chapter 6

>>Consumer and Producer Surplus

Section 1: Consumer Surplus and the Demand Curve

The market in used textbooks is not a big business in terms of dollars and cents. But it is a convenient starting point for developing the concepts of consumer and producer surplus.

So let's look at the market for used textbooks, starting with the buyers. The key point, as we'll see in a minute, is that the demand curve is derived from their preferences—and that those same preferences also determine how much they gain from the opportunity to buy used books.

Willingness to Pay and the Demand Curve

A consumer's **willingness to pay** for a good is the maximum price at which he or she would buy that good.

A used book is not as good as a new book—it will be battered and coffee-stained, may include someone else's highlighting, and may not be completely up to date. How much this bothers you depends on your own preferences. Some potential buyers would prefer to buy the used book if it is only slightly cheaper, while others would buy the used book only if it is considerably cheaper than a new book. Let's define a potential buyer's **willingness to pay** as the maximum price at which he or she would buy a good, in this case a used textbook. An individual won't buy the book if it costs more than this amount but is eager to do so if it costs less. If the price is just equal to an individual's willingness to pay, he or she is indifferent between buying and not buying.

The table in Figure 6-1 shows five potential buyers of a used book that costs $100 new, listed in order of their willingness to pay. At one extreme is Aleisha, who will buy a second-hand book even if the price is as high as $59. Brad is less willing to have a used book and will buy one only if the price is $45 or less. Claudia is willing to pay only $35, Darren only $25. And Edwina, who really doesn't like the idea of a used book, will buy one only if it costs no more than $10.

How many of these five students will actually buy a used book? It depends on the price. If the price of a used book is $55, only Aleisha buys one; if the price is $40, Aleisha and Brad both buy used books, and so on. So the information in the table on willingness to pay also defines the *demand schedule* for used textbooks.

As we saw in Chapter 3, we can use this demand schedule to derive the market demand curve shown in Figure 6-1. Because we are considering only a small number of consumers, this curve doesn't look like the smooth demand curves of earlier chapters, where markets contained hundreds or thousands of consumers. This demand curve is step-shaped, with alternating horizontal and vertical segments. Each horizontal segment—each step—corresponds to one potential buyer's willingness to pay. However, we'll see shortly that for the analysis of consumer surplus it doesn't matter whether the demand curve is stepped, as in this figure, or whether there are many consumers, making the curve smooth.

Figure 6-1 **The Demand Curve for Used Textbooks**

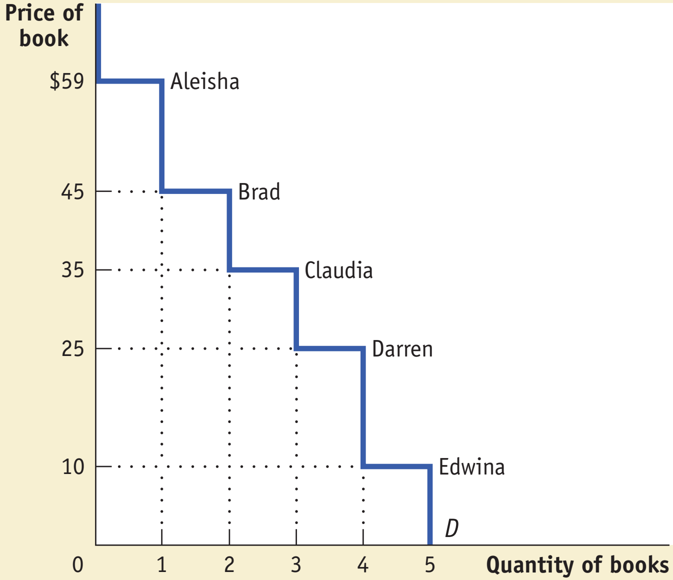

Potential buyers	Willingness to pay
Aleisha	$59
Brad	45
Claudia	35
Darren	25
Edwina	10

With only five potential consumers in this market, the demand curve is step-shaped. Each step represents one consumer, and its height indicates that consumer's willingness to pay, the maximum price at which each student will buy a used textbook, as indicated in the table. Aleisha has the highest willingness to pay at $59, Brad has the next highest at $45, and so on down to Edwina with the lowest at $10. At a price of $59 the quantity demanded is one (Aleisha); at a price of $45 the quantity demanded is two (Aleisha and Brad), and so on until you reach a price of $10 at which all five students are willing to purchase a book.

Willingness to Pay and Consumer Surplus

Suppose that the campus bookstore makes used textbooks available at a price of $30. In that case Aleisha, Brad, and Claudia will buy books. Do they gain from their purchases, and if so, how much?

The answer, shown in Table 6-1, is that each student who purchases a book does achieve a net gain but that the amount of the gain differs among students.

Aleisha would have been willing to pay $59, so her net gain is $59 – $30 = $29. Brad would have been willing to pay $45, so his net gain is $45 – $30 = $15. Claudia would have been willing to pay $35, so her net gain is $35 – $30 = $5. Darren and Edwina, however, won't be willing to buy a used book at a price of $30, so they neither gain nor lose.

Individual consumer surplus is the net gain to an individual buyer from the purchase of a good. It is equal to the difference between the buyer's willingness to pay and the price paid.

The net gain that a buyer achieves from the purchase of a good is called that buyer's **individual consumer surplus**. What we learn from this example is that every buyer of a good achieves some individual consumer surplus.

TABLE 6-1

Consumer Surplus When the Price of a Used Textbook Is $30

Potential buyer	Willingness to pay	Price paid	Individual consumer surplus = willingness to pay − price paid
Aleisha	$59	$30	$29
Brad	45	30	15
Claudia	35	30	5
Darren	25	—	—
Edwina	10	—	—

Total consumer surplus: $49

Total consumer surplus is the sum of the individual consumer surpluses of all the buyers of a good.

The term **consumer surplus** is often used to refer to both individual and to total consumer surplus.

The sum of the individual consumer surpluses achieved by all the buyers of a good is known as the **total consumer surplus** achieved in the market. In Table 6-1, the total consumer surplus is the sum of the individual consumer surpluses achieved by Aleisha, Brad, and Claudia: $29 + $15 + $5 = $49.

Economists often use the term **consumer surplus** to refer to both individual and total consumer surplus. We will follow this practice; it will always be clear in context whether we are referring to the consumer surplus achieved by an individual or by all buyers.

Total consumer surplus can be represented graphically. Figure 6-2 reproduces the demand curve from Figure 6-1. Each step in that demand curve is one book wide and represents one consumer. For example, the height of Aleisha's step is $59, her willingness to pay. This step forms the top of a rectangle, with $30—the price she actually pays for a book—forming the bottom. The area of Aleisha's rectangle, ($59 – $30) × 1 = $29, is her consumer surplus from purchasing a book at $30. So the individual consumer surplus Aleisha gains is the *area of the dark blue rectangle* shown in Figure 6-2.

In addition to Aleisha, Brad and Claudia will also buy books when the price is $30. Like Aleisha, they benefit from their purchases, though not as much, because they each have a lower willingness to pay. Figure 6-2 also shows the consumer surplus gained by Brad and Claudia; again, this can be measured by the areas of the appropriate rectangles. Darren and Edwina, because they do not buy books at a price of $30, receive no consumer surplus.

The total consumer surplus achieved in this market is just the sum of the individual consumer surpluses received by Aleisha, Brad, and Claudia. So total consumer surplus is equal to the combined area of the three rectangles—the entire shaded area in Figure 6-2. Another way to say this is that total consumer surplus is equal to the area that is under the demand curve but above the price.

This illustrates the following general principle: *The total consumer surplus generated by purchases of a good at a given price is equal to the area below the demand curve but above that price*. The same principle applies regardless of the number of consumers.

When we consider large markets, this graphical representation becomes extremely helpful. Consider, for example, the sales of personal computers to millions of potential buyers. Each potential buyer has a maximum price that he or she is willing to pay. With so many potential buyers, the demand curve will be smooth, like the one shown in Figure 6-3.

Figure 6-2

Consumer Surplus in the Used-Textbook Market

At a market price of $30, Aleisha, Brad, and Claudia each buy a book but Darren and Edwina do not. Aleisha, Brad, and Claudia get individual consumer surpluses equal to the difference between their willingness to pay and the market price, illustrated by the areas of the shaded rectangles. Both Darren and Edwina have a willingness to pay less than $30, so they are unwilling to buy a book in this market; they receive zero consumer surplus. The total consumer surplus is given by the entire shaded area—the sum of the individual consumer surpluses of Aleisha, Brad, and Claudia—equal to $29 + $15 + $5 = $49.

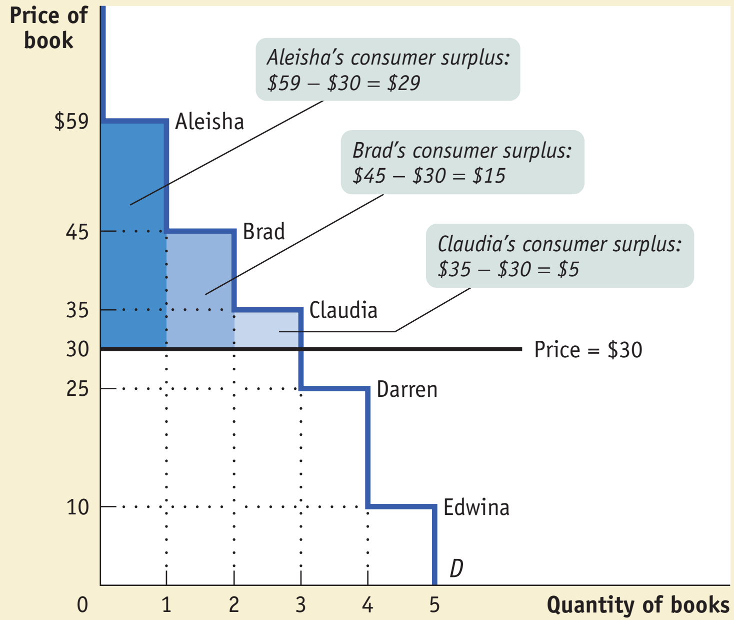

Suppose that at a price of $1,500, a total of 1 million computers are purchased. How much do consumers gain from being able to buy those 1 million computers? We could answer that question by calculating the consumer surplus of each individual buyer and then adding these numbers up to arrive at a total. But it is much easier just to look at Figure 6-3 and use the fact that the total consumer surplus is equal to the

Figure 6-3

Consumer Surplus

The demand curve for computers is smooth because there are many potential buyers of computers. At a price of $1,500, 1 million computers are demanded. The consumer surplus at this price is equal to the shaded area: the area below the demand curve but above the price. This is the total gain to consumers generated from consuming computers when the price is $1,500.

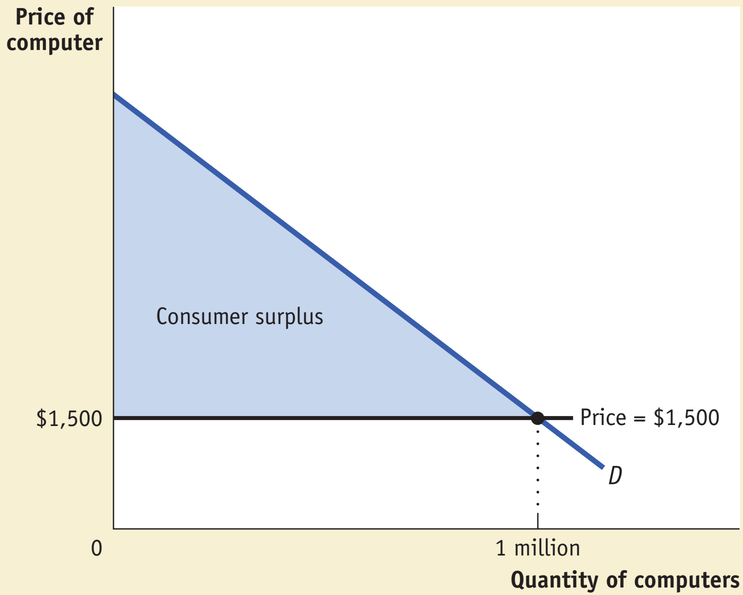

shaded area. As in our original example, consumer surplus is equal to the area below the demand curve but above the price.

How Changing Prices Affect Consumer Surplus

It is often important to know how much consumer surplus *changes* when the price changes. For example, we may want to know how much consumers are hurt if a frost in Florida drives up orange prices or how much consumers gain if the introduction of fish farming makes salmon less expensive. The same approach we have used to derive consumer surplus can be used to answer questions about how changes in prices affect consumers.

Let's return to the example of the market for used textbooks. Suppose that the bookstore decided to sell used textbooks for $20 instead of $30. How much would this increase consumer surplus?

The answer is illustrated in Figure 6-4. As shown in the figure, there are two parts to the increase in consumer surplus. The first part, shaded dark blue, is the gain of those who would have bought books even at the higher price. Each of the students who would have bought books at $30—Aleisha, Brad, and Claudia—pays $10 less, and therefore each gains $10 in consumer surplus from the fall in price to $20. So the dark blue area represents the $30 increase in consumer surplus to those three buyers. The second part, shaded light blue, is the gain of those who would not have bought a book at $30 but are willing to pay more than $20. In this case that means Darren, who would not have bought a book at $30 but does buy one at $20. He gains $5—the difference between the new price of $20 and his willingness to pay, $25. So the light blue area represents a further $5 gain in consumer surplus. The total increase in consumer surplus is the sum of the shaded areas, $35. Likewise, a rise in market price from $20 to $30 would decrease consumer surplus by an amount equal to the sum of the shaded areas.

Figure 6-4

Consumer Surplus and a Fall in the Market Price of Used Textbooks

There are two parts to the increase in consumer surplus generated by a fall in market price from \$30 to \$20. The first is given by the dark blue rectangle: each person who would have bought at the original price of \$30—Aleisha, Brad, and Claudia—receives an increase in consumer surplus equal to the total fall in price, \$10. So the area of the dark blue rectangle corresponds to an amount equal to 3 × \$10 = \$30. The second part is given by the light blue rectangle: the increase in consumer surplus for those who would *not* have bought at the original price of \$30 but who buy at the new price of \$20—namely, Darren. Darren's willingness to pay is \$25, so he now receives a consumer surplus of \$5. The total increase in consumer surplus is 3 × \$10 + \$5 = \$35, represented by the sum of the shaded areas. Likewise, a rise in market price from \$20 to \$30 would decrease consumer surplus by an amount equal to the sum of the shaded areas.

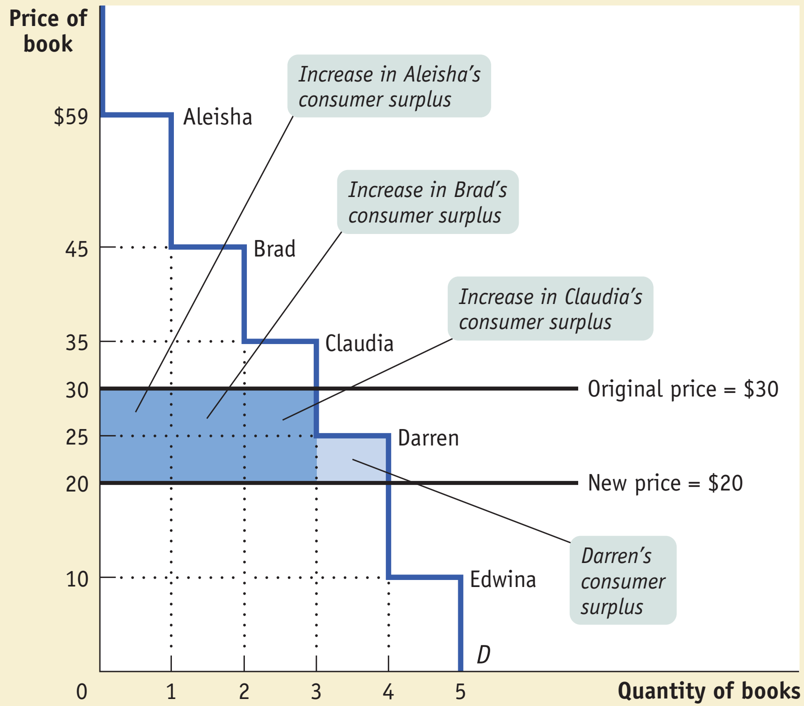

Figure 6-4 illustrates that when the price of a good falls, the area under the demand curve but above the price—which we have seen is equal to the total consumer surplus—increases. Figure 6-5 shows the same result for the case of a smooth demand curve, the demand for personal computers. Here we assume that the price of computers falls from $5,000 to $1,500, leading to an increase in the quantity demanded from 200,000 to 1 million units. As in the used-textbook example, we divide the gain in consumer surplus into two parts. The dark blue rectangle in Figure 6-5 corresponds to the dark blue area in Figure 6-4: it is the gain to the 200,000 people who would have bought computers even at the higher price of $5,000. As a result of the price fall, each receives additional surplus of $3,500. The light blue triangle in Figure 6-5 corresponds to the light blue area in Figure 6-4: it is the gain to people who would not have bought the good at the higher price but are willing to do so at a price of $1,500. For example, the light blue triangle includes the gain to someone who would have been willing to pay $2,000 for a computer and therefore gains $500 in consumer surplus when he or she is able to buy a computer for only $1,500. As before, the total gain in consumer surplus is the sum of the shaded areas, the increase in the area under the demand curve but above the market price.

What would happen if the price of a good were to rise instead of fall? We would do the same analysis in reverse. Suppose, for example, that for some reason the price of computers rises from $1,500 to $5,000. This would lead to a fall in consumer surplus, equal to the shaded area in Figure 6-5. This loss consists of two parts. The dark blue rectangle represents the losses to consumers who would still buy a computer, even at a price of $5,000. The light blue triangle represents the loss to consumers who decide not to buy a computer at the higher price.

Figure 6-5

A Fall in the Market Price Increases Consumer Surplus

A fall in the market price of a computer from $5,000 to $1,500 leads to an increase in the quantity demanded and an increase in consumer surplus. The change in the total consumer surplus is given by the sum of the shaded areas: the total area below the demand curve but between the old and new prices. Here, the dark blue area represents the increase in consumer surplus for the 200,000 consumers who would have bought a computer at the original price of $5,000; they each receive an increase in consumer surplus of $3,500. The light blue area represents the increase in consumer surplus for those willing to buy at a price equal to or greater than $1,500 but less than $5,000. Similarly, a rise in the market price of a computer from $1,500 to $5,000 generates a decrease in consumer surplus equal to the sum of the two shaded areas.

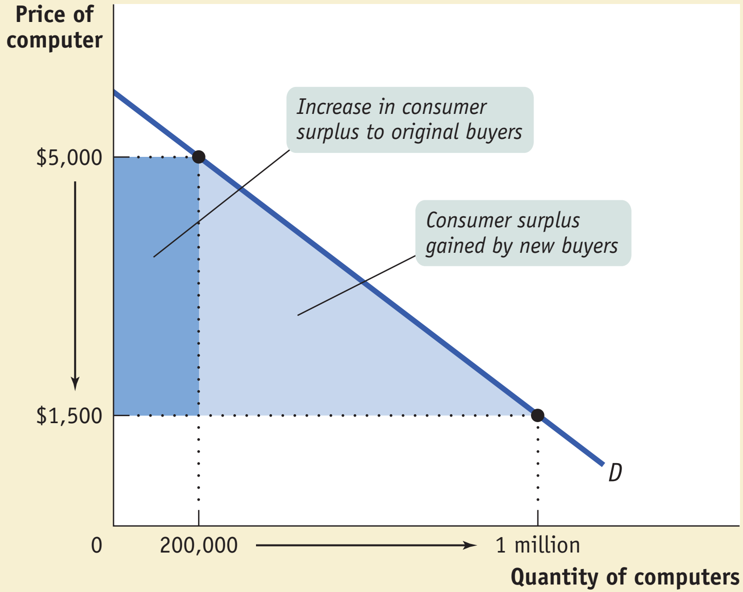

chapter 6

>>Consumer and Producer Surplus

Section 2: Producer Surplus and the Supply Curve

Just as buyers of a good would have been willing to pay more for their purchase than the price they actually pay, sellers of a good would have been willing to sell it for less than the price they actually receive. We can therefore carry out an analysis of producer surplus and the supply curve that is almost exactly parallel to that of consumer surplus and the demand curve.

Cost and Producer Surplus

Consider a group of students who are potential sellers of used textbooks. Because they have different preferences, the various potential sellers differ in the price at which they are willing to sell their books. The table in Figure 6-6 shows the prices at which several different students would be willing to sell. Andrew is willing to sell the book as long as he can get anything more than $5; Betty won't sell unless she can get at least $15; Carlos, unless he can get $25; Donna, unless she can get $35; Engelbert, unless he can get $45.

A potential seller's **cost** is the lowest price at which he or she is willing to sell a good.

The lowest price at which a potential seller is willing to sell has a special name in economics: it is called the seller's **cost**. So Andrew's cost is $5, Betty's is $15, and so on.

Figure 6-6 The Supply Curve for Used Textbooks

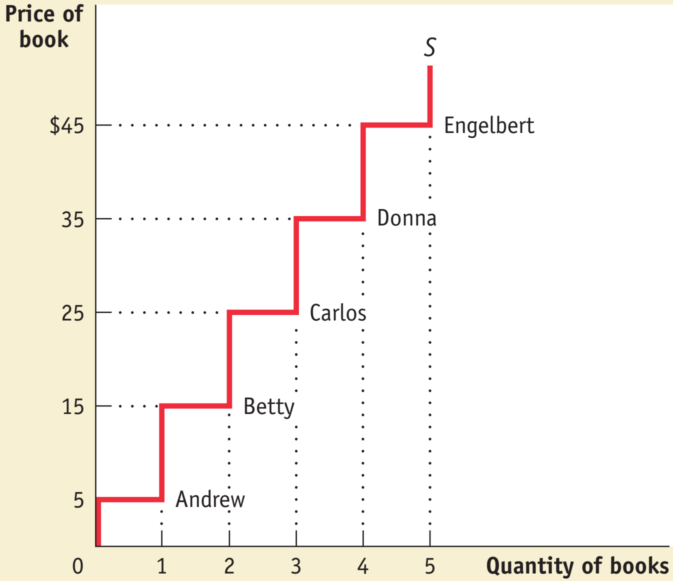

Potential sellers	Cost
Andrew	$5
Betty	15
Carlos	25
Donna	35
Engelbert	45

The supply curve illustrates sellers' cost, the lowest price at which a potential seller is willing to sell the good, and quantity supplied at that price. Each of the five students here has one book to sell and each has a different cost, as indicated in the accompanying table. At a price of $5 the quantity supplied is one (Andrew), at $15 it is two (Andrew and Betty), and so on until you reach $45, the price at which all five students are willing to sell.

Using the term *cost*, which people normally associate with the monetary cost of producing a good, may sound a little strange when applied to sellers of used textbooks. The students don't have to manufacture the books, so it doesn't cost the student who sells a book anything to make that book available for sale, does it?

Yes, it does. A student who sells a book won't have it later, as part of a personal collection. So there is an *opportunity cost* to selling a textbook, even if the owner has completed the course for which it was required. And remember that one of the basic principles of economics is that the true measure of the cost of doing anything is always its opportunity cost—the real cost of something is what you give up to get it.

So it is good economics to talk of the minimum price at which someone will sell a good as the "cost" of selling that good, even if he or she doesn't spend any money to make the good available for sale. Of course, in most real-world markets the sellers are also those who produce the good—and therefore *do* expend money to make the good available for sale. In this case the cost of making the good available for sale *includes* monetary costs—but it may also include other opportunity costs.

Getting back to the example, suppose that Andrew sells his book for $30. Clearly he has gained from the transaction: he would have been willing to sell for only $5, so he has gained $25. This gain, the difference between the price he actually gets and his cost—the minimum price at which he would have been willing to sell—is known as his **individual producer surplus.**

Individual producer surplus is the net gain to a seller from selling a good. It is equal to the difference between the price received and the seller's cost.

Total producer surplus in a market is the sum of the individual producer surpluses of all the sellers of a good. Economists use the term **producer surplus** to refer both to individual and to total producer surplus.

Just as we derived the demand curve from the willingness to pay of different consumers, we can derive the supply curve from the cost of different producers. The step-shaped curve in Figure 6-6 shows the supply curve implied by the costs shown in the accompanying table. At a price less than $5, none of the students are willing to sell; at a price between $5 and $15, only Andrew is willing to sell, and so on.

As in the case of consumer surplus, we can add the individual producer surpluses of sellers to calculate the **total producer surplus**, the total gains to sellers in the market. Economists use the term **producer surplus** to refer to either total or indi-

TABLE 6-2

Producer Surplus When the Price of a Used Textbook Is $30

Potential seller	Cost	Price received	Individual producer surplus = price received − cost
Andrew	$5	$30	$25
Betty	15	30	15
Carlos	25	30	5
Donna	35	—	—
Engelbert	45	—	—
Total producer surplus: $45			

vidual producer surplus. Table 6-2 shows the net gain to each of the students who would sell a used book at a price of $30: $25 for Andrew, $15 for Betty, and $5 for Carlos. The total producer surplus is $25 + $15 + $5 = $45.

As with consumer surplus, the producer surplus gained by those who sell books can be represented graphically. Figure 6-7 reproduces the supply curve from Figure 6-6. Each step in that supply curve is one book wide and represents one seller. The height of Andrew's step is $5, his cost. This forms the bottom of a rectangle, with $30, the price he actually receives for his book, forming the top. The area of this rectangle, ($30 – $5) × 1 = $25, is his producer surplus. So the producer surplus Andrew gains from selling his book is the *area of the dark red rectangle* shown in the figure.

Let's assume that the campus bookstore is willing to buy all the used copies of this book that students are willing to sell at a price of $30. Then, in addition to Andrew, Betty and Carlos will also sell their books. They will also benefit from their sales, though not as much as Andrew, because they have higher costs. Andrew, as we have seen, gains $25. Betty gains a smaller amount: since her cost is $15, she gains only $15. Carlos gains even less, only $5.

Again, as with consumer surplus, we have a general rule for determining the total producer surplus from sales of a good: *The total producer surplus from sales of a good at a given price is the area above the supply curve but below that price.*

Figure 6-7

Producer Surplus in the Used-Textbook Market

At a market price of $30, Andrew, Betty, and Carlos each sell a book but Donna and Engelbert do not. Andrew, Betty, and Carlos get individual producer surpluses equal to the difference between the market price and their cost, illustrated here by the shaded rectangles. Donna and Engelbert each have a cost that is greater than the market price of $30, so they are unwilling to sell a book and therefore receive zero producer surplus. The total producer surplus is given by the entire shaded area, the sum of the individual producer surpluses of Andrew, Betty, and Carlos, equal to $25 + $15 + $5 = $45.

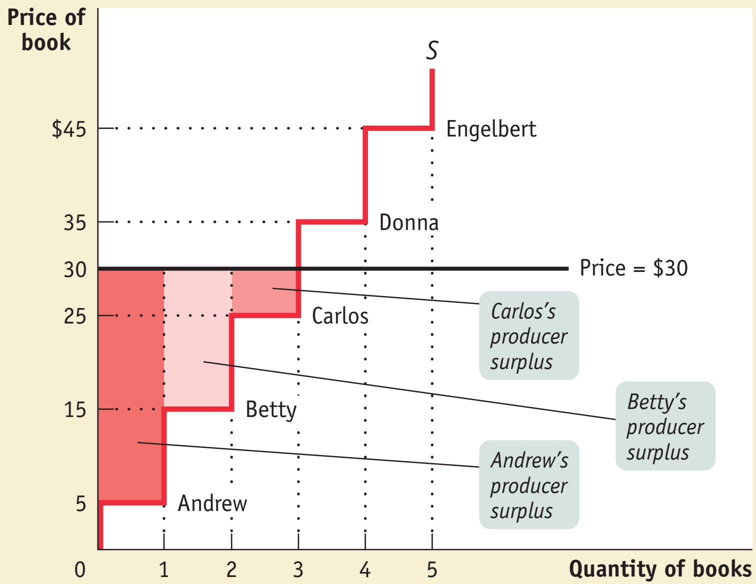

This rule applies both to examples like the one shown in Figure 6-7, where there are a small number of producers and a step-shaped supply curve, and to more realistic examples where there are many producers and the supply curve is more or less smooth.

Consider, for example, the supply of wheat. Figure 6-8 shows how the producer surplus depends on the price per bushel. Suppose that, as shown in the figure, the price is $5 per bushel and farmers supply 1 million bushels. What is the benefit to the farmers

Figure 6-8

Producer Surplus

Here is the supply curve for wheat. At a price of $5 per bushel, farmers supply 1 million bushels. The producer surplus at this price is equal to the shaded area: the area above the supply curve but below the price. This is the total gain to producers—farmers in this case—from supplying their product when the price is $5.

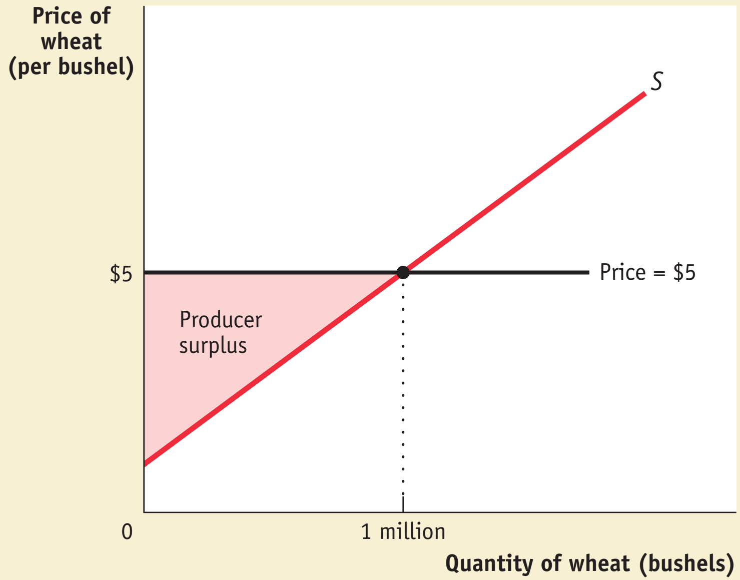

from selling their wheat at a price of $5? Their producer surplus is equal to the shaded area in the figure—the area above the supply curve but below the price of $5 per bushel.

Changes in Producer Surplus

If the price of a good rises, producers of the good will experience an increase in producer surplus, though not all producers gain the same amount. Some producers would have produced the good even at the original price; they will gain the entire price increase on every unit they produce. Other producers will enter the market because of the higher price; they will gain only the difference between the new market price and their cost.

Figure 6-9 is the supply counterpart of Figure 6-5 in "Section 1: Consumer Surplus and the Demand Curve." It shows the effect on producer surplus of a rise in the price of wheat from $5 to $7 per bushel. The increase in producer surplus is the entire shaded area, which consists of two parts. First, there is a red rectangle corresponding to the gains of those farmers who would have supplied wheat even at the original $5 price. Second, there is an additional pink triangle that corresponds to the gains of those farmers who would not have supplied wheat at the original price but are drawn into the market by the higher price.

If the price were to fall from $7 to $5 per bushel, the story would run in reverse. The whole shaded area would now be the decline in producer surplus, the fall in the area above the supply curve but below the price. The loss would consist of two parts, the loss to farmers who would still grow wheat at a price of $5 (the red rectangle) and the loss to farmers who decide not to grow wheat because of the lower price (the pink triangle).

Figure 6-9

A Rise in the Market Price Increases Producer Surplus

A rise in the market price of wheat from \$5 to \$7 leads to an increase in the quantity supplied and an increase in producer surplus. The change in the total producer surplus is given by the sum of the shaded areas: the total area above the supply curve but between the old and new prices. The red area represents the gain to the farmers who would have supplied 1 million bushels at the original price of \$5; they each receive an increase in producer surplus of \$2 for each of those bushels. The triangular pink area represents the increase in producer surplus achieved by the farmers who supply the additional 500,000 bushels because of the higher price. Similarly, a fall in the market price of wheat generates a decrease in producer surplus equal to the shaded areas.

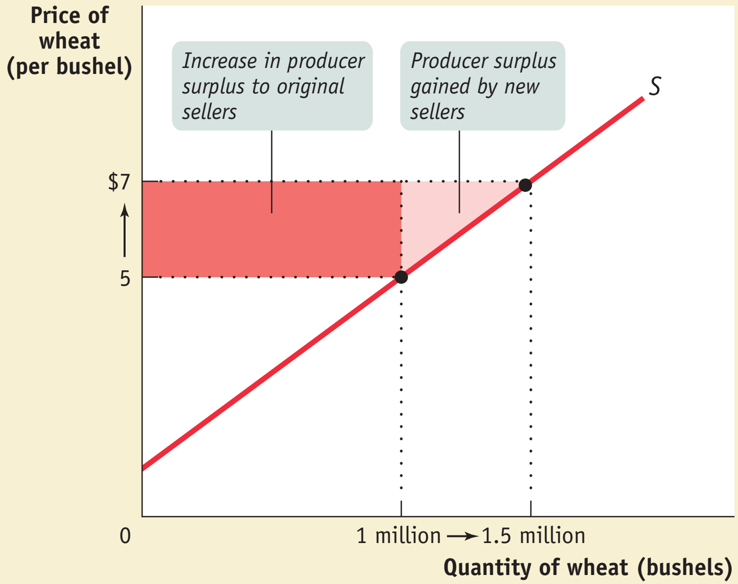

chapter 6

>>Consumer and Producer Surplus

Section 3: Consumer Surplus, Producer Surplus, and the Gains from Trade

One of the nine core principles of economics we introduced in Chapter 1 is that markets are a remarkably effective way to organize economic activity: they generally make society as well off as possible given the available resources. The concepts of consumer surplus and producer surplus can help us deepen our understanding of why this is so.

The Gains from Trade

Let's go back to the market in used textbooks but now consider a much bigger market—say, one at a large state university—where there are many potential buyers and sellers. Let's line up incoming students—who are potential buyers of the book—in order of their willingness to pay, so that the entering student with the highest willingness to pay is potential buyer number 1, the student with the next highest willingness to pay is number 2, and so on. Then we can use their willingness to pay to derive a demand curve like the one in Figure 6-10. Similarly, we can line up outgoing students, who are potential sellers of the book, in order of their cost, starting with the student with the lowest cost, then the student with the next lowest cost, and so on, to derive a supply curve like the one shown in the same figure.

As we have drawn the curves, the market reaches equilibrium at a price of $30 per book, and 1,000 books are bought and sold at that price. The two shaded triangles show the consumer surplus (blue) and the producer surplus (red) generated by this

Figure 6-10

Total Surplus

In the market for used textbooks, the equilibrium price is $30 and the equilibrium quantity is 1,000 books. Consumer surplus is given by the blue area, the area below the demand curve but above the market price. Producer surplus is given by the red area, the area above the supply curve but below the market price. The sum of the blue and the red areas is total surplus, the total benefit to society from the production and consumption of the good.

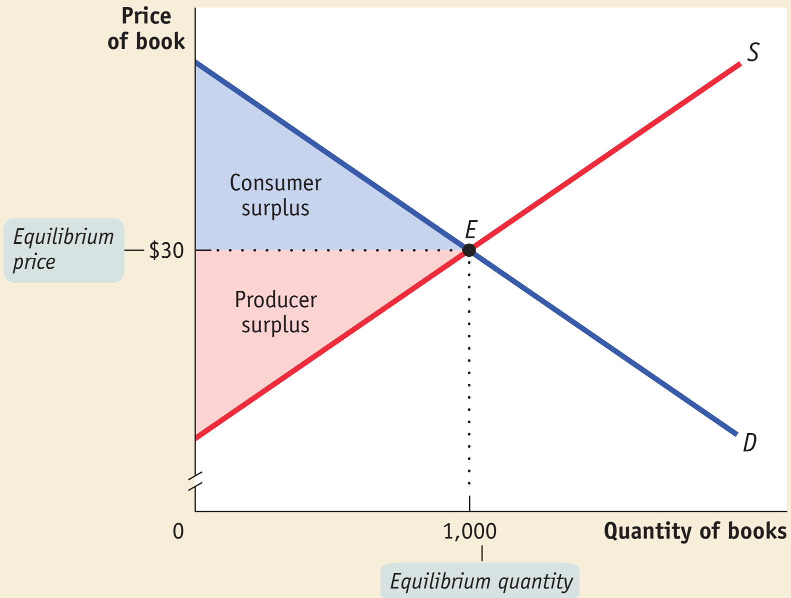

The **total surplus** generated in a market is the total net gain to consumers and producers from trading in the market. It is the sum of the producer and the consumer surplus.

market. The sum of consumer and producer surplus is known as the **total surplus** generated in a market.

The striking thing about this picture is that both consumers and producers gain—that is, both consumers and producers are better off because there is a market in this good. But this should come as no surprise—it illustrates another core principle of economics: there are *gains from trade*. These gains from trade are the reason everyone is better off participating in a market economy than they would be if each individual tried to be self-sufficient.

But are we as well off as we could be? This brings us to the question of the efficiency of markets.

The Efficiency of Markets: A Preliminary View

Markets produce gains from trade, but in Chapter 1 we made a bigger claim: that markets are usually *efficient*. That is, we claimed that once the market has produced its gains from trade, there is usually no way to make anyone better off without making someone else worse off (with some well-defined exceptions).

We're not yet ready to carry out a full discussion of the efficiency of markets—that will have to wait until we've looked in more detail at the behavior of producers and consumers. However, we can get an intuitive sense of the efficiency of markets by noticing a key feature of the market equilibrium shown in Figure 6-10: the maximum possible total surplus is achieved at market equilibrium. That is, the market equilibrium allocates the consumption of the good among potential consumers and sales of the good among potential sellers in a way that achieves the highest possible gain to society.

How do we know this? By comparing the total surplus generated by the consumption and production choices in the market equilibrium to the surplus generated by a different set of production and consumption choices. We can show that any change from the market equilibrium reduces total surplus.

Let's consider three ways in which you might try to increase the total surplus:

1. *Reallocate consumption among consumers*—take the good away from buyers who would have purchased the good in the market equilibrium, and instead give it to potential consumers who would not have bought it in equilibrium.
2. *Reallocate sales among sellers*—take sales away from sellers who would have sold the good in the market equilibrium, and instead compel potential sellers who would not have sold the good in equilibrium to sell it.
3. *Change the quantity traded*—compel consumers and producers to transact either more or less than the equilibrium quantity.

It turns out that each of these actions will not only fail to increase the total surplus; in fact, each will reduce the total surplus.

Figure 6-11 shows why reallocating consumption of the good among consumers will reduce the total surplus. Points *A* and *B* show the positions on the demand curve of two potential buyers of a used book, Ana and Bob. As we can see from the figure, Ana is willing to pay $35 for a book, but Bob is willing to pay only $25. Since the equilibrium price is $30, Ana buys a book and Bob does not.

Now suppose that we try to reallocate consumption. This would mean taking a book away from somebody who *would* have bought one at the equilibrium price of $30, like Ana, and giving that book to someone who would *not* have bought at that price, like Bob. But since the book is worth $35 to Ana, but only $25 to Bob, this would *reduce total consumer surplus* by $35 – $25 = $10.

This result doesn't depend on which two students we pick. Every student who buys a book in equilibrium has a willingness to pay that is *more* than $30, and every student who doesn't buy a book has a willingness to pay that is *less* than $30. So reallocating the good among consumers always means taking a book away from a student who values it more and giving it to a student who values it less, which necessarily reduces consumer surplus.

A similar argument, illustrated by Figure 6-12, holds for producer surplus. Here points *X* and *Y* show the positions on the supply curve of Xavier, who has a cost of $25, and Yvonne, who has a cost of $35. At the equilibrium price of $30, Xavier would sell his book but Yvonne would not. If we reallocated sales, forcing Xavier to keep his book and forcing Yvonne to give up hers, total producer surplus would be reduced by $35 – $25 = $10. Again, it doesn't matter which two students we choose. Any student who

Figure 6-11

Reallocating Consumption Lowers Consumer Surplus

Ana (point *A*) has a willingness to pay of $35. Bob (point *B*) has a willingness to pay of only $25. At the market equilibrium price of $30, Ana purchases a book but Bob does not. If we rearrange consumption by taking a book from Ana and giving it to Bob, consumer surplus declines by $10 and, as a result, total surplus declines by $10. The market equilibrium generates the highest possible consumer surplus by ensuring that those who consume the good are those who value it the most.

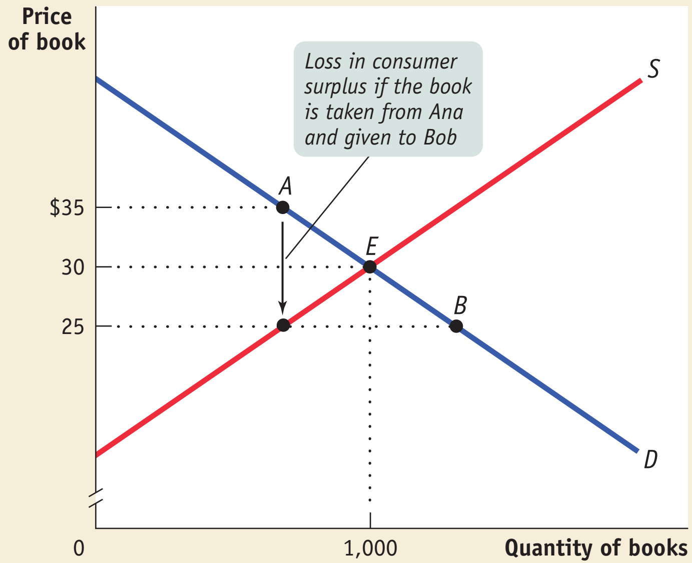

sells a book in equilibrium has a lower cost than any student who does not, so reallocating sales among sellers necessarily increases total cost and reduces producer surplus. In this way the market equilibrium generates the highest possible producer surplus: it ensures that those who sell their books are those who most value the right to sell them.

Finally, changing the quantity bought and sold reduces the sum of producer and consumer surplus. Figure 6-13 shows all four students: potential buyers Ana and

Figure 6-12

Reallocating Sales Lowers Producer Surplus

Yvonne (point *Y*) has a cost of $35, $10 more than Xavier (point *X*) who has a cost of $25. At the market equilibrium price of $30, Xavier sells a book, but Yvonne does not. If we rearrange sales by preventing Xavier from selling his book and compelling Yvonne to sell hers, producer surplus declines by $10 and, as a result, total surplus declines by $10. The market equilibrium generates the highest possible producer surplus by assuring that those who sell the good are those who value the right to sell it the most.

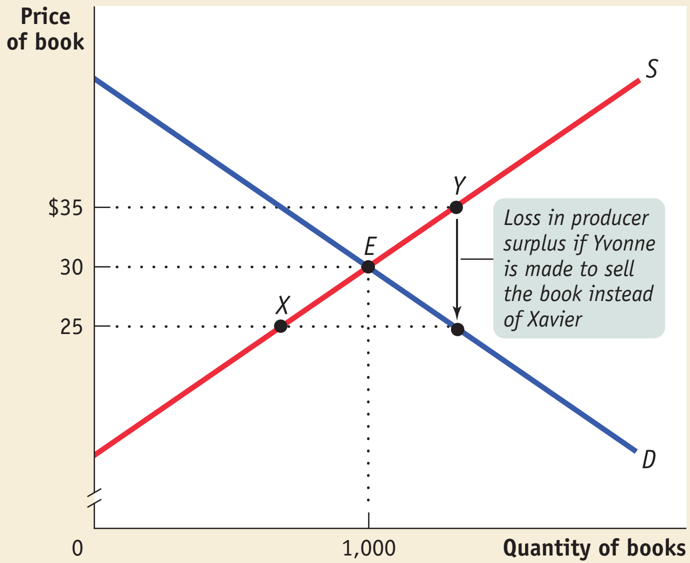

Bob, potential sellers Xavier and Yvonne. To reduce sales, we would have to prevent someone like Xavier, who would have sold the book in equilibrium, from making the sale; and the book would then not be made available to someone like Ana who would have bought it in equilibrium. As we've seen, however, Ana would be willing to pay $35, but Xavier's cost is only $25. So preventing this sale would reduce total surplus

Figure 6-13

Changing the Quantity Lowers Total Surplus

If Xavier (point *X*) were prevented from selling his book to someone like Ana (point *A*), total surplus would fall by $10, the difference between Ana's willingness to pay ($35) and Xavier's cost ($25). This means that total surplus falls whenever fewer than 1,000 books—the equilibrium quantity—are transacted. Likewise, if Yvonne (point *Y*) were compelled to sell her book to someone like Bob (point *B*), total surplus would also fall by $10, the difference between Yvonne's cost ($35) and Bob's willingness to pay ($25). This means that total surplus falls whenever more than 1,000 books are transacted. These two examples show that at market equilibrium, all beneficial transactions—and only beneficial transactions—occur.

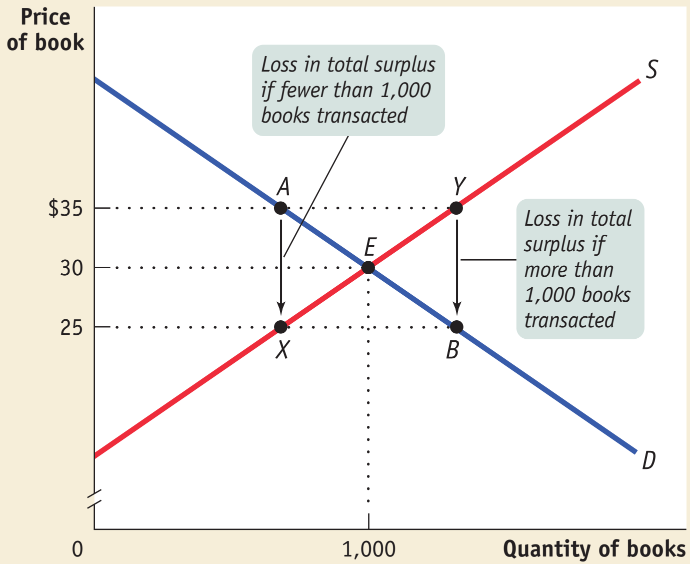

by $35 – $25 = $10. Once again, this result doesn't depend on which two students we pick: any student who would have sold the book in equilibrium has a cost of *less* than $30, and any student who would have purchased the book at equilibrium would be willing to pay *more* than $30, so preventing any sale that would have taken place in equilibrium reduces total surplus.

Finally, to increase sales would mean forcing someone like Yvonne, who would not have sold her book in equilibrium, to sell it, and giving it to someone like Bob, who would not have bought a book in equilibrium. Because Yvonne's cost is $35 but Bob is only willing to pay $25, this reduces total surplus by $10. And once again it doesn't matter which two students we pick—anyone who wouldn't have bought the book is willing to pay less than $30, and anyone who wouldn't have sold has a cost of more than $30.

What we have shown is that the market equilibrium maximizes total surplus—the sum of producer and consumer surplus. It does this because the market performs four important functions:

1. It allocates consumption of the good to the potential buyers who value it the most, as indicated by the fact that they have the highest willingness to pay.
2. It allocates sales to the potential sellers who most value the right to sell the good, as indicated by the fact that they have the lowest cost.
3. It ensures that every consumer who makes a purchase values the good more than every seller who makes a sale, so that all transactions are mutually beneficial.
4. It ensures that every potential buyer who doesn't make a purchase values the good less than every potential seller who doesn't make a sale, so that no mutually beneficial transactions are missed.

A caveat: it's important to realize that although the market equilibrium maximizes the total surplus, this does not mean that it is the best outcome for every individual consumer and producer. Other things being equal, each buyer would like to pay less and each seller would like to receive more. So some people would benefit from the price controls discussed in Chapter 4. A price ceiling that held down the market price would leave some consumers—those who managed to make a purchase—better off than they would be at equilibrium. A price floor that kept the price up would benefit some sellers—those who managed to make a sale.

But in the market equilibrium there is no way to make some people better off without making others worse off—and that's the definition of efficiency.

A Few Words of Caution

Markets are an amazingly effective way to organize economic activity; we've just demonstrated that, under certain conditions, a market is actually efficient—there is literally no way to make anyone better off without making someone else worse off.

But how secure is this result? Are markets really that good?

Market failure occurs when a market fails to be efficient.

The answer is "not always." As we discussed briefly in Chapter 1 in our ninth and final principle of economics (*when markets don't achieve efficiency, government intervention can improve society's welfare*), markets can fail to be efficient for a number of reasons. When a market is not efficient, we have what is known as a case of **market failure**. We will examine various causes of *market failure* in depth in later chapters; for now, let's review the three main reasons why markets sometimes fall short of efficiency in reality.

First, markets can fail when, in an attempt to capture more resources, one party prevents mutually beneficial trades from occurring. This situation arises, for instance, when a market contains only a single seller of a good, known as a *monopolist*. In this case, the assumption we have relied on in supply and demand analysis—that no individual buyer and seller can have a noticeable effect on the market price—is no longer valid; the monopolist can determine the market price. As we'll see in Chapter 14, this gives rise to inefficiency as a monopolist manipulates the market price in order to increase profits, thereby preventing mutually beneficial trades from occurring.

Second, actions of individuals sometimes have *side effects* on the welfare of other individuals that markets don't take into account. The best-known example of such an *externality* is pollution. We'll see in Chapter 19 that pollution and other externalities also give rise to inefficiency.

Third, markets for some goods can fail because these goods, by their very nature, are unsuited for efficient management by markets. In Chapter 18 we will analyze goods that fall into this category because of problems of *private information*—information about a good that some people possess but others don't. In Chapter 20, we will encounter other types of goods that fall into this category—*public goods*, *common resources*, and *artificially scarce goods*. These are goods for which markets fail because of problems in limiting people's access to and consumption of the good. And in Chapter 22 we will learn about *information goods:* goods like a downloaded tune, that are costly to create but, once created, cost nothing to consume.

But even with these caveats, it's remarkable how well markets work at maximizing the gains from trade. ■

chapter 6

>>Consumer and Producer Surplus

Section 4: Applying Consumer and Producer Surplus: The Efficiency Costs of a Tax

The concepts of consumer and producer surplus are extremely useful in many economic applications. Among the most important of these is assessing the efficiency cost of taxation.

In Chapter 4 we introduced the concept of an *excise tax*, a tax on the purchase or sale of a good. We saw that such a tax drives a *wedge* between the price paid by consumers and that received by producers: the price paid by consumers rises, and the price received by producers falls, with the difference equal to the tax per unit. The *incidence* of the tax—how much of the burden falls on consumers, how much on producers—does not depend on who actually writes the check to the government. Instead, as we saw in Chapter 5, the burden of the tax depends on the price elasticities of supply and demand: the higher the price elasticity of demand, the greater the burden on producers; the higher the price elasticity of supply, the greater the burden on consumers.

We also learned that there is an additional cost of a tax, over and above the money actually paid to the government. A tax causes a *deadweight loss* to society, because less of the good is produced and consumed than in the absence of the tax. As a result, some mutually beneficial trades between producers and consumers do not take place.

Now we can complete the picture, because the concepts of consumer and producer surplus are what we need to pin down precisely the deadweight loss that an excise tax imposes.

Figure 6-14 shows the effects of an excise tax on consumer and producer surplus. In the absence of the tax, the equilibrium is at *E*, and the equilibrium price and quantity are P_E and Q_E, respectively. An excise tax drives a wedge equal to the amount of the tax between the price received by producers and the price paid by consumers, reducing the quantity bought and sold. In this case, where the tax is *T* dollars per unit, the quantity bought and sold falls to Q_T. The price paid by consumers rises to P_C, the demand price of the reduced quantity, and the price received by producers falls to P_P, the supply price of that quantity. The difference between these prices, $P_C - P_P$, is equal to the excise tax, *T*.

What we can now do, using the concepts of producer and consumer surplus, is show exactly how much surplus producers and consumers lose as a result of the tax.

We saw earlier, in Figure 6-5 in Section 6.1: Consumer Surplus and the Demand Curve, that a fall in the price of a good generates a gain in consumer surplus that is equal to the sum of the areas of a rectangle and a triangle. A price increase causes a loss to consumers that looks exactly the same. In the case of an excise tax, the rise in the price paid by consumers causes a loss equal to the sum of the area of the dark blue rectangle labeled *A* and the area of the light blue triangle labeled *B* in Figure 6-14.

Meanwhile, the fall in the price received by producers causes a fall in producer surplus. This, too, is the sum of the areas of a rectangle and a triangle. The loss in producer surplus is the sum of the areas of the red rectangle labeled *C* and the pink triangle labeled *F* in Figure 6-14.

Of course, although consumers and producers are hurt by the tax, the government gains revenue. The revenue the government collects is equal to the tax per unit sold,

T, multiplied by the quantity sold, Q_T. This revenue is equal to the area of a rectangle Q_T wide and T high. And we already have that rectangle in the figure: it is the sum of rectangles A and C. So the government gains part of what consumers and producers lose from an excise tax.

But there is a part of the loss to producers and consumers from the tax that is not offset by a gain to the government—specifically, the two triangles B and F. The dead-

Figure 6-14

A Tax Reduces Consumer and Producer Surplus

Before the tax, the equilibrium price and quantity are P_E and Q_E, respectively. After an excise tax of T per unit is imposed, the price to consumers rises to P_C and consumer surplus falls by the sum of the dark blue rectangle, labeled A, and the light blue triangle, labeled B. The tax also causes the price to producers to fall to P_P; producer surplus falls by the sum of the red rectangle, labeled C, and the pink triangle, labeled F. The government receives revenue from the tax, $Q_T \times T$, which is given by the sum of the areas A and C. Areas B and F represent the losses to consumer and producer surplus that are not collected by the government as revenue; they are the deadweight loss to society of the tax.

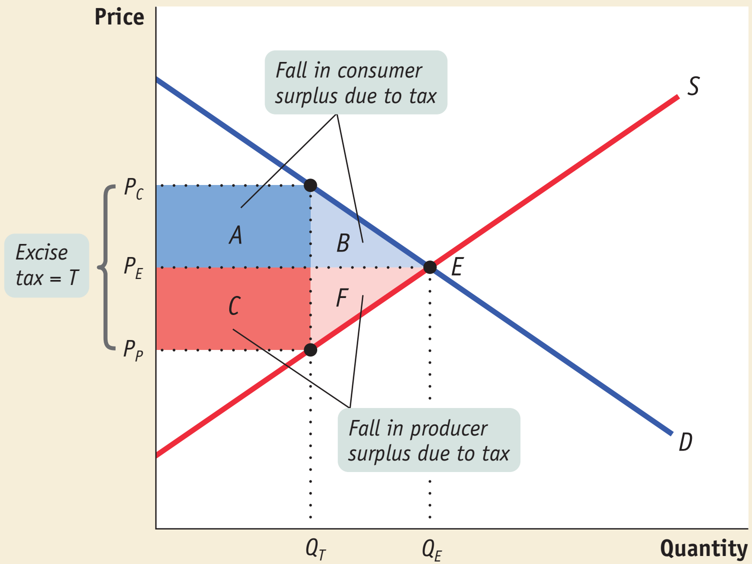

weight loss caused by the tax is equal to the combined area of these triangles. It represents the total surplus that would have been generated by transactions that do not take place because of the tax.

Figure 6-15 is a version of the same picture, leaving out the shaded rectangles—which represent money shifted from consumers and producers to the government—and showing only the deadweight loss, this time as a triangle shaded yellow. The base of that triangle is the tax wedge, *T*; the height of the triangle is the reduction the tax causes in the quantity sold, $Q_E - Q_T$. Notice that if the excise tax *didn't* reduce the quantity bought and sold in this market—if Q_T weren't less than Q_E—the deadweight loss represented by the yellow triangle would disappear. This observation ties in with the explanation given in Chapter 4 of why an excise tax generates a deadweight loss to society: the tax causes inefficiency because it discourages mutually beneficial transactions between buyers and sellers.

The idea that deadweight losses can be measured by the area of a triangle recurs in many economic applications. Deadweight-loss triangles are produced not only by excise taxes but also by other types of taxation. They are also produced by other kinds of distortions of markets, such as monopoly. And triangles are often used to evaluate other public policies besides taxation—for example, decisions about whether to build new highways.

The general rule for economic policy is that other things equal, you want to choose the policy that produces the smallest deadweight loss. This principle gives valuable guidance on everything from the design of the tax system to environmental policy. But how can we predict the size of the deadweight loss associated with a given policy? For the answer to that question, we return to a familiar concept: elasticity.

Deadweight Loss and Elasticities

The deadweight loss from an excise tax arises because it prevents some mutually beneficial transactions from occurring. In particular, the producer and consumer surplus that is forgone from these missing transactions is equal to the size of the deadweight loss itself. This means that the larger the number of transactions that are impeded by the tax, the larger the deadweight loss.

Figure 6-15

The Deadweight Loss of a Tax

A tax leads to a deadweight loss because it creates inefficiency: some mutually beneficial transactions never take place because of the tax, namely the transactions $Q_E - Q_T$. The yellow area here represents the value of the deadweight loss: it is the total surplus that would have been gained from the $Q_E - Q_T$ transactions. If the tax had not discouraged transactions—had the number of transactions remained at Q_E—no deadweight loss would have been incurred.

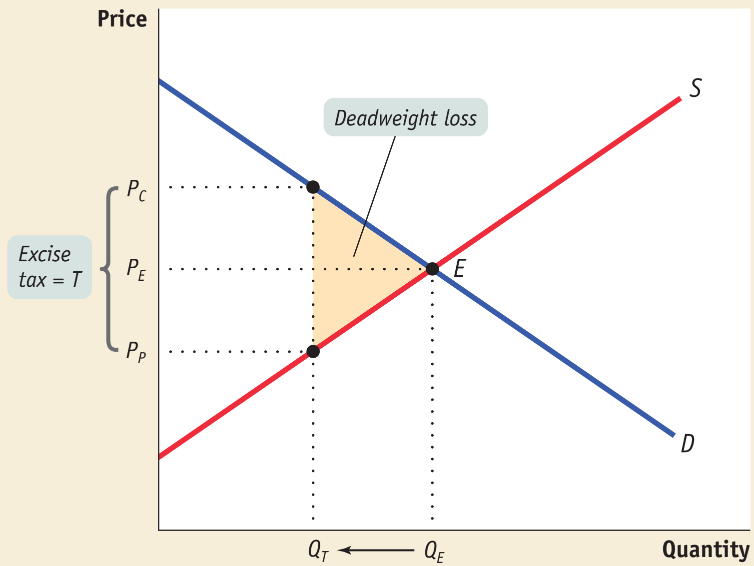

This gives us an important clue in understanding the relationship between elasticity and the size of deadweight loss from a tax. Recall that when demand or supply is elastic, it means that the quantity demanded or the quantity supplied is relatively responsive to price. So a tax imposed on a good for which either demand or supply, or both, is elastic will cause a relatively large decrease in the quantity transacted and a large deadweight loss. And when we say that demand or supply is inelastic, we mean that the quantity demanded or the quantity supplied is relatively unresponsive to price. As a result, a tax imposed when demand or supply, or both, is inelastic will cause a relatively small decrease in quantity transacted and a small deadweight loss.

The four panels of Figure 6-16 illustrate the positive relationship between price elasticity of either demand or supply and the deadweight loss of taxation. In each panel, the size of the deadweight loss is given by the area of the shaded triangle. In panel (a), the deadweight-loss triangle is large because demand is relatively elastic—a large number of transactions fail to occur because of the tax. In panel (b), the same supply curve is drawn as in panel (a), but demand is now relatively inelastic; as a result, the triangle is small because only a small number of transactions are forgone. Likewise, panels (c) and (d) contain the same demand curve but different supply curves. In panel (c), an elastic supply curve gives rise to a large deadweight-loss triangle, but in panel (d) an inelastic supply curve gives rise to a small deadweight-loss triangle.

As the following story illustrates, the implication of this result is clear: if you want to lessen the efficiency costs of taxation, you should devise taxes to fall on goods for which either demand or supply, or both, is relatively inelastic. And this lesson carries a flip-side: using a tax to purposely decrease the amount of a harmful activity, such as underage drinking, will have the most impact when that activity is elastically demanded or supplied. In the extreme case in which demand is perfectly inelastic (a vertical demand curve), the quantity demanded is unchanged by the imposition of the tax. As a result, the tax imposes no deadweight loss. Similarly, if supply is perfectly inelastic, (a vertical supply curve), the quantity supplied is unchanged by the tax and there is also no deadweight loss.

Figure 6-16 Deadweight Loss and Elasticities

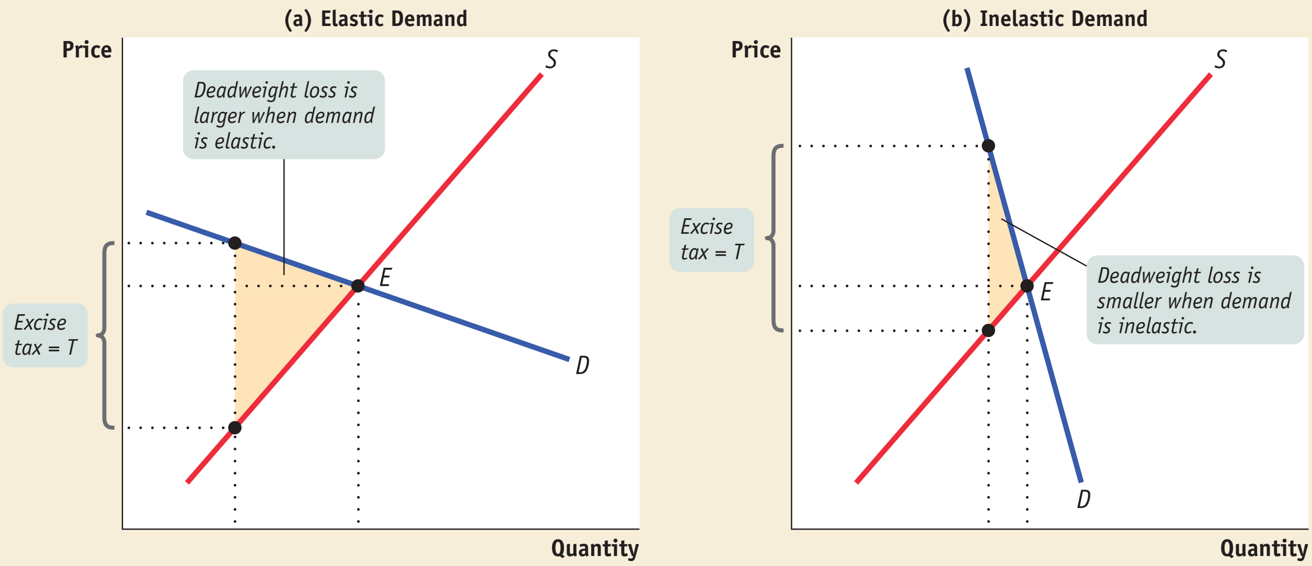

Figure 6-16 Deadweight Loss and Elasticities (continued)

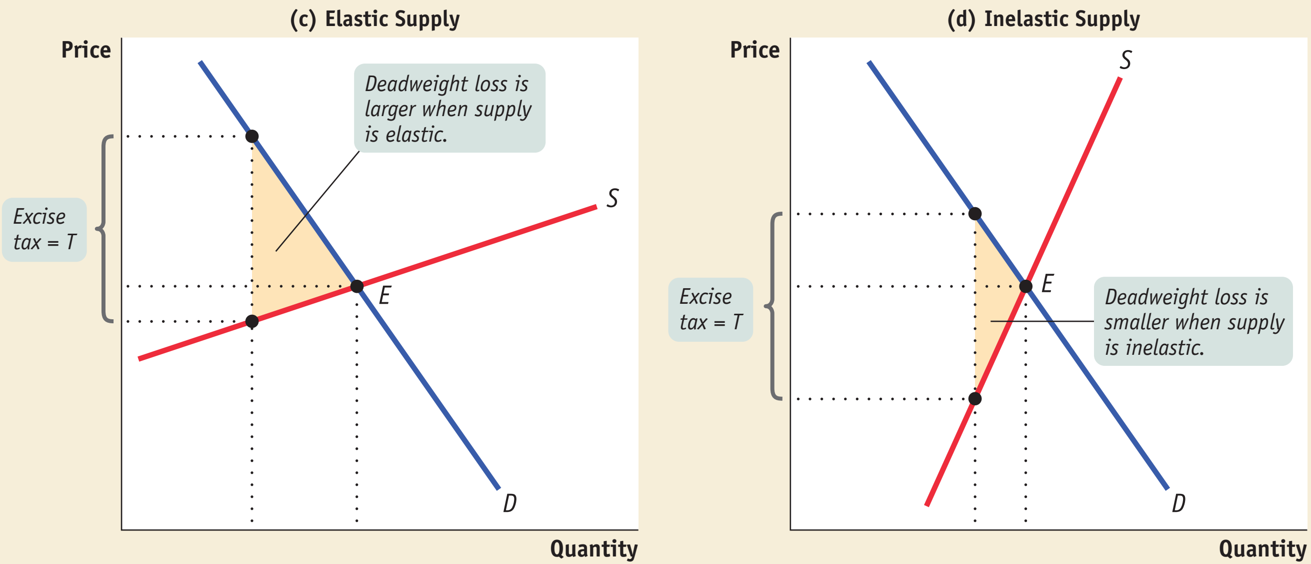

Demand is elastic in panel (a) and inelastic in panel (b), but the supply curves are the same. Supply is elastic in panel (c) and inelastic in panel (d), but the demand curves are the same. The deadweight losses are larger in panels (a) and (c) than in panels (b) and (d) because the greater the elasticity of demand or supply, the greater the tax-induced fall in the quantity transacted. In contrast, when demand or supply is inelastic, the smaller the tax-induced fall in the quantity transacted, and the smaller the deadweight loss.