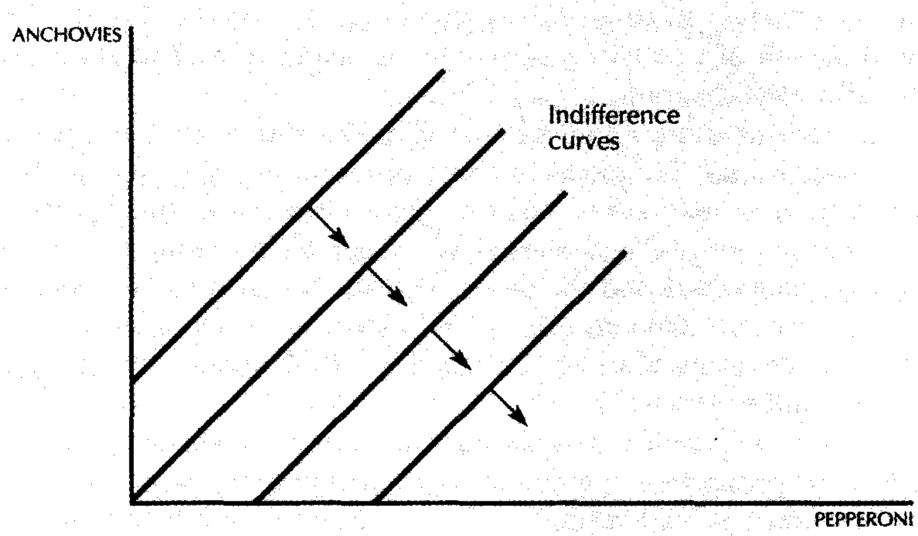


**Bads.** Here anchovies are a "bad," and pepperoni is a "good" for this consumer. Thus the indifference curves have a positive slope.

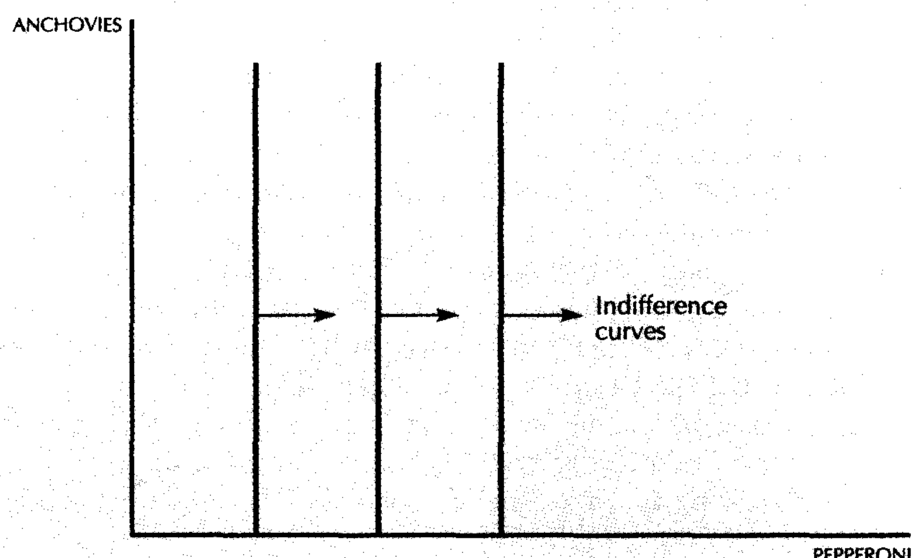


**A neutral good.** The consumer likes pepperoni but is neutral about anchovies, so the indifference curves are vertical lines.

He only cares about the amount of pepperoni he has and doesn't care at all about how many anchovies he has. The more pepperoni the better, but adding more anchovies doesn't affect him one way or the other.

## Satiation

We sometimes want to consider a situation involving **satiation**, where there is some overall best bundle for the consumer, and the "closer" he is to that best bundle, the better off he is in terms of his own preferences. For example, suppose that the consumer has some most preferred bundle of goods $(\bar{x}_1, \bar{x}_2)$, and the farther away he is from that bundle, the worse off he is. In this case we say that $(\bar{x}_1, \bar{x}_2)$ is a **satiation** point, or a **bliss** point. The indifference curves for the consumer look like those depicted in Figure 3.7. The best point is $(\bar{x}_1, \bar{x}_2)$ and points farther away from this bliss point lie on "lower" indifference curves.

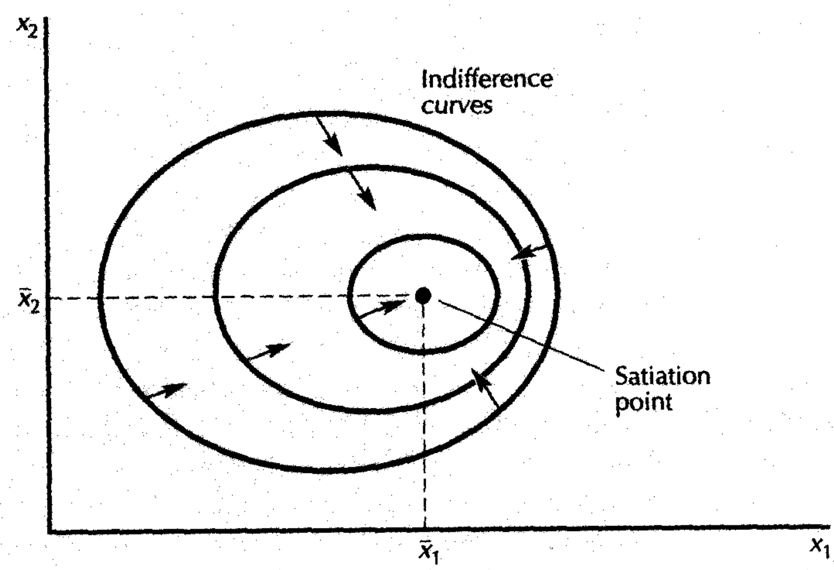


**Satiated preferences.** The bundle $(\bar{x}_1, \bar{x}_2)$ is the satiation point or bliss point, and the indifference curves surround this point.

In this case the indifference curves have a negative slope when the consumer has "too little" or "too much" of both goods, and a positive slope when he has "too much" of one of the goods. When he has too much of one of the goods, it becomes a bad—reducing the consumption of the bad good moves him closer to his "bliss point." If he has too much of both goods, they both are bads, so reducing the consumption of each moves him closer to the bliss point.

Suppose, for example, that the two goods are chocolate cake and ice cream. There might well be some optimal amount of chocolate cake and ice cream that you would want to eat per week. Any less than that amount would make you worse off, but any more than that amount would also make you worse off.

If you think about it, most goods are like chocolate cake and ice cream in this respect—you can have too much of nearly anything. But people would generally not voluntarily *choose* to have too much of the goods they consume. Why would you choose to have more than you want of something? Thus the interesting region from the viewpoint of economic choice is where you have *less* than you want of most goods. The choices that people actually care about are choices of this sort, and these are the choices with which we will be concerned.

### Discrete Goods

Usually we think of measuring goods in units where fractional amounts make sense—you might on average consume 12.43 gallons of milk a month even though you buy it a quart at a time. But sometimes we want to examine preferences over goods that naturally come in discrete units.

For example, consider a consumer's demand for automobiles. We could define the demand for automobiles in terms of the time spent using an automobile, so that we would have a continuous variable, but for many purposes it is the actual number of cars demanded that is of interest.

There is no difficulty in using preferences to describe choice behavior for this kind of discrete good. Suppose that $x_2$ is money to be spent on other goods and $x_1$ is a **discrete good** that is only available in integer amounts. We have illustrated the appearance of indifference "curves" and a weakly preferred set for this kind of good in Figure 3.8. In this case the bundles indifferent to a given bundle will be a set of discrete points. The set of bundles at least as good as a particular bundle will be a set of line segments.

The choice of whether to emphasize the discrete nature of a good or not will depend on our application. If the consumer chooses only one or two units of the good during the time period of our analysis, recognizing the discrete nature of the choice may be important. But if the consumer is choosing 30 or 40 units of the good, then it will probably be convenient to think of this as a continuous good.

## 3.5 Well-Behaved Preferences

We've now seen some examples of indifference curves. As we've seen, many kinds of preferences, reasonable or unreasonable, can be described by these simple diagrams. But if we want to describe preferences in general, it will be convenient to focus on a few general shapes of indifference curves. In

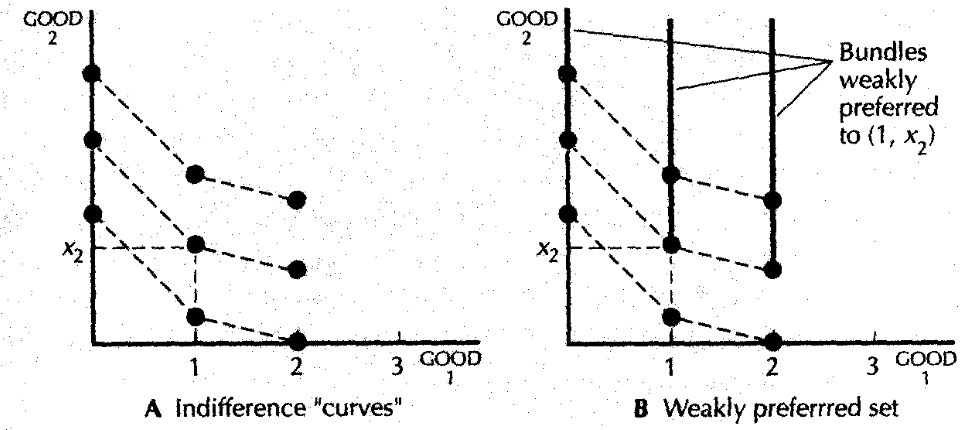


**A discrete good.** Here good 1 is only available in integer amounts. In panel A the dashed lines connect together the bundles that are indifferent, and in panel B the vertical lines represent bundles that are at least as good as the indicated bundle.

Figure 3.8

this section we will describe some more general assumptions that we will typically make about preferences and the implications of these assumptions for the shapes of the associated indifference curves. These assumptions are not the only possible ones; in some situations you might want to use different assumptions. But we will take them as the defining features for **well-behaved indifference curves**.

First we will typically assume that more is better, that is, that we are talking about *goods*, not bads. More precisely, if $(x_1, x_2)$ is a bundle of goods and $(y_1, y_2)$ is a bundle of goods with at least as much of both goods and more of one, then $(y_1, y_2) \succ (x_1, x_2)$. This assumption is sometimes called **monotonicity** of preferences. As we suggested in our discussion of satiation, more is better would probably only hold up to a point. Thus the assumption of monotonicity is saying only that we are going to examine situations *before* that point is reached—before any satiation sets in—while more *still is* better. Economics would not be a very interesting subject in a world where everyone was satiated in their consumption of every good.

What does monotonicity imply about the shape of indifference curves? It implies that they have a *negative* slope. Consider Figure 3.9. If we start at a bundle $(x_1, x_2)$ and move anywhere up and to the right, we must be moving to a preferred position. If we move down and to the left we must be moving to a worse position. So if we are moving to an *indifferent* position, we must be moving either left and up or right and down: the indifference curve must have a negative slope.

Second, we are going to assume that averages are preferred to extremes. That is, if we take two bundles of goods $(x_1, x_2)$ and $(y_1, y_2)$ on the same indifference curve and take a weighted average of the two bundles such as

$$\left(\frac{1}{2}x_1 + \frac{1}{2}y_1, \frac{1}{2}x_2 + \frac{1}{2}y_2\right),$$

then the average bundle will be at least as good as or strictly preferred to each of the two extreme bundles. This weighted-average bundle has the average amount of good 1 and the average amount of good 2 that is present in the two bundles. It therefore lies halfway along the straight line connecting the x–bundle and the y–bundle.

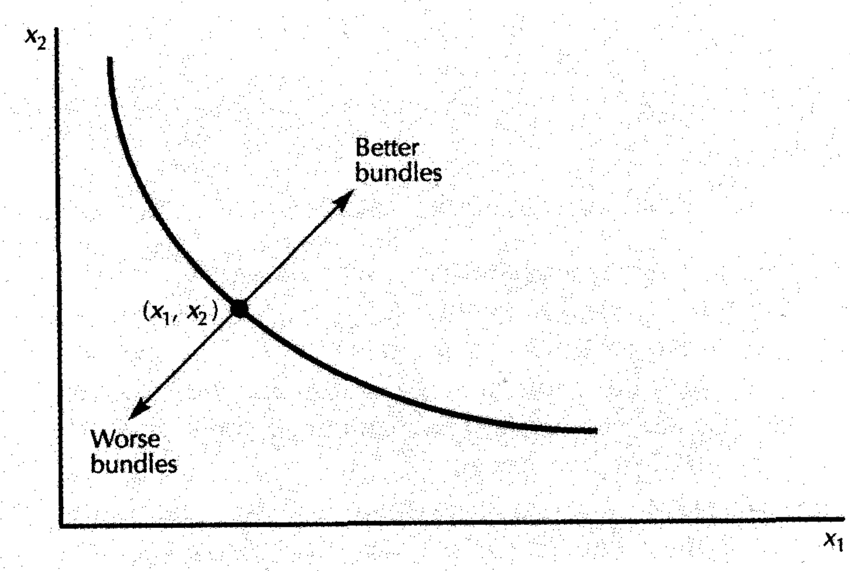


**Monotonic preferences.** More of both goods is a better bundle for this consumer; less of both goods represents a worse bundle.

Actually, we're going to assume this for any weight $t$ between 0 and 1, not just 1/2. Thus we are assuming that if $(x_1, x_2) \sim (y_1, y_2)$, then

$$(tx_1 + (1-t)y_1, tx_2 + (1-t)y_2) \succeq (x_1, x_2)$$

for any $t$ such that $0 \leq t \leq 1$. This weighted average of the two bundles gives a weight of $t$ to the x-bundle and a weight of $1 - t$ to the y-bundle. Therefore, the distance from the x-bundle to the average bundle is just a fraction $t$ of the distance from the x-bundle to the y-bundle, along the straight line connecting the two bundles.

What does this assumption about preferences mean geometrically? It means that the set of bundles weakly preferred to $(x_1, x_2)$ is a **convex set**. For suppose that $(y_1, y_2)$ and $(x_1, x_2)$ are indifferent bundles. Then, if averages are preferred to extremes, all of the weighted averages of $(x_1, x_2)$ and $(y_1, y_2)$ are weakly preferred to $(x_1, x_2)$ and $(y_1, y_2)$. A convex set has the property that if you take *any* two points in the set and draw the line segment connecting those two points, that line segment lies entirely in the set.

Figure 3.10A depicts an example of convex preferences, while Figures 3.10B and 3.10C show two examples of nonconvex preferences. Figure 3.10C presents preferences that are so nonconvex that we might want to call them "concave preferences."

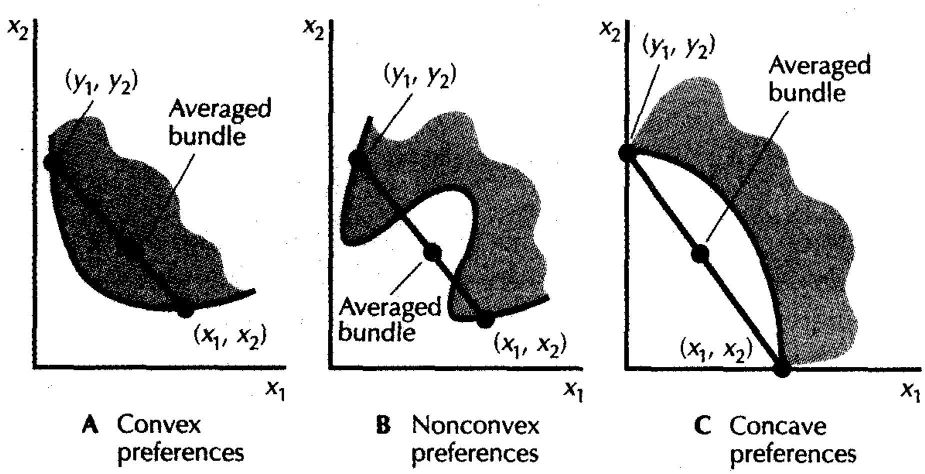


**Various kinds of preferences.** Panel A depicts convex preferences, panel B depicts nonconvex preferences, and panel C depicts "concave" preferences.

Can you think of preferences that are not convex? One possibility might be something like my preferences for ice cream and olives. I like ice cream and I like olives ... but I don't like to have them together! In considering my consumption in the next hour, I might be indifferent between consuming 8 ounces of ice cream and 2 ounces of olives, or 8 ounces of olives and 2 ounces of ice cream. But either one of these bundles would be better than consuming 5 ounces of each! These are the kind of preferences depicted in Figure 3.10C.

Why do we want to assume that well-behaved preferences are convex? Because, for the most part, goods are consumed together. The kinds of preferences depicted in Figures 3.10B and 3.10C imply that the consumer would prefer to specialize, at least to some degree, and to consume only one of the goods. However, the normal case is where the consumer would want to trade some of one good for the other and end up consuming some of each, rather than specializing in consuming only one of the two goods.

In fact, if we look at my preferences for *monthly* consumption of ice cream and olives, rather than at my immediate consumption, they would tend to look much more like Figure 3.10A than Figure 3.10C. Each month I would prefer having some ice cream and some olives—albeit at different times—to specializing in consuming either one for the entire month.

Finally, one extension of the assumption of convexity is the assumption of **strict convexity**. This means that the weighted average of two indifferent bundles is *strictly* preferred to the two extreme bundles. Convex preferences may have flat spots, while *strictly* convex preferences must have indifferences curves that are "rounded." The preferences for two goods that are perfect substitutes are convex, but not strictly convex.

## 3.6 The Marginal Rate of Substitution

We will often find it useful to refer to the slope of an indifference curve at a particular point. This idea is so useful that it even has a name: the slope of an indifference curve is known as the **marginal rate of substitution (MRS)**. The name comes from the fact that the MRS measures the rate at which the consumer is just willing to substitute one good for the other.

Suppose that we take a little of good 1, $\Delta x_1$, away from the consumer. Then we give him $\Delta x_2$, an amount that is just sufficient to put him back on his indifference curve, so that he is just as well off after this substitution of $x_2$ for $x_1$ as he was before. We think of the ratio $\Delta x_2/\Delta x_1$ as being the *rate* at which the consumer is willing to substitute good 2 for good 1.

Now think of $\Delta x_1$ as being a very small change—a marginal change. Then the rate $\Delta x_2/\Delta x_1$ measures the *marginal* rate of substitution of good 2 for good 1. As $\Delta x_1$ gets smaller, $\Delta x_2/\Delta x_1$ approaches the slope of the indifference curve, as can be seen in Figure 3.11.

When we write the ratio $\Delta x_2/\Delta x_1$, we will always think of both the numerator and the denominator as being small numbers—as describing *marginal* changes from the original consumption bundle. Thus the ratio defining the MRS will always describe the slope of the indifference curve: the rate at which the consumer is just willing to substitute a little more consumption of good 2 for a little less consumption of good 1.

One slightly confusing thing about the MRS is that it is typically a *negative* number. We've already seen that monotonic preferences imply that indifference curves must have a negative slope. Since the MRS is the numerical measure of the slope of an indifference curve, it will naturally be a negative number.

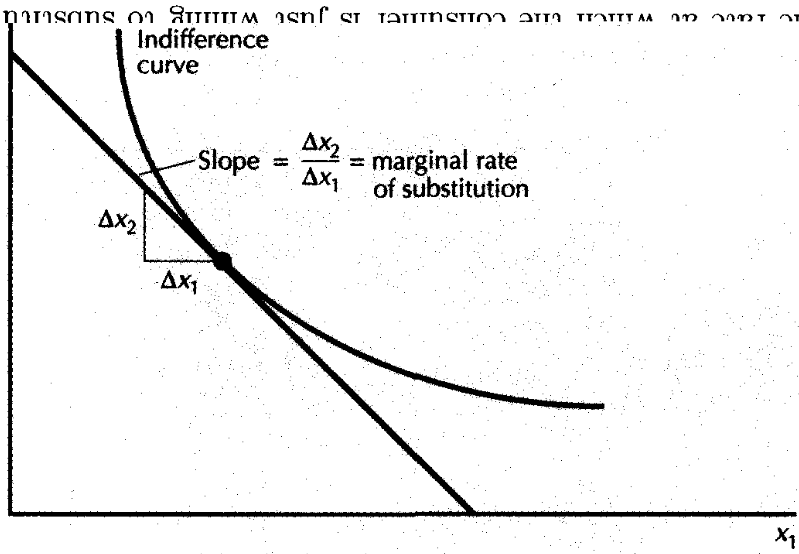


**The marginal rate of substitution (MRS).** The marginal rate of substitution measures the slope of the indifference curve.

The marginal rate of substitution measures an interesting aspect of the consumer's behavior. Suppose that the consumer has well-behaved preferences, that is, preferences that are monotonic and convex, and that he is currently consuming some bundle $(x_1, x_2)$. We now will offer him a trade: he can exchange good 1 for 2, or good 2 for 1, in any amount at a "rate of exchange" of $E$.

That is, if the consumer gives up $\Delta x_1$ units of good 1, he can get $E\Delta x_1$ units of good 2 in exchange. Or, conversely, if he gives up $\Delta x_2$ units of good 2, he can get $\Delta x_2/E$ units of good 1. Geometrically, we are offering the consumer an opportunity to move to any point along a line with slope $-E$ that passes through $(x_1, x_2)$, as depicted in Figure 3.12. Moving up and to the left from $(x_1, x_2)$ involves exchanging good 1 for good 2, and moving down and to the right involves exchanging good 2 for good 1. In either movement, the exchange rate is $E$. Since exchange always involves giving up one good in exchange for another, the exchange *rate* $E$ corresponds to a *slope* of $-E$.

We can now ask what would the rate of exchange have to be in order for the consumer to want to stay put at $(x_1, x_2)$? To answer this question, we simply note that any time the exchange line *crosses* the indifference curve, there will be some points on that line that are preferred to $(x_1, x_2)$—that lie above the indifference curve. Thus, if there is to be no movement from

$(x_1, x_2)$, the exchange line must be tangent to the indifference curve. That is, the slope of the exchange line, $-E$, must be the slope of the indifference curve at $(x_1, x_2)$. At any other rate of exchange, the exchange line would cut the indifference curve and thus allow the consumer to move to a more preferred point.

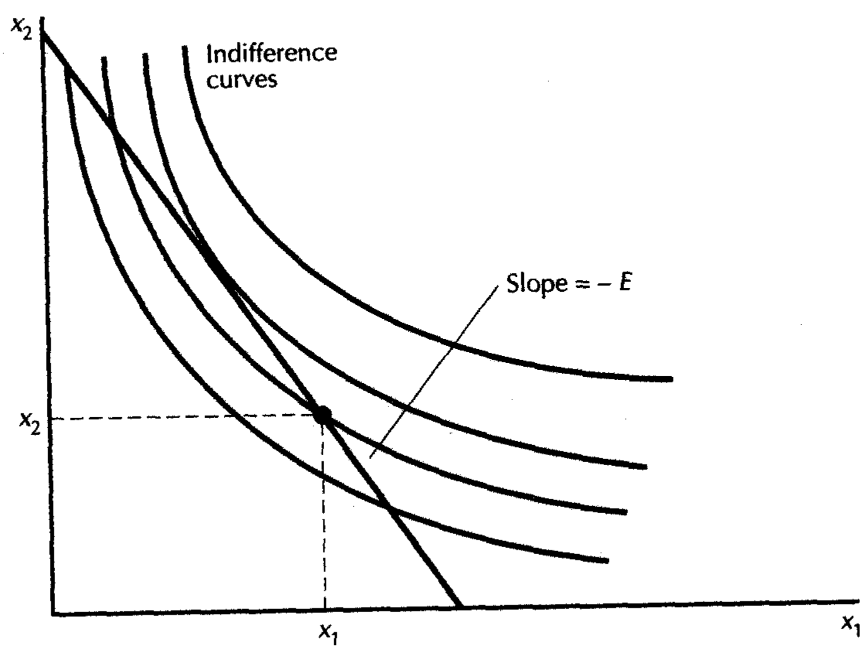


Figure 3.12 **Trading at an exchange rate.** Here we are allowing the consumer to trade the goods at an exchange rate $E$, which implies that the consumer can move along a line with slope $-E$.

Thus the slope of the indifference curve, the marginal rate of substitution, measures the rate at which the consumer is just on the margin of trading or not trading. At any rate of exchange other than the MRS, the consumer would want to trade one good for the other. But if the rate of exchange equals the MRS, the consumer wants to stay put.

## 3.7 Other Interpretations of the MRS

We have said that the MRS measures the rate at which the consumer is just on the margin of being willing to substitute good 1 for good 2. We could also say that the consumer is just on the margin of being willing to "pay" some of good 1 in order to buy some more of good 2. So sometimes you hear people say that the slope of the indifference curve measures the **marginal willingness to pay**.

If good 2 represents the consumption of "all other goods," and it is measured in dollars that you can spend on other goods, then the marginal-willingness-to-pay interpretation is very natural. The marginal rate of substitution of good 2 for good 1 is how many dollars you would just be willing to give up spending on other goods in order to consume a little bit more of good 1. Thus the MRS measures the marginal willingness to give up dollars in order to consume a small amount more of good 1. But giving up those dollars is just like paying dollars in order to consume a little more of good 1.

If you use the marginal-willingness-to-pay interpretation of the MRS, you should be careful to emphasize both the "marginal" and the "willingness" aspects. The MRS measures the amount of good 2 that one is *willing* to pay for a *marginal* amount of extra consumption of good 1. How much you actually *have* to pay for some given amount of extra consumption may be different than the amount you are willing to pay. How much you have to pay will depend on the price of the good in question. How much you are willing to pay doesn't depend on the price—it is determined by your preferences.

Similarly, how much you may be willing to pay for a large change in consumption may be different from how much you are willing to pay for a marginal change. How much you actually end up buying of a good will depend on your preferences for that good and the prices that you face. How much you would be willing to pay for a small amount extra of the good is a feature only of your preferences.

## 3.8 Behavior of the MRS

It is sometimes useful to describe the shapes of indifference curves by describing the behavior of the marginal rate of substitution. For example, the "perfect substitutes" indifference curves are characterized by the fact that the MRS is constant at $-1$. The "neutrals" case is characterized by the fact that the MRS is everywhere infinite. The preferences for "perfect complements" are characterized by the fact that the MRS is either zero or infinity, and nothing in between.

We've already pointed out that the assumption of monotonicity implies that indifference curves must have a negative slope, so the MRS always involves reducing the consumption of one good in order to get more of another for monotonic preferences.

The case of convex indifference curves exhibits yet another kind of behavior for the MRS. For strictly convex indifference curves, the MRS—the slope of the indifference curve—decreases (in absolute value) as we increase $x_1$. Thus the indifference curves exhibit a **diminishing marginal rate of**

**substitution**. This means that the amount of good 1 that the person is willing to give up for an additional amount of good 2 increases the amount of good 1 increases. Stated in this way, convexity of indifference curves seems very natural: it says that the more you have of one good, the more willing you are to give some of it up in exchange for the other good. (But remember the ice cream and olives example—for some pairs of goods this assumption might not hold!)

## Summary

1. Economists assume that a consumer can rank various consumption possibilities. The way in which the consumer ranks the consumption bundles describes the consumer's preferences.

2. Indifference curves can be used to depict different kinds of preferences.

3. Well-behaved preferences are monotonic (meaning more is better) and convex (meaning averages are preferred to extremes).

4. The marginal rate of substitution (MRS) measures the slope of the indifference curve. This can be interpreted as how much the consumer is willing to give up of good 2 to acquire more of good 1.

## REVIEW QUESTIONS

1. If we observe a consumer choosing $(x_1, x_2)$ when $(y_1, y_2)$ is available one time, are we justified in concluding that $(x_1, x_2) \succ (y_1, y_2)$?

2. Consider a group of people A, B, C and the relation "at least as tall as," as in "A is at least as tall as B." Is this relation transitive? Is it complete?

3. Take the same group of people and consider the relation "strictly taller than." Is this relation transitive? Is it reflexive? Is it complete?

4. A college football coach says that given any two linemen A and B, he always prefers the one who is bigger and faster. Is this preference relation transitive? Is it complete?

5. Can an indifference curve cross itself? For example, could Figure 3.2 depict a single indifference curve?

6. Could Figure 3.2 be a single indifference curve if preferences are monotonic?

7. If both pepperoni and anchovies are bads, will the indifference curve have a positive or a negative slope?

8. Explain why convex preferences means that "averages are preferred to extremes."

9. What is your marginal rate of substitution of \$1 bills for \$5 bills?

10. If good 1 is a "neutral," what is its marginal rate of substitution for good 2?

11. Think of some other goods for which your preferences might be concave.