

Welfare Economics and the Gains from Trade

Should some people be taxed to pay for other people's health care and education? If so, how much? Should the government subsidize purchases of home insulation? Should foreign goods be taxed to protect American workers and manufacturers?

Should governments pass laws to lower the prices of some goods and raise the prices of others? How do we know which policies are best for the economy?

The answer is that there is no answer, because there is no such thing as what's "best for the economy." Every policy you can imagine is good for some people and bad for others.

How then can you know which policies to support? Perhaps you're guided partly by self-interest. Or perhaps you prefer to balance your self-interest against what you think is fair and just for your fellow citizens.

But when a tax or a subsidy or a new law benefits some of those fellow citizens at the expense of others, how can you decide whether that policy is on net good or bad? What you need is a *normative criterion*—a way of balancing the benefits that accrue to some people against the costs that are imposed on others.

Economic theory can't tell you what your normative criterion should be. It can, however, suggest some candidates for normative criteria and help you understand what it would mean to adopt one criterion or another. In this chapter, we'll make a particularly detailed analysis of one candidate, called the *efficiency criterion*.

But before we can talk about weighing benefits against costs, we need a way of *measuring* benefits and costs. We'll start by measuring the gains from trade. When a consumer purchases a dozen eggs from a farmer, each is better off (or at least not worse off)—otherwise no trade would have occurred in the first place. The question we will address is: How *much* better off are they?

Once we know how to measure the gains from trade, we can ask how these gains are affected by various changes in market conditions. Such changes include taxes, price controls, subsidies, quotas, rationing, and so forth. We will be able to see who gains and who loses from such policies and to evaluate the size of these gains and losses.

Finally, we will learn one of the most remarkable facts in economics: In a competitive equilibrium, the sum of all the gains to all the market participants is as large as possible. This fact, called the *invisible hand theorem*, suggests one normative standard by which market outcomes can be judged. In the appendix to the chapter, we will compare this normative standard with a variety of alternatives.

8.1 Measuring the Gains from Trade

When a consumer buys eggs from a farmer, each one gains from the trade. Our first task is to measure the extent of these gains.

Consumers' and Producers' Surplus

First we consider the gains to the consumer. We begin by developing a geometric measure of the value that the consumer places on his purchases.

Marginal Value and Demand

Suppose you're so hungry that you're willing to pay up to \$15 for an egg. Does it follow that you'd be willing to pay up to \$30 for 2 eggs? Probably not. Once you've eaten your first egg, you might be willing to pay only \$10 for a second, or, in other words, \$25 total for 2 eggs.

In this case we say that (to you) the **marginal value** of the first egg is \$15 and the marginal value of the second is \$10. The **total value** (again to you) of 1 egg is \$15 and the total value of 2 eggs is \$25.

In general, we expect the marginal value of your second egg to be less than the marginal value of your first, and the marginal value of your third egg to be even lower. Why? Because when you have only 1 egg, you put it to the most valuable use you can think of. Depending on your tastes, that might mean frying it for breakfast. When you have a second egg, you put it to the second most valuable use you can think of—maybe by making egg salad for lunch. Even if you combine your 2 eggs to make an omelet, it's reasonable to think that the second half of that omelet is worth less to you than the first half.

As you acquire more eggs, their marginal value continues to fall. Consider a consumer whose marginal values are given by Table A in Exhibit 8.1. If the market price is \$7 per egg, how many eggs does this consumer buy? He certainly buys a first egg: He values it at \$15 and can get it for \$7. He also buys a second egg, which he values at \$13 and can also get for \$7. Likewise, he buys a third egg. The fourth egg, which he values at \$7 and can buy for \$7, is a matter of indifference; we will assume that the consumer buys this egg as well. The fifth egg would be a bad buy for our consumer; it provides only \$5 worth of additional value and costs \$7 to acquire. He buys 4 eggs.

Exercise 8.1 Add to Table A in Exhibit 8.1 a “Net Gain” column displaying the difference between total value and total cost. Verify that the consumer is best off when he buys 4 eggs.

Exercise 8.2 How many eggs does the consumer buy when the market price is \$5 per egg? Explain why.

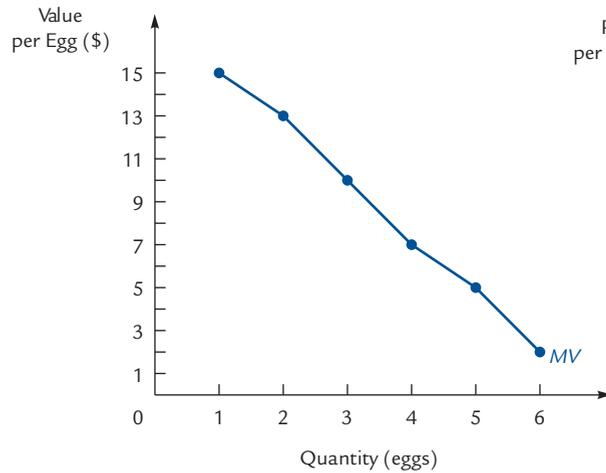
There is nothing new in this reasoning; it is just an application of the equimarginal principle. The consumer buys eggs as long as the marginal value of an egg exceeds its price and stops when the two become equal. In other words, he chooses that quantity at which price equals marginal value. In Table B of Exhibit 8.1 we record the number of eggs the consumer will purchase at each price. Table B is the consumer's demand

Marginal value
The maximum amount a consumer would be willing to pay to acquire one additional item.

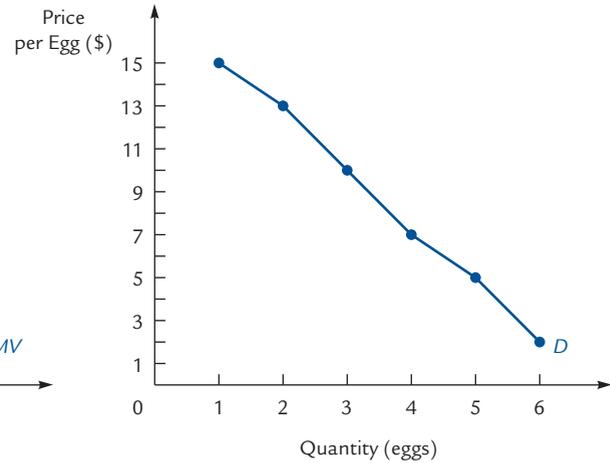
Total value
The maximum amount a consumer would be willing to pay to acquire a given quantity of items.

EXHIBIT 8.1

Demand and Marginal Value



A



B

Table A. Total and Marginal Value

Quantity	Total Value	Marginal Value
1	\$15	\$15/egg
2	28	13
3	38	10
4	45	7
5	50	5
6	52	2

Table B. Demand

Price	Quantity
\$15/egg	1
13	2
10	3
7	4
5	5
2	6

At a given market price the consumer will choose a quantity that equates price with marginal value. As a result, his demand curve for eggs is identical with his marginal value curve.

schedule, and the corresponding graph is a picture of his demand curve for eggs.¹ (Compare this reasoning with the derivation of Farmer Vickers's supply curve in Exhibit 7.4.)

The graphs in Exhibit 8.1 display both the consumer's marginal value curve and his demand curve for eggs. The curves are identical, although they differ conceptually. To read the marginal value curve, take a given quantity and read the corresponding marginal value off the vertical axis. To read the demand curve, take a given price and read the corresponding quantity off the horizontal axis.

Again, what we have learned is not new. The marginal value of an egg, measured in dollars, is the same thing as the consumer's marginal rate of substitution between

¹ More precisely, the graph is a picture of his *compensated* demand curve. When we talk about "willingness to pay" for an additional egg, we are asking what number of dollars the consumer could sacrifice for that egg and remain equally happy. The points on the marginal value curve all represent points on the same indifference curve for the consumer.

All of the demand curves in this chapter are really compensated demand curves. However, the compensated and uncompensated demand curves coincide when income effects are small, so measurements using the ordinary (uncompensated) demand curve are good approximations for most purposes.

eggs and dollars: It is the number of dollars for which he would be just willing to trade an egg. In an indifference curve diagram between eggs and dollars, the marginal value is the slope of an indifference curve and the price is the slope of the budget line. We saw in Section 3.2 that the consumer's optimum occurs at a point where the marginal value is equal to the price and that this is the source of the consumer's demand curve.

Total Value as an Area

Suppose the consumer of Exhibit 8.1 acquires 4 eggs. We would like to depict geometrically their total value. We begin by depicting the \$15 in value represented by the first egg. This \$15 is the area of rectangle 1 in panel A of Exhibit 8.2. The height of the rectangle is 15, and the width of the rectangle (which stretches from a quantity of 0 to a quantity of 1) is 1. Thus, the area is $15 \times 1 = 15$. The \$13 in value that the consumer receives from the second egg is represented by rectangle 2 in the same graph. The height of this rectangle is 13 and its width is 1, so its area is $13 \times 1 = 13$.

The area of rectangle 3 is the marginal value of the third egg, and the area of rectangle 4 is the marginal value of the fourth egg. The total value of the 4 eggs is the sum of the 4 marginal values, or the total area of the 4 rectangles. That is, the total value is $\$(15 + 13 + 10 + 7) = \45 .

Actually, what we have done is only approximately correct. That is because the marginal value table in Exhibit 8.1 omits some information. It does not show the value of $11/2$ eggs or $31/4$ eggs, for example.² In order to consider such quantities, we might make our measurements not in *eggs*, but in *quarter-eggs*. If we do so, the quantity of quarter-eggs bought is 16, and the four rectangles of panel A of Exhibit 8.2 are replaced by the 16 rectangles of panel B, each one-quarter as wide as the original ones. Refining things even further, we could measure quantities in hundredth-eggs, making 400 rectangles. As our fundamental units get smaller, our approximation to the total value of 4 eggs gets better. The total value of the consumer's 4 eggs is exactly equal to the shaded area in panel C.

The total value of the consumer's purchases is equal to the area under the demand curve out to the quantity demanded.³



Dangerous
Curve

The total value of 4 eggs is completely independent of their market price. Imagine offering the consumer a choice of living in two worlds, both identical except for the fact that in one world he has no eggs and in the other he has 4 eggs. Ask him what is the most he would be willing to pay to live in the second world rather than the first. His answer to that question is the total value that he places on 4 eggs.

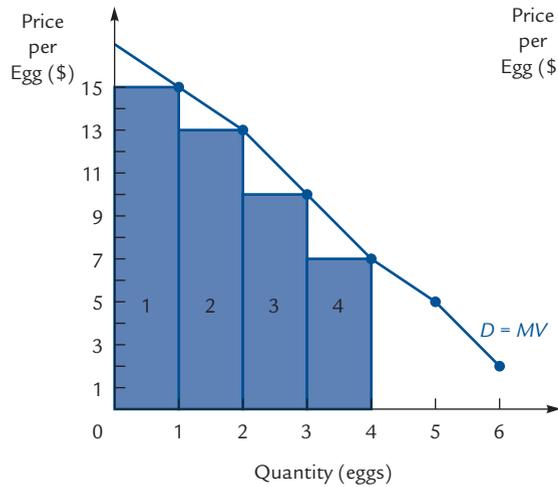
² You might think it is impossible to buy just one-quarter of an egg, but this is not so. Remember that every demand curve has a unit of time implicitly associated with it. If our demand curves are *per week*, then the way to buy exactly one-quarter of an egg per week is to buy one every four weeks.

³ If you have had a course in calculus, you might be interested to know that we have just "proven" the fundamental theorem of calculus! Think of total value as a function (where quantity is the variable). The marginal value is the addition to total value when quantity is increased by one small unit. In other words, marginal value is the derivative of total value. The area under the marginal value curve out to a given quantity is the integral of marginal value from zero out to that quantity. We have argued that this integral is equal to the total value associated with that quantity. In other words, integrating the derivative brings you back to the original function.

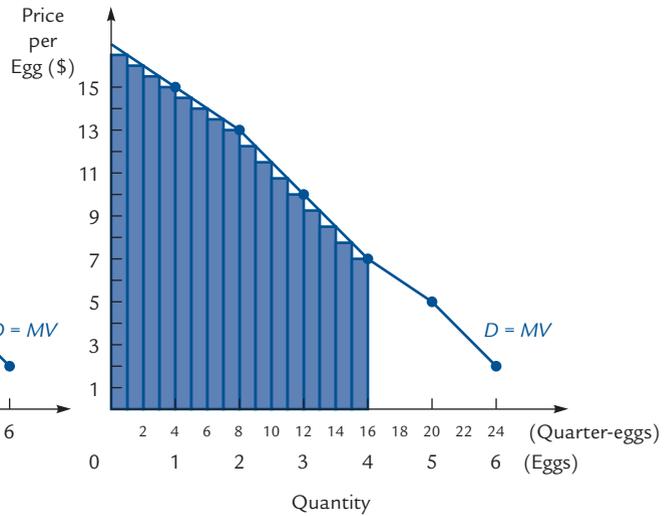
Perhaps you knew the fundamental theorem of calculus but always accepted it as a mysterious fact of nature. If so, thinking about the economics of total and marginal values should give you some real insight into why the fundamental theorem is true.

EXHIBIT 8.2

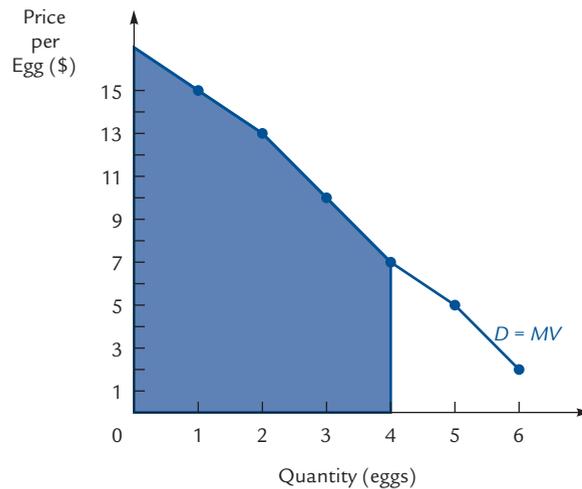
Total Value



A



B



C

When the consumer buys 4 eggs, their marginal values (\$15, \$13, \$10, and \$7) can be read off the demand curve. Their values are represented by the areas of rectangles 1 through 4 in panel A. Therefore, their total value is the sum of the areas of the rectangles, or \$45.

We can get a more accurate estimate of total value if we measure eggs in smaller units. Panel B shows the calculation of total value when we measure by the quarter-egg instead of by the whole egg. As we take smaller and smaller units, we approach the shaded area in panel C, which is the exact measure of total value when the consumer buys 4 eggs.

The Consumer's Surplus

Suppose the market price of an egg is \$7. At this price, the consumer of Exhibit 8.1 buys 4 eggs, with a total value of (approximately) \$45, which is represented by the entire shaded area in Exhibit 8.3.⁴

When the consumer buys those eggs, his total expenditure is only $4 \times \$7 = \28 —a bargain, considering that he'd have been willing to pay up to \$45. That \$28 is represented in Exhibit 8.3 by area *B*, which is a rectangle with height 7 and width 4. The extent of the bargain is measured by the difference $\$45 - \$28 = \$17$, which is area *A*. We call that area the **consumer's surplus** in the market for eggs. It is the total value (to him) of the eggs he buys, minus what he actually pays for them. In summary, we have:

Consumer's surplus

The consumer's gain from trade; the amount by which the value of his purchases exceeds what he actually pays for them.

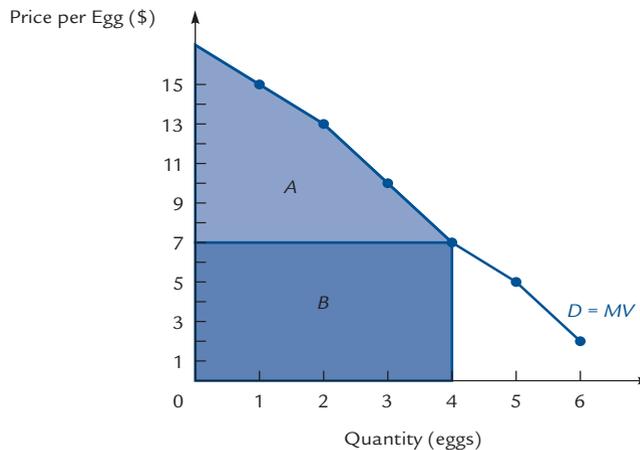
Total Value	=	$A + B$	=	\$45
Expenditure	=	B	=	\$28
Consumer's Surplus	=	A	=	\$17

This consumer would be willing to pay up to \$17 for a ticket to enter a grocery store where he can buy eggs. If the store lets him in for free, it's as if the consumer has received a gift (i.e., a free admission ticket) that he valued at \$17. You can think of the consumer's surplus as the value of that gift. Geometrically, we have seen that

The consumer's surplus is the area under the demand curve down to the price paid and out to the quantity demanded.

EXHIBIT 8.3

The Consumer's Surplus



In order to acquire 4 eggs, the consumer would be willing to pay up to the entire shaded area, $A + B$. At a price of \$7 per egg, his actual expenditure for 4 eggs is \$28, which is area *B*. The difference, area *A*, is his consumer's surplus.

⁴ \$45 is the area of the four rectangles in Exhibit 8.2A, which is approximately the same as the shaded regions in Exhibits 8.2C and 8.3. The approximation is good enough that from now on, we will say that the area of the shaded region is \$45.

Notice that the consumer's surplus is measured in units of *dollars*. In general, horizontal distances represent *quantities*, vertical distances represent *prices* (in units of, for example, dollars per egg), and areas represent numbers of *dollars*.



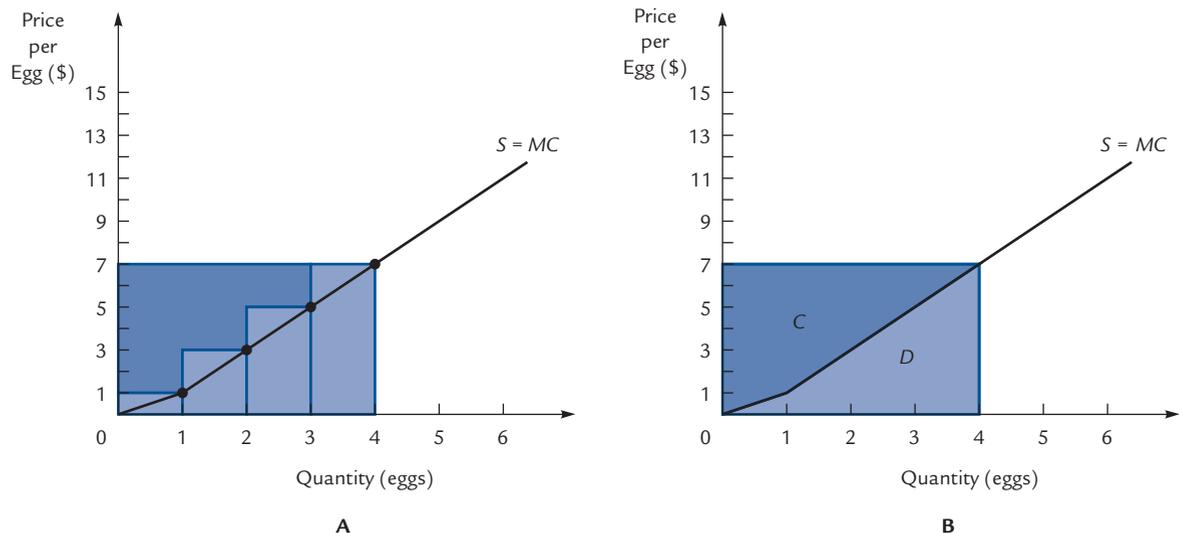
Dangerous
Curve

The Producer's Surplus

The consumer is not the only party to a transaction and not the only one to gain from it. We can also calculate the *producer's* gains from trade. Imagine a producer with the marginal cost curve shown in Exhibit 8.4. Suppose that this producer supplies 4 eggs to the marketplace. What is the cost of supplying these 4 eggs? It is the sum of the marginal cost of supplying the first egg (\$1), the marginal cost of the second (\$3), the marginal cost of the third (\$5), and the marginal cost of the fourth (\$7). These numbers are represented by the 4 rectangles in panel A of Exhibit 8.4. Their heights are 1, 3, 5, and 7, and they each have width 1.

As with the consumer's total value, we must realize that the rectangles of panel A provide only an approximation, because we are making the faulty assumption that eggs can be produced only in whole-number quantities. A more accurate picture would include very thin rectangles, and the sum of their areas would be the area labeled *D* in the second panel. This is the cost of providing 4 eggs.⁵ Area *D* is approximately equal to $\$(1 + 3 + 5 + 7) = \16 .

EXHIBIT 8.4 The Producer's Surplus



If the producer supplies 4 eggs, his cost is the sum of the 4 marginal costs, which are represented by the rectangles in panel A. If we measure eggs in very small units, we find that an exact measure of his cost is area *D* in panel B. At a market price of \$7, revenue is $\$7 \times 4 = \28 , which is the area of rectangle *C* + *D*. Thus, the producer's surplus is area *C*.

⁵ By adding up the producer's marginal costs, we are excluding any fixed costs that the producer might have. This is because we are considering only how the producer is affected by trade, whereas the producer would incur the fixed costs even without trading. This makes the fixed costs irrelevant to the discussion.

Next we depict the producer's total revenue. This is easy: He sells 4 eggs at \$7 apiece, so his revenue is $4 \times \$7 = \28 , which is the area of the rectangle $C + D$ in Exhibit 8.4.

Now we can compute the producer's gains from trade: Total revenues are $C + D$ and production costs are D . The difference, area C , is called the **producer's surplus** and represents the gains to the producer as a result of his participation in the marketplace. In this example, we have

Producer's surplus

The producer's gain from trade; the amount by which his revenue exceeds his variable production costs.

$$\begin{array}{rclclcl} \text{Total Revenue} & = & C + D & = & \$28 \\ - \text{Production Costs} & = & D & = & \$16 \\ \hline \text{Producer's Surplus} & = & C & = & \$12 \end{array}$$

This producer would be willing to pay up to \$12 for a license to sell eggs. If no license is required, it's as if the producer has received a gift (i.e., a free license) that he valued at \$12. You can think of the producer's surplus as the value of that gift.

If the producer is competitive, his marginal cost curve can be identified with his supply curve. Therefore:

The producer's surplus is the area above the supply curve up to the price received and out to the quantity supplied.

For a noncompetitive producer, we would want to change *supply curve* to *marginal cost curve* in the preceding sentence, but for a competitive producer these are the same thing.

Social Gain

In panel A of Exhibit 8.5 we have drawn both the supply and the demand curve on the same graph. The consumer's surplus is taken from Exhibit 8.3 and the producer's surplus is taken from Exhibit 8.4.

Social gain or welfare gain

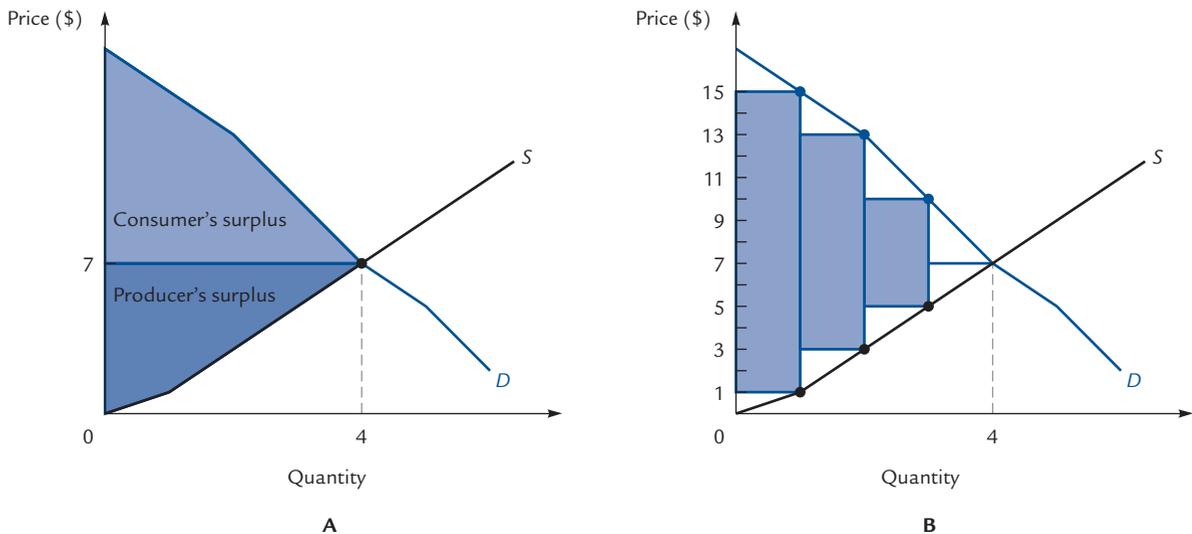
The sum of the gains from trade to all participants.

The consumer's and producer's surpluses depicted in Exhibit 8.5 provide a measure of the gains to both parties. Their sum is called the **social gain**, or **welfare gain**, due to the existence of the market. Students sometimes want to know where these gains are coming from: If the consumer and the producer have both gained, then who has lost? The answer is *nobody*. The process of trade creates welfare gains, which simply did not exist before the trading took place. The fact that the world as a whole can be made better off should not strike you as surprising: Imagine the total value of all the goods in the world 100 years ago and compare it with the value of what you see around you today. In a very real sense the difference can be thought of as the sum of all the little triangles of surplus that have been created by consumers and producers over the passage of time.

There is another way to measure the welfare gains created by the marketplace. Rather than separately computing a consumer's surplus and a producer's surplus, we can calculate the total welfare gain created by each egg. This is shown in panel B of Exhibit 8.5. The first rectangle represents the difference between the marginal value of the first egg and the marginal cost of producing it, which is precisely the welfare gain due to that egg. The height of the rectangle is $15 - 1 = 14$, and its width is 1, giving an area of 14. The second rectangle has a height of $13 - 3 = 10$ and a width of 1, giving an area of 10, which is the welfare gain from the second egg. The welfare gain due to the exchange of 4 eggs is the sum of the 4 rectangles (the fourth "rectangle" has height zero!). As usual, our focus on whole numbers has forced us to approximate: The total welfare gain is actually the entire shaded area between the supply and demand curves out to the equilibrium point.

EXHIBIT 8.5

Welfare Gains



Panel A shows the consumer's surplus and the producer's surplus when 4 eggs are sold at a price of \$7. The sum of these areas is the total welfare gain. The second panel shows another way to calculate the welfare gain. The first egg creates a gain equal to the area of the first rectangle, the second creates a gain equal to the area of the second rectangle, and so on. When units are taken to be small, the sum of these areas is the shaded region, which is the sum of the consumer's and producer's surpluses.

Notice that the total welfare gain (shown in panel B) is the sum of the consumer's and producer's surpluses (shown in panel A). This is as it should be: All of the gains have to go somewhere, and there are only the consumer and the producer to collect them.

Social Gains and Markets

Next we want to consider markets with more than one consumer and with more than one producer. It turns out that consumers' and producers' surpluses can again be computed in exactly the same way.

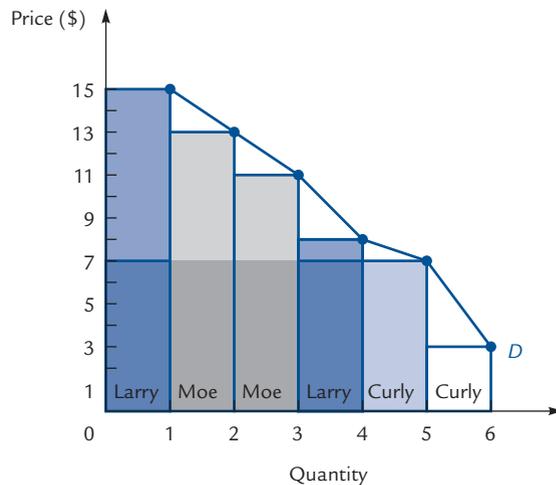
Imagine a world with three consumers: Larry, Moe, and Curly. Exhibit 8.6 displays each man's marginal value schedule for eggs. In this world, when the price is \$15, Larry buys 1 egg and Moe and Curly each buy 0 eggs. The total quantity demanded is 1. At a price of \$13, Larry and Moe buy 1 each and Curly buys 0; the quantity demanded is 2. At a price of \$11, Larry buys 1, Moe buys 2, and Curly buys 0 for a total of 3, and so on. The resulting demand curve is also shown in Exhibit 8.6.

The rectangles below the demand curve represent the marginal values of the eggs that are bought. Each rectangle is labeled with the name of the man who consumes the corresponding egg: The first egg sold is bought by Larry, the second and third by Moe, the fourth by Larry, the fifth and sixth by Curly.

Now suppose that the price of eggs is \$7. How many eggs are sold, and what is their total value? Larry buys 2 (the first and fourth), Moe buys 2 (the second and third), and Curly buys 1 (the fifth). The values of these eggs are given by the areas of the

EXHIBIT 8.6

Consumers' Surplus in the Market



Larry		Moe		Curly	
Quantity	Marginal Value	Quantity	Marginal Value	Quantity	Marginal Value
1	\$15/egg	1	\$13/egg	1	\$7/egg
2	8	2	11	2	3

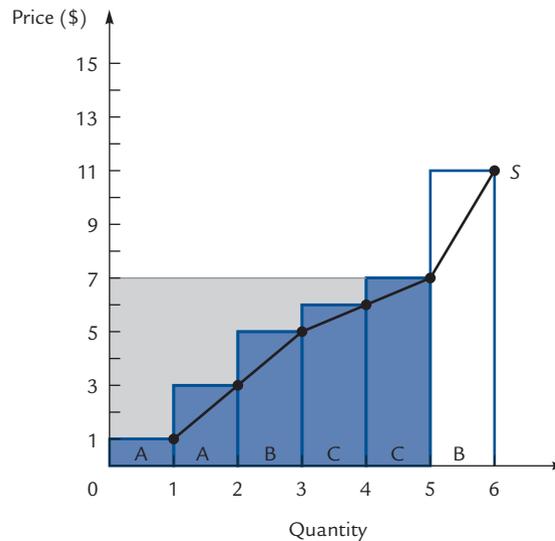
The demand curve is constructed from the marginal value curves of the three individuals. At a price of \$7, Larry buys 2 eggs that he values at \$15 and \$8, Moe buys 2 that he values at \$13 and \$11, and Curly buys 1 that he values at \$7. These marginal values are represented by the first 5 rectangles under the demand curve, each labeled with the appropriate consumer's name. The total value of the 5 eggs to the consumers is the sum of the areas of the first 5 rectangles. The cost to the consumers is the darker area. The consumers' surplus is what remains; it is the area under the demand curve down to the price paid and out to the quantity purchased.

corresponding rectangles, and the total value to the consumers is the sum of the 5 areas, which are shaded in Exhibit 8.6. From this must be subtracted the total amount that the consumers pay for the 5 eggs, which is represented by the darker, lower portions of the rectangles. The remaining portion, above the \$7 price line, is the consumers' surplus. The consumers' surplus is composed of many rectangles, and each consumer receives some of these rectangles as his share of the welfare gain. But, just as before, the total consumers' surplus is represented by the area under the demand curve down to the price paid and out to the quantity demanded.

An analogous statement holds for producers' surplus. Suppose that three different firms have the marginal cost schedules shown in Exhibit 8.7. The total supply curve is given by the graph. The colored rectangles corresponding to individual eggs are labeled with the names of the firms that produce them. Producers' surplus is given by total revenue (the entire shaded region) *minus* the sum of the areas of these rectangles, out to the quantity produced. That is, the producers' surplus is the gray part of the shaded region in the exhibit. This surplus is divided up among the producers, but the total of all the producers' surplus is still given by the area above the supply curve up to the price received and out to the quantity supplied.

EXHIBIT 8.7

Producers' Surplus in the Market



Firm A		Firm B		Firm C	
Quantity	Marginal Value	Quantity	Marginal Value	Quantity	Marginal Cost
1	\$1	1	\$5	1	\$6
2	3	2	11	2	7

The market supply curve is the industry's marginal cost curve. When the price is \$7, Firm A produces 2 items at marginal costs of \$1 and \$3, Firm B produces 1 at a marginal cost of \$5, and Firm C produces 2 at marginal costs of \$6 and \$7. These costs are represented by the colored rectangles below the supply curve. The revenue earned by producers is the entire shaded region. The gray portion of that region above the supply curve is the producers' surplus.

8.2 The Efficiency Criterion

Suppose the government decides to impose a sales tax on coffee and give away the tax revenue (say, as welfare payments or Social Security payments). Is that a good or a bad policy?

Both coffee drinkers and coffee sellers will tend to oppose this policy, because a sales tax simultaneously raises the price to demanders and lowers the price to suppliers. On the other hand, the citizens who are slated to receive the tax revenue will tend to favor the policy. How should we weigh the interests of one group against those of another?

A **normative criterion** is a general method for making this sort of decision. One example of a normative criterion is *majority rule*: Every citizen gets one vote to cast for or against the tax, and we bow to the will of the majority. In this case, the tax will probably be defeated if the coffee buyers and coffee sellers outnumber the tax recipients, and the tax will probably pass if the tax recipients outnumber the buyers and sellers.

One problem with the majority rule criterion is that it allows the slight preference of a majority to overrule the strong preference of a minority. For example, suppose

Normative criterion

A general method for choosing among alternative social policies.

that you and nine of your fellow students vote to burn down your economics professor's house for amusement. By the majority rule criterion, your 10 votes in favor of this activity outweigh the professor's 1 vote against. Nevertheless, most people would agree that burning down the house is a bad thing to do. So there is apparently something wrong with unrestricted majority rule.

Efficiency criterion

A normative criterion according to which your votes are weighted according to your willingness to pay for your preferred outcome.

An alternative to majority rule is the **efficiency criterion**. According to the efficiency criterion, everyone is permitted to cast a number of votes proportional to his stake in the outcome, where your stake in the outcome is measured by how much you'd be willing to pay to get your way. So, for example, if 10 students each think it would be worth \$10 to watch the professor's house go up in flames, while the professor thinks it would be worth \$1,000 to prevent that outcome, then each of the students gets 10 votes and the professor gets 1,000 votes. The house burning is defeated by a vote of 1,000 to 100.

Is that the right outcome? Most people seem to think so, for several reasons. Here's one of those reasons: The house burning essentially takes \$1,000 from the professor in order to give the students \$100 worth of enjoyment. But as long as you're willing to take \$1,000 from the professor, wouldn't it make more sense to take it in cash, hand out \$100 to the students (making them just as happy as the house burning would), and then have \$900 left over to do more good? You could, for example, give even more money to the students, give some back to the professor, or give it all to someone else, or any combination of those things.

The house burning is *inefficient* in the sense that it takes \$1,000 worth of house, converts it into \$100 worth of pleasure, and effectively throws \$900 away. The efficiency criterion says precisely that this kind of inefficiency is a bad thing.

One advantage of the efficiency criterion is that when it is applied consistently, you'll have the most influence on the issues you care about the most. In the appendix to this chapter, we will consider several alternatives to the efficiency criterion. In this section, we will explore the consequences of accepting the efficiency criterion and applying it to evaluate public policies. This will enable us to judge various policies—such as the sales tax on coffee—to be either “good” or “bad” *as judged by the efficiency criterion*. Of course, it does not follow that those policies are necessarily either good or bad in a larger sense. The efficiency criterion is one possible method of choosing among policies, and it is a method that you might come either to approve or disapprove.

To help you decide whether you like the efficiency criterion, it will be useful to see what it recommends in a variety of specific circumstances. That's what we'll do in this section.

Consumers' Surplus and the Efficiency Criterion

Suppose that we are deciding whether it should be legal to produce, sell, and buy eggs. Among the parties who will be interested in the outcome of this debate are the people who like to eat eggs for breakfast. They'll want to vote for legal egg sales. How many votes should we give them?

The answer, according to the efficiency criterion, is that they should receive votes in proportion to their willingness to pay for the right to buy eggs. That willingness to pay is measured by the consumers' surplus. For example, consider the consumer depicted back in Exhibit 8.3; this consumer receives a number of votes proportional to area *A*.

If we want to know how many votes should be allocated to *all* egg consumers (as opposed to the *single* egg consumer of Exhibit 8.3), we can use the market demand

curve to measure the total consumers' surplus. For example, suppose the area under the market demand curve, out to the quantity of eggs consumed and down to the market price of eggs, is equal to \$10,000. Then we know that egg consumers *as a group* should receive 10,000 pro-egg votes. (In other words, each consumer receives a number of votes proportional to his individual consumer's surplus, and we know that the sum of all these numbers is 10,000.)

Likewise, we can use the producers' surplus to compute the number of pro-egg votes cast by egg producers. If there are any anti-egg votes (say, from people who hate living next door to chicken farms), their number is a bit harder to calculate in practice. But in principle, the farmer's neighbors get a number of votes proportional to what they would be willing to pay to make the chickens go away. (The easiest case is the case where there are no unhappy neighbors; then there might be *zero* votes in favor of banning egg production.)

The Effect of a Sales Tax

Now let's return to the issue of a sales tax on coffee and evaluate that policy according to the efficiency criterion.

Panel A of Exhibit 8.8 shows the supply and demand for coffee. Panel B shows the same market after a 5¢-per-cup sales tax is placed on consumers. As we know from Chapter 1, this has the effect of lowering the demand curve vertically a distance of 5¢.

Before the sales tax is imposed, the consumers' and producers' surpluses are as shown in panel A. The sum of these is the total welfare gained by all members of society, and we will refer to it as the *social gain*. In terms of the areas in panel B, we have:

$$\begin{array}{rcl}
 \text{Consumers' Surplus} & = & A + B + C + D + E \\
 \text{Producers' Surplus} & = & F + G + H + I \\
 \hline
 \text{Social Gain} & = & A + B + C + D + E + F + G + H + I
 \end{array}$$

Once the sales tax is imposed, we need to recompute the consumers' and producers' surpluses. The consumers' surplus is the area below the demand curve down to the price paid and out to the quantity demanded. The question now arises: Which demand curve? The answer is: The original demand curve, because this is the curve that reflects the consumers' true marginal values. Which price? The price paid by demanders: P_d . Which quantity? The quantity that is bought when the tax is in effect: Q' . The consumers' surplus is area $A + B$.

In other words, the sales tax causes the consumers' surplus to fall by the amount $C + D + E$. Thus, consumers would collectively be willing to pay up to $C + D + E$ to prevent the tax, and we will eventually allow them to cast $C + D + E$ votes against it.

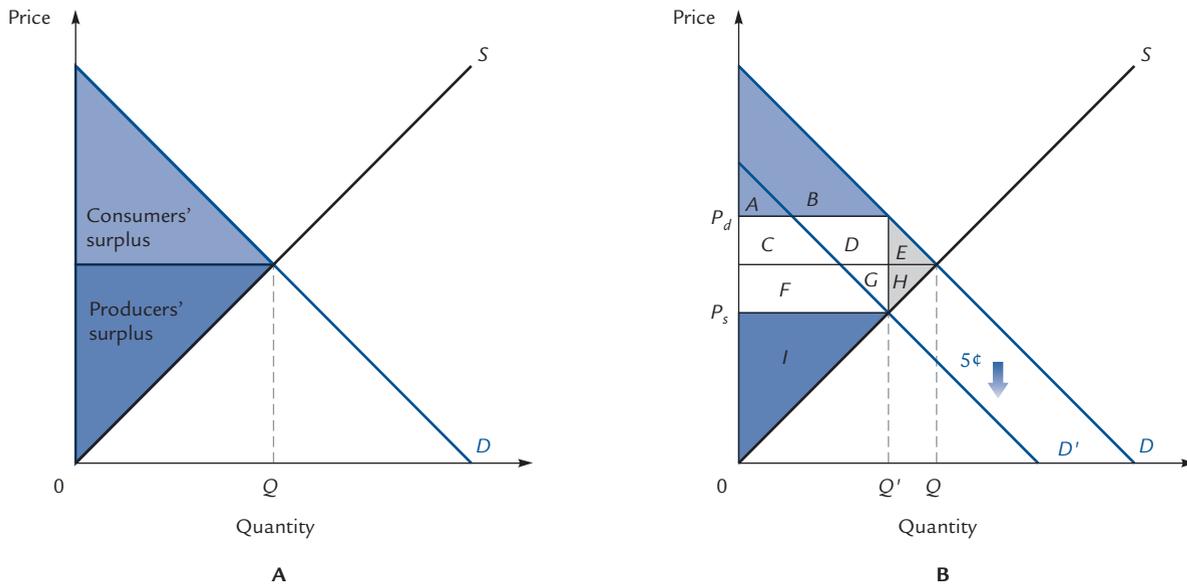
What about producers' surplus? We need to look at the area above the supply curve up to the price received and out to the quantity supplied. The relevant price to suppliers is P_s , and the relevant quantity is the quantity being sold in the presence of the sales tax: Q' . The producers' surplus is I . The tax costs producers $F + G + H$, so we will allow them to cast $F + G + H$ votes against the tax.

We can now make the following tabulation:

	Before Sales Tax	After Sales Tax
Consumers' Surplus	$A + B + C + D + E$	$A + B$
Producers' Surplus	$F + G + H + I$	I
Social Gain	$A + B + C + D + E$ $+ F + G + H + I$?

EXHIBIT 8.8

The Effect of a Sales Tax



	Before Sales Tax	After Sales Tax
Consumers' Surplus	$A + B + C + D + E$	$A + B$
Producers' Surplus	$F + G + H + I$	I
Tax Revenue	—	$C + D + F + G$
Social Gain	$A + B + C + D + E + F + G + H + I$	$A + B + C + D + F + G + I$
Deadweight Loss	—	$E + H$

Before the sales tax is imposed, consumers' and producers' surpluses are as shown in panel A. The first column of the chart shows these surpluses in terms of the labels in panel B. The second column shows the gains to consumers and producers after the imposition of the sales tax and includes a row for the gains to the recipients of the tax revenue. The total social gain after the tax is less than the social gain before the tax. The difference between the two is area $E + H$, the deadweight loss.

What about the social gain after the sales tax is imposed? Can't we find it by simply adding the consumers' and producers' surpluses? The answer is no, because there is now an additional component to consider. We must ask what becomes of the tax revenue that is collected by the government. The simplest assumption is that it is given to somebody (perhaps as a welfare or Social Security payment). Alternatively, it might be spent to purchase goods and services that are then given to somebody. In some form or another, some individual (or group of individuals) ultimately collects the tax revenue, and that individual is part of society. The revenue that the recipients collect is welfare gained.

How much tax revenue is there? The answer: It is equal to the tax per cup (5¢) times the number of cups sold (Q'). Because the vertical distance between the two demand curves is 5¢, the amount of this revenue is equal to the area of the rectangle $C + D + F + G$ (height = 5¢, width = Q'). The recipients of the tax revenue gain $C + D + F + G$

as a result of the tax, and so will be allowed to cast $C + D + F + G$ votes in its favor. The final version of our table is this:

	Before Sales Tax	After Sales Tax
Consumers' Surplus	$A + B + C + D + E$	$A + B$
Producers' Surplus	$F + G + H + I$	I
Tax Revenue	—	$C + D + F + G$
Social Gain	$A + B + C + D + E$ $+ F + G + H + I$	$A + B + C + D + F$ $+ G + I$

The social gain entry is obtained by adding the entries in the preceding three rows. Even after the tax revenue is taken into account, the total gain to society is still less after the tax than it was before. The reduction in total gain is called the **deadweight loss** due to the tax. In this example, the deadweight loss is equal to the area $E + H$. Other terms for the deadweight loss are *social loss*, *welfare loss*, and *efficiency loss*.

Let's tabulate the votes for and against this tax. Consumers cast $C + D + E$ votes against; producers cast $F + G + H$ votes against, and the recipients of tax revenue cast $C + D + F + G$ votes in favor. The tax is defeated by a margin of $E + H$ votes, so, according to the efficiency criterion, the tax is a bad thing.

It is no coincidence that the margin of defeat ($E + H$) is equal to the deadweight loss. The efficiency criterion always recommends the policy that creates the greatest social gain. If an alternative policy creates a smaller social gain, the difference is equal to the deadweight loss from that policy *and* to the margin by which that policy loses in the election prescribed by the efficiency criterion.

In doing the computations, we have considered three separate groups: consumers, producers, and the recipients of tax revenue. Some individuals might belong to two or even all three of these groups. A seller of coffee might also be a drinker of coffee; a drinker of coffee might be one of the group of people to whom the government gives the tax proceeds. Such an individual receives shares of more than one of the areas in the graph. Someone who both supplies and demands coffee will get a piece of the producers' surplus in his role as a producer and a piece of the consumers' surplus in his role as a consumer. Nevertheless, we keep track of the consumers' and producers' surpluses separately.

Deadweight loss

A reduction in social gain.



Dangerous
Curve

The Hidden and Nonhidden Assumptions

Our rejection of the sales tax is based on several hidden assumptions. First, we assumed that in the absence of the sales tax, the market price would be determined by the intersection of supply and demand. (We used this assumption when we computed the consumers' and producers' surpluses in the "no tax" column.) Although that assumption holds in competitive markets, we will see in Chapter 10 that it need *not* hold when there are firms with monopoly power.

Second, we assumed that the government simply gives away the tax revenue, as opposed to using it for some purpose that is even more valuable. We used this assumption when we entered the value of tax revenue at $C + D + F + G$. That's the amount of revenue *collected*, and it's certainly still the value of the revenue if it's simply given away. But if, for example, $C + D + F + G = \$100$, and if the government uses that \$100 to construct a post office that has a value of \$300 (measured by people's willingness to pay

for the post office), then our calculation of social gain in the “after sales tax” column is off by \$200. In Chapter 14, we will discuss the circumstances in which governments might be able to spend money more efficiently than individuals can.

Third, we assumed that the production and consumption of coffee does not affect anyone but the producers and consumers. But suppose that coffee producers use heavy machinery that keeps their neighbors awake at night or that coffee drinkers use styro-foam cups that they throw by the roadside when they’re done. Then there should be additional rows in our chart to reflect the concerns of sleep-deprived neighbors and Sunday motorists who prefer not to confront other people’s litter. By omitting these rows, we assumed that there are no significant concerns of this kind. In Chapter 13, we will discuss how to incorporate such concerns in the analysis.

In addition to these hidden assumptions, we have made the nonhidden assumption that the efficiency criterion is an appropriate way to judge a policy. If any one of these assumptions is violated, we might need to reconsider the desirability of the sales tax on coffee.

Understanding Deadweight Loss

Exhibit 8.9 presents another view of the deadweight loss. The prices and quantities are the same as in panel B of Exhibit 8.8. At the original equilibrium quantity Q , the social gain is the sum of all the rectangles. At Q' , which is the quantity with the tax, the social gain consists of only the color rectangles. The next cup of coffee after Q' would increase welfare if it were produced, because the marginal value it provides (read off the demand curve) exceeds the marginal cost of producing it (read off the supply curve). However, that cup is not produced and an opportunity to add to welfare is lost.

EXHIBIT 8.9

Deadweight Loss



If the market operates at the equilibrium quantity Q , all of the rectangles are included in the social gain. If for any reason the market operates at the quantity Q' (e.g., because of a tax), then only the color rectangles are included. The units of output that could create the gray rectangles are never produced, and those rectangles of gain are never created. The gray rectangles, representing gains that could have been created but weren't, constitute the deadweight loss.

The deadweight loss calculated in Exhibit 8.9 is the same as the deadweight loss calculated in Exhibit 8.8, where it corresponds to the area $E + H$.

If we think of the social gain as a pie divided among various groups, then a tax has two effects: It changes the way the pie is distributed, and it simultaneously changes the size of the entire pie. Thus, in Exhibit 8.8, the pie originally consists of all the lettered areas. The tax reduces the consumers' and producers' pieces. On the other hand, the recipients of the tax revenue, who get nothing in the absence of the tax, now receive a piece of the pie. After adding up everyone's pieces, we find that the total pie has shrunk; the losses to the losers exceed the gains to the winners. The shrinkage in the pie is the deadweight loss.

The deadweight loss is *not* due in any way to the costs of collecting the tax. We have been assuming that these costs are negligible. If, in fact, it is necessary to hire tax collectors, to provide them with office space, and to buy them computers, or if it is costly for citizens to compute their taxes or to deliver them to the government, these are additional losses that are not included in our computation of the deadweight loss.



Dangerous
Curve

Exercise 8.3 In Exhibit 8.8 how much does each group of losers lose? How much does each group of winners win? Is the excess of losses over gains equal to the deadweight loss?

A moral of this story is that “taxes are bad”—though not in the sense you might think. You might think that taxes are bad because paying them makes you poorer. True, but collecting them makes somebody else richer. In Exhibit 8.8 the areas $C + D + F + G$ that are paid in taxes do end up in somebody's pocket. Whether this is a good thing or a bad thing depends on whose pocket you care about most. The aspect of the tax that is unambiguously “bad” is the deadweight loss. This is a loss to consumers and producers that is not offset by a gain to anybody.

Exercise 8.4 Work out the effects of an excise tax of 5¢ per cup of coffee. (*Hint:* We already know that an excise tax has exactly the same effects as a sales tax, so you will know your answer is right if it gives exactly the same results as in Exhibit 8.8.)

Whenever a policy creates a deadweight loss, it is possible to imagine an alternative policy that would be better for everybody. Exhibit 8.10 revisits the effect of a 5¢ sales tax. The tax costs consumers the amount $C + D + E$, costs producers the amount $F + G + H$, and delivers $C + D + F + G$ to the tax recipients.

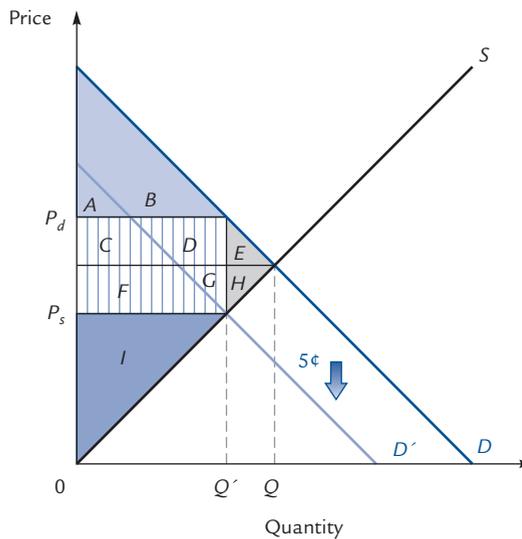
Suppose that *instead* of the 5¢ sales tax, we adopt the following plan: One night, without warning, the tax collector breaks into the homes of the consumers and steals a total of $C + D + E$ dollars from their dresser drawers; then he breaks into the homes of the producers and steals a total of $F + G + H$. As far as the consumers and producers are concerned, this is neither better nor worse than being taxed. The tax collector then gives $C + D + F + G$ dollars to the tax recipients—the same amount they'd have received in tax revenue. The tax collector now has $E + H$ dollars left over to do some additional good with. The collector can return part of it to the consumers and producers, give part of it to the tax recipients, give part of it to charity, keep part of it, or do any combination of these things. If he wants to, he can give *everyone* a small sliver of that $E + H$ and make *everyone* better off.

By eliminating the deadweight loss of $E + H$ dollars, the tax collector can do exactly an additional $E + H$ dollars worth of good. This isn't too surprising: When there's more surplus to go around, we can always find a way to increase everyone's share. Or in other words: When the pie is bigger, you can always give everyone a bigger piece.

EXHIBIT 8.10

Eliminating Deadweight Loss

	No Tax	With Sales Tax
Consumers' Surplus	$A + B + C + D + E$	$A + B$
Producer's Surplus	$F + G + H + I$	I
Tax Revenue	—	$C + D + F + G$
Social Gain	$A + B + C + D + E + F + G + H + I$	$A + B + C + D + F + G + I$
Deadweight Loss	—	$E + H$



The sales tax costs consumers $C + D + E$ and producers $F + G + H$. If, instead of imposing a sales tax, the collector simply steals these amounts from the consumers and the producers, he can give $C + D + F + G$ to the tax recipients (just as they'd receive under the sales tax) and still have $E + H$ left over to do some additional good.



Dangerous
Curve

An important feature of our alternative policy is that it is totally unexpected and nobody can do anything to avoid it. If people know in advance, for example, that the tax collector will be stealing from all producers of coffee, the producers will react to this as they would to a tax and produce less. Their exact reaction will depend on the collector's exact policy: If he steals more from those who produce more, he is effectively imposing an excise tax, which causes each firm to reduce its quantity. If he steals equally from all producers, the main effect will be to drive some producers out of the industry altogether.

When people anticipate the collector's actions, they will take steps to avoid them. These steps will include producing and consuming less coffee, and this will create a deadweight loss. The only way to avoid a deadweight loss is for the market to produce the equilibrium quantity of coffee, and this happens only if nobody is given a chance to alter his or her behavior in order to reduce the tax burden.

As with the tax policy, we are ignoring any costs involved with implementing the alternative policy (such as the collector's expenditure on burglar tools or the value of his time). Any such costs would lessen the social gain.



Dangerous
Curve

It would be enormously impractical to subject consumers and producers to random unexpected thefts. That's one reason we still have taxes, despite the deadweight loss. When a tax creates a reasonably small deadweight loss, we might be willing to live with it; when another tax creates a much larger deadweight loss, we might be more inclined to think hard about alternatives.



Dangerous
Curve

Other Normative Criteria

The simplest of all normative criteria is the **Pareto criterion**, according to which one policy is “better” than another when it is preferred unanimously. In Exhibit 8.10, this means that the alternative policy is better than the sales tax, because everyone—consumers, producers, *and* recipients of tax revenue—agrees on this assessment. But according to the Pareto criterion, there is no way to decide between the “no tax” policy in the first column and the “sales tax” policy in the second column. Consumers and producers prefer the first, while tax recipients prefer the second. There is no unanimity; therefore, the Pareto criterion remains silent.

The great advantage of the Pareto criterion is that its recommendations, when it makes them, are extremely noncontroversial. Who can disagree with the outcome of a unanimous election? The offsetting disadvantage is that the Pareto criterion usually makes no recommendation at all, because unanimity is rarely found.

One modification of the Pareto criterion is the **potential Pareto criterion**, according to which any proposal that *could* be unanimously defeated should be rejected—even if the proposal that defeats it is not really in the running. For example, suppose in Exhibit 8.10 that we are asked to choose between the “no tax” proposal in the first column and the “sales tax” proposal in the second. According to the potential Pareto criterion, we should reject the sales tax because it loses unanimously to the alternative proposal in the third column—and that's enough to disqualify it, even if the alternative policy is not under serious consideration.

In all of our examples, the potential Pareto criterion and the efficiency criterion will make identical recommendations. It's easy to see why if you return to the pie analogy: The efficiency criterion says that we should always try to make the total “pie” of social gain as big as possible. The potential Pareto criterion says that if there's a way to make everyone's piece of pie bigger, you're not doing things right. But to say that everyone's piece could be made bigger is the same thing as saying that the pie could be made bigger—so whatever the potential Pareto criterion rejects, the efficiency criterion will reject as well.

Many economists regard the potential Pareto criterion and the efficiency criterion as good rough guides to policy choices, though few would defend them as the sole basis on which to make such decisions. Regardless of your feelings on this issue, calculations of social gains and deadweight losses can still be useful in understanding the consequences of various alternatives. If a policy causes a large deadweight loss, it is at least worth considering whether there is some good way to revise the policy so that the loss can be made smaller.

Pareto criterion

A normative criterion according to which one policy is better than another when it is preferred unanimously.

Potential Pareto criterion

A normative criterion according to which any proposal that can be unanimously defeated—even by a candidate not under consideration—should be rejected.

8.3 Examples and Applications

The machinery of consumers' and producers' surpluses is widely applicable, as the following sequence of examples will illustrate. All of them use just one basic procedure, which is summarized in Exhibit 8.11.

Subsidies

Suppose that the government institutes a new program whereby buyers of home insulation receive a rebate of \$50 for every unit of insulation they purchase. This has the effect of shifting the demand curve upward a vertical distance \$50, from D to D' in Exhibit 8.12.

With the subsidy, the quantity sold is Q' , at a market price of P_s . This is the price suppliers receive for insulation. However, consumers actually pay less, because they receive a payment of \$50 from the government, so that the consumer's actual cost is $P_s - \$50 = P_d$.

To calculate consumers' and producers' surpluses before the subsidy, we use the equilibrium price and quantity. This is shown in the first column of the table in Exhibit 8.12.

After the subsidy, consumers purchase quantity Q' at a price to them of P_d . Their consumers' surplus is the area under the *original* demand curve D out to this quantity and down to this price. We use the original demand curve because it is this curve that represents the true marginal value of insulation to consumers. The intrinsic value of home insulation is not changed by the subsidy. Therefore, the consumers' surplus is the area $A + C + F + G$, as recorded in the second column of the table.

To calculate producers' surplus, we use the quantity Q' and the producers' price P_s . This yields the area $C + D + F + H$, which is also recorded in the table.

We are still not finished. The subsidy being paid to consumers must come from somewhere, presumably from tax revenues. This represents a cost to taxpayers equal to the number of units of insulation sold times \$50 per unit. Geometrically, this is

EXHIBIT 8.11

Calculating the Consumers' and Producers' Surpluses

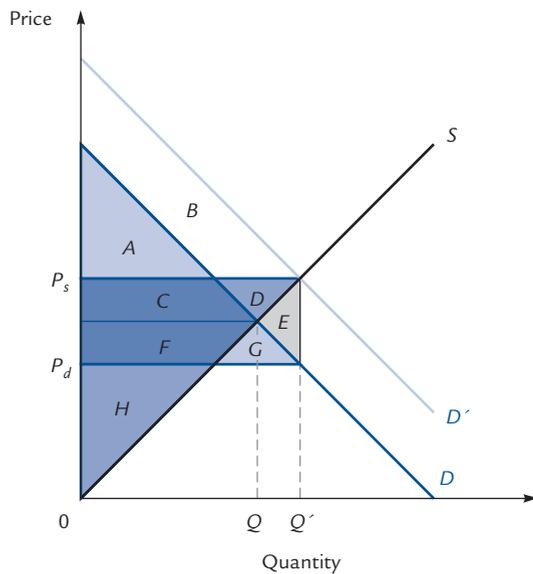
You will often be asked to calculate the effects of governmental policies on consumers' and producers' surpluses. Here are some rules to help you:

1. Begin by drawing a supply and demand diagram showing equilibrium both before and after the policy is imposed. Draw horizontal and vertical lines from the interesting points in your diagram to the axes. After a while you will get a feel for which lines to draw and which to omit. It never hurts to draw more than you need.
2. Before you proceed, label every area that is even possibly relevant.
3. When calculating consumers' surplus, use only the demand curve and prices and quantities that are relevant to the consumer. When calculating producers' surplus, use only the supply curve and prices and quantities relevant to the producer.
4. Remember that the demand and supply curves are relevant only because they are equal to the marginal value and marginal cost curves. If for some reason the demand curve should separate from the marginal value curve, continue to use the marginal value for calculating consumers' surplus. Do likewise if the supply curve should separate from the marginal cost curve.
5. Check your work with a picture like Exhibit 8.9: Calculate the social gain directly by drawing rectangles of "welfare gains" for each item actually produced and by summing the areas of these rectangles. The sum should equal the total of the gains to all of the individuals involved.

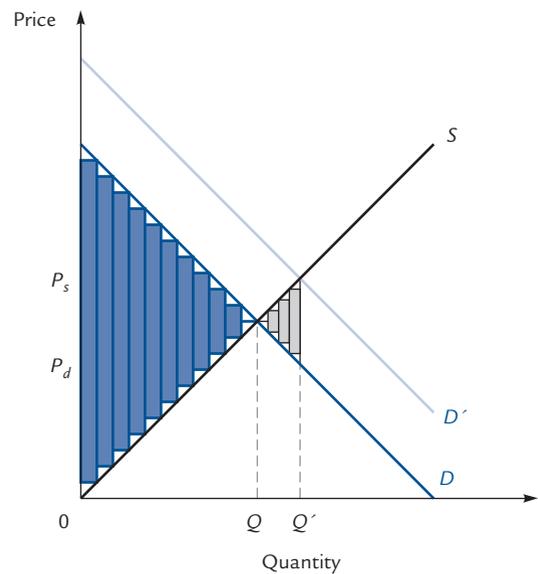
EXHIBIT 8.12

The Effect of a Subsidy

	Before Subsidy	After Subsidy
Consumers' Surplus	$A + C$	$A + C + F + G$
Producer's Surplus	$F + H$	$C + D + F + H$
Cost to Taxpayers	—	$-(C + D + E + F + G)$
Social Gain	$A + C + F + H$	$A + C + F + H - E$
Deadweight Loss	—	E



A



B

The table shows the gains to consumers and producers before and after the institution of a \$50-per-unit government subsidy to home insulation. With the subsidy in effect, there is a cost to taxpayers that must be subtracted when we calculate the social gain. We find that the social gain with the subsidy is lower by E than the social gain without the subsidy. E is the deadweight loss.

To check our work, we can consider the social gain created by each individual unit of insulation, shown in panel B. Each unit up to the equilibrium quantity Q creates a rectangle of social gain. After Q units have been produced, we enter a region where marginal cost exceeds marginal value. Each unit produced in this region creates a social loss equal to the excess of marginal cost over marginal value; these losses are represented by the gray rectangles, which stop at the quantity Q' that is actually produced. The social gain is equal to the sum of the colored rectangles minus the sum of the gray ones. Because the social gain without the subsidy is just the sum of the colored rectangles, the gray rectangles represent the deadweight loss.

represented by the rectangle $C + D + E + F + G$. This cost is a loss to the taxpayers and so must be subtracted in the computation of social gain. The deadweight loss of E is the difference between social gain before and after the subsidy.

According to the efficiency criterion, the subsidy should be rejected: It gathers $F + G$ votes in favor from consumers and $C + D$ votes in favor from producers, but

$C + D + E + F + G$ votes opposed from taxpayers. Thus, it loses by a margin of E , which (noncoincidentally) is the deadweight loss.

Exercise 8.5 Verify the calculation of social gain in Exhibit 8.12.



Dangerous
Curve

Students often want to know how areas C and F can be part of both the consumers' surplus and the producers' surplus. The answer is that surplus is not an area at all—the area is just a *measure* of surplus. The fact that you have 12 yards of carpet and your friend has 12 yards of carpet does not mean that you both own the same “yards,” only that each of you owns carpeting that can be measured by the same yardstick. The areas of surplus are yardsticks with which we measure different individuals' gains from trade.

Panel B in Exhibit 8.12 provides a way to check our work. The colored rectangles to the left of equilibrium represent gains to social welfare just as in Exhibit 8.9. In this case, however, *more* than the equilibrium quantity is produced. Consider the first item produced after equilibrium. The marginal value of this item to consumers (read off the original demand curve) is *less* than the marginal cost of producing it. The difference between the two is the area of the first gray rectangle. This area therefore represents a net welfare *loss* to society. Similarly, the next item produced represents a welfare loss in the amount of the area of the second gray rectangle, and so on out to the quantity Q' . The total welfare loss is the sum of these rectangles, which is equal to the area E in panel A. Therefore, area E should be the deadweight loss, and the calculation in the table is confirmed.

An alternative way to calculate the consumers' surplus is shown in Exhibit 8.13. For most purposes, it suffices to use either the method of Exhibit 8.12 or that of Exhibit 8.13. Because both always lead to the same answers, you need to master only one of them. However, there will be a few occasions later on in this book where you will find it much easier to use the alternative method of Exhibit 8.13.

Price Ceilings

Price ceiling

A maximum price at which a product can be legally sold.

Effective price ceiling

A price ceiling set below the equilibrium price.

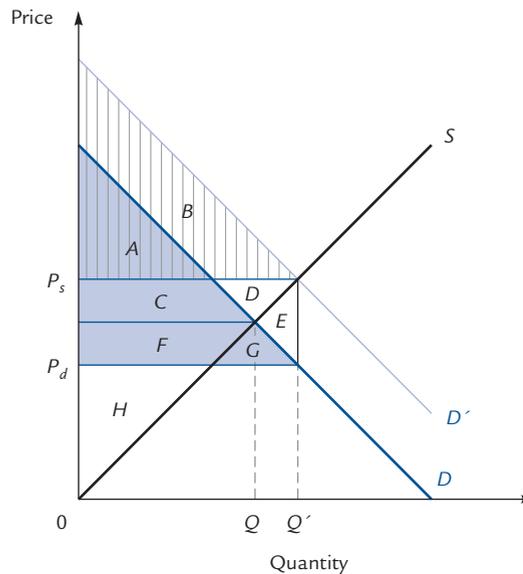
A **price ceiling** is a legally mandated maximum price at which a good may be sold. The effect of a price ceiling depends on its level. If the legal maximum is above the equilibrium price that prevails anyway, then the price ceiling has no effect (a law forbidding any piece of bubble gum to sell for more than \$2,000 will not change anyone's behavior). An **effective price ceiling** is one set below the equilibrium price, like the price P_0 in Exhibit 8.14.

At the price P_0 , producers want to sell the quantity Q_s and consumers want to buy the quantity Q_d . What quantity actually gets traded? The answer is Q_s , because as soon as Q_s units are sold, the sellers pack up and go home. When buyers and sellers disagree about quantity, the group wanting to trade fewer items always wins, because trading stops as soon as either party loses interest.

Another, and very real, possibility must be considered: Because buyers are frustrated, they will be willing to offer prices higher than P_0 , and sellers may accept these prices in violation of the law. For purposes of our simple analysis, we will assume that the law is perfectly enforced and this does not occur. We will also assume that the

EXHIBIT 8.13

Another Way to Do It



The method that we have been using to calculate social gain is adequate for most problems. However, there is an alternate method that may occasionally be more convenient.

We illustrate with the “subsidy” example from Exhibit 8.12. The graph here is identical to panel A in that exhibit. There are two ways of viewing a \$50 rebate for home insulation. The first way, which we adopted in Exhibit 8.12, is to say that the rebate does not alter the value of home insulation. Consumers benefit from the subsidy by being able to buy insulation at a lower price. This is why we use the *old* demand curve (the true marginal value curve) and the price paid by consumers (P_d) as boundaries for the area of consumer surplus. This gives the shaded area $A + C + F + G$.

An alternative and *equally valid* point of view is to say that a subsidy is like a \$50 bill taped to each unit of insulation. This raises the marginal value of the insulation by \$50. However, we now have to view the insulation as being purchased at the market price P_s . If we said that the insulation has increased in value *and* that the consumer is paying less than market price for his insulation, we would be wrongly double-counting the \$50 rebate.

From the alternative point of view, the consumer surplus is the area under the *new* demand curve down to the market price P_s and out to the quantity Q' . That is, the striped area $A + B$.

If both points of view are equally valid, how can they give different answers? The answer is: They don't. In fact, area $A + B$ is equal to area $A + C + F + G$. They have to be equal, because each represents the consumers' surplus calculated correctly, and there can be only one consumers' surplus. If you find that argument unconvincing, try proving directly that the two areas are equal. This is an exercise in high school geometry if you assume all curves are straight lines; it is an exercise in calculus otherwise.

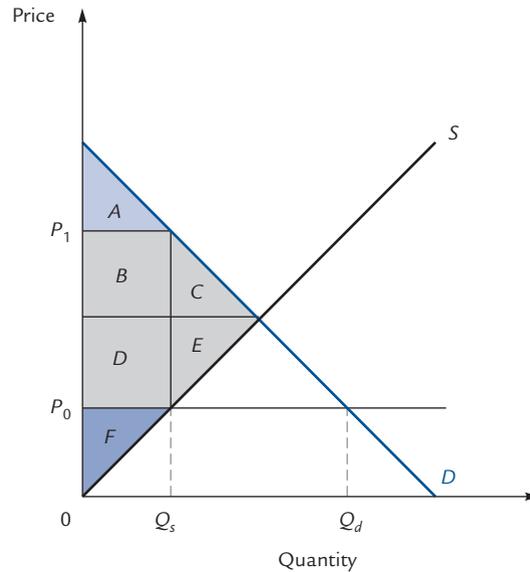
enforcement is costless (otherwise, the cost of enforcement would have to be subtracted from social gain).⁶

The quantity sold is Q_s . What price do consumers pay? You may think the answer is obviously P_0 , but this is incorrect. At a price of P_0 , consumers want to buy more goods than are available. Therefore, *they compete with each other* to acquire the limited supply.

⁶ Here is an interesting puzzle. Why is it that in “victimless crimes” like prostitution and the sale of drugs, both parties are held criminally liable, whereas in the equally “victimless” crime of violating a price control, only the seller faces legal consequences? For an interesting discussion of this puzzle, see J. Lott and R. Roberts, “Why Comply: One-Sided Enforcement of Price Controls and Victimless Crime Laws,” *Journal of Legal Studies* 18 (1989).

EXHIBIT 8.14

A Price Ceiling



	Before Ceiling	After Ceiling
Consumers' Surplus	$A + B + C$	A
Producers' Surplus	$D + E + F$	F
Social Gain	$A + B + C + D + E + F$	$A + F$
Deadweight Loss	—	$B + C + D + E$

At a maximum legal price of P_0 , demanders want to buy more than suppliers want to sell. Therefore, they compete against each other for the available supply, by waiting in line, advertising, and so forth. This increases the actual price to consumers. The full price to consumers must be bid all the way up to P_1 , because at any lower price the quantity demanded still exceeds the quantity supplied, leading to increases in the lengths of waiting lines.

The deadweight loss comes about for two reasons. First, there is the reduction in quantity from equilibrium to Q_s . This loss is the area $C + E$. Second, there is the value of the consumers' time spent waiting in line. This is equal to $P_1 - P_0$ times the quantity of items purchased, which is the rectangle $B + D$.

Depending on the nature of the good, this may take the form of standing in line, searching from store to store, advertising, or any of a number of other possibilities. All of these activities are costly, in time, gasoline, energy, and other currency, and these costs must be added to the "price" that consumers actually pay for the item.

How high does the price go? It must go to exactly P_1 in Exhibit 8.14. At any lower price the quantity demanded still exceeds Q_s , and consumers intensify their efforts. Only when the "price" reaches P_1 does the market equilibrate.

Of course, even though P_1 is the price paid by consumers, the price received by suppliers is still P_0 . Therefore, we use P_1 to calculate consumers' surplus and P_0 to calculate producers' surplus. In each case, the appropriate quantity is Q_s , the quantity actually traded. The computations are shown in Exhibit 8.14.

Exercise 8.6 Verify the correctness of the table in Exhibit 8.14.

The deadweight loss calculated in Exhibit 8.14 comes about for two reasons. First, there is the reduction in quantity from Q to Q_s , which leads to a social loss of $C + E$, just as in the case of a tax. However, now there is another sort of loss as well. The value of the time people spend waiting in lines is equal to the value of the time-per-unit-purchased ($P_1 - P_0$) times the quantity of units purchased (Q_s), which is the rectangle $B + D$. Taken together, these effects account for the entire deadweight loss.

Notice that from a social point of view there is a great difference between a *price control* that drives the demanders' price up to P_1 and a *tax* that drives the demanders' price up to P_1 . Because the revenue from a tax is wealth transferred from one individual to another, it is neither a gain nor a loss to society as a whole. But the value of the time spent waiting in lines is wealth lost and never recovered by anyone.

Some of the deadweight loss can be avoided if there is a class of people whose time is relatively inexpensive. Those people will offer their services as "searchers" or "line-standers" and consumers will pay them up to $P_1 - P_0$ per item for their services. The income to the line-standers, minus the value of their time, is a gain that offsets part of the lost area $B + D$.

Of course, some consumers whose time has low value might stand in line to make their own purchases. We view these consumers as having purchased line-standing services from themselves at the going price of $P_1 - P_0$. Such a consumer earns part of area A as a consumer and part of area $B + D$ as a line-stander.



Dangerous
Curve

The reduction in deadweight loss through the use of line-standers doesn't work if too many people have low time values. In that case, all of those people attempt to become line-standers and the lines get longer, so that the value of the time each one spends waiting gets bid back up to $P_1 - P_0$.



Dangerous
Curve

Tariffs

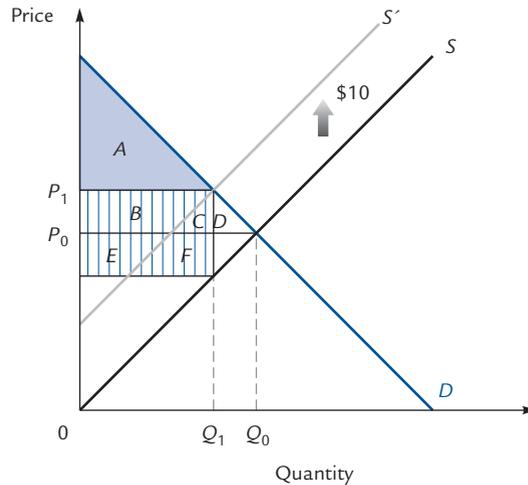
Suppose that Americans buy all of their cameras from Japanese companies. It is proposed that a tariff of \$10 per camera be imposed on all such imports and that the proceeds be distributed to Americans chosen at random. What areas must we measure to see whether the tariff makes Americans as a whole better off?

Exhibit 8.15 shows the market for cameras, with both the original and post-tariff supply curves. The table shows the gains to Americans before and after the tariff. These gains are calculated using the pretariff price and quantity of P_0 and Q_0 and the posttariff price and quantity of P_1 and Q_1 . Notice that we do not include the producers' surplus, because this is earned by the Japanese companies and the question asks only about the welfare of Americans. If we had been asked about the welfare of the entire world, we would have included producers' surplus in our calculations.

Exercise 8.7 Calculate the social gains to the entire world before and after the tariff is imposed.

EXHIBIT 8.15

A Tax on Imported Cameras



	Before Tariff	After Tariff
Consumers' Surplus	$A + B + C + D$	A
Tariff Revenue	—	$B + C + E + F$
Social Gain	$A + B + C + D$	$A + B + C + E + F$

If cameras are supplied by foreigners and purchased by Americans, then a tariff affects Americans through the consumers' surplus and through the tax revenue that it generates.

Now we return to the question: What areas must we measure? The answer is evidently that one must compare area D with area $E + F$. If $E + F$ is bigger, the tariff improves the welfare of Americans; otherwise it reduces their welfare.

In practice, these areas can be estimated if the supply and demand curves can be estimated, and, as we remarked in Chapter 1, there are econometric methods available for this. Therefore, an economist can contribute meaningfully to a debate about tariffs by computing the relevant areas and reporting which policy is better—*provided* that the goal is to maximize Americans' welfare.

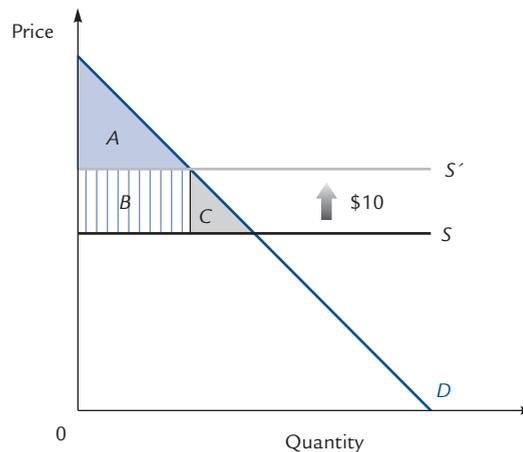
It is often a reasonable assumption that a country faces flat supply curves for imported items. The reason for this is that Japanese firms sell cameras in many foreign countries, and the United States is only a small part of their market. Thus, changes in quantity that appear big (from our point of view) may in fact correspond only to very small movements along the Japanese supply curves and hence to small changes in price. Exhibit 8.16 shows the analysis of a tariff when the supply curve is flat. In this case, you can see that the tariff always reduces Americans' welfare.

Tariffs and Domestic Industries

A more interesting example involves tariffs on a product that is produced both domestically and abroad. Suppose that Americans buy cars from Japan subject to a flat Japanese supply curve at a price P_0 , and that domestic car manufacturers have

EXHIBIT 8.16

A Tariff on Imported Cameras That Are Elastically Supplied



	Before Tariff	After Tariff
Consumers' Surplus	$A + B + C$	A
Tariff Revenue	—	B
Social Gain	$A + B + C$	$A + B$
Deadweight Loss	—	C

If the United States is a small part of the market to which the Japanese sell cameras, then Americans will face a flat supply curve. In this case, a tariff always reduces the welfare of Americans.

the upward-sloping supply curve shown in panel A of Exhibit 8.17. Assuming that all cars are identical, no consumer will be willing to pay more than P_0 for a domestic car, because the consumer can always buy an import instead. Therefore, all cars sell at a price of P_0 . At this price, domestic manufacturers produce Q_0 cars and domestic consumers buy Q_1 . The difference, $Q_1 - Q_0$, is the number of imports. Table A in Exhibit 8.17 shows the consumers' and producers' surpluses.

Now suppose that we impose a tariff of \$500 on each imported car. This raises the foreign supply curve \$500 to a level of $P_0 + \$500$. The price of cars goes up to $P_0 + \$500$, the quantity supplied domestically goes up to Q_0' (in panel B of Exhibit 8.17), and the quantity demanded falls to Q_1' . The quantity imported falls to $Q_1 - Q_0'$.

In Exhibit 8.17 Table B shows the consumers' and producers' surpluses both before and after the tariff. (The "before" column, of course, simply repeats the calculation from Table A.) What about revenue from the tariff? The number of imported cars is $Q_1' - Q_0'$, and the tariff is \$500 on each of these. Thus, the tariff revenue (which ends up in American pockets) is $(Q_1' - Q_0') \times \$500$, and this is the area of rectangle I . This is recorded in Table B, along with a comparison of social gains.

We can see that even when there is a domestic industry that benefits from the tariff, and even though the tariff revenue is a gain to the country, tariffs still cause a deadweight loss (we say that they are *inefficient*) because consumers lose more than all other groups gain.

EXHIBIT 8.17

A Tariff When There Is a Domestic Industry

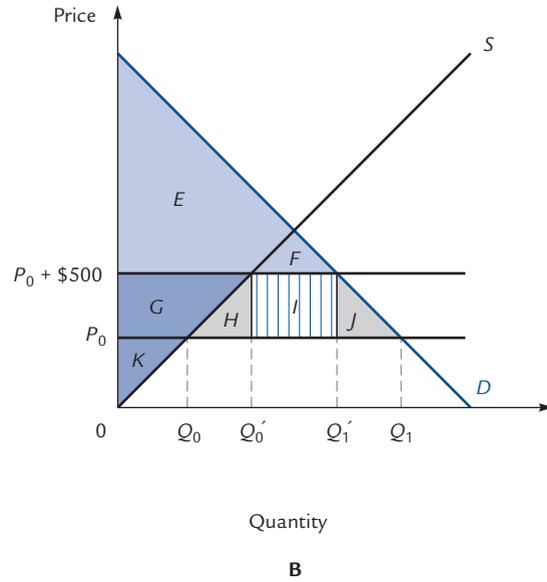
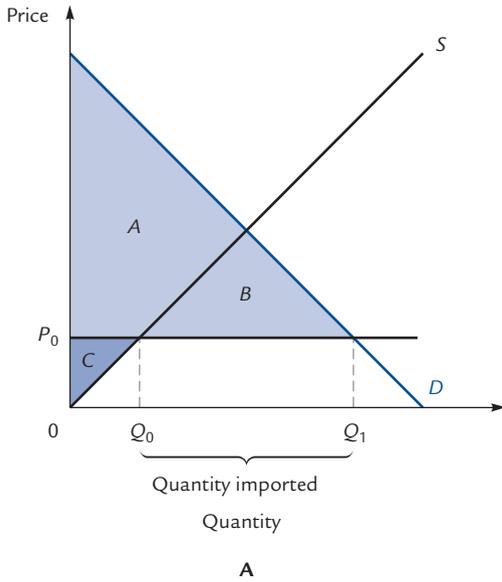


Table A

Consumers' Surplus	$A + B$
Producers' Surplus	C
Social Gain	$A + B + C$

Table B

	Before Tariff	After Tariff
Consumers' Surplus	$E + F + G + H + I + J$	$E + F$
Producers' Surplus	K	$G + K$
Tax Revenue	—	I
Social Gain	$E + F + G + H + I + J + K$	$E + F + G + K + I$
Deadweight Loss	—	$H + J$

We assume that Americans can buy any number of cars from Japan at the price P_0 . The supply curve S shows how many cars American manufacturers will provide at each price. At the price P_0 , American producers supply Q_0 cars and American consumers purchase Q_1 cars. The difference, $Q_1 - Q_0$, is the number of cars imported. Table A shows the gains to Americans.

In panel B we see the effect of a \$500 tariff on imported cars. The price of a foreign car rises to $P_0 + \$500$, and the number of imports falls to $Q'_1 - Q'_0$. Table B compares gains before and after the tariff. Note that the first column of Table B is identical to Table A except that it uses the labels from panel B rather than panel A. The tariff revenue is computed by observing that the area of rectangle I is $(Q'_1 - Q'_0) \times \$500$.

Exercise 8.8 Suppose that the government wants to benefit domestic auto producers and the recipients of tax revenue at the expense of car buyers. Devise an efficient (though perhaps impractical) way of doing this that makes everybody happier than a tariff does.

Robbery

From the point of view of economic efficiency (i.e., the maximization of the total gains to all members of society), a loss to one group that is exactly offset by a gain to another group is a “wash.” To one who is interested only in maximizing social gain,

such a transfer is neither a good thing nor a bad thing. How, then, should such a one feel about *robbery*?

Many people think that robbery constitutes a social loss equal to the value of what is stolen. Their reasoning is simple but faulty: They notice the loss to the victim without noticing the offsetting gain to the robber. A more sophisticated answer would be that robbery is a matter of indifference, because stolen goods do not disappear from society; they only change ownership.

However, this more sophisticated answer is also wrong. There *is* a social cost to robbery. It is the opportunity cost of the robber's time and energy. The robber who steals your bicycle could, perhaps, with the same expenditure of energy, be building a bicycle of his own. If he did, society would have two bicycles; when he steals yours instead, society has only one. The option to steal costs society a bicycle.

This shows that robbery is socially costly; we still have to ask: *How costly?* To answer this, it is reasonable to treat robbery as a competitive industry: Robbers continue to rob until the marginal cost (in time, energy, and so on) of committing an additional crime is equal to the marginal revenue (in loot). The cost is what interests us, the loot is observable, and we know that the two are equal. So, at the margin, we can reckon the cost of a robbery as approximately equal to the value of what is stolen.

This tells us that the amount stolen is a correct measure of the cost of the last robbery committed. In Exhibit 8.18 we calculate the total social cost of all robberies. Suppose that a robber can expect to earn $\$R$ each time he commits a robbery. Then robbers steal until the marginal cost of stealing is equal to $\$R$; that is, they commit Q robberies. The amount stolen is $\$R \times Q$, the area $A + B$. However, the robbers' total costs are given by the area under the supply curve, A . This cost to the robbers is society's cost as well. Therefore, the total social cost of all robberies (A) is less than the value of what is taken ($A + B$).

This analysis ignores the very real possibility that people will take costly steps to protect themselves from robbery—installing burglar alarms, deadbolt locks, and the like. These additional costs are also due to the existence of robbery and must be added to area A in order to calculate the full social cost of robberies.



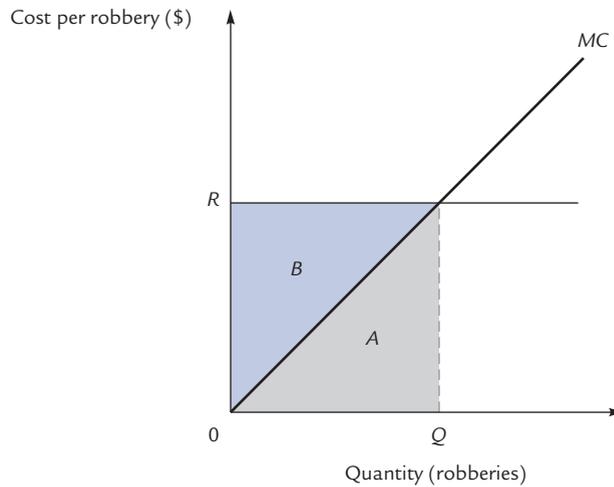
Dangerous
Curve

The more general lesson of this example is that effort expended in nonproductive activity is social loss. Accountants devising new methods of tax avoidance, lawyers in litigation, lobbyists seeking laws to transfer wealth to their clients, and all of the resources that they employ (secretaries, file clerks, photocopy machines, telephone services, and so on) are often unproductive from a social point of view. Whatever they win for their clients is a loss for their adversaries. In the absence of this activity, all of these resources could be employed elsewhere, making society richer.

On the other hand, some of this seemingly unproductive activity serves hidden and valuable purposes. Suppose that a law is passed requiring that all owners of apple orchards donate \$5,000 each to the president's brother. The owners of apple orchards might hire a lobbyist to assist them in having this law overturned. If the effort is successful, apple growers win only what the president's brother loses, and so at one level of analysis the lobbyist's time contributes nothing to the welfare of society. On the other hand, if all orange growers were made very nervous by this law and planned to burn down their orange trees as a precaution against their being next, then the lobbyist saves a lot of valuable orange trees through his efforts. Insofar as redistributing income affects the incentives to engage in productive activities, it can indirectly affect society's welfare.

EXHIBIT 8.18

The Social Cost of Robbery



Producers' Surplus (Earned by Robber)	B
Loss to Victim	$A + B$
Social Loss	A

We suppose that a robber can expect to earn $\$R$ for each robbery he commits. Then robberies will take place until the robbers' marginal cost (the opportunity cost of their time, energy, and so on) equals $\$R$. The number of robberies committed is Q , and robbers earn a producers' surplus of B . However, victims lose the amount stolen, which is $A + B$. There is a net social loss of A . If society pursues economic efficiency, A is the maximum amount it would be willing to spend to prevent all robberies.

Theories of Value

We have defined value in terms of consumers' willingness to pay, and we have discovered that the price of an item is equal to its marginal value. Other theories of value have arisen in the history of economics, only to be abandoned when careful analysis revealed them as erroneous. Because such errors are still common in much discussion by noneconomists, it is worth examining them to see why they should be avoided.

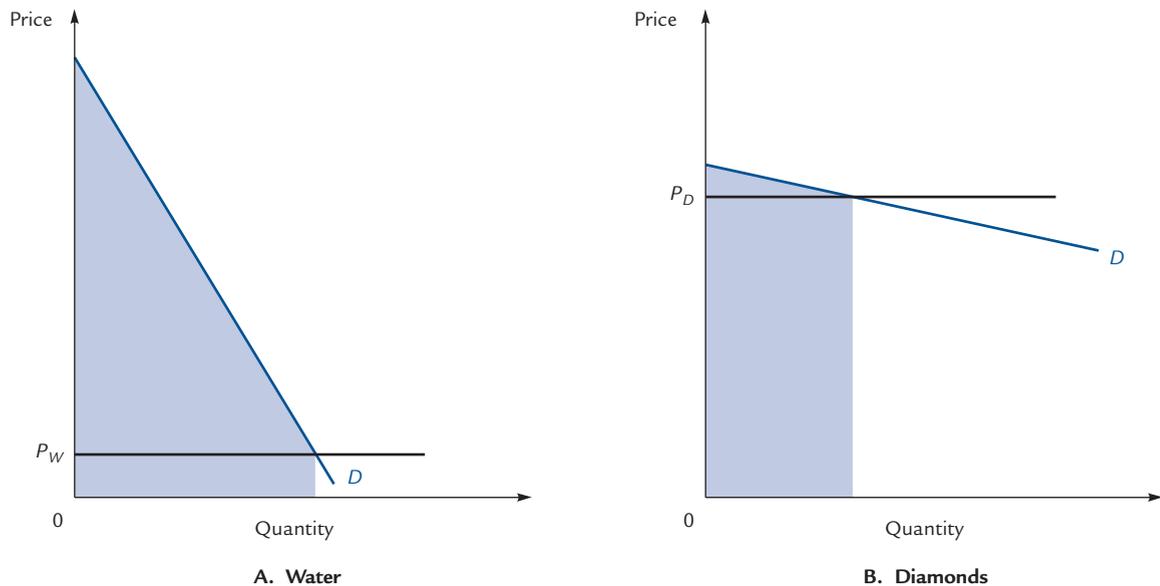
The Diamond-Water Paradox

Many classical economists were puzzled by the so-called diamond-water paradox. How can it be that water, which is essential for life and therefore as "valuable" a thing as can be imagined, is so inexpensive relative to diamonds, which are used primarily for decoration and the production of nonessential goods? If price reflects value, shouldn't a gallon of water be worth innumerable diamonds?

The paradox is resolved when you realize that price reflects not *total* value, but *marginal* value. Exhibit 8.19 depicts the demand curves for water and for diamonds, together with their market prices and the corresponding consumer's surpluses. The marginal value of your first gallon of water is indeed much higher than the marginal value of your first diamond, and this is reflected by the heights of the demand curves

EXHIBIT 8.19

The Diamond–Water Paradox Resolved



If you had no water and no diamonds, you would be willing to pay far more for a first bucket of water than for a first diamond. Therefore, your marginal value (= demand) curve for water starts out much higher than your marginal value curve for diamonds. At the market prices P_W and P_D , you consume Q_W buckets of water and Q_D diamonds, so that the marginal value of a bucket of water (P_W) is much less than the marginal value of a diamond (P_D).

It is true that the total value of all your water (the shaded area in panel A) is greater than the total value of all your diamonds (the shaded area in panel B). The graphs show that this is perfectly consistent with a low marginal value for water and a high marginal value for diamonds. The price, which is equal to the marginal value, should not be expected to reflect the total value.

at low quantities. But this has nothing to do with the *price* of water; the price is equal to the marginal value of the last bucket consumed, and this may be very low if you consume many gallons.

Notice that the total value (the colored area) in the market for water is much higher than in the market for diamonds: If you lost all of your water and all of your diamonds, you would be willing to pay more to retrieve the water than to retrieve the diamonds. In consequence, the consumers' surplus is much higher in the market for water than in the market for diamonds. Exhibit 8.19 shows that there is nothing paradoxical about a low price and a large consumers' surplus existing simultaneously.

The Labor Theory of Value

The **labor theory of value** is an error that deceived such diverse economists as Adam Smith and Karl Marx. In its simplest form it says that the price of an item is determined by the amount of labor used in its production.⁷ In this form it is clearly false: You can

Labor theory of value

The assertion that the value of an object is determined by the amount of labor involved in its production.

⁷ Of course, this is a simpler form than economists have ever believed; typical versions restrict attention to "socially necessary" labor and include the labor of previous generations who built machines used in current production.

expend an enormous quantity of labor digging a gigantic hole in your backyard, and the price that hole commands in the marketplace may be far less than the price of a short story produced by a good writer in an afternoon, sitting at a word processor in an air-conditioned house sipping lemonade.⁸

For a theory so evidently false, the labor theory of value (even in this simple form) is remarkably pervasive. You will hear it argued that doctors “ought to” earn high salaries because of all the effort involved in earning their medical degrees, or that people in occupation A “ought to” earn as much as people in occupation B because they work equally hard. Such arguments ignore the fact that value is determined not by the cost of inputs, but by demand—the consumer’s willingness to pay for the good or service being offered.

Another common belief that embodies the labor theory fallacy is that a meaning can be attached to the “book value” of a firm. A firm’s “book value” is a measure of what it would cost to produce the actual physical assets of the firm. It is computed, for example, by adding up the cost of the bricks used to build the firm’s plants and office buildings, the desks and chairs in the executive offices, the machines along the assembly line, and the letterhead stationery in the cabinets. This book value can be compared to the actual price at which one could acquire the entire firm (say, by purchasing all its stock). It sometimes happens that a firm can be acquired for less than book value, and it is widely believed that this represents a bargain.

Not so. The fact that a factory is built from \$1 million worth of bricks does not make that factory worth \$1 million, any more than your application of \$1 million worth of labor would make a hole in your backyard worth \$1 million. If your labor is devoted to the production of something that nobody wants, or if the bricks are glued together to form a factory that produces nothing useful, this will be reflected in the price. What we have here is a *brick theory of value*, different perhaps from the labor theory of value, but perfectly analogous and just as false.

A final example illustrates both the diamond–water and the labor theory paradoxes. It is sometimes argued that something must be wrong with society’s values when a baseball player (for example) earns a seven-figure salary for playing a game that (1) he enjoys anyway and (2) produces little social value compared with something like teaching elementary school, which is far less lucrative. The first point is the labor theory of value again. It errs by assuming that how hard the baseball player works determines the value of what he produces. The second point uses the erroneous reasoning that underlies the diamond–water paradox. It may very well be that teachers (like water) produce far more total social value than star baseball players. But it can be simultaneously true that *one additional teacher* produces less social value than one additional star baseball player. This can be the case, for example, if there are many teachers and few star baseball players. We should not expect the price of a teacher or a baseball player to tell us anything about the total value to society of the two professions.

8.4 General Equilibrium and the Invisible Hand

Based on the examples in Section 8.3, you might have begun to suspect that any deviation from competitive equilibrium leads to a reduction in social gain. In this section, we will see that this is, in fact, the case.

⁸ It is true, of course, that in a competitive market, price equals marginal cost (and a competitive producer will not choose to dig a hole in his backyard for sale in the marketplace). But marginal cost is not labor cost. Some labor costs may be sunk (and therefore irrelevant), and many relevant costs have nothing to do with labor. The relevant costs, as always, are the opportunity costs—the writer could be writing a movie script instead.

The Fundamental Theorem of Welfare Economics

Exhibit 8.20 shows the competitive market for potatoes. We can ask two questions about this market ostensibly as different as questions can be:

1. What is the quantity of potatoes actually produced and sold?
2. Suppose you were a benevolent dictator, concerned only with maximizing the total welfare gains to all of society. What quantity of potatoes would you *order* produced and sold?

Note well the dissimilarity between these questions. One is a question about what *is*; the other is a question about what *ought* to be.

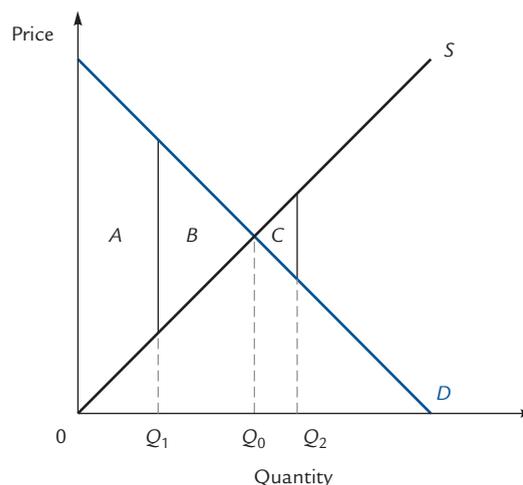
We know the answers to each of these questions. They are:

1. The quantity of potatoes produced and sold is at Q_0 , where supply equals demand. We have seen that individual suppliers and demanders, seeking to maximize their own profits and their own happiness, choose to operate at this point.
2. To maximize social gains, you would continue ordering potatoes to be produced as long as their marginal value exceeds the marginal cost of producing them. You would stop when marginal cost equals marginal value, that is, at Q_0 .

The choice of Q_0 yields a social gain of $A + B$ in Exhibit 8.20. A despot who made the mistake of ordering only Q_1 potatoes produced would limit social gain to area A . If the same benevolent dictator made the mistake of ordering Q_2 , area C would be subtracted from the social gain, because it is made up of rectangles whose areas represent an excess of marginal cost over marginal value.

EXHIBIT 8.20

The Invisible Hand



Under competition, the quantity produced is Q_0 , where supply equals demand. A benevolent dictator who wanted to maximize social gain would employ the equimarginal principle and order potatoes to be produced to a quantity where marginal cost equals marginal value. This also occurs at quantity Q_0 .

If the dictator ordered Q_1 potatoes produced, social gain would be area A ; if he ordered Q_2 , social gain would be $A + B - C$. The maximum social gain, at Q_0 , is $A + B$.

It is astounding that the two questions have identical answers. The coincidence results from the prior coincidences of the supply curve with the marginal cost curve, and of the demand curve with the marginal value curve.

It is not only astounding that the two answers are identical but it is fortunate. It means that people living in a competitive world achieve the maximum possible social gain without any need of a benevolent despot. The market alone achieves an outcome that is economically efficient. To say the same thing in different words, competitive equilibrium is Pareto-optimal.

The eighteenth-century economist Adam Smith was so struck by this observation that he described it with one of the world's most enduring metaphors. Of the individual participant in the marketplace, he said: "He intends only his own gain, and he is . . . led by an invisible hand to promote an end which was no part of his intention."⁹

Noneconomists frequently misunderstand what Smith meant by the *invisible hand*. Some think it is a metaphor for an ideology or a philosophical point of view; the notion has even been described as a theological one! In fact, the *invisible hand* expresses what is at bottom a mathematical truth. The point of equilibrium (where competitive suppliers operate "intending only their own gain") is also the point of maximum social gain (an end that is no part of any individual participant's intention).

The Idea of a General Equilibrium

The preceding analysis is striking, but it is incomplete. By participating in the potato market, people change conditions in other markets as well. When he grows more or fewer potatoes, a farmer consequently grows less or more of something else. The amount of labor that he hires changes. When a consumer changes his potato consumption, he probably also changes his consumption of rice, and of butter. At one further remove, any change in the potato market affects the potato farmer's income, which affects his purchases of shoes, which affects the market for leather, which affects the market for something else, ad infinitum. If we really want to understand the welfare consequences of competitive equilibrium in the potato industry, we need to consider its effects in all of these other markets as well. Could it be that by maximizing welfare gains in one market, we are imposing a net welfare *loss* in the totality of all other markets?

It was not until the 1950s, nearly 200 years after Adam Smith, that economists developed the mathematical tools necessary to deal fully with this complicated question. In that decade, economists such as Kenneth Arrow, Gerard Debreu, and Lionel McKenzie devised techniques that make it possible to study all the markets in the economy at one time. In this they were advancing a subject called **general equilibrium analysis**, first invented by the nineteenth-century economist Léon Walras. One of the great and powerful results of general equilibrium theory is that even in view of the effects of all markets on all other markets, competitive equilibrium is still Pareto-optimal. This discovery is usually called the *first fundamental theorem of welfare economics*, or the *invisible hand theorem*.

The invisible hand theorem says, in essence, that in competitive markets, people who selfishly pursue their own interests end up achieving an outcome that is socially desirable. Outside of competitive markets, such good fortune is not to be expected. The governor of Colorado recently told of walking down a suburban street where each homeowner was out blowing leaves onto his neighbor's lawn. Each homeowner acted

General equilibrium analysis

A way of modeling the economy so as to take account of all markets at once and of all the interactions among them.

⁹ From Book 4 of Smith's monumental work *The Wealth of Nations*, first published in 1776.

selfishly, and the outcome was highly undesirable. If the homeowners had all agreed to spend the afternoon watching football, they would have enjoyed themselves more and had the same number of leaves on their lawns at the end of the day. Because the decision to blow leaves takes place outside of the market system, there is no reason to expect it to yield outcomes that are in any sense desirable. In Chapters 10 through 14 we will see many more such examples. The fact that the invisible hand theorem fails so easily in so many contexts makes it utterly remarkable that it succeeds in the particular context of competitive markets.

The Pareto optimality of competitive equilibrium is a deep and wondrous fact about the price system. No analogous statement is true in the absence of competition or in the absence of prices. The invisible hand theorem is a remarkable truth.

An Edgeworth Box Economy

The invisible hand theorem is true in very complex economies with many participants and many markets, but we will illustrate it (and the basic ideas of general equilibrium analysis) only in the simplest possible case. Assume a world with two people (Aline and Bob) and two goods (food and clothing). We will simplify further by assuming that there is no production in this world; Aline and Bob can only trade the goods that already exist. These assumptions will enable us to present a complete general equilibrium model and to illustrate the invisible hand theorem.

Because there is no production in this world, there is only a fixed, unchangeable amount of food and clothing. In panel A of Exhibit 8.21 we draw a box that has a width equal to the amount of food in existence, and a height equal to the amount of clothing. Such a box is called an **Edgeworth box**.¹⁰ Using the lower left-hand corner as the origin, we draw Aline's indifference curves between food and clothing. We also mark one point of special interest: It is Aline's **endowment point**, O , representing the basket of food and clothing that she owns at the beginning of the story.

In panel B we do a strange thing: We turn the entire page upside down, and we draw Bob's (black) indifference curves in the same box. For him, the food axis is the line that Aline views as the top of the box, and the clothing axis is the line that Aline views as the right side of the box.

To plot Bob's endowment point, remember that the width of the box is equal to the sum of Bob's and Aline's food endowments and that the height is equal to the sum of their clothing endowments. A moment's reflection should convince you that Bob's endowment point (measured along *his* axes) is the same as Aline's endowment point (measured along *her* axes).

Panel C shows a piece of panel B: All but two indifference curves have been eliminated. We have retained only those indifference curves (one of Aline's and one of Bob's) that pass through the endowment point.

Now suppose that Bob and Aline discuss the possibility of trade. Aline vetoes any trade that moves her into region A , C , or E , because these all represent moves to lower indifference curves from her point of view. Similarly, Bob vetoes any trade that moves him into region A , B , or E . (Hold the book upside down for help in seeing this!) However, a movement anywhere inside region D benefits both Aline and Bob. For this reason, region D is called the **region of mutual advantage**, and Aline and Bob can arrange a trade that moves them into this area.

Edgeworth box

A certain diagrammatic representation of an economy with two individuals, two goods, and no production.

Endowment point

The point representing the initial holdings of an individual in an Edgeworth box.

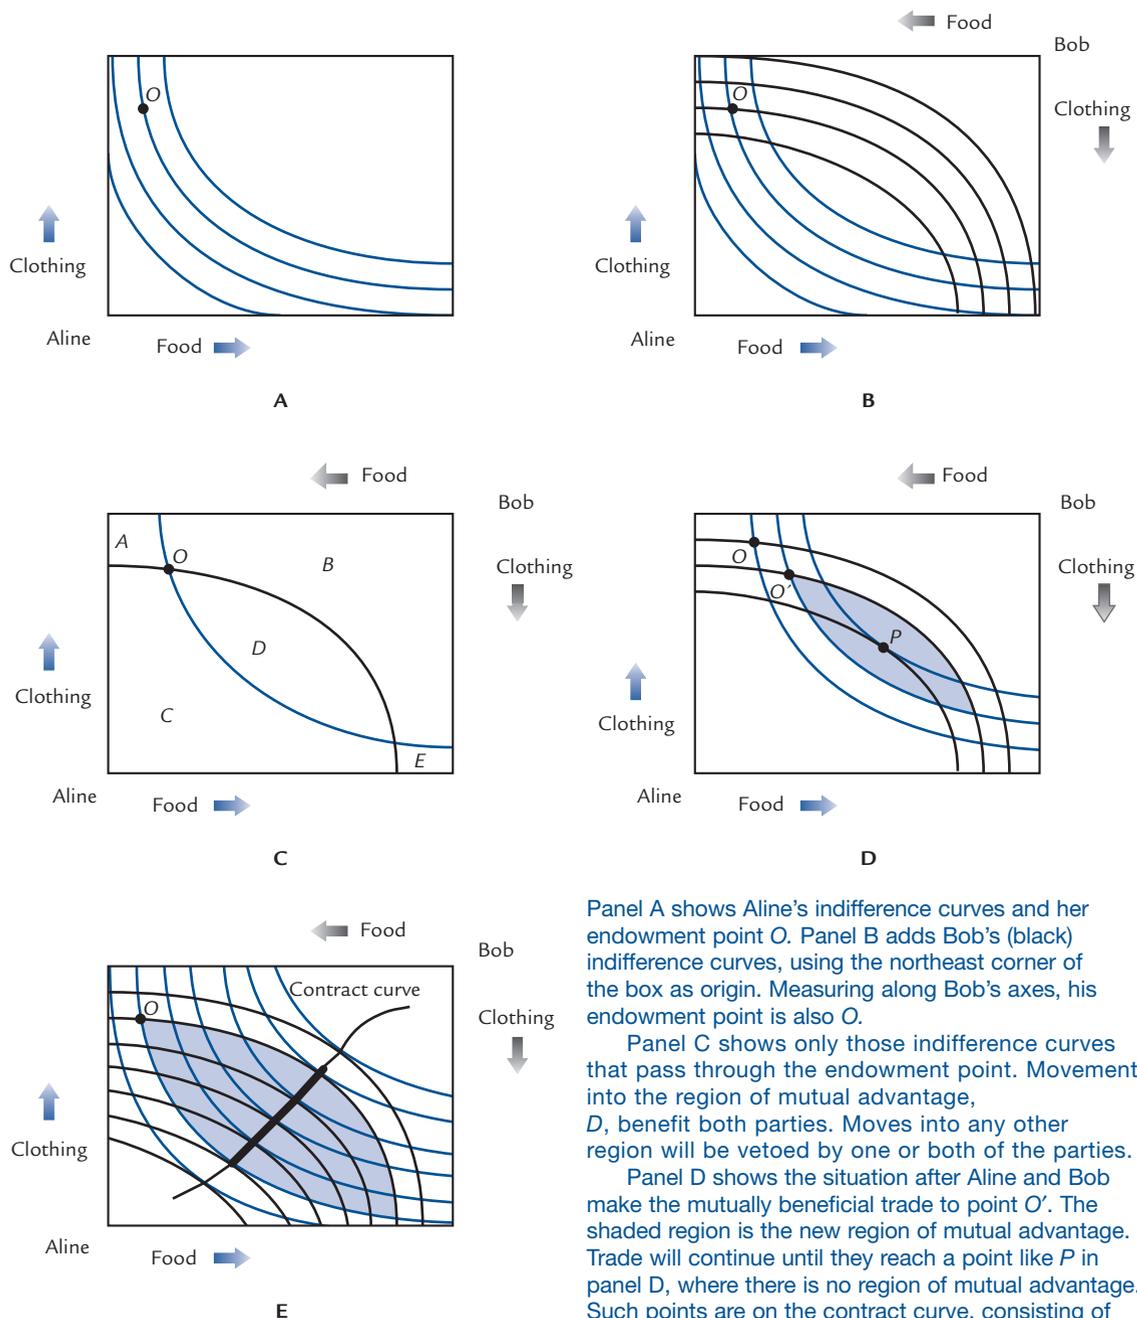
Region of mutual advantage

The set of points that are considered at least as good as the initial endowment.

¹⁰ The Edgeworth box is named after the nineteenth-century British economist F. Y. Edgeworth.

EXHIBIT 8.21

Trade in an Edgeworth Box Economy



Panel A shows Aline's indifference curves and her endowment point O . Panel B adds Bob's (black) indifference curves, using the northeast corner of the box as origin. Measuring along Bob's axes, his endowment point is also O .

Panel C shows only those indifference curves that pass through the endowment point. Movements into the region of mutual advantage, D , benefit both parties. Moves into any other region will be vetoed by one or both of the parties.

Panel D shows the situation after Aline and Bob make the mutually beneficial trade to point O' . The shaded region is the new region of mutual advantage. Trade will continue until they reach a point like P in panel D, where there is no region of mutual advantage. Such points are on the contract curve, consisting of the tangencies between Aline's and Bob's indifference curves. The points on the contract curve are precisely those that are Pareto-optimal.

In panel E the shaded region is the original region of mutual advantage. Trade leads to the choice of a point on the contract curve in this region. The darker segment of the contract curve is the set of possible outcomes.

After moving to a new point, O' , inside the region of mutual advantage, Aline and Bob face a new, smaller region of mutual advantage, as shown in panel D. They will move to a new point in this new region and will continue this process until no region of mutual advantage remains. This occurs precisely when they reach a point where their indifference curves are tangent to each other, such as the point P in panel D.

A point of tangency between Aline's and Bob's indifference curves is a point from which no further mutually beneficial trade is possible. In other words, such a point is Pareto-optimal; from that point no change can improve both parties' welfare. The collection of all Pareto-optimal points forms a curve, which is called the **contract curve** and is illustrated in panel E.

We do not know in advance exactly what point Aline and Bob will reach through the trading process. We know only that it will be somewhere within the original region of mutual advantage, and that it will be on the contract curve. The set of possible outcomes is the darker segment of the contract curve shown in panel E.

Contract curve

The set of Pareto-optimal points.

Competitive Equilibrium in the Edgeworth Box

Our analysis has revealed an infinite variety of possible outcomes for the bargaining process. Next we ask what can happen if Aline and Bob play according to a far more restrictive set of rules. Instead of letting them bargain in whatever way they choose, we require them to bargain through the mechanism of a price system.

The new rules of the game work this way: Aline and Bob decide on a relative price for food and clothing. At this price, each decides how much of each commodity he or she would like to buy or sell. If their desires are compatible (i.e., if Aline wants to buy just as much food as Bob wants to sell), they carry out the transaction. If their desires are not compatible, they decide on a new relative price and try again. This process continues until they find a relative price that “clears the market” in the sense that quantities demanded equal quantities supplied.

Why would Aline and Bob ever agree to such a strange and restricted set of rules? They wouldn't, because two people can bargain far more effectively without introducing the artifice of market-clearing prices. But our interest in Aline and Bob is not personal; we are concerned with them only because we are interested in the workings of much larger markets, and such markets *do* operate through a price mechanism. So we shall force Aline and Bob to behave the way people in large markets behave, hoping that their responses will teach us something about those large markets.

Suggesting a relative price is equivalent to suggesting a slope for Aline's budget line. Once we know this slope, we know her entire budget line. This is because her budget line must pass through her endowment point, in view of the fact that she can always achieve this point by refusing to trade. Bob's budget line (viewed from his upside-down perspective) is the same as Aline's. In panel A of Exhibit 8.22 a relative price has been suggested that leads Aline to choose point X and Bob to choose point Y . The total quantity of food demanded is more than exists in the world; the total quantity of clothing demanded fails to exhaust the available supply. The market has not cleared and a new relative price must be tried. In view of the outcome at the current price, it seems sensible to raise the relative price of food. That is, we try a steeper budget line, as in panel B. This time Aline and Bob both choose the same point Z and the market clears.

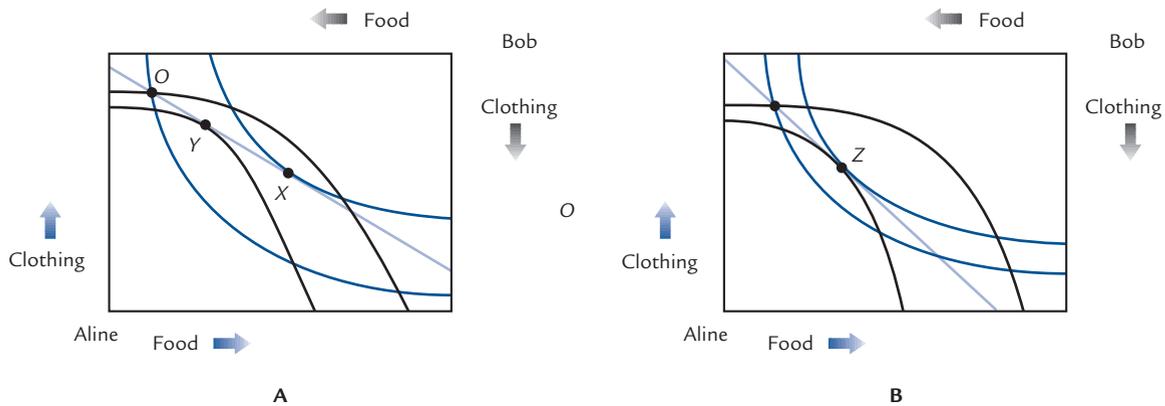
The mutually acceptable point Z in panel B is called a **competitive equilibrium** for this economy. It requires finding a budget line that goes through the original endowment point and leads to the same optimum point for Aline that it does for Bob.

Competitive equilibrium

A point that everyone will choose to trade to, for some appropriate market prices.

EXHIBIT 8.22

Competitive Equilibrium in an Edgeworth Box Economy



In panel A a relative price has been suggested that leads to the budget line pictured. (This is Aline's budget line from her perspective and Bob's budget line from his.) Aline chooses point X and Bob chooses point Y. But these points are not the same; the quantities that Aline wants to buy and sell are not the same quantities that Bob wants to sell and buy.

In panel B a different relative price has been suggested. At this price Aline's desires are compatible with Bob's. Point Z is a competitive equilibrium.

It is not immediately obvious that a competitive equilibrium should even exist, but it turns out to be possible to prove this.¹¹

The Invisible Hand in the Edgeworth Box

At the competitive equilibrium Z of Exhibit 8.22, Aline's indifference curve is tangent to the budget line, and Bob's indifference curve is tangent to the same budget line. It follows that Aline's and Bob's indifference curves are tangent to each other. This, in turn, means that the competitive equilibrium is a point on the contract curve—that is, it is Pareto-optimal.

This reasoning shows that in an Edgeworth box economy, any competitive equilibrium is Pareto-optimal. That is, the invisible hand theorem is true.

We began this section by noticing that competitive equilibrium is Pareto-optimal in the context of a single market. We have just seen that the same is true in the context of an entire economy (albeit an extraordinarily simple economy in which no production takes place). The same is also true in far more complex models involving many markets and incorporating production, though this requires advanced mathematics to prove.

General Equilibrium with Production

In the Edgeworth box economy there is no production. Next we will study general equilibrium in an economy where production is possible.

¹¹ In fact, you can prove it, if you have had a course in calculus. Define the *aggregate excess demand* for food as the sum of the quantities demanded by Bob and Aline, minus the world supply of food. At a price of zero, draw the budget line and compute the aggregate excess demand. Do the same at an infinite price. Now use the intermediate value theorem to complete the proof.

Robinson Crusoe

Robinson Crusoe lives alone on an island where the only foods he can produce are tomatoes and fish. He grows the tomatoes and catches the fish. Because each activity takes time, he can have more of one only by accepting less of the other.

Exhibit 8.23 shows the various combinations of tomatoes and fish that Robinson could produce in a week. If he grows no tomatoes, he can catch 15 fish. If he catches no fish, he can grow 18 tomatoes. The curve displaying all of his options is called Robinson's **production possibility curve**.

If Robinson starts at point E in the diagram and gives up a single tomato, he can catch ΔF additional fish. We can think of ΔF as the relative price of tomatoes in terms of fish. ΔF is also the slope of the production possibility curve at E . Therefore, the slope of the production possibility curve is equal to the relative price of tomatoes in terms of fish.

At point E , Robinson grows a lot of tomatoes. Because of diminishing marginal returns to farming on a fixed quantity of land, it takes a lot of effort to grow one more tomato. By giving up his last tomato, Robinson frees up a lot of time and catches a large number (ΔF) of fish. By contrast, if Robinson started out at a point near the northwest corner of the production possibility curve, the marginal tomato would require less effort.

Giving it up would only free a small amount of time; moreover, diminishing marginal returns to fishing render that time relatively unproductive. (Notice that Robinson is already catching a lot of fish.) In consequence, the price of a tomato in terms of fish is very low near the northwest corner, just as it is very high near the southeast corner. Remembering that price equals slope, this tells us that

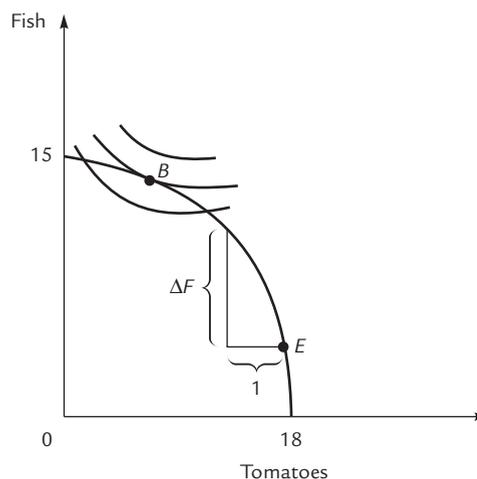
The production possibility curve bows outward from the origin.

Production possibility curve

The curve displaying all baskets that can be produced.

EXHIBIT 8.23

The Production Possibility Curve



The curve shows the various combinations of tomatoes and fish that Robinson can produce. Its slope shows how many fish he can have in exchange for one tomato and can therefore be thought of as the relative price of tomatoes. At point E that relative price is the distance ΔF . Robinson chooses a point of tangency with an indifference curve; that is, he chooses point B .

To complete the analysis, we must bring Robinson's indifference curves into the picture. Robinson chooses his favorite point on his production possibility curve, which is the tangency B . At this point, Robinson equates the relative price of tomatoes (the slope of the production possibility curve) with the marginal rate of substitution between tomatoes and fish (the slope of the indifference curve).

The Open Economy

Now suppose that Robinson establishes contact with the natives of a large nearby island. His own island is transformed into an **open economy**, one that can trade with outsiders at prices determined in world markets. The going price of a tomato on this other island is P fish dinners.

Robinson now faces two separate choices. First, how should he allocate his time between farming and fishing? Second, how should he allocate his consumption between tomatoes and fish?

We know how to answer the second question. Robinson chooses the tangency between his budget line and an indifference curve. What is his budget line? It is a line with absolute slope P (P being the relative price of tomatoes) and passing through the point representing Robinson's production. Why must it pass through that point? Because Robinson can always consume at that point by simply not trading with his neighbors. Since that point is available to him, it must be on his budget line.

Panel A of Exhibit 8.24 shows several lines with absolute slope P . If Robinson produces either basket A or basket E , his budget line is the lightest of these. If he produces B or D , his budget line is the middle one. If he produces C , the dark line is his budget line. It is best to have a budget line as far from the origin as possible, so C is Robinson's best choice. That is,

Production occurs at the point where the production possibility curve is tangent to a line of slope P . The line of tangency becomes the budget line.

Panel B of Exhibit 8.24 shows Robinson's consumption choice. Having produced basket C , he has the budget line shown; along this budget line he selects basket X . Notice that X is superior to the basket B that Robinson would consume in the absence of trade. Robinson gains from trade with his neighbors. We can go on to ask: How much does he gain?

To answer this question, we must compare two different prices. One is the **autarkic relative price** that would prevail on Robinson's island if there were no trade. With no trade, Robinson would choose point B in Exhibit 8.25 and would have the budget line shown in color. The slope of that line is the autarkic relative price of tomatoes.

The second interesting price is the **world relative price** at which Robinson can trade with his neighbors. Suppose first that the world relative price happens by chance to equal the autarkic relative price. In that case, Robinson's budget line must be tangent to the production possibility curve and parallel to the colored line; that is, his budget line is the colored line itself. He produces at the point B and consumes at the point B . But this is exactly the same point that Robinson chose in Exhibit 8.23, when there was no opportunity to trade. In other words,

If the autarkic and world relative prices are equal, then there is no gain from trade.

Suppose, alternatively, that the world relative price is given by the slope of the black line in Exhibit 8.25. Then Robinson produces at C and consumes at X , which makes him happier than if he were to consume at B . In this case, he gains from trade.

Open economy

An economy that trades with outsiders at prices determined in world markets.

Autarkic relative price

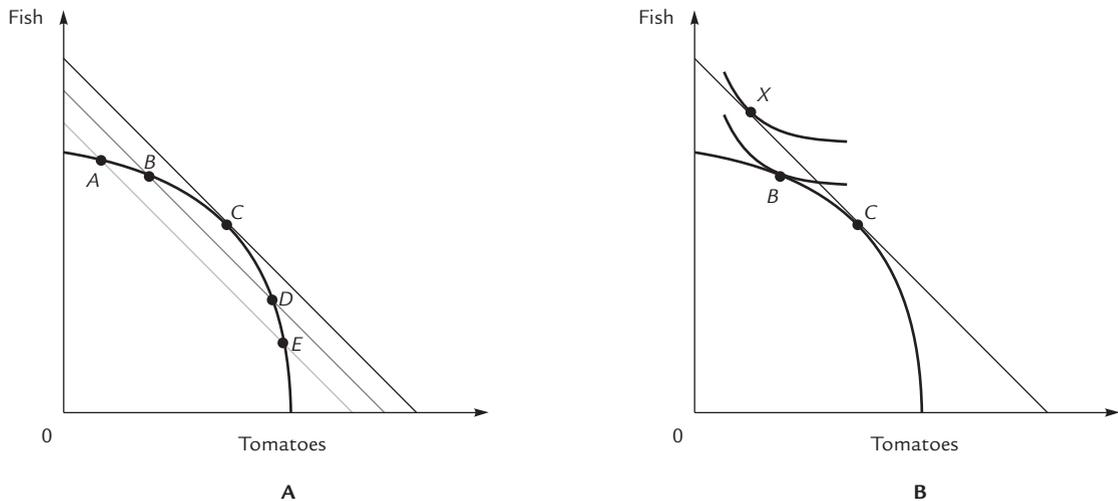
The relative price that would prevail if there were no trade with foreigners.

World relative price

The relative price that prevails in the presence of trade with foreigners.

EXHIBIT 8.24

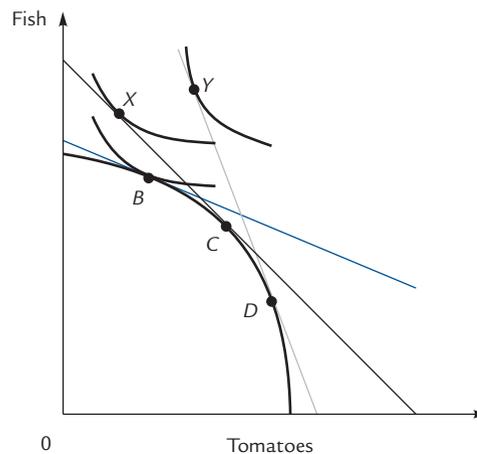
Production and Consumption with Foreign Trade



When Robinson can trade with his neighbors at a relative price of P fish per tomato, he faces a budget line of absolute slope P . All of the lines in panel A have that slope. By choosing a basket to produce, Robinson can choose his budget line from among the lines pictured. If he produces basket A or basket E, he has the light budget line; if he produces basket B or basket D, he has the middle budget line; if he produces basket C, he has the dark budget line. The dark budget line is the best one to have, so Robinson produces basket C. He then trades along the budget line to his optimal basket X, shown in panel B. Without trade, Robinson would choose basket B. Since basket X is preferred to basket B, Robinson gains from trade.

EXHIBIT 8.25

Autarkic versus World Relative Prices



The slope of the brown line represents the autarkic relative price on Robinson's island. If the world relative price is the same as the autarkic relative price, then Robinson both produces and consumes basket B, just as he would with no opportunity to trade.

If, instead, the world relative price is given by the slope of the black line, then Robinson produces basket C and consumes basket X, which is an improvement over basket B. If the world relative price goes up to the slope of the gray line, then Robinson produces basket D and consumes basket Y, which is a further improvement.

Next, suppose that the world relative price differs even more from the autarkic relative price, being given by the slope of the gray line in Exhibit 8.25. Then Robinson produces D and consumes Y , which is better even than X .

The more the world relative price differs from the autarkic relative price, the more Robinson gains from trade.

Exercise 8.9 The black and gray lines in Exhibit 8.25 represent world relative prices that are greater than the autarkic relative price. Draw some budget lines that result when world relative prices are *less* than the autarkic relative price. Check that it remains true that the gains from trade are greater when the world relative price is further from the autarkic relative price.

What determines the world relative price? The answer is: supply and demand by everyone in the world, including Robinson. Thus, the world price is a sort of average of the autarkic relative prices on all of the various islands in Robinson's trading group. If Robinson's supply and demand constitute a large percentage of the world's supply and demand, then his own autarkic relative price counts quite heavily in this average, bringing the world relative price closer to the autarkic one. This, in turn, reduces Robinson's gains from trade.

If, on the other hand, Robinson is an insignificant player in the world market, then there is a greater chance that the world relative price differs substantially from his autarkic one. In this case, Robinson's gains from trade are greater.

All of this serves to illustrate a point we made back in Chapter 2: To gain from trade, it pays to be different from the world. Small countries are more likely to be different from the world than large countries are. Therefore, small countries have more to gain from international trade than large ones do. For many goods, world relative prices do not differ significantly from U.S. relative prices, so the United States has relatively little to gain from trade in these goods. But New Zealand, for example, where the autarkic relative price of wool is quite low, benefits greatly from being able to trade its wool for other goods at the comparatively high world relative price.

The World Economy

We have seen how Robinson Crusoe reacts to world prices, and we have asserted that these world prices are determined by supply and demand. To complete the picture of the world economy, we have only to understand exactly how the world supply and demand curves are determined.

To derive a point on the supply curve for tomatoes, we imagine a price and ask what quantity Robinson supplies. Referring again to Exhibit 8.25, suppose that the colored budget line has absolute slope P . Then Robinson produces basket B , and the quantity of tomatoes he supplies is the horizontal coordinate of this point (whether he supplies them to himself or to someone else is not relevant here). The price P corresponds to this quantity on the supply curve.

To get another point on the supply curve, suppose the black budget line has absolute slope P' . At this price, Robinson produces at point C , and the corresponding quantity of tomatoes is paired with price P' on his supply curve.

A similar procedure generates points on Robinson's demand curve. When the price is P , he has the colored budget line and demands a quantity of tomatoes given by the horizontal coordinate of point B . When the price is P' , he has the black budget line and demands a quantity given by the horizontal coordinate of point X .

In this way, we can generate Robinson's supply and demand curves for tomatoes. We can do the same for all his trading partners. We get world supply and demand curves by adding the individual supply and demand curves, and these determine a world equilibrium price.

Summary

Consumers and producers both gain from trade. Consumers' and producers' surpluses are measures of the extent of their gains.

When the consumer buys a good X, the total value of his purchase is given by the area under his demand curve out to the quantity. This area is the most that he would be willing to pay in exchange for that quantity of X. After we subtract the total cost to the consumer, we are left with the area under his demand curve down to the price paid and out to the quantity consumed. This area is his consumer's surplus. It is the amount that the consumer would be willing to pay in exchange for being allowed to purchase good X.

The producer's surplus is the excess of the producer's revenues over his costs. It is measured by the area above the supply curve up to the price received and out to the quantity supplied.

When there is more than one consumer or more than one producer, the total surplus to all consumers is given by the area under the market demand curve down to the price paid and out to the quantity demanded. The total surplus to all producers is given by the area above the market supply curve up to the price received and out to the quantity sold.

Policies such as taxes or price controls can change prices and quantities and consequently change the consumers' and producers' surpluses. They also sometimes generate tax revenue (which is a gain to somebody) or impose a cost on taxpayers (which is a loss). Social gain is the sum of consumers' and producers' surpluses, plus any other gains, minus any losses. If a policy reduces social gain below what it might have been, the amount of the reduction is known as a deadweight loss.

Whenever there is deadweight loss, it is possible to devise an alternative policy that is Pareto-preferred (i.e., preferred by everybody) to the current policy. A policy is said to be Pareto-optimal, or efficient, if no other policy is Pareto-preferred.

The efficiency criterion is a normative criterion asserting that we should prefer policies that maximize social gain, or, equivalently, minimize deadweight loss. Few (if any) would argue that the efficiency criterion should be the sole guide to policy, but many economists consider it reasonable to use it as a rough guideline. When a policy creates large deadweight losses, there may be a Pareto-preferred policy that is actually possible to implement.

The invisible hand theorem states that competitive equilibrium is Pareto-optimal. That is, in a competitive market where each individual seeks only his own personal gains, it turns out to be the case that social gains are maximized. This is true in individual markets and remains true when the entire economy is taken into account. The Edgeworth box presents an example of a complete economy that can be used to illustrate the workings of the invisible hand.

It is also possible to study general equilibrium in economies with production. The opportunity to trade with outsiders confers benefits on the members of such an economy. The more world prices differ from autarkic relative prices, the greater those benefits tend to be.

Author Commentary

www.cengage.com/economics/landsburg

- AC1.** Read about the conflict between the Pareto criterion and individual freedom. Also read about how good voting systems are hard to find, both in politics and in sports.

Review Questions

- R1.** Explain why a consumer's demand curve is identical to his marginal value curve.
- R2.** What geometric areas represent the value of the goods that a consumer purchases and the cost of producing those goods? What geometric area represents the social gain from the goods' production, and why?
- R3.** What geometric areas represent the consumers' and producers' surpluses, and why?
- R4.** Analyze the effect on social welfare of a sales tax.
- R5.** Analyze the effect on social welfare of a subsidy.
- R6.** Analyze the effect on social welfare of a price ceiling.
- R7.** Analyze the effect on social welfare of a tariff, assuming that the country imposing the tariff constitutes a small part of the entire market. First answer assuming that the good in question is available only from abroad, then repeat your answer assuming that there is a domestic industry.
- R8.** "The fact that secretaries are paid less than corporate executives shows that society values secretarial services less than it values the work of executives." Comment.
- R9.** State the invisible hand theorem. Illustrate its meaning using supply and demand curves.
- R10.** Explain the difference between the allocation of resources and the distribution of income. With which is the efficiency criterion concerned?
- R11.** Using an Edgeworth box, illustrate the region of mutual advantage and the contract curve. Explain why trade will always lead to a point that is both in the region and on the curve.
- R12.** Using an Edgeworth box, illustrate the competitive equilibrium. Explain how you know that the competitive equilibrium is on the contract curve. How does this illustrate the invisible hand theorem?
- R13.** Show how Robinson Crusoe chooses his consumption point when he is unable to trade. Show how he chooses his production and consumption points when trade becomes an option.

Problem Set

1. **True or False:** If consumers buy 1,000 heads of lettuce per week, and if the price of lettuce falls by 10¢ per head, then the consumer's surplus will increase by \$100.
2. Suppose that your demand curves for gadgets and widgets are both straight lines but your demand curve for gadgets is much more elastic than your demand curve for widgets. Each is selling at a market price of \$10, and at that price you choose to buy exactly 30 gadgets and 30 widgets.
 - a. From which transaction do you gain more surplus?
 - b. If forced at gunpoint to buy either an extra gadget or an extra widget, which would you buy?
 - c. Illustrate the change in your consumer's surplus as a result of the forced transaction of part (b).
3. Adam and Eve consume only apples. Of the following allocations of apples, which are preferred to which others according to (a) the Pareto criterion, and (b) the efficiency criterion?
 - a. Adam has 12 apples and Eve has 0 apples.
 - b. Adam has 9 apples and Eve has 3 apples.
 - c. Adam has 6 apples and Eve has 6 apples.
 - d. Adam has 0 apples and Eve has 12 apples.
 - e. Adam has 5 apples and Eve has 5 apples.
4. **True or False:** Cheap foreign goods hurt American producers and are therefore bad according to the efficiency criterion.
5. **True or False:** If there is a fixed amount of land in Wyoming, then a sales tax on Wyoming land will have no effect on social welfare.
6. Home insulation is currently subsidized. Draw a graph (as in Exhibit 8.12) that shows the gains and losses to all relevant groups. Explain how you could, in principle, make *everyone* happier by eliminating the subsidy and instead transferring income from some people to others. Be explicit about exactly who you'd take income from, how much you would take, and how you would distribute it.
7. The demand and supply curves for gasoline are the same in Upper Slobbovia as in Lower Slobbovia. However, in Upper Slobbovia everybody's time is worth just \$1 per hour, while in Lower Slobbovia everybody's time is worth \$10 per hour.

True or False: If both countries impose a price ceiling on gasoline, the value of time wasted in waiting lines will be higher in Lower Slobbovia than in Upper Slobbovia.
8. In the preceding problem, suppose that there is also a country of Middle Slobbovia, where the value of various people's time ranges between \$1 and \$10. If Middle Slobbovia imposes a price ceiling on gasoline, how will the value of time wasted in waiting lines compare to the time wasted in Upper and Lower Slobbovia?
9. Suppose the equilibrium price of potatoes is \$5 per pound, but the government imposes a price ceiling of \$2 per pound, creating a deadweight loss. **True or False:** If the government imposes an *excise tax* of \$1 per pound of potatoes (while continuing to maintain the price ceiling), then the deadweight loss will get even larger.

10. Suppose the equilibrium price of potatoes is \$5 per pound, but the government imposes a price ceiling of \$2 per pound, creating a deadweight loss. **True or False:** If the government imposes a *sales tax* of \$1 per pound of potatoes (while continuing to maintain the price ceiling), then the deadweight loss will get even larger.
11. Suppose there is a federal excise tax on gasoline of 4.3¢ per gallon. A United States senator has proposed eliminating this tax, but requiring oil companies to pass all of the savings on to the consumer (by maintaining a new price at the pump that is 4.3¢ lower than the current price). Show the deadweight loss under the current tax and under the senator's plan. Can you tell which is bigger?
12. **True or False:** A price ceiling on wheat would cause the price of bread to fall.
13. In equilibrium, 500 pounds of potatoes are sold each week. However, the government prints up and randomly distributes 250 non-reusable ration tickets each week, and requires that buyers present one ration ticket for each pound of potatoes they buy. Therefore consumers can purchase only 250 pounds of potatoes per week. Ration tickets can be freely bought and sold. Draw a graph illustrating the market for potatoes, and on your graph indicate:
- The price of a pound of potatoes.
 - The price of a ration ticket.
 - The consumer and producer surplus in the potato market.
 - The value of the ration tickets to the citizens who are randomly chosen to receive them.
 - The deadweight loss.
14. The American supply and demand curves for potatoes cross at \$5 a pound, but potatoes are available in any quantity from abroad at \$2 a pound. Each week, American sellers produce 500 pounds of potatoes and American buyers purchase a total of 1,200 pounds. Suppose the government prints and randomly distributes 250 non-reusable ration tickets each week, and requires that buyers present one ration ticket for each pound of *imported* potatoes that they buy. (You don't need a ration ticket to buy an American potato.) Ration tickets can be freely bought and sold. Draw a graph illustrating the market for potatoes, and on your graph indicate the price Americans now pay for potatoes, the price of a ration ticket, the gains and losses to all relevant groups, and the deadweight loss.
15. Widgets are produced by a competitive industry and sold for \$5 apiece. The government requires each widget firm to have a license, and charges the highest license fee firms are willing to pay. If the government were to impose an excise tax of \$3 per widget, how much would the license fee have to change? (Illustrate your answer as an area on a graph). Would total government revenue (i.e., excise tax revenue plus revenue from license fees) rise or fall as a result of the excise tax?
16. In equilibrium, 2,000 pounds of potatoes are sold each month. A new law requires sellers to buy permits before they can sell potatoes. One permit allows you to sell one pound, and permits can't be reused. The government creates only 1,000 permits and sells them to the highest bidders. Use a graph to show the new price of potatoes, the price of a permit, the gains and losses to all relevant groups, and the deadweight loss.
17. Toys are produced by a competitive industry. Santa Claus gives away one million free toys each year. Illustrate Santa's effect on a) the price of toys, b) the consumer

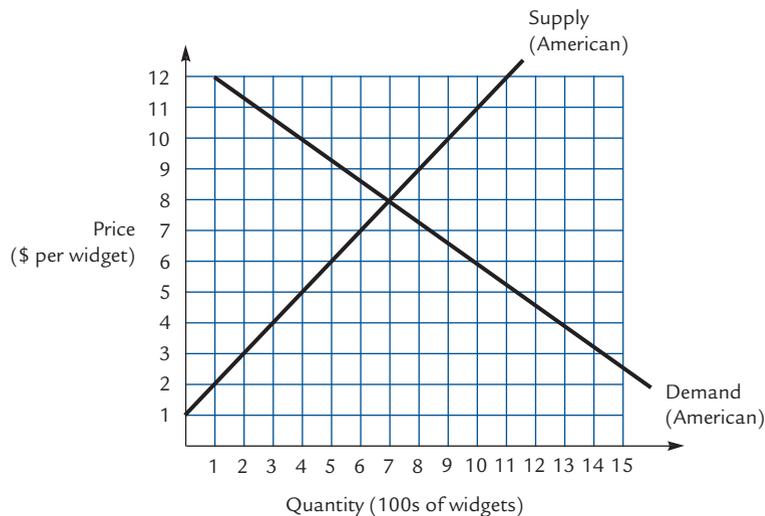
surplus, c) the producer surplus earned by commercial toy manufacturers, and d) social gain. (Don't worry about gains or losses to Santa.) (*Hint*: Remember that the toys Santa distributes are free.)

- 18.** Suppose the government sets an effective price *floor* (i.e., a price above equilibrium) in the market for oranges and agrees to buy all oranges that go unsold at that price. The oranges purchased by the government are discarded. Illustrate the number of oranges purchased by the government. Illustrate the gains and losses to all relevant groups of Americans and the deadweight loss.
- 19.** The American demand and supply curves for oranges cross at a price of \$8, but all Americans are free to buy or sell oranges on the world market at a price of \$5. One day, the U.S. government announces that it will pay \$6 apiece for American oranges and will buy as many oranges as Americans want to sell at that price. The government then takes these oranges and resells them on the world market at \$5 apiece. Illustrate the gains and losses to all relevant groups of Americans, and illustrate the deadweight loss.
- 20.** In the tariff example of Exhibit 8.17, divide the two triangles of deadweight loss into individual rectangles of loss, as in Exhibit 8.9. Give an intuitive explanation of the loss that each of those rectangles represents.
- 21.** Suppose the U.S. supply and demand curves for automobiles cross at a price of \$15,000 but (identical) automobiles can be purchased from abroad for \$10,000. Now suppose the government imposes a \$2,000 excise tax on every car produced in the United States (regardless of whether the car is sold in the United States or abroad).
- What price must Americans pay for cars before the tax is imposed? What price must Americans pay for cars after the tax is imposed? (*Hint*: Americans can always buy cars on the world market and so will never pay more than the world price for a car.) What prices do U.S. producers receive for their cars before and after the tax is imposed?
 - Before and after the tax is imposed, calculate the gains to all relevant groups of Americans. What is the deadweight loss due to the tax?
- 22.** The equilibrium price of an apple is 25¢. The government sets a price ceiling of 15¢ per apple, and requires sellers to provide as many apples as buyers want to buy at that price. Draw a graph to illustrate the price ceiling and answer the following questions in terms of areas on your graph:
- What areas would you measure to determine whether sellers continue selling apples in the short run? (To answer this, you can assume that all sellers are identical.)
 - If sellers continue selling apples, what area represents the deadweight loss and why?
 - If sellers don't continue selling apples, what area represents the deadweight loss and why?
- 23.** Suppose the U.S. supply and demand curves for automobiles cross at a price of \$15,000 but (identical) automobiles can be purchased from abroad for \$10,000. Now suppose the government imposes a \$2,000 sales tax on every American who buys a car (regardless of whether the car is produced domestically or abroad).
- What price must Americans pay for cars before the tax is imposed? What price must Americans pay for cars after the tax is imposed? (*Hint*: American

- suppliers can always sell cars abroad for \$10,000 and so will never sell cars for less.) What prices do U.S. producers receive for their cars before and after the tax is imposed?
- b. Before and after the tax is imposed, calculate the gains to all relevant groups of Americans. What is the deadweight loss due to the tax?
- 24.** Suppose the U.S. supply and demand curves for automobiles cross at a price of \$15,000 but (identical) automobiles can be purchased from abroad for \$10,000. Now suppose the government offers a subsidy of \$2,000 to each American who buys an imported car. Buyers of domestic cars receive no subsidy.
- a. What price do Americans pay for domestic cars before the subsidy is offered? What is the most an American will pay for a domestic car after the subsidy is offered?
 - b. Given your answer to part (a), and given that anyone can buy or sell cars abroad at the world price of \$10,000, how many cars will U.S. producers want to sell in the United States?
 - c. Before and after the subsidy is offered, calculate the gains to all relevant groups of Americans. What is the deadweight loss due to the subsidy?
 - d. How does your answer change if U.S. producers are prohibited from selling cars abroad?
- 25.** Suppose the U.S. supply and demand curves for automobiles cross at a price of \$15,000 but (identical) automobiles can be purchased from abroad for \$10,000. Now suppose the government offers U.S. producers a \$2,000 subsidy for every car they produce (regardless of whether the car is sold in the United States or abroad).
- a. What prices must Americans pay for cars before and after the subsidy is offered? What prices do U.S. producers feel they are receiving before and after the subsidy is offered?
 - b. Before and after the subsidy is offered, calculate the gains to all relevant groups of Americans. What is the deadweight loss due to the subsidy?
- 26.** Suppose the U.S. supply and demand curves for automobiles cross at a price of \$15,000 and that (identical) automobiles can be purchased from abroad for \$10,000. Now suppose the government offers a \$2,000 subsidy to every American who buys a car (regardless of whether the car is foreign or domestic).
- a. At what prices do U.S. producers sell their cars before and after the subsidy is offered? What prices do U.S. consumers feel like they are paying before and after the subsidy is offered?
 - b. Before and after the subsidy is offered, calculate the gains to all relevant groups of Americans. What is the deadweight loss due to the subsidy?
- 27.** The American supply and demand curves for cars cross at \$15,000. Foreigners will sell us any number of cars at the world price of \$10,000. Now the government announces two new taxes: a sales tax of \$1,000 on each American car, and a sales tax (i.e., a tariff) of \$3,000 on each foreign car. Illustrate the gains and/or losses to all relevant groups of Americans as a result of the combined tax program, and illustrate the deadweight loss.
- 28.** There is currently a sales tax on all cars, foreign and domestic. In order to help the American car industry, the government is thinking of eliminating the sales tax.

Plan A is to eliminate the tax for domestic cars only; Plan B is to eliminate the tax for both domestic and foreign cars.

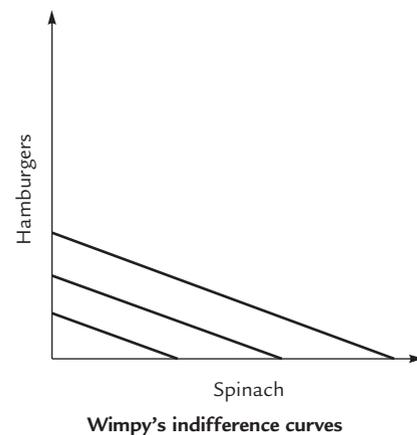
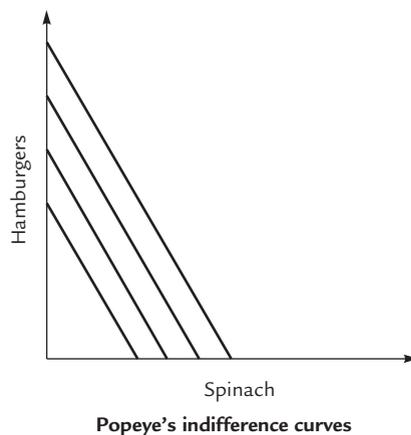
- a. Which plan is better for domestic car makers?
 - b. **True or False:** Plan B, combined with an appropriate redistribution of income, can make *everybody* happier than Plan A. If your answer is true, explicitly describe the appropriate redistribution of income. If your answer is false, explain carefully why no such redistribution is possible.
- 29.** The American supply and demand curves for bananas cross at \$5. Foreigners will buy as many bananas as Americans want to sell at \$10. The government subsidizes exports by giving sellers \$2 for each banana they sell abroad. Bananas cannot be imported into the United States.
- a. Illustrate the deadweight loss from the subsidy.
 - b. Devise a program that *everyone*—buyers, sellers, and the recipients of the revenue—prefers to the subsidy.
- 30.** The American supply and demand curves for widgets are illustrated at the top of the next page. Foreign widget makers will sell any quantity of widgets to Americans at a price of \$2 apiece.



Suppose the government distributes 500 ration coupons, which can be freely bought and sold. To buy a foreign-made widget, you must present a ration coupon. (An American widget requires no coupon.) What is the price of a ration coupon? (Your answer should be a number.)

- 31.** Suppose that the government successfully maintains a price P_0 for wheat that is above the equilibrium price. At this price, consumers want to purchase Q_d bushels of wheat and farmers want to produce Q_s . The way that the government maintains the price P_0 is by offering farmers a cash reward for limiting their production.
- a. By how much must farmers agree to cut back production in order for the program to be successful?
 - b. Show on a graph the minimum payment that the government must make to farmers in order for them to agree to the deal.

- c. Assuming that the government makes this minimum payment, use your graph to show the gains and losses to consumers, producers, and taxpayers from this arrangement. Calculate the deadweight loss.
- 32.** The American demand and supply curves for labor cross at a wage rate of \$25 per hour. However, American firms can hire as many foreign workers as they want to at a wage of \$15 per hour. (Assume that foreign workers are exactly as productive as American workers.) A new law requires American firms to pay \$25 an hour to Americans and to hire every American who wants to work at that wage. Firms may still hire any number of foreigners at \$15 per hour.
- Before the law is enacted, what wage do American workers earn? Illustrate the consumers' and producers' surpluses earned by American workers and American firms.
 - After the law is enacted, illustrate the number of Americans hired and the number of foreigners hired, the consumers' and producers' surpluses earned by American workers and American firms, and the deadweight loss.
- 33.** Widgets are produced by a constant-cost industry. Suppose the government decides to institute an annual subsidy of \$8,000 per year to every firm that produces widgets.
- Explain why, in the long run, each firm's producer surplus must fall by \$8,000.
 - Suppose the subsidy causes the price of widgets to fall by \$1. With the subsidy in place, does each firm produce more than, fewer than, or exactly 8,000 widgets a year?
 - Suppose the government replaces the per-firm subsidy with a per-widget subsidy of \$1 per widget produced. In the long run, is this change good or bad for consumers? Is it good or bad for producers? (*Hint:* Remember the zero-profit condition!) Is it good or bad for taxpayers? Is it good or bad according to the efficiency criterion?
- 34. True or False:** If all thieves are identical, then the social cost of robbery is equal to the value of the stolen goods.
- 35.** Popeye and Wimpy trade only with each other. Popeye has 8 hamburgers and 2 cans of spinach, and Wimpy has 2 hamburgers and 8 cans of spinach. Their indifference, somewhat unusually, are all straight lines, Popeye's being much steeper than Wimpy's:



In an Edgeworth box, show the initial endowment, the region of mutual advantage, the contract curve, and the competitive equilibrium.

- 36.** Robinson Crusoe lives alone on an island, producing nuts and berries and trading with people on other islands. If his production possibility curve is a straight line, what can you conclude about the quantities of nuts and berries he will produce?
- 37.** Robinson Crusoe lives alone on an island, producing nuts and berries and trading with people on other islands. **True or False:** If nuts are an inferior good for Robinson, then his supply curve for nuts must be upward sloping.

Appendix to Chapter 8

Normative Criteria

Suppose that by ordering the execution of one innocent man you could save the lives of five others, equally innocent. Should you do it?

Consequentialist moral theories assert that the correctness of an act depends only on its consequences. A simple consequentialist position might be that one lost life is less bad than five lost lives, so the execution should proceed.

Other views are possible. You might argue that if the one man to be executed is happy and fulfilled, while the other five lead barely tolerable lives, then it would be better to spare the one and sacrifice the five. This position is still consequentialist, because it judges an action by its consequences: The sacrifice of one happy life versus the sacrifice of five unhappy lives.

There are also moral theories that are not consequentialist. Some are based on natural rights. One could argue that a man has a natural right to live and that there can be no justification for depriving him of this right, regardless of the consequences. There can be no execution, even if it would save a hundred innocent lives.

In the heated public debate about abortion, both sides have tended to make arguments that go beyond consequentialism. One side defends a “right to choose,” while the other defends a “right to life.” A strictly consequentialist view would discard any discussion of “rights” and judge the desirability of legalized abortion strictly on the basis of its implications for human happiness. This is not enough to settle the issue; one must still face extraordinarily difficult questions about how to trade off different people’s happiness and potential happiness. Consequentialism, like natural rights doctrine, accommodates many precepts and conclusions.

The efficiency criterion is an example of a consequentialist normative theory. Which kind of world is better: One with 10 people, each earning \$50,000 per year, or one with 10 people, of whom 3 earn \$30,000 and 7 earn \$100,000? According to a strict application of the efficiency criterion, the second world is better, because total income is \$790,000 instead of \$500,000. The world with more wealth is the better world.

There are many other possible viewpoints, some consequentialist and some not. One might argue that certain people are more deserving of high income than others, and so there is no way to choose between the two income distributions without knowing more about the characteristics of the people involved. Such a position introduces criteria other than the ultimate consequences for human happiness, and so can be characterized as nonconsequentialist.

Judging the desirability of outcomes requires a normative theory. Economics can help us understand the implications of various theories, and perhaps help us choose among them. For the most part, economic analysis tends to focus on the various

Consequentialist moral theories

Moral theories that assert that the correctness of an act can be judged by its consequences.

consequentialist theories. This is not because natural rights doctrines are uninteresting; it just seems to be the case that (so far) economics has less to say about them.

Some Normative Criteria

Here are a few of the normative criteria that economists have thought about.

Majority Rule

According to this simple criterion, the better of two outcomes is the one that most people prefer.

A number of objections can be raised. One is that majority rule does not provide a coherent basis for choosing among *three* or more possible outcomes. Sharon, Lois, and Bram plan to order a pizza with one topping. Their preferences are shown in the following table:

	Sharon	Lois	Bram
First Choice	Peppers	Anchovies	Onions
Second Choice	Anchovies	Onions	Peppers
Third Choice	Onions	Peppers	Anchovies

A majority (Sharon and Bram) prefers peppers to anchovies, a different majority (Sharon and Lois) prefers anchovies to onions, and a third (Lois and Bram) prefers onions to peppers. No matter what topping is chosen, there is some majority that prefers a different one.

A more fundamental objection to majority rule is that it forces us to accept outcomes that almost all people agree are undesirable. If 60% of the people vote to torture and maim the other 40% for their own amusement, a true believer in majority rule is forced to admit the legitimacy of their decision.

A less flamboyant example is a proposed tax policy that would have the effect of increasing 51% of all household incomes by \$1 per year while decreasing 49% of all household incomes by \$10,000 per year. The majority supports the proposal. Do you think it should be implemented?

The Kaldor–Hicks Potential Compensation Criterion

The British economists Nicholas Kaldor and Sir John Hicks suggest a normative criterion under which a change is a good thing if it would be possible in principle for the winners to compensate the losers for their losses and still remain winners.

If a policy increases Jack's income by \$10, reduces Jill's by \$5, and has no other effects, should it be implemented? According to Kaldor–Hicks, the answer is yes, because Jack could in principle reimburse Jill for her loss and still come out ahead. On the other hand, a policy that increases Jack's income by \$10 while reducing Jill's by \$15 is a bad thing, because there is no way for Jack to reimburse Jill out of his winnings.

In applications like this, the Kaldor–Hicks criterion and the efficiency criterion amount to the same thing. When Jack gains \$10 and Jill loses \$5, social gains increase by \$5, so the policy is a good one. When Jack gains \$10 and Jill loses \$15, there is a deadweight loss of \$5, so the policy is bad.

However, there are potential subtleties that we did not address when we discussed the efficiency criterion in Section 8.1. Suppose that Jack has a stamp collection that he values very highly. Aside from his stamp collection, he owns nothing of great value and

in fact barely gets enough to eat. Nevertheless, he would be unwilling to sell his stamp collection for anything less than \$100,000. On the other hand, if the collection were taken from him, he would be willing to pay only \$100 to get it back; any higher payment would mean starvation.

Jill values Jack's stamp collection at \$50,000, regardless of who currently owns it. Should the collection be taken from Jack and given to Jill? If so, she would gain \$50,000 in surplus, which is not enough to compensate Jack for his \$100,000 loss. The Kaldor–Hicks criterion opposes such a move.

On the other hand, suppose that the stamp collection has already found its way into Jill's hands. If it is restored to Jack, he gains something that he values at \$100, not enough to compensate Jill for her \$50,000 loss. Kaldor–Hicks opposes this move also.

Thus, we get the somewhat paradoxical result that Jack gets to keep his stamp collection, unless it accidentally finds its way into Jill's hands, in which case Jack is not allowed to get it back.

Such paradoxes did not arise when we applied the efficiency criterion in Section 8.1. Why not? When we experimented with changing government policies, making some people better off and others worse off, we implicitly assumed that there were no resulting income effects on demand. If there are income effects, they cause the demand curve to shift at the moment when the policy is implemented. (In the present example, Jack's demand for stamps shifts dramatically depending on whether he already owns them or not.) This makes welfare analysis ambiguous: Should we calculate surplus using the old demand curve or the new one?

However, when the changes being contemplated do not affect large fractions of people's income, the Kaldor–Hicks criterion becomes unambiguous and equivalent to the efficiency criterion we have already studied.

The Veil of Ignorance

Let us repeat an earlier question. Is it better for everyone to earn \$50,000 or for 70% of us to earn \$100,000 while the rest earn \$30,000?

The philosopher John Rawls has popularized a way to think about such problems.¹ Imagine two planets: On Planet X everyone earns \$50,000; on Planet Y 70% earn \$100,000 and the rest earn \$30,000. On which planet would you rather be born? Your honest answer reveals which income distribution is morally preferable.

When you choose where to be born, it is important that you not know who you will be. If you knew that you'd be rich on Planet Y, you would presumably choose Y; if you knew you'd be poor on Y you would presumably choose X. But Rawls insists that we imagine making the decision from behind a *veil of ignorance*, deprived of any knowledge of whose life we will live.

A potential problem with the “veil of ignorance” criterion is that there might be honest disagreements about which is the better world. But Rawls contends that such disagreements arise because of different circumstances in our present lives. If we take seriously the presumption of the veil, that we have not yet lived and are all equally likely to live one life as another, then the reasons for disagreement will vanish and we will achieve unanimity. The unanimous decision is the right decision.

Suppose that a potential change in policy would enrich one billionaire by \$10,000 while costing eight impoverished people \$1,000 each. The efficiency criterion

¹ See J. Rawls, *A Theory of Justice* (Harvard University Press, 1971). A similar idea had appeared in J. C. Harsanyi; “Cardinal Utility in Welfare Economics and the Theory of Risk Taking,” *Journal of Political Economy* 61 (1953): 434–435.

pronounces such a policy a good one. Rawls's criterion probably would not. If you did not yet know whether you were going to be the billionaire or one of the impoverished, it seems likely that you would oppose this policy, on a variety of grounds. First, \$10,000 is unlikely to make much difference in a billionaire's life, while a loss of \$1,000 can be devastating if you are very poor. Second, it is 8 times more likely that you will be poor than rich. Rawls would argue that behind the veil of ignorance, the vote against this policy would be unanimous. Therefore, the policy is bad.

The veil of ignorance can be used to justify various forms of social insurance, in which income is redistributed from the more to the less fortunate. Some misfortunes do not usually strike until late in life, and we can buy insurance against them at our leisure. But other misfortunes are evident from birth, making insurance impossible. You can't insure against being born into poverty, or with below-average intelligence. There is a plausible case that behind the veil, we *would* insure ourselves, by agreeing that those born into the best circumstances will transfer income to those born into the worst. Because everyone behind the veil would want this agreement, it is a good thing and should be enforced.

The Maximin Criterion

The maximin criterion says that we should always prefer that outcome which maximizes the welfare of the worst-off member of society. Taken to the extreme, this means that a world in which everyone is a millionaire, except for one man who has only \$200, is not as good as a world in which everyone has only \$300 except for one man who has \$201.

Perhaps nobody would want to apply the maximin criterion in a circumstance quite so extreme as this. But John Rawls believes that for the most part, souls living behind the veil of ignorance would want the maximin criterion to be applied. This is because people abhor risk and worry about the prospect of being born unlucky. Therefore, while still behind the veil, their primary concern is to improve the lot of the least fortunate members of society.

According to Rawls, then, the maximin criterion is not really a new criterion at all, but instead prescribes essentially the same outcomes that the veil of ignorance criterion prescribes.²

The Ideal Participant Criterion

This is a slight variant on the veil of ignorance criterion, developed by Professor Tyler Cowen for the purpose of thinking about the problem of population but applicable more generally. (We will briefly address the population problem later in this appendix.) According to this criterion, we should imagine living many lives in succession, one each in the circumstances of every person on earth. The right outcomes are the ones we would choose before setting out on this long journey.

In comparing the ideal participant criterion with the more standard veil of ignorance criterion, you might want to consider two critical questions. First, in what

² A complete statement of Rawls's position would have to incorporate at least two additional subtleties. First, Rawls believes that from behind the veil, people's first priority would be to design social institutions that guarantee individual liberty. Having narrowed down to this set of institutions, they would then choose among them according to the maximin criterion. Second, Rawls does not want to apply the maximin criterion to particular details of the income distribution or human interactions. He wants to apply it instead to the design of social institutions. Thus, a Rawlsian might focus not on designing the ideal income distribution but rather on designing an ideal tax structure, from which the income distribution would arise. Rawls seeks that tax structure, among all of those that are consistent with individual liberty, which maximizes benefits to the least well-off members of society.

circumstances would these criteria lead to the same choices and in what circumstances would they disagree? Second, is there some more fundamental moral principle from which we can deduce a preference for one of the two criteria over the other? So far, economists have not found much to say about either of these issues.

Utilitarianism

Utilitarianism, a creation of the philosopher Jeremy Bentham, asserts that it is meaningful to measure each person's *utility*, or happiness, by a number. This makes it possible to make meaningful comparisons across people: If your utility is 4 and mine is 3, then you are happier than I am. (By contrast, many modern economists deny that any precise meaning can be attached to the statement "Person X is happier than Person Y.")³

Starting from the assertion that utilities are meaningful, utilitarians argue that the best outcome is the one that maximizes the sum of everybody's utilities. By this criterion, it is often better to augment the income of a poor man than a rich man, because an extra dollar contributes more to the poor man's utility than to the rich man's. This conclusion need not follow, however. One can imagine that the poor man has for some reason a much lower *capacity for happiness* than the rich man has, so that additional income contributes little to his enjoyment of life.

A generalized form of utilitarianism proposes that we assign a *weight* to each person and maximize the *weighted* sum of their utilities. If Jack has weight 2 and Jill has weight 3, then we choose the outcome that maximizes twice Jack's utility plus 3 times Jill's. The source of the weights themselves is left open, or is determined by any of various auxiliary theories.⁴

Under quite general circumstances, it is possible to prove that utilitarianism, with any choice of weights, always leads to a Pareto-optimal outcome and that utilitarian criteria are the only criteria that always lead to Pareto-optimal outcomes. This is so even if we drop the assumption that it is meaningful to compare different people's utilities.

Fairness

Economists have attempted to formalize the notion of fairness in a variety of ways, usually in the context of allocating fixed supplies of more than one good. In a world with 6 apples and 6 oranges, it seems absurd to insist that Jack and Jill each end up with 3 of each fruit; after all Jack might have a strong preference for apples and Jill for oranges. On the other hand, it seems quite unfair for either Jack or Jill to have all of the food while the other one starves. What precisely distinguishes those allocations that we think are equitable?

A widely studied criterion is that allocations should be **envy-free**, which means that no person would prefer somebody else's basket of goods to his own. Any allocation of apples and oranges is envy-free if neither Jack nor Jill would want to trade places with the other, given the choice.

Utilitarianism

The belief that utility, or happiness, can be meaningfully measured and that it is desirable to maximize the sum of everyone's utility.

Envy-free allocation

An outcome in which nobody would prefer to trade baskets with anybody else.

³ Utilitarians are not the only ones who believe that they can compare different people's happiness. In order to apply the maximin criterion, for example, it is necessary to make sense of the notion of the "least well-off" member of society.

⁴ The primary proponents of utilitarianism among economists were H. Sidgwick and F. Y. Edgeworth (the same Edgeworth of the Edgeworth box). For a very interesting attempt to reconstruct the weights that Sidgwick and Edgeworth had in mind, see M. Yaari, 'Rawls, Edgeworth, Shapley, Nash: Theories of Distributive Justice Re-examined,' *Journal of Economic Theory* 24 (1981): 1–39.

In an Edgeworth box economy, it is possible to show that if each trading partner starts with equal shares of everything (3 apples and 3 oranges each), then any competitive equilibrium is envy-free. This is an important result, because we already know that any competitive equilibrium is efficient as well; that is, it satisfies the efficiency criterion. This implies that in such an economy, it is always possible to achieve an outcome that is simultaneously efficient and envy-free, satisfying two criteria at once.

Optimal Population

What is the right number of people?⁵ If large populations imply crowding and unpleasantness, then how much is too much? Would it be better if there were only 10 people, each deliriously happy, or if there were 1 billion people, each slightly less happy? Where should we draw the line?

It should first be noted that the implied premise is at least debatable. A 10% increase in the current world population would change a lot of things, some for the better and some for the worse. The new arrivals would consume resources (which is bad for the rest of us) and produce output (which is good for the rest of us); it is unclear whether we'd be better or worse off on balance.

Still, it is probable that beyond some point—though it might be very far beyond the point we're at now—increases in population will make life less pleasant for everyone. At what point does the population become 'too big'? The population problem tends to confound the usual normative criteria, which are designed to address the problem of allocating resources among a fixed number of people.

We could adopt the utilitarian prescription, attempting to maximize total utility. A world of 1 billion reasonably happy people is better than a world with 100 extremely happy people, because total utility is higher in the first of these worlds. But the same criterion dictates that a world of 10 trillion people, each leading a barely tolerable existence, can be superior to the world of 1 billion who are reasonably happy. To some economists, this conclusion is self-evidently absurd. Professor Derek Parfit has endowed it with a proper name: He calls it the Repugnant Conclusion.⁶ To Parfit and others, any moral theory that entails the Repugnant Conclusion must be rejected. There are others, though, who think that the repugnance of the Repugnant Conclusion is far from evident.

An alternative is to maximize average (as opposed to total) utility. In practice, people are probably happier on average when the population is reasonably large (so that there is greater efficiency in production, a wider range of consumer goods, and a better chance of finding love). Therefore, a world of 1 billion might lead to higher average utility than a world of 100, even though a grossly overcrowded world of 1 trillion is worse than either. An objection to the average utility criterion is that it always implies that the world would be a better place if everyone with below-average utility were removed.

Alternatively, we can step behind the veil of ignorance and ask how many of us should be born. The trade-off is this: If the population is too large, the world is an unpleasant place, but if it is too small, most of us never get a chance to live. The conceptual problem here is to decide exactly how many souls there are behind the veil. Is there one for every person who *might* be born? Is that an infinite number? If so, then

⁵ This entire section owes much to Tyler Cowen's paper "Normative Population Theory," *Social Choice and Welfare* 6 (1989): 33–43.

⁶ D. Parfit, *Reasons and Persons* (Oxford University Press, 1984).

each has effectively zero chance of being among the finitely many lucky ones who do get born, rendering each indifferent to what the world is like. If instead there is a large, finite number of souls behind the veil, what determines that large, finite number?

Tyler Cowen has raised an additional objection to the veil of ignorance criterion. He asks a form of the following question: Suppose that you were offered a bet, whereby there is a 1% chance that 100 duplicate copies of earth will be created and a 99% chance that all human life will disappear. Would you take the bet? Behind the veil you would, because it actually increases the chance of your birth without changing the average quality of human life. Yet, Cowen argues, the bet is obviously a bad one.⁷ Because the veil criterion leads us to choose a bad bet, it must be a bad criterion.

Cowen has argued that the Ideal Participant Criterion is the ideal criterion for considering problems of population. You can read his arguments in the paper cited in the footnote at the beginning of this section. But the issue is very far from settled. In a world where we can't agree on what the speed limit should be, a consensus on population size will probably be a long time coming.

Author Commentary

www.cengage.com/economics/landsburg

- AC1 A-1.** Good voting systems are hard to find, both in politics and in sports.
- AC1 A-2.** If we take the veil of ignorance seriously, what does it dictate? For some back-of-the-envelope calculations, read this article.
- AC1 A-3.** Read about one point of view on various normative criteria applied to tax policy.
- AC1 A-4.** What does utilitarianism say about the case for environmental conservation?
- AC1 A-5.** For more on optimal population, see this article.
- AC1 A-6.** This is an article on optimal population.

⁷ It should perhaps be mentioned that what is obvious to Cowen is not obvious to everyone, among them the author of your textbook.

