CHAPTER 16

EQUILIBRIUM

In preceding chapters we have seen how to construct individual demand curves by using information about preferences and prices. In Chapter 15 we added up these individual demand curves to construct market demand curves. In this chapter we will describe how to use these market demand curves to determine the equilibrium market price.

In Chapter 1 we said that there were two fundamental principles of microeconomic analysis. These were the optimization principle and the equilibrium principle. Up until now we have been studying examples of the optimization principle: what follows from the assumption that people choose their consumption optimally from their budget sets. In later chapters we will continue to use optimization analysis to study the profit-maximization behavior of firms. Finally, we combine the behavior of consumers and firms to study the equilibrium outcomes of their interaction in the market.

But before undertaking that study in detail it seems worthwhile at this point to give some examples of equilibrium analysis—how the prices adjust so as to make the demand and supply decisions of economic agents compatible. In order to do so, we will have to briefly consider the other side of the market—the supply side.

16.1 Supply

We have already seen a few examples of supply curves. In Chapter 1 we looked at a vertical supply curve for apartments. In Chapter 9 we considered situations where consumers would choose to be net suppliers or demanders of goods that they owned, and we analyzed labor-supply decisions.

In all of these cases the supply curve simply measured how much the consumer was willing to supply of a good at each possible market price. Indeed, this is the definition of the supply curve: for each p, we determine how much of the good will be supplied, $S(p)$. In the next few chapters we will discuss the supply behavior of firms. However, for many purposes, it is not really necessary to know where the supply curve or the demand curve comes from in terms of the optimizing behavior that generates the curves. For many problems the fact that there is a functional relationship between the price and the quantity that consumers want to demand or supply at that price is enough to highlight important insights.

16.2 Market Equilibrium

Suppose that we have a number of consumers of a good. Given their individual demand curves we can add them up to get a market demand curve. Similarly, if we have a number of independent suppliers of this good, we can add up their individual supply curves to get the **market supply curve**.

The individual demanders and suppliers are assumed to take prices as given—outside of their control—and simply determine their best response given those market prices. A market where each economic agent takes the market price as outside of his or her control is called a **competitive market.**

The usual justification for the competitive-market assumption is that each consumer or producer is a small part of the market as a whole and thus has a negligible effect on the market price. For example, each supplier of wheat takes the market price to be more or less independent of his actions when he determines how much wheat he wants to produce and supply to the market.

Although the market price may be independent of any *one* agent's actions in a competitive market, it is the actions of all the agents together that determine the market price. The **equilibrium price** of a good is that price where the supply of the good equals the demand. Geometrically, this is the price where the demand and the supply curves cross.

If we let $D(p)$ be the market demand curve and $S(p)$ the market supply curve, the equilibrium price is the price p^* that solves the equation

$$D(p^*) = S(p^*).$$

The solution to this equation, p^*, is the price where market demand equals market supply.

Why should this be an equilibrium price? An economic equilibrium is a situation where all agents are choosing the best possible action for themselves and each person's behavior is consistent with that of the others. At any price other than an equilibrium price, some agents' behaviors would be infeasible, and there would therefore be a reason for their behavior to change. Thus a price that is not an equilibrium price cannot be expected to persist since at least some agents would have an incentive to change their behavior.

The demand and supply curves represent the optimal choices of the agents involved, and the fact that they are equal at some price p^* indicates that the behaviors of the demanders and suppliers are compatible. At any price *other* than the price where demand equals supply these two conditions will *not* be met.

For example, suppose that we consider some price $p' < p^*$ where demand is greater than supply. Then some suppliers will realize that they can sell their goods at more than the going price p' to the disappointed demanders. As more and more suppliers realize this, the market price will be pushed up to the point where demand and supply are equal.

Similarly if $p' > p^*$, so that demand is less than supply, then some suppliers will not be able to sell the amount that they expected to sell. The only way in which they will be able to sell more output will be to offer it at a lower price. But if all suppliers are selling the identical goods, and if some supplier offers to sell at a lower price, the other suppliers must match that price. Thus excess supply exerts a downward pressure on the market price. Only when the amount that people want to buy at a given price equals the amount that people want to sell at that price will the market be in equilibrium.

16.3 Two Special Cases

There are two special cases of market equilibrium that are worth mentioning since they come up fairly often. The first is the case of fixed supply. Here the amount supplied is some given number and is independent of price; that is, the supply curve is vertical. In this case the equilibrium *quantity* is determined entirely by the supply conditions and the equilibrium *price* is determined entirely by demand conditions.

The opposite case is the case where the supply curve is completely horizontal. If an industry has a perfectly horizontal supply curve, it means that the industry will supply any amount of a good at a constant price. In this situation the equilibrium *price* is determined by the supply conditions, while the equilibrium *quantity* is determined by the demand curve.

The two cases are depicted in Figure 16.1. In these two special cases the determination of price and quantity can be separated; but in the general case the equilibrium price and the equilibrium quantity are jointly determined by the demand and supply curves.

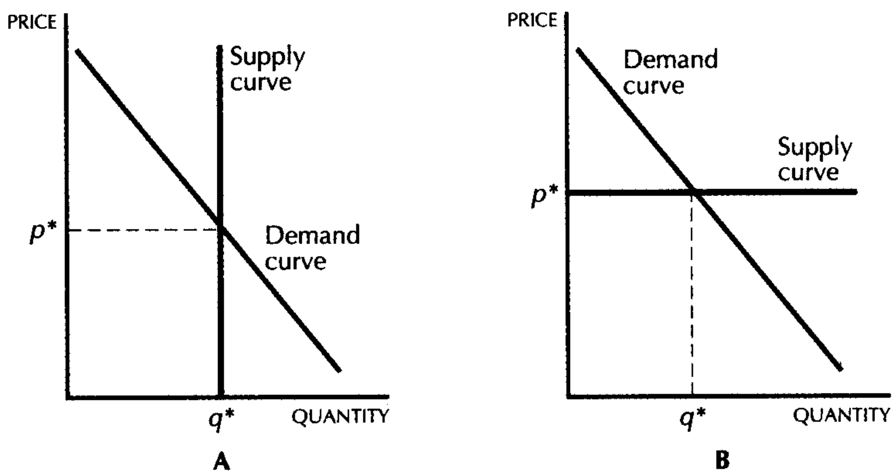

Special cases of equilibrium. Case A shows a vertical supply curve where the equilibrium price is determined solely by the demand curve. Case B depicts a horizontal supply curve where the equilibrium price is determined solely by the supply curve.

16.4 Inverse Demand and Supply Curves

We can look at market equilibrium in a slightly different way that is often useful. As indicated earlier, individual demand curves are normally viewed as giving the optimal quantities demanded as a function of the price charged. But we can also view them as inverse demand functions that measure the *price* that someone is willing to pay in order to acquire some given amount of a good. The same thing holds for supply curves. They can be viewed as measuring the quantity supplied as a function of the price. But we can also view them as measuring the *price* that must prevail in order to generate a given amount of supply.

These same constructions can be used with *market* demand and *market* supply curves, and the interpretations are just those given above. In this framework an equilibrium price is determined by finding that quantity at which the amount the demanders are willing to pay to consume that quantity is the same as the price that suppliers must receive in order to supply that quantity.

Thus, if we let $P_S(q)$ be the inverse supply function and $P_D(q)$ be the inverse demand function, equilibrium is determined by the condition

$$P_S(q^*) = P_D(q^*).$$

EXAMPLE: Equilibrium with Linear Curves

Suppose that both the demand and the supply curves are linear:

$$D(p) = a - bp$$

$$S(p) = c + dp.$$

The coefficients (a, b, c, d) are the parameters that determine the intercepts and slopes of these linear curves. The equilibrium price can be found by solving the following equation:

$$D(p) = a - bp = c + dp = S(p).$$

The answer is

$$p^* = \frac{a-c}{d+b}.$$

The equilibrium quantity demanded (and supplied) is

$$\begin{aligned} D(p^*) &= a - bp^* \\ &= a - b\frac{a-c}{b+d} \\ &= \frac{ad+bc}{b+d}. \end{aligned}$$

We can also solve this problem by using the inverse demand and supply curves. First we need to find the inverse demand curve. At what price is some quantity q demanded? Simply substitute q for $D(p)$ and solve for p. We have

$$q = a - bp,$$

so

$$P_D(q) = \frac{a-q}{b}.$$

In the same manner we find

$$P_S(q) = \frac{q-c}{d}.$$

Setting the demand price equal to the supply price and solving for the equilibrium quantity we have

$$P_D(q) = \frac{a-q}{b} = \frac{q-c}{d} = P_S(q)$$

$$q^* = \frac{ad+bc}{b+d}.$$

Note that this gives the same answer as in the original problem for both the equilibrium price and the equilibrium quantity.

16.5 Comparative Statics

After we have found an equilibrium by using the demand equals supply condition (or the demand price equals the supply price condition), we can see how it will change as the demand and supply curves change. For example, it is easy to see that if the demand curve shifts to the right in a parallel way—some fixed amount more is demanded at every price—the equilibrium price and quantity must both rise. On the other hand, if the supply curve shifts to the right, the equilibrium quantity rises, but the equilibrium price must fall.

What if both curves shift to the right? Then the quantity will definitely increase while the change in price is ambiguous—it could increase or it could decrease.

EXAMPLE: Shifting Both Curves

Question: Consider the competitive market for apartments described in Chapter 1. Let the equilibrium price in that market be p^* and the equilibrium quantity be q^*. Suppose that a developer converts m of the apartments to condominiums, which are bought by the people who are currently living in the apartments. What happens to the equilibrium price?

Answer: The situation is depicted in Figure 16.2. The demand and supply curves both shift to the left by the *same* amount. Hence the price is unchanged and the quantity sold simply drops by m.

Algebraically the new equilibrium price is determined by

$$D(p) - m = S(p) - m,$$

which clearly has the same solution as the original demand equals supply condition.

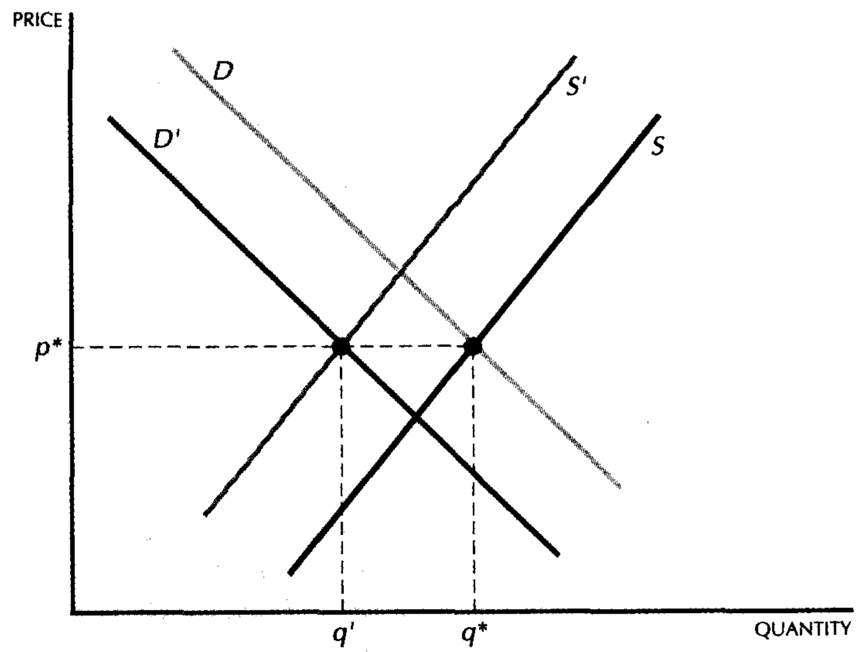

Shifting both curves. Both demand and supply curves shift to the left by the same amount, which implies the equilibrium price will remain unchanged.

16.6 Taxes

Describing a market before and after taxes are imposed presents a very nice exercise in comparative statics, as well as being of considerable interest in the conduct of economic policy. Let us see how it is done.

The fundamental thing to understand about taxes is that when a tax is present in a market, there are *two* prices of interest: the price the demander pays and the price the supplier gets. These two prices—the demand price and the supply price—differ by the amount of the tax.

There are several different kinds of taxes that one might impose. Two examples we will consider here are **quantity taxes** and **value** taxes (also called **ad valorem** taxes).

A quantity tax is a tax levied per unit of quantity bought or sold. Gasoline taxes are a good example of this. The gasoline tax is roughly 12 cents a gallon. If the demander is paying $P_D = \$1.50$ per gallon of gasoline, the supplier is getting $P_S = \$1.50 - .12 = \$1.38$ per gallon. In general, if t is the amount of the quantity tax per unit sold, then

$$P_D = P_S + t.$$

A value tax is a tax expressed in percentage units. State sales taxes are the most common example of value taxes. If your state has a 5 percent sales tax, then when you pay \$1.05 for something (including the tax), the supplier gets \$1.00. In general, if the tax rate is given by τ, then

$$P_D = (1 + \tau)P_S.$$

Let us consider what happens in a market when a quantity tax is imposed. For our first case we suppose that the supplier is required to pay the tax, as in the case of the gasoline tax. Then the amount supplied will depend on the supply price—the amount the supplier actually gets after paying the tax—and the amount demanded will depend on the demand price—the amount that the demander pays. The amount that the supplier gets will be the amount the demander pays minus the amount of the tax. This gives us two equations:

$$D(P_D) = S(P_S)$$

$$P_S = P_D - t.$$

Substituting the second equation into the first, we have the equilibrium condition:

$$D(P_D) = S(P_D - t).$$

Alternatively we could also rearrange the second equation to get $P_D = P_S + t$ and then substitute to find

$$D(P_S + t) = S(P_S).$$

Either way is equally valid; which one you use will depends on convenience in a particular case.

Now suppose that instead of the supplier paying the tax, the demander has to pay the tax. Then we write

$$P_D - t = P_S,$$

which says that the amount paid by the demander minus the tax equals the price received by the supplier. Substituting this into the demand equals supply condition we find

$$D(P_D) = S(P_D - t).$$

Note that this is the same equation as in the case where the supplier pays the tax. As far as the equilibrium price facing the demanders and the suppliers is concerned, it really doesn't matter who is responsible for paying the tax—it just matters that the tax must be paid by someone.

This really isn't so mysterious. Think of the gasoline tax. There the tax is included in the posted price. But if the price were instead listed as the before-tax price and the gasoline tax were added on as a separate item to be paid by the demanders, then do you think that the amount of gasoline demanded would change? After all, the final price to the consumers would be the same whichever way the tax was charged. Insofar as the consumers can recognize the net cost to them of goods they purchase, it really doesn't matter which way the tax is levied.

There is an even simpler way to show this using the inverse demand and supply functions. The equilibrium quantity traded is that quantity q^* such that the demand price at q^* *minus the tax being paid* is just equal to the supply price at q^*. In symbols:

$$P_D(q^*) - t = P_S(q^*).$$

If the tax is being imposed on the suppliers, then the condition is that the supply price *plus the amount of the tax* must equal the demand price:

$$P_D(q^*) = P_S(q^*) + t.$$

But these are the same equations, so the same equilibrium prices and quantities must result.

Finally, we consider the geometry of the situation. This is most easily seen by using the inverse demand and supply curves discussed above. We want to find the quantity where the curve $P_D(q) - t$ crosses the curve $P_S(q)$. In order to locate this point we simply shift the demand curve down by t and see where this shifted demand curve intersects the original supply curve. Alternatively we can find the quantity where $P_D(q)$ equals $P_S(q) + t$. To do this, we simply shift the supply curve up by the amount of the tax. Either way gives us the correct answer for the equilibrium quantity. The picture is given in Figure 16.3.

From this diagram we can easily see the qualitative effects of the tax. The quantity sold must decrease, the price paid by the demanders must go up, and the price received by the suppliers must go down.

Figure 16.4 depicts another way to determine the impact of a tax. Think about the definition of equilibrium in this market. We want to find a quantity q^* such that when the supplier faces the price p_s and the demander faces the price $p_d = p_s + t$, the quantity q^* is demanded by the demander and supplied by the supplier. Let us represent the tax t by a vertical line segment and slide it along the supply curve until it just touches the demand curve. That point is our equilibrium quantity!

EXAMPLE: Taxation with Linear Demand and Supply

Suppose that the demand and supply curves are both linear. Then if we impose a tax in this market, the equilibrium is determined by the equations

$$a - bp_D = c + dp_S$$

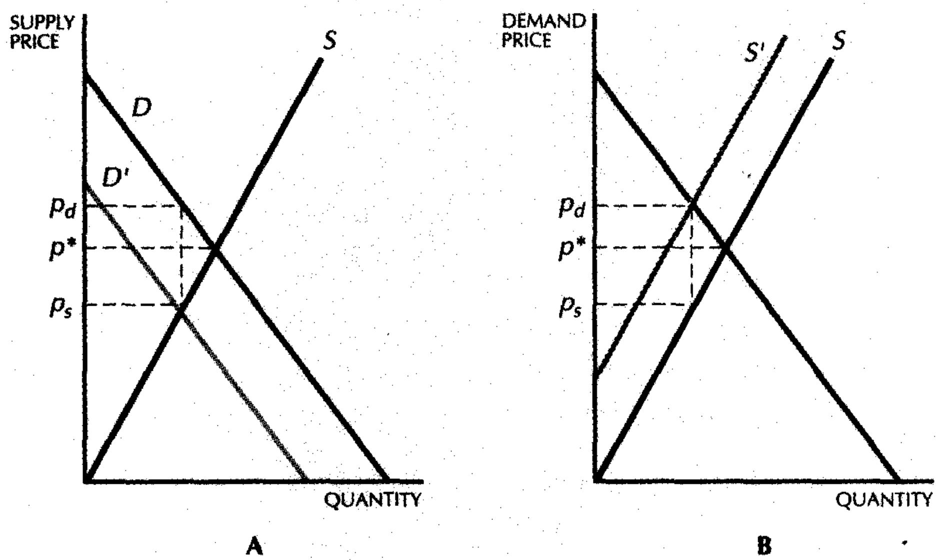

The imposition of a tax. In order to study the impact of a tax, we can either shift the demand curve down, as in panel A, or shift the supply curve up, as in panel B. The equilibrium prices paid by the demanders and received by the suppliers will be the same either way.

and

$$p_D = p_S + t.$$

Substituting from the second equation into the first, we have

$$a - b(p_S + t) = c + dp_S.$$

Solving for the equilibrium supply price, p_S^*, gives

$$p_S^* = \frac{a - c - bt}{d + b}.$$

The equilibrium demand price, p_D^*, is then given by $p_S^* + t$:

$$\begin{aligned} p_D^* &= \frac{a - c - bt}{d + b} + t \\ &= \frac{a - c + dt}{d + b}. \end{aligned}$$

Note that the price paid by the demander increases and the price received by the supplier decreases. The amount of the price change depends on the slope of the demand and supply curves.

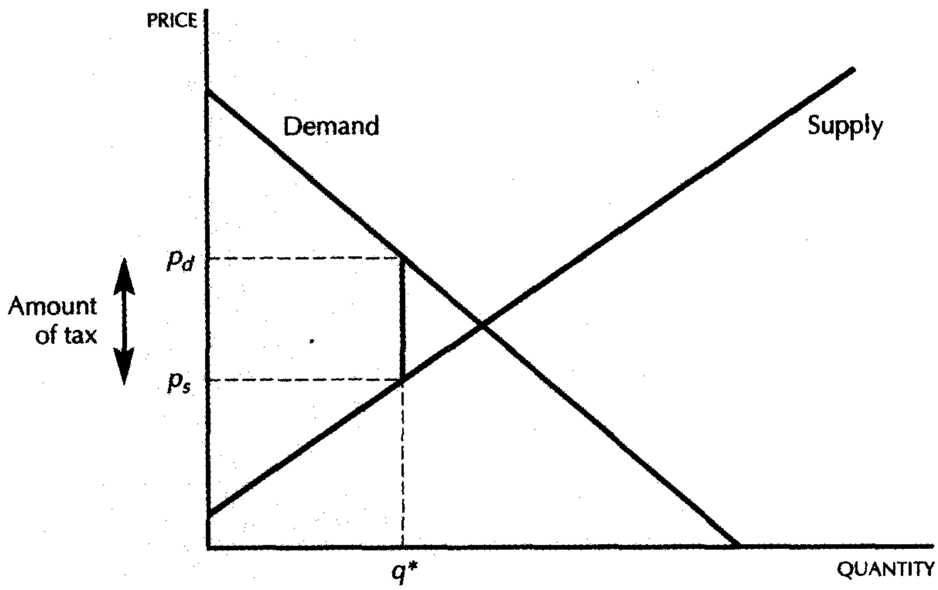

Another way to determine the impact of a tax. Slide the line segment along the supply curve until it hits the demand curve.

16.7 Passing Along a Tax

One often hears about how a tax on producers doesn't hurt profits, since firms can simply pass along a tax to consumers. As we've seen above, a tax really shouldn't be regarded as a tax on firms or on consumers. Rather, taxes are on transactions *between* firms and consumers. In general, a tax will both raise the price paid by consumers and lower the price received by firms. How much of a tax gets passed along will therefore depend on the characteristics of demand and supply.

This is easiest to see in the extreme cases: when we have a perfectly horizontal supply curve or a perfectly vertical supply curve. These are also known as the case of **perfectly elastic** and **perfectly inelastic** supply.

We've already encountered these two special cases earlier in this chapter. If an industry has a horizontal supply curve, it means that the industry will supply any amount desired of the good at some given price, and zero units of the good at any lower price. In this case the price is entirely determined by the supply curve and the quantity sold is determined by demand. If an industry has a vertical supply curve, it means that the quantity of the good is fixed. The equilibrium price of the good is determined entirely by demand.

Let's consider the imposition of a tax in a market with a perfectly elastic supply curve. As we've seen above, imposing a tax is just like shifting the

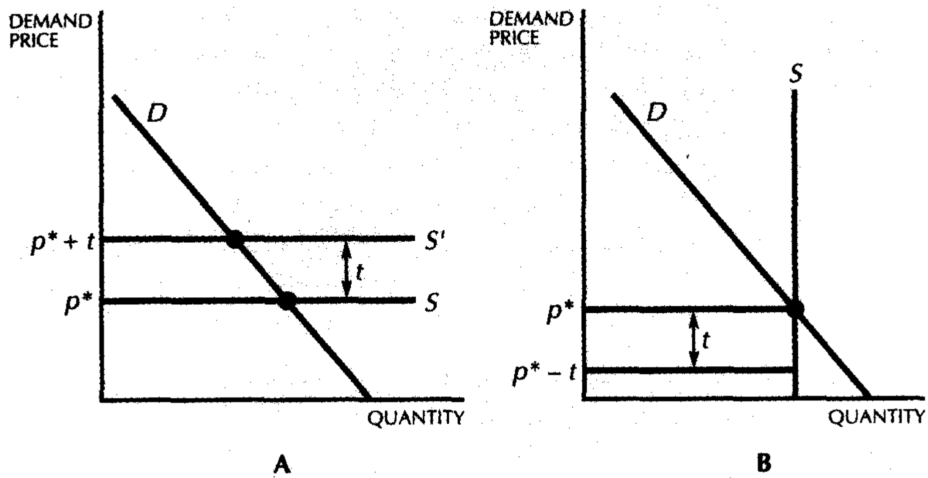

Special cases of taxation. (A) In the case of a perfectly elastic supply curve the tax gets completely passed along to the consumers. (B) In the case of a perfectly inelastic supply none of the tax gets passed along.

supply curve up by the amount of the tax, as illustrated in Figure 16.5A.

In the case of a perfectly elastic supply curve it is easy to see that the price to the consumers goes up by exactly the amount of the tax. The supply price is exactly the same as it was before the tax, and the demanders end up paying the entire tax. When you think about the meaning of the horizontal supply curve, this is not hard to understand. The horizontal supply curve means that the industry is willing to supply any amount of the good at some particular price, p^*, and zero amount at any lower price. Thus, if any amount of the good is going to be sold at all in equilibrium, the suppliers must receive p^* for selling it. This effectively determines the equilibrium supply price, and the demand price is $p^* + t$.

The opposite case is illustrated in Figure 16.5B. If the supply curve is vertical and we "shift the supply curve up," we don't change anything in the diagram. The supply curve just slides along itself, and we still have the same amount of the good supplied, with or without the tax. In this case, the demanders determine the equilibrium price of the good, and they are willing to pay a certain amount, p^*, for the supply of the good that is available, tax or no tax. Thus they end up paying p^*, and the suppliers end up receiving $p^* - t$. The entire amount of the tax is paid by the suppliers.

This case often strikes people as paradoxical, but it really isn't. If the suppliers could raise their prices after the tax is imposed and still sell their entire fixed supply, they would have raised their prices before the tax was imposed and made more money! If the demand curve doesn't move, then the only way the price can increase is if the supply is reduced. If a policy doesn't change either supply or demand, it certainly can't affect price.

Now that we understand the special cases, we can examine the in-between case where the supply curve has an upward slope but is not perfectly vertical. In this situation, the amount of the tax that gets passed along will depend on the steepness of the supply curve relative to the demand curve. If the supply curve is nearly horizontal, nearly all of the tax gets passed along to the consumers, while if the supply curve is nearly vertical, almost none of the tax gets passed along. See Figure 16.6 for some examples.

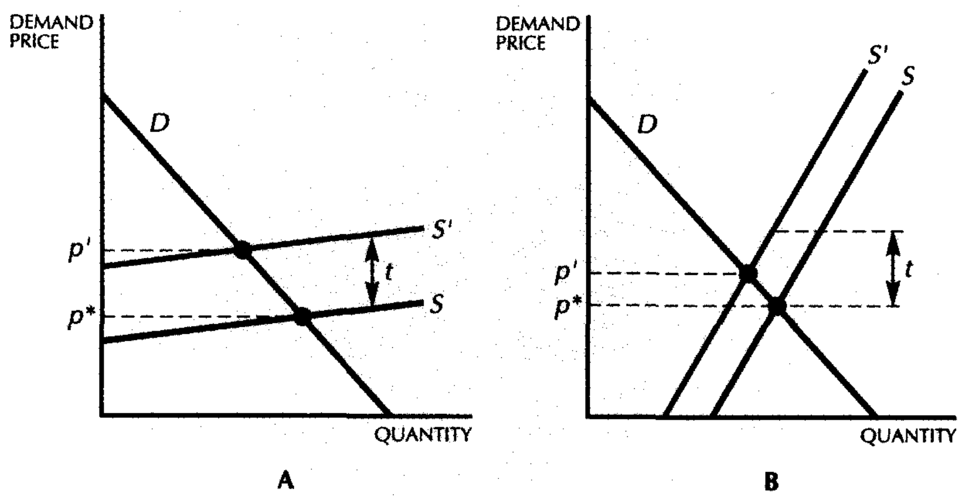

Passing along a tax. (A) If the supply curve is nearly horizontal, much of the tax can be passed along. (B) If it is nearly vertical, very little of the tax can be passed along.

16.8 The Deadweight Loss of a Tax

We've seen that taxing a good will typically increase the price paid by the demanders and decrease the price received by the suppliers. This certainly represents a cost to the demanders and suppliers, but from the economist's viewpoint, the real cost of the tax is that the output has been reduced.

The lost output is the social cost of the tax. Let us explore the social cost of a tax using the consumers' and producers' surplus tools developed in Chapter 14. We start with the diagram given in Figure 16.7. This depicts the equilibrium demand price and supply price after a tax, t, has been imposed.

Output has been decreased by this tax, and we can use the tools of consumers' and producers' surplus to value the social loss. The loss in consumers' surplus is given by the areas $A + B$, and the loss in producers' surplus is given in areas $C + D$. These are the same kind of losses that we examined in Chapter 14.

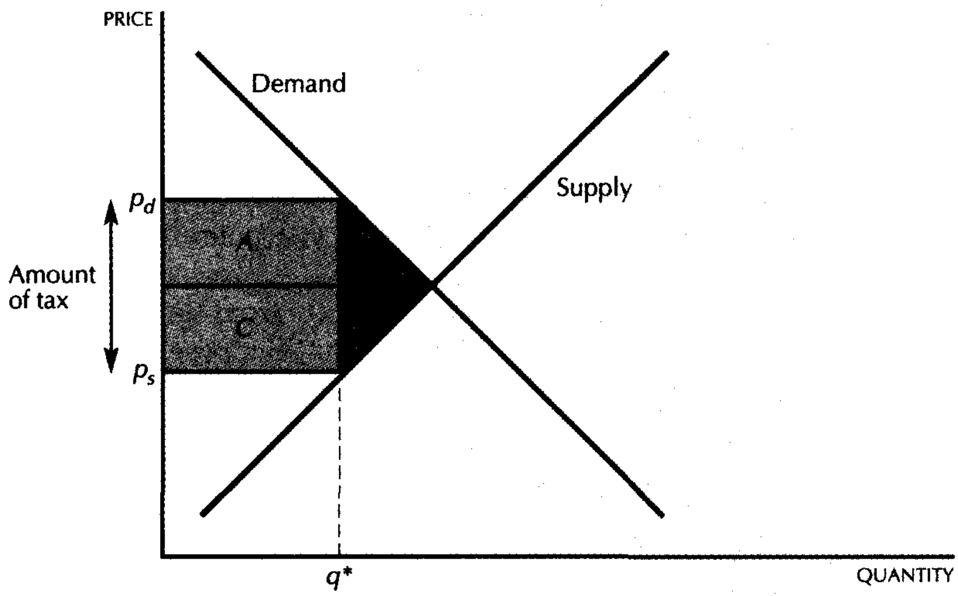

The deadweight loss of a tax. The area $B + D$ measures the deadweight loss of the tax.

Since we're after an expression for the social cost of the tax, it seems sensible to add the areas $A+B$ and $C+D$ to each other to get the total loss to the consumers and to the producers of the good in question. However, we've still left out one party—namely, the government.

The government *gains* revenue from the tax. And, of course, the consumers who benefit from the government services provided with these tax revenues also gain from the tax. We can't really say how much they gain until we know what the tax revenues will be spent on.

Let us make the assumption that the tax revenues will just be handed back to the consumers and the producers, or equivalently that the services provided by the government revenues will be just equal in value to the revenues spent on them.

Then the net benefit to the government is the area $A + C$—the total revenue from the tax. Since the loss of producers' and consumers' surpluses are net costs, and the tax revenue to the government is a net benefit, the total net cost of the tax is the algebraic sum of these areas: the loss in consumers' surplus, $-(A+B)$, the loss in producers' surplus, $-(C+D)$, and the gain in government revenue, $+(A+C)$.

The net result is the area $-(B+D)$. This area is known as the **deadweight loss** of the tax or the **excess burden** of the tax. This latter phrase is especially descriptive.

Recall the interpretation of the loss of consumers' surplus. It is how much the consumers would pay to avoid the tax. In terms of this diagram the consumers are willing to pay $A+B$ to avoid the tax. Similarly, the producers are willing to pay $C+D$ to avoid the tax. Together they are willing to pay $A+B+C+D$ to avoid a tax that raises $A+C$ dollars of revenue. The *excess burden* of the tax is therefore $B+D$.

What is the source of this excess burden? Basically it is the lost value to the consumers and producers due to the reduction in the sales of the good. You can't tax what isn't there.[1] So the government doesn't get any revenue on the reduction in sales of the good. From the viewpoint of society, it is a pure loss—a deadweight loss.

We could also derive the deadweight loss directly from its definition, by just measuring the social value of the lost output. Suppose that we start at the old equilibrium and start moving to the left. The first unit lost was one where the price that someone was willing to pay for it was just equal to the price that someone was willing to sell it for. Here there is hardly any social loss since this unit was the marginal unit that was sold.

Now move a little farther to the left. The demand price measures how much someone was willing to pay to receive the good, and the supply price measures the price at which someone was willing to supply the good. The difference is the lost value on that unit of the good. If we add this up over the units of the good that are not produced and consumed because of the presence of the tax, we get the deadweight loss.

EXAMPLE: The Market for Loans

The amount of borrowing or lending in an economy is influenced to a large degree by the interest rate charged. The interest rate serves as a price in the market for loans.

We can let $D(r)$ be the demand for loans by borrowers and $S(r)$ be the supply of loans by lenders. The equilibrium interest rate, r^*, is then determined by the condition that demand equal supply:

$$D(r^*) = S(r^*). \tag{16.1}$$

Suppose we consider adding taxes to this model. What will happen to the equilibrium interest rate?

[1] At least the government hasn't figured out how to do this yet. But they're working on it.

In the U.S. economy individuals have to pay income tax on the interest they earn from lending money. If everyone is in the same tax bracket, t, the after-tax interest rate facing lenders will be $(1-t)r$. Thus the supply of loans, which depends on the after-tax interest rate, will be $S((1-t)r)$.

On the other hand, the Internal Revenue Service code allows many borrowers to deduct their interest charges, so if the borrowers are in the same tax bracket as the lenders, the after-tax interest rate they pay will be $(1-t)r$. Hence the demand for loans will be $D((1-t)r)$. The equation for interest rate determination with taxes present is then

$$D((1-t)r') = S((1-t)r'). \tag{16.2}$$

Now observe that if r^* solves equation (16.1), then $r^* = (1-t)r'$ must solve equation (16.2) so that

$$r^* = (1-t)r',$$

or

$$r' = \frac{r^*}{(1-t)}.$$

Thus the interest rate in the presence of the tax will be higher by $1/(1-t)$. The *after-tax* interest rate $(1-t)r'$ will be r^*, just as it was before the tax was imposed!

Figure 16.8 may make things clearer. Making interest income taxable will tilt the supply curve for loans up by a factor of $1/(1-t)$; but making interest payments tax deductible will also tilt the demand curve for loans up by $1/(1-t)$. The net result is that the market interest rate rises by precisely $1/(1-t)$.

Inverse demand and supply functions provide another way to look at this problem. Let $r_b(q)$ be the inverse demand function for borrowers. This tells us what the after-tax interest rate would have to be to induce people to borrow q. Similarly, let $r_l(q)$ be the inverse supply function for lenders. The equilibrium amount lent will then be determined by the condition

$$r_b(q^*) = r_l(q^*). \tag{16.3}$$

Now introduce taxes into the situation. To make things more interesting, we'll allow borrowers and lenders to be in different tax brackets, denoted by t_b and t_l. If the market interest rate is r, then the after-tax rate facing borrowers will be $(1-t_b)r$, and the quantity they choose to borrow will be determined by the equation

$$(1-t_b)r = r_b(q)$$

or

$$r = \frac{r_b(q)}{1-t_b}. \tag{16.4}$$

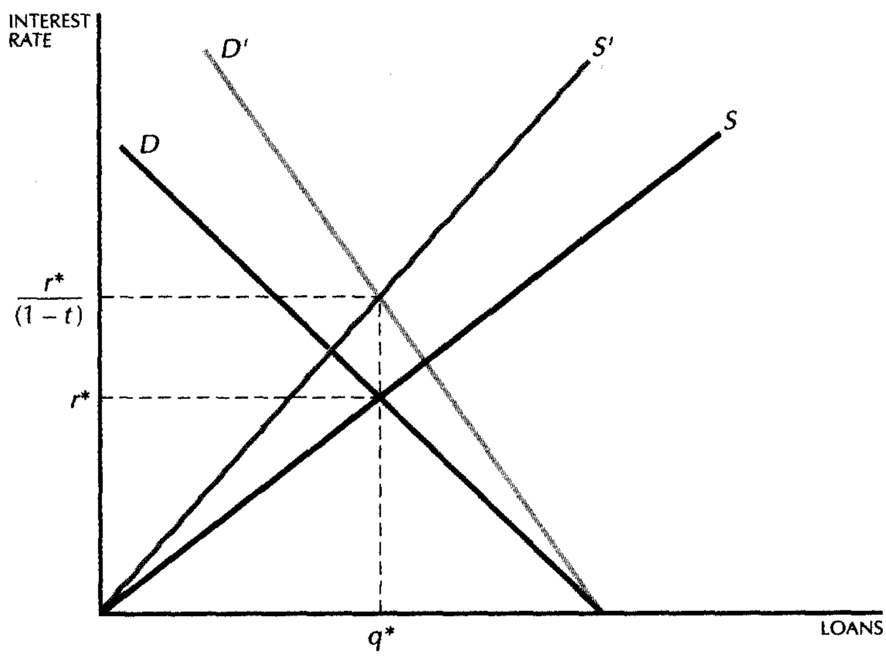

Equilibrium in the loan market. If borrowers and lenders are in the same tax bracket, the after-tax interest rate and the amount borrowed are unchanged.

Similarly, the after-tax rate facing lenders will be $(1 - t_l)r$, and the amount they choose to lend will be determined by the equation

$$(1 - t_l)r = r_l(q)$$

or

$$r = \frac{r_l(q)}{1 - t_l}. \tag{16.5}$$

Combining equations (16.4) and (16.5) gives the equilibrium condition:

$$r = \frac{r_b(\hat{q})}{1 - t_b} = \frac{r_l(\hat{q})}{1 - t_l}. \tag{16.6}$$

From this equation it is easy to see that if borrowers and lenders are in the same tax bracket, so that $t_b = t_l$, then $\hat{q} = q^*$. What if they are in different tax brackets? It is not hard to see that the tax law is subsidizing borrowers and taxing lenders, but what is the net effect? If the borrowers face a higher price than the lenders, then the system is a net tax on borrowing, but if the borrowers face a lower price than the lenders, then it is a net subsidy. Rewriting the equilibrium condition, equation (16.6), we have

$$r_b(\hat{q}) = \frac{1 - t_b}{1 - t_l} r_l(\hat{q}).$$

Thus borrowers will face a higher price than lenders if

$$\frac{1-t_b}{1-t_l} > 1,$$

which means that $t_l > t_b$. So if the tax bracket of lenders is greater than the tax bracket of borrowers, the system is a net tax on borrowing, but if $t_l < t_b$, it is a net subsidy.

EXAMPLE: Food Subsidies

In years when there were bad harvests in nineteenth-century England the rich would provide charitable assistance to the poor by buying up the harvest, *consuming a fixed amount of the grain, and selling the remainder to* the poor at half the price they paid for it. At first thought this seems like it would provide significant benefits to the poor, but on second thought, doubts begin to arise.

The only way that the poor can be made better off is if they end up consuming more grain. But there is a fixed amount of grain available after the harvest. So how can the poor be better off because of this policy?

As a matter of fact they are not; the poor end up paying exactly the same price for the grain with or without the policy. To see why, we will model the equilibrium with and without this program. Let $D(p)$ be the demand curve for the poor, K the amount demanded by the rich, and S the fixed amount supplied in a year with a bad harvest. By assumption the supply of grain and the demand by the rich are fixed. Without the charity provided by the rich, the equilibrium price is determined by total demand equals total supply:

$$D(p^*) + K = S.$$

With the program in place, the equilibrium price is determined by

$$D(\hat{p}/2) + K = S.$$

But now observe: if p^* solves the first equation, then $\hat{p} = 2p^*$ solves the second equation. So when the rich offer to buy the grain and distribute it to the poor, the market price is simply bid up to twice the original price—and the poor pay the same price they did before!

When you think about it this isn't too surprising. If the demand of the rich is fixed and the supply of grain is fixed, then the amount that the poor can consume is fixed. Thus the equilibrium price facing the poor is determined entirely by their own demand curve; the equilibrium price will be the same, regardless of how the grain is provided to the poor.

EXAMPLE: Subsidies in Iraq

Even subsidies that are put in place "for a good reason" can be extremely difficult to dislodge. Why? Because they create a political constituency that comes to rely on them. This is true in every country, but Iraq represents a particularly egregious case. As of 2005, fuel and food subsidies in Iraq consumed nearly one third of the government's budget.[2]

Almost all of the Iraqi government's budget comes from oil exports. There is very little refining capacity in the country, so Iraq imports gasoline at 30 to 35 cents a liter, which it then sells to the public at 1.5 cents. A substantial amount of this gasoline is sold on the black market and smuggled into Turkey, where gas is about one dollar a liter.

Food and fuel oil are also highly subsidized. Politicians are reluctant to remove these subsidies due to the politically unstable environment. When similar subsidies were removed in Yemen, there was rioting in the streets, with dozens of people dying. A World Bank study concluded that more than half of the GDP in Iraq was spent on subsidies. According to the finance minister, Ali Abdulameer Allawi, "They've reached the point where they've become insane. They distort the economy in a grotesque way, and create the worst incentives you can think of."

16.9 Pareto Efficiency

An economic situation is **Pareto efficient** if there is no way to make any person better off without hurting anybody else. Pareto efficiency is a desirable thing—if there is some way to make some group of people better off, why not do it?—but efficiency is not the only goal of economic policy. For example, efficiency has almost nothing to say about income distribution or economic justice.

However, efficiency is an important goal, and it is worth asking how well a competitive market does in achieving Pareto efficiency. A competitive market, or any economic mechanism, has to determine two things. First, how much is produced, and second, who gets it. A competitive market determines how much is produced based on how much people are willing to pay to purchase the good as compared to how much people must be paid to supply the good.

Consider Figure 16.9. At any amount of output less than the competitive amount q^*, there is someone who is willing to supply an extra unit of the

[2] James Glanz, "Despite Crushing Costs, Iraqi Cabinet Lets Big Subsidies Stand," *New York Times*, August 11, 2005.

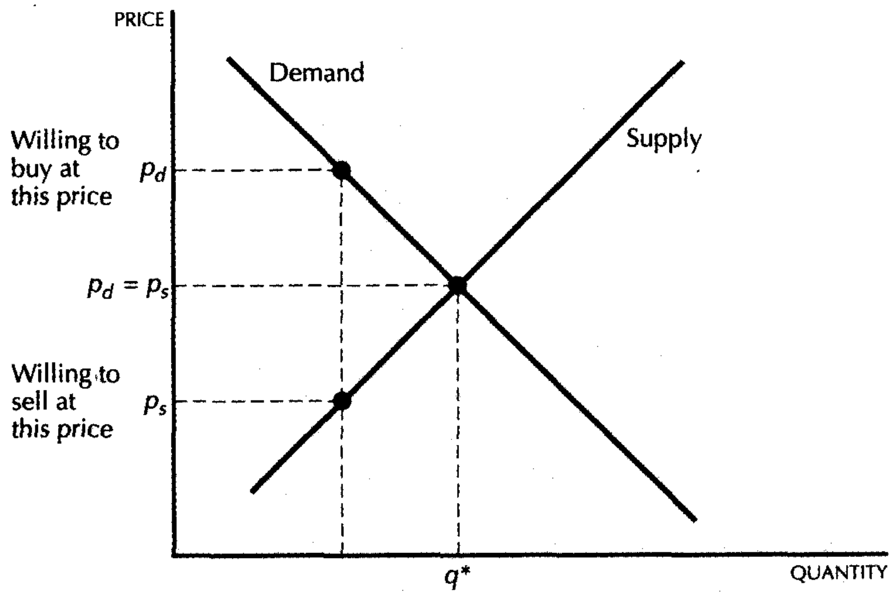

Pareto efficiency. The competitive market determines a Pareto efficient amount of output because at q^* the price that someone is willing to pay to buy an extra unit of the good is equal to the price that someone must be paid to sell an extra unit of the good.

good at a price that is less than the price that someone is willing to pay for an extra unit of the good.

If the good were produced and exchanged between these two people at any price between the demand price and the supply price, they would both be made better off. Thus any amount less than the equilibrium amount cannot be Pareto efficient, since there will be at least two people who could be made better off.

Similarly, at any output larger than q^*, the amount someone would be willing to pay for an extra unit of the good is less than the price that it would take to get it supplied. Only at the market equilibrium q^* would we have a Pareto efficient amount of output supplied—an amount such that the willingness to pay for an extra unit is just equal to the willingness to be paid to supply an extra unit.

Thus the competitive market produces a Pareto efficient amount of output. What about the way in which the good is allocated among the consumers? In a competitive market everyone pays the same price for a good—the marginal rate of substitution between the good and "all other goods" is equal to the price of the good. Everyone who is willing to pay this price is able to purchase the good, and everyone who is not willing to pay this price cannot purchase the good.

What would happen if there were an allocation of the good where the marginal rates of substitution between the good and "all other goods" were not the same? Then there must be at least two people who value a marginal unit of the good differently. Maybe one values a marginal unit at \$5 and one values it at \$4. Then if the one with the lower value sells a bit of the good to the one with the higher value at any price between \$4 and \$5, both people would be made better off. Thus any allocation with different marginal rates of substitution cannot be Pareto efficient.

EXAMPLE: Waiting in Line

One commonly used way to allocate resources is by making people wait in line. We can analyze this mechanism for resource allocation using the same tools that we have developed for analyzing the market mechanism. Let us look at a concrete example: suppose that your university is going to distribute tickets to the championship basketball game. Each person who waits in line can get one ticket for free.

The cost of a ticket will then simply be the cost of waiting in line. People who want to see the basketball game very much will camp out outside the ticket office so as to be sure to get a ticket. People who don't care very much about the game may drop by a few minutes before the ticket window opens on the off chance that some tickets will be left. The willingness to pay for a ticket should no longer be measured in dollars but rather in waiting time, since tickets will be allocated according to willingness to wait.

Will waiting in line result in a Pareto efficient allocation of tickets? Ask yourself whether it is possible that someone who waited for a ticket might be willing to sell it to someone who didn't wait in line. Often this will be the case, simply because willingness to wait and willingness to pay differ across the population. If someone is willing to wait in line to buy a ticket and then sell it to someone else, allocating tickets by willingness to wait does not exhaust all the gains to trade—some people would generally still be willing to trade the tickets after the tickets have been allocated. Since waiting in line does not exhaust all of the gains from trade, it does not in general result in a Pareto efficient outcome.

If you allocate a good using a price set in dollars, then the dollars paid by the demanders provide benefits to the suppliers of the good. If you allocate a good using waiting time, the hours spent in line don't benefit anybody. The waiting time imposes a cost on the buyers of the good and provide no benefits at all to the suppliers. Waiting in line is a form of **deadweight loss**—the people who wait in line pay a "price" but no one else receives any benefits from the price they pay.

Summary

1. The supply curve measures how much people will be willing to supply of some good at each price.

2. An equilibrium price is one where the quantity that people are willing to supply equals the quantity that people are willing to demand.

3. The study of how the equilibrium price and quantity change when the underlying demand and supply curves change is another example of comparative statics.

4. When a good is taxed, there will always be two prices: the price paid by the demanders and the price received by the suppliers. The difference between the two represents the amount of the tax.

5. How much of a tax gets passed along to consumers depends on the relative steepness of the demand and supply curves. If the supply curve is horizontal, all of the tax gets passed along to consumers; if the supply curve is vertical, none of the tax gets passed along.

6. The deadweight loss of a tax is the net loss in consumers' surplus plus producers' surplus that arises from imposing the tax. It measures the value of the output that is not sold due to the presence of the tax.

7. A situation is Pareto efficient if there is no way to make some group of people better off without making some other group worse off.

8. The Pareto efficient amount of output to supply in a single market is that amount where the demand and supply curves cross, since this is the only point where the amount that demanders are willing to pay for an extra unit of output equals the price at which suppliers are willing to supply an extra unit of output.

REVIEW QUESTIONS

1. What is the effect of a subsidy in a market with a horizontal supply curve? With a vertical supply curve?

2. Suppose that the demand curve is vertical while the supply curve slopes upward. If a tax is imposed in this market who ends up paying it?

3. Suppose that all consumers view red pencils and blue pencils as perfect substitutes. Suppose that the supply curve for red pencils is upward sloping. Let the price of red pencils and blue pencils be p_r and p_b. What would happen if the government put a tax only on red pencils?

4. The United States imports about half of its petroleum needs. Suppose that the rest of the oil producers are willing to supply as much oil as the United States wants at a constant price of $25 a barrel. What would happen to the price of domestic oil if a tax of $5 a barrel were placed on foreign oil?

5. Suppose that the supply curve is vertical. What is the deadweight loss of a tax in this market?

6. Consider the tax treatment of borrowing and lending described in the text. How much revenue does this tax system raise if borrowers and lenders are in the same tax bracket?

7. Does such a tax system raise a positive or negative amount of revenue when $t_l < t_b$?