ECONOMICS of the PUBLIC SECTOR

THIRD EDITION

JOSEPH E. STIGLITZ

W.W. NORTON & COMPANY
NEW YORK / LONDON

19 Taxation and Economic Efficiency

FOCUS QUESTIONS

1 How is the efficiency loss associated with taxation measured? On what does its magnitude depend?

2 What is meant by the income effect and substitution effects of a tax? Why do they normally reinforce each other for taxes on commodities, but work against each other for taxes on wages and interest?

3 How large are the efficiency losses associated with taxes on labor and savings?

All taxes affect economic behavior. They transfer resources from individuals to the government. As a result, individuals must alter their behavior in some way. If they do not adjust the amount of work they do, they must reduce their consumption. They may work more, enjoying less leisure; by working more, they need reduce their consumption less.

No matter how individuals adjust, an increase in taxes must make them worse off.[1] But some taxes reduce individuals' welfare less, for each dollar of revenue raised, than do other taxes. Tax policy is concerned with designing tax structures which minimize welfare loss for any given amount of revenue raised—while still attaining the other objectives of tax policy discussed in Chapter 17. This chapter analyzes the determinants of welfare loss; Chapter 20 then uses the results to describe the basic principles of optimal taxation.

This chapter is divided into six sections. The first analyzes the effects of a tax on a consumption good, such as beer. After describing the effects qualitatively, the second section shows how the distortions can be quantified. The third section analyzes inefficiencies associated with taxes on producers. The fourth and fifth sections show how the same principles may be applied to taxes on the return to savings and on wages. The final section discusses various attempts to quantify the effects of taxation on labor supply.

[1] This ignores, of course, the benefits that may accrue from the increased government expenditures that result from the increased taxes. In a sense, this chapter looks at the "costs" of government programs, that are associated with the taxes to finance them, while earlier chapters in the book looked at the benefits. An overall assessment requires balancing the two. Throughout this chapter we also ignore general equilibrium effects: before-tax wages and prices will be assumed to be unaffected by the imposition of a tax.

EFFECT OF TAXES BORNE BY CONSUMERS

We begin the analysis with the simplest case, that of a tax borne fully by consumers. Assume that an individual's income is fixed, and he can choose between purchasing two commodities, soda and beer. His budget constraint is the line *SB* in Figure 19.1. This gives the various combinations of soda and beer that the individual can purchase. If he spent all his income on soda, he could purchase the amount *S*; if he spent all his income on beer, he could purchase the amount *B*.

Suppose that the government imposes a tax on beer. What will be the effect? (Throughout this section, we will assume that the consumer price rises by the full amount of the tax; that is, consumers bear the full burden of

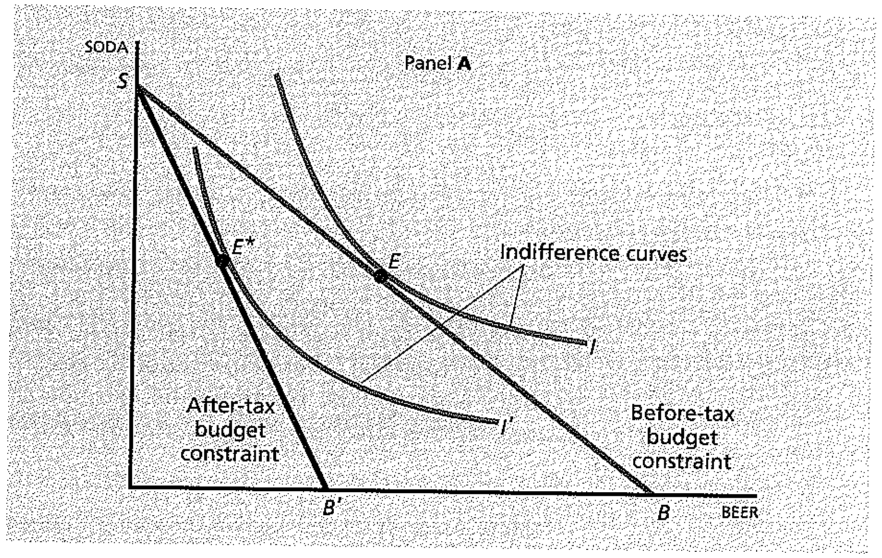

FIGURE 19.1 **Equilibrium After the Imposition of a Tax on Beer** The effect of the tax is to shift the budget constraint down and, thus, the equilibrium changes from *E* to *E**.

the tax. This will happen if the supply curves for beer and soda are infinitely elastic, as we showed in Chapter 18.) The tax on beer shifts the budget constraint in to SB'. The individual can still, if he wishes, spend all his income on soda, in which case he obtains S units of soda. But beer is now more expensive, so he can purchase less of it with his income.

Initially, the individual allocated his income by choosing point E on this budget constraint. This is the point of tangency between the budget constraint and the indifference curve. After the imposition of the tax, there is a new equilibrium, at point E*. At E^*, the individual consumes less beer than at E.

SUBSTITUTION AND INCOME EFFECTS

The tax decreases the individual's consumption of beer, for two reasons. First, the tax—like any tax or loss of income—makes him worse off, by leaving him with less money to spend. Normally, when an individual is worse off, he consumes less of all goods. The amount by which his consumption of the taxed good is reduced because he is worse off is called the **income effect** of the tax. Second, the tax makes beer more expensive than other goods. When a good becomes relatively more expensive, individuals find substitutes for it. The extent to which consumption of the taxed good is reduced because of the increased *relative price* is the **substitution effect.**

Figure 19.2 shows how to decompose the movement from E to E^*—the reduction in beer consumption—into income and substitution effects. We first ask, how would consumption of beer have been reduced if we had taken away income from the individual, to put him on the new, lower indifference curve, but had not at the same time changed relative prices? This change is reflected in the budget constraint $\hat{S}\hat{B}$, which is parallel to the original budget constraint (implying the same prices), but tangent to the indifference curve I', at $\hat{E}$. The corresponding reduction in beer consumption is the income effect.

The movement from $\hat{E}$ to E^*, and the corresponding reduction in beer consumption, is the substitution effect. It represents the reduction in consumption *due solely to changes in relative prices.*

Income and substitution effects work in the same direction in the case of a beer tax: beer consumption drops continually as we move from E to $\hat{E}$ to E^*.

DETERMINING THE SIZE OF THE SUBSTITUTION EFFECT The magnitude of the substitution effect depends on how easy it is to substitute other goods for the taxed good. This is reflected in the shape of the indifference curves. If they are relatively flat, then substitution is easy, and the substitution effect is large.[2] Figure 19.2B illustrates the extreme case where indifference curves are L-shaped and there is no substitution effect.

[2] More precisely, it depends on the *elasticity of substitution,* which is defined as the percentage change in relative quantities consumed from a percentage change in relative prices. The L-shaped indifference curve of Figure 19.2B has a zero elasticity of substitution. The other extreme case is a straight-line indifference curve, in which case the elasticity of substitution is said to be infinite.

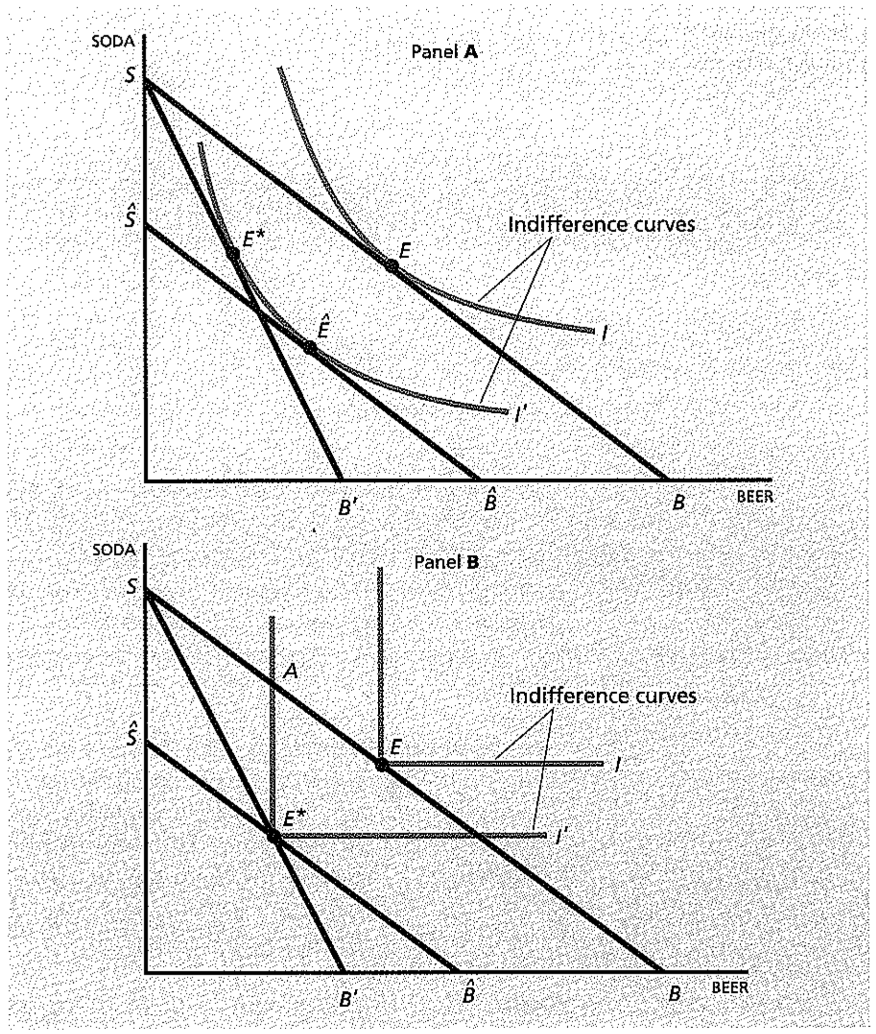

FIGURE 19.2 **Income and Substitution Effects of a Tax on Beer Consumption** Panel A decomposes the movement from E to E^* into income and substitution effects. The movement from E to $\hat{E}$ is the income effect, and the movement from $\hat{E}$ to E^* is the substituton effect. Panel B, where indifference curves are L-shaped, represents the case in which there is no substitution effect.

QUANTIFYING THE DISTORTIONS

Any tax must have effects on consumption. After all, the purpose of a tax is to transfer purchasing power from the individual to the government. Individuals have to reduce their consumption of something. An efficient tax minimizes the welfare loss per unit revenue raised. Chapter 17 introduced the concept of a lump-sum tax, a tax which the individual must pay regardless on what he does. Such a tax simply moves the budget constraint in a parallel

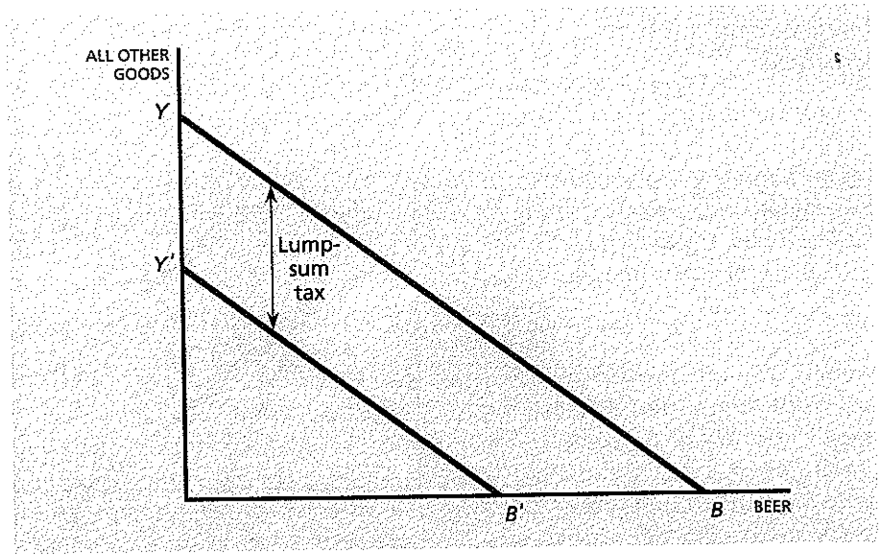

FIGURE 19.3 **Lump-Sum Tax** The vertical distance between the two budget constraints measures the magnitude of the lump-sum tax.

way, as illustrated in Figure 19.3. In the figure, we have put expenditures on beer on the horizontal axis and expenditures on all other goods on the vertical axis. Thus, point *Y*, where the individual consumes no beer, measures his income before tax; and point *Y'* measures his income after tax; the vertical distance, *YY'*, measures the lump-sum tax. The budget constraint is

Expenditures on beer + expenditures on all other goods = income − lump-sum taxes,

where expenditures on beer = $p_B B$, the price of beer times the quantity of beer purchased.

We compare the effect of any tax—such as a tax on beer—with the effect of a lump-sum tax by asking: For the same revenue, how much worse off are individuals with the tax on beer than they would have been with the lump-sum tax? The extra loss in welfare is called the **deadweight loss.** Equivalently, we can ask: For the same effect on individual welfare, how much *extra* revenue would a lump-sum tax have raised, or how much *less* revenue does the beer tax raise? The difference in revenue is how we measure the deadweight loss of the tax.

MEASURING DEADWEIGHT LOSS USING INDIFFERENCE CURVES

Figure 19.4 contrasts the effect of a tax on beer with a lump-sum tax. The beer tax rotates the individual's budget constraint down, from *YB* to *YB'*. The income raised by the tax is the vertical distance between the before-tax budget constraint and the after-tax budget constraint. Clearly, when no beer is consumed (point *Y*), no revenue is raised. The more

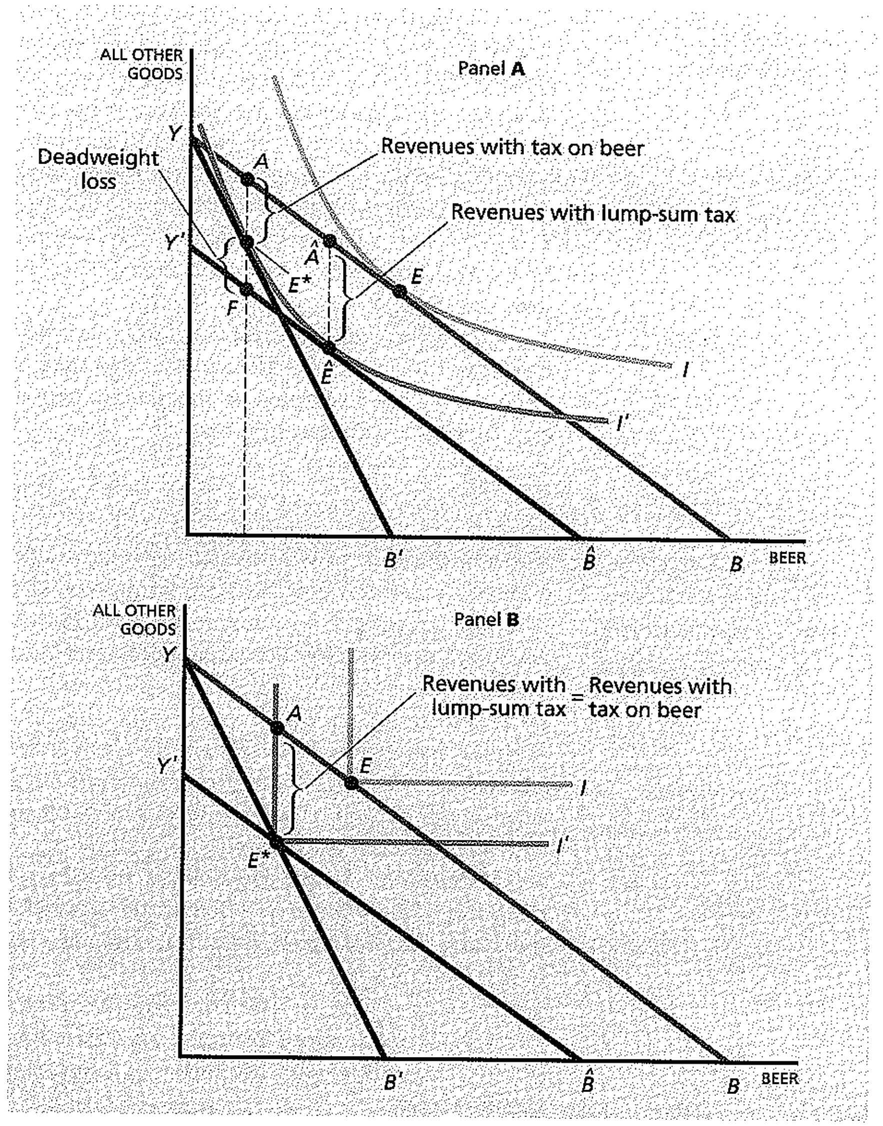

FIGURE 19.4 **Measuring the Deadweight Loss Using Indifference Curves** Individuals choose the amount of beer to consume by the tangency between their indifference curve and the budget constraint. The beer tax rotates the budget constraint. The lump-sum tax moves the budget constraint down parallel. Panel A: The extra revenue raised by the lump-sum tax is *E*F*. Panel B: When there is no substitution effect, the beer tax has no deadweight loss; a lump-sum tax and a tax on beer raise the same revenue.

beer that is consumed, the greater the tax revenue. The revenue raised is *AE*.*

The lump-sum tax with the same effect on utility moves the budget constraint from YB to $Y'\hat{B}$ and the equilibrium is now $\hat{E}$. The revenue raised is again the vertical difference between the new and the old budget constraints—this represents the amount of income that had to be taken away to leave the individual on the same indifference curve. Since the new and old budget constraints are parallel, that vertical distance $\hat{A}\hat{E}$ is exactly equal to AF. (The vertical distance between parallel lines is the same at any location.) Thus, *the lump-sum tax with the same effect on utility raises an additional revenue in the amount of E^*F.* E^*F is the measure of the deadweight loss associated with the tax.

The *magnitude of the deadweight loss depends on the substitution effect.* This is illustrated in Figure 19.4B, which is identical to 19.4A, except now indifference curves are L-shaped, so there is no substitution effect, and, it is apparent, there is no deadweight loss.

MEASURING DEADWEIGHT LOSS USING COMPENSATED DEMAND CURVES

There is another way of measuring deadweight loss, making use of the concepts of consumer surplus and compensated demand curves introduced in Chapter 5. Assume we have imposed a tax of 30 cents per bottle of beer, and, with the tax, the individual consumes ten bottles a week. We ask the individual how much he would be willing to give to the government if the tax were eliminated. In other words, what lump-sum tax would leave him at the same utility level he reached when he was subject to the 30-cent tax on beer? Clearly, he would be willing to pay at least 30 cents × 10 per week. Any extra revenue that such a tax would generate is the deadweight loss associated with the use of a distortionary tax system.

We now show how to calculate the deadweight loss using a consumer's *compensated demand curve.* The compensated demand curve gives the individual's demand for beer, assuming that as the price is lowered, income is being taken away from him in such a way as to leave him on the same indifference curve. We use the compensated demand curve because we wish to know how much more revenue we could have achieved with a nondistortionary tax, still leaving the individual just as well off as he was with the distortionary tax.

Assume initially the price of a bottle of beer is $1.50, including the 30-cent tax, and the individual consumes ten bottles a week. We then ask him how much extra he would be willing to pay to consume eleven bottles a week. He is willing to pay only $1.40. The total amount that the individual would be willing to pay us as a lump-sum tax if we lowered the tax from 30 cents to 20 cents (and lowered the price of beer from $1.50 to $1.40) is 10 cents × the 10 bottles he previously purchased, or $1.00 (the area *FGCD* in Figure 19.5A).

We now ask him to assume he is in a situation where we levied a $1.00 lump-sum tax and charged $1.40 each for eleven bottles of beer. How much *extra* would he be willing to pay for one extra bottle? Assume the individual said $1.30. We can now calculate the total lump-sum tax that an individual

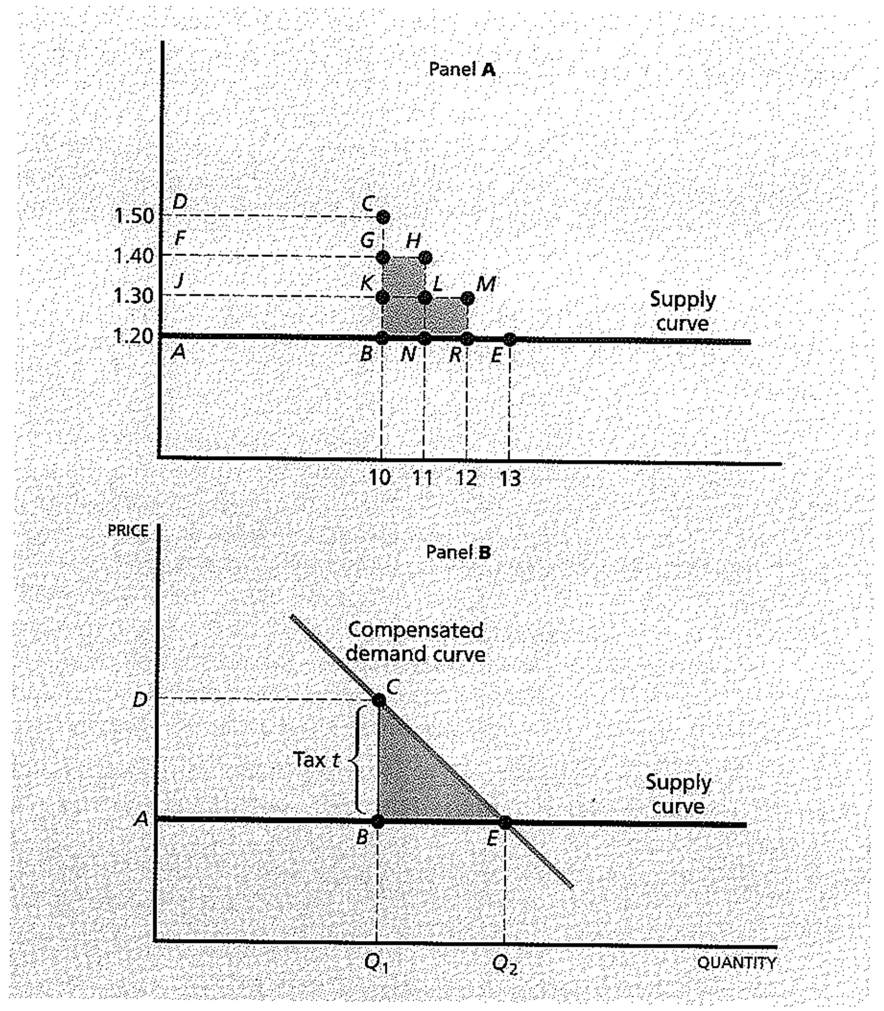

FIGURE 19.5 **Using Compensated Demand Curves to Measure Deadweight Loss** Government revenue is area *ABCD*. Panel A shows how much the individual would be willing to pay to have the price of beer reduced from $1.50 to $1.20 (keeping him at the same level of utility). The difference between this and the tax revenue raised (the area *ABCD*) is the deadweight loss (the shaded area). Panel B illustrates the case where the level of consumption can be varied in very small increments.

would be willing to pay if the price were reduced from $1.50 to $1.30. He would be willing to pay 20 cents a bottle for the first ten bottles (the area *JKCD*) and 10 cents for the next (the area *GKLH*), for a total of $2.10.

Finally, we ask him to assume he is in a situation where we levied a $2.10 lump-sum tax and charged $1.30 each for twelve bottles. How much extra would he be willing to pay for one extra bottle? Assume the individual said \$1.20. We could now calculate the total lump-sum tax that an individual would be willing to pay for the elimination of the 30-cent tax. He would be willing to pay 30 cents on the first ten bottles (the area *ABCD*), 20 cents on the next bottle (the area *BNHG*), 10 cents on the twelfth bottle (the area *NRML*), for a total of \$3.30. The tax revenue from the tax was \$3.00 (the area *ABCD*). The deadweight loss is 30 cents (the shaded area).

More generally, the amount that an individual would be willing to pay to have the price reduced by 1 cent is just 1 cent times the quantity consumed. As we lower the price, the quantity consumed increases. In Figure 19.5B the total the individual would be willing to pay to have the price reduced from *D* to *A* is the area *AECD*, which takes account of the change in the after-tax quantity consumed as the price is reduced. But of that, *ABCD* is the tax revenue (the tax *AD*—which equals *BC*—times the quantity consumed, *AB*). Hence the deadweight loss, the difference between the two, is just the triangle *BCE*. Figure 19.6 shows that as we double the tax rate, we more than double the deadweight loss.

Figure 19.7 shows that, for a given tax rate, the deadweight loss is greater the flatter (or more precisely, the more elastic) the demand curve. (Remember that the elasticity of the demand curve gives the percentage change in demand as a result of a 1 percent change in price.)

We now make these insights more precise.

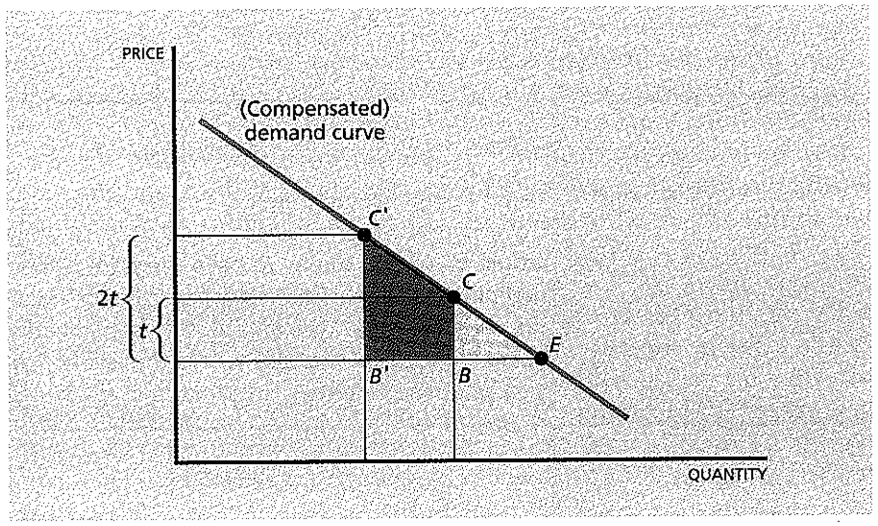

FIGURE 19.6 **Effect of an Increase in Tax Rate on the Deadweight Loss** A doubling of the tax rate more than doubles the deadweight loss. (The area *B'C'E* is four times the area *BCE*.)

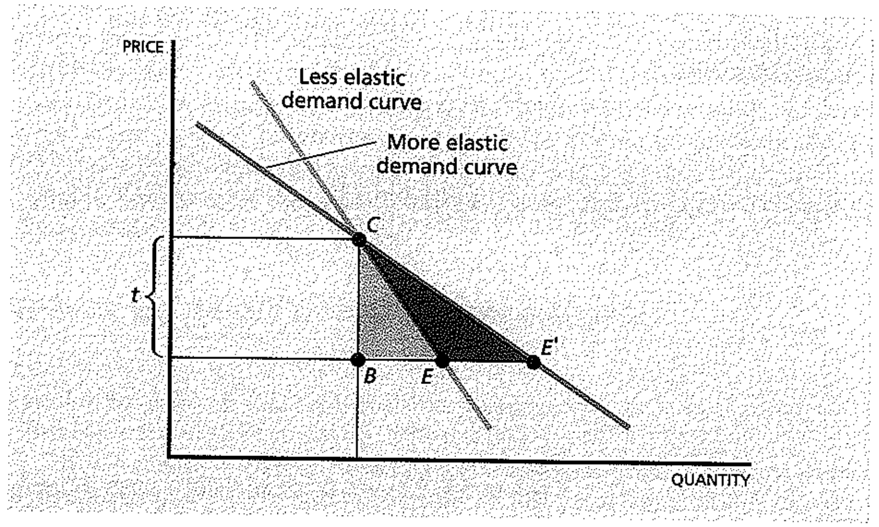

FIGURE 19.7 **Effect of an Increase in the (Compensated) Elasticity of Demand on Deadweight Loss** An increase in the elasticity of the (compensated) demand curve increases the deadweight loss. (*BEC* is deadweight loss from the less elastic demand curve, *BE'C* from the more elastic demand curve.)

CALCULATING THE DEADWEIGHT LOSS

Return to Figure 19.5, in which we used compensated demand curves to measure deadweight loss. The height of the triangle, *BC*, is equal to the tax, *t*. *BE* is the change in quantity as a result of the tax. Recall that the elasticity of demand gives the percentage change in quantity as a result of a 1 percent change in price, that is

$$\eta = \frac{\Delta Q/Q}{\Delta p/p}.$$

where the symbol ΔQ represents the change in quantity and the symbol Δp represents the change in price. The symbol Δ is the (capital) Greek letter "Delta" and is conventionally used to represent a change. The symbol η is the Greek letter "eta" and is conventionally used to represent the elasticity of demand. Rearranging, we can write the change in quantity as.

$$\Delta Q = \frac{\Delta p}{p} Q\eta.$$

This equation has the natural interpretation that the change in quantity will be larger, the larger the change in price and the larger the elasticity of demand. But the change in price is just the per unit tax, *t*. Thus, substituting, we obtain

$$BE = \frac{t}{p} Q\eta.$$

Now the area of the triangle *BCE* is just

$$\frac{t \cdot BE}{2} = \frac{1}{2}\frac{t^2}{p}Q\eta$$

$$= \frac{1}{2}\left[\frac{t}{p}\right]\left[\frac{t}{p}\right]pQ\eta$$

$$= \frac{1}{2}\hat{t}^2 pQ\eta,$$

where $\hat{t} \equiv \frac{t}{p}$ is the tax *rate*, the ratio of the tax to the price.

THE DETERMINANTS OF DEADWEIGHT LOSS The above formula identifies two of the primary determinants of deadweight loss. Deadweight loss increases with the *square* of the tax rate. High tax rates are far more distortionary than low tax rates.

And deadweight loss increases with the *elasticity of the compensated demand curve.*[3] The latter is precisely the substitution effect identified earlier as the critical determinant of deadweight loss. When indifference curves are very flat, the elasticity of the compensated demand curve is large—that is, a small percentage change in price leads to a large change in consumption. (Remember, the compensated demand curve simply describes a movement along an indifference curve, since, by definition, individuals are being compensated to keep them on the same indifference curve.) Many of the goods on which excise taxes are imposed have relatively low elasticities of demand,

DEADWEIGHT LOSS OF A TAX

The deadweight loss of a tax increases with the magnitude of the substitution effect (or the elasticity of the compensated demand curve) and increases with the *square* of the tax rate.

[3] Recall from Chapter 5 that the compensated demand curve is closely related to the ordinary demand curve. When price rises (say, as a result of the tax), individuals are worse off. If the individual previously purchased 100 bottles a beer a year, a 10-cent price increase makes him worse off; if we gave him $10, he would be fully compensated. The effect of a compensated price increase is just the ordinary direct effect, plus the effect of giving an individual an extra $10. If Bill spends only .1 percent of his income on beer, then the extra income induces an additional beer expenditure of 10 cents: there is little difference between the impact of a compensated and an uncompensated change.

so that the deadweight loss is relatively small. For instance, the 10 percent airline ticket tax is estimated to have a deadweight loss equal to 2.5 percent of the revenue raised (on the basis of an estimated .5 price elasticity of demand), an 8 percent beer tax generates a deadweight loss equal to 1.2 percent of the revenue raised (on the basis of an estimated price elasticity of .3), and a 15 percent cigarette tax is estimated to lead to a deadweight loss equal to 3 percent of the revenue raised (on the basis of a price elasticity of demand of .4).

EFFECT OF TAXES BORNE BY PRODUCERS

Up to now, this chapter has focused on the distortionary effects of a tax on a consumption good. We assumed that supply curves were horizontal, so the entire burden of the tax was on consumers.

But, at least in the short run, most supply curves are upward-sloping. This means that part of the burden of any tax on a consumption good will fall on producers. Will this cause an excess burden on producers, above and beyond the direct burden of the tax revenue? The answer is yes, except in the special case where the supply curve is vertical (that is, the elasticity of supply is zero).

Consider how a supply schedule (curve) is constructed. At each price, firms produce up to the point where price equals marginal cost. If the supply schedule is upward-sloping, the marginal cost rises as production rises. The area between the supply curve and price measures the **producer surplus,** which is just the difference between revenues and total variable costs. Changes in this area thus measure changes in profits.

Consider the example illustrated in panel A of Figure 19.8. What happens to profits as price increases from 1 to 4 and output increases from 1 to 4? The first unit of output costs $1; the next, $2; the third, $3; and the fourth, $4. If we pay the firm $4 for each unit, so it produces four units, the firm gets $3 more than marginal costs for producing the first unit, $2 more than marginal costs for producing the second unit, and $1 more than marginal costs for producing the third unit. The total profits are $3 + 2 + 1 = $6. Imposing a tax which lowers the price received by the producer to $3 lowers profit to $2 + $1 = $3. But if the tax is $1 per unit, tax collection will be $2, so the deadweight loss is $1.

This can be seen more generally in panel B of Figure 19.8. Assume initially the producer is receiving the price p. Then a tax is imposed that lowers the amount he receives to $p - t$. In the initial situation, his total profits are given by the area *DBC*.[4] Now, his profits are reduced to *DGE*. The change in his profits area is *EGBC*. But of this change, part accrues to the government as tax revenue—the rectangle *EGHC*. The tax on producers has resulted in

[4] More accurately, the shaded area measures the difference between revenues and total variable costs. To calculate profits, we need to subtract fixed costs. (Fixed costs are costs that are incurred as long as the firm operates; they do not depend on the scale of production.)

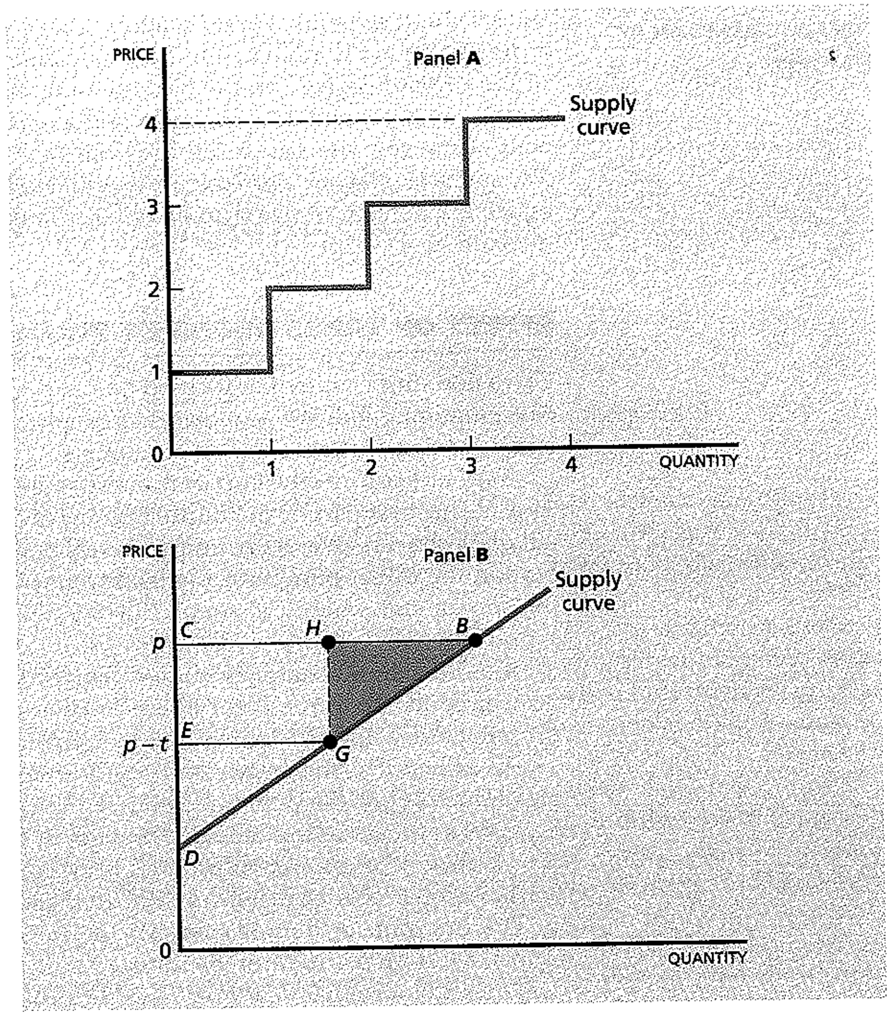

FIGURE 19.8 **The Deadweight Loss of a Tax on Production** *BGH* measures the deadweight loss of a tax on production.

producers' profits being reduced by more than government revenue has increased. The difference between the two is the deadweight loss associated with the tax. It is simply the shaded area *BGH*. To put it another way, the government could have imposed a lump-sum tax on the firm, which left price at p and which left the firm at the same level of profits as it had with a price of $p - t$. That lump-sum tax would have generated higher revenues, by the amount *BGH*, than the tax on the output of the firm.

It is clear that the steeper—the more inelastic—the supply schedule, the smaller the deadweight loss. In particular, we can, as before, show that for small taxes the deadweight loss increases with the square of the tax rate and with the supply elasticity.

EFFECT OF TAXES BORNE BY PRODUCERS

A similar analysis applies to taxes on goods that are used in production. For instance, assume we had a tax on some input, such as steel, into an industry (automobiles). We can ask what lump-sum tax we could impose on the industry that would have the same effect on profits as the tax on steel.[5] The difference in revenues raised by the lump-sum tax and the tax on steel is the deadweight loss from the tax. The magnitude of the deadweight loss will depend on the possibilities of substitution. If the firm cannot substitute any other input for steel (even partially), the tax on steel is no different from a tax on output. There is no distortionary effect on the input mix and, hence, no deadweight loss associated with a change in the input mix.

EFFECTS OF TAXES BORNE PARTLY BY CONSUMERS, PARTLY BY PRODUCERS

It is straightforward to combine our analysis of producer deadweight loss with consumer deadweight loss. Figure 19.9 illustrates the case of a tax that is borne partly by producers (the price they receive falls from p to p_s) and partly by consumers (the price they pay rises from p to p_c). The change in market demand can be decomposed into two parts, just as before. The

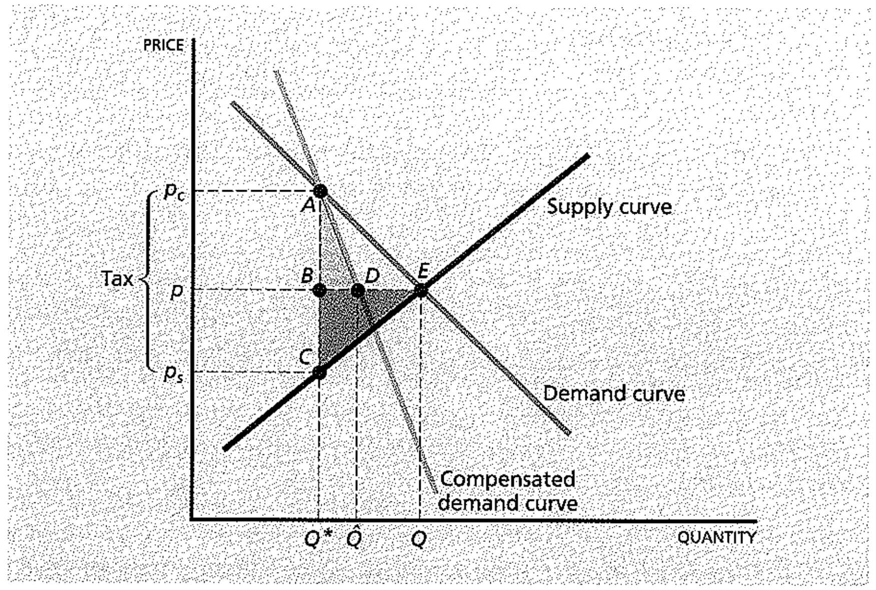

FIGURE 19.9 **Deadweight Loss from a Tax That Is Borne Partly by Consumers and Partly by Producers** The consumers' deadweight loss is the triangle *ABD;* the producers' is the triangle *BCE*. If the compensated and uncompensated demand schedules coincide, as they will if the demand curve is not sensitive to small changes in income, then the total deadweight loss is the big triangle *ACE*.

[5] This is not, of course, the only deadweight loss arising from the input tax. Since it increases the marginal cost of production, the input tax will result in an increase in the price consumers pay, and there will be a deadweight loss to consumers.

movement from Q to $\hat{Q}$ is the income effect of the tax; the movement from $\hat{Q}$ to Q^* is the substitution effect, as consumers substitute away from the taxed good along the compensated demand curve. That is, in the new equilibrium, at the price p_c, consumers are clearly worse off than they were at the original equilibrium price, p. If we ask how much they would have consumed, at the original (non-tax-distorted) price p, but at the new lower level of welfare, the answer is $\hat{Q}$, the point along the compensated demand schedule through A at the price p. The deadweight loss is associated with the movement along the compensated demand schedule, with the reduction of consumption from $\hat{Q}$ to Q^*, and is given by the triangle ABD.

What matters for producers, however, is the total change in quantity, from Q to Q^*, so their deadweight loss is the triangle BCE. The total deadweight loss is the sum of these two triangles, and depends, as before, on the elasticities of demand and supply.

TAXATION OF SAVINGS

The individual's allocation of her income between consumption this period and consumption in the future is very much like her decision about allocating her income between two different commodities.

By giving up one dollar of consumption today, the individual can obtain $(1 + r)$ of extra consumption dollars next period, where r is the interest rate. That is, if the individual saves the dollar and deposits it in a bank, she gets back at the end of the period her dollar plus the interest it has earned. Thus, $1/1 + r$ is the price of consumption tomorrow, relative to consumption today.

If the individual neither borrowed nor saved money, she would consume whatever her wages were in the two periods. We denote the wages in the initial period by w_0 and wages in the next by w_1. Suppose that w_0 and w_1 correspond to point W in Figure 19.10. By borrowing, the individual can consume more today, but at the expense of consuming less next period. By saving, the individual can consume more next period, but at the expense of consuming less this period.

The individual thus faces a budget constraint. She can either have $\bar{C}$ units of consumption today, or $(1 + r)\bar{C}$ units of consumption tomorrow, or any point on the straight line joining the two points, as depicted in Figure 19.10. The individual has a set of indifference curves between present consumption and future consumption,just as the individual has between beer and soda; each indifference curve gives those combinations of current and future consumption that leave her at the same level of utility. The individual is willing to consume less today in return for more future consumption. As her present consumption gets smaller and smaller, she becomes less willing to give up more; and as her future consumption gets larger and larger, the extra benefit she gets from each additional unit of future consumption gets smaller and smaller. Thus, the amount of increased consumption next period required to compensate the individual for a reduction by one unit in current consumption becomes larger and larger. That is why the indiffer-

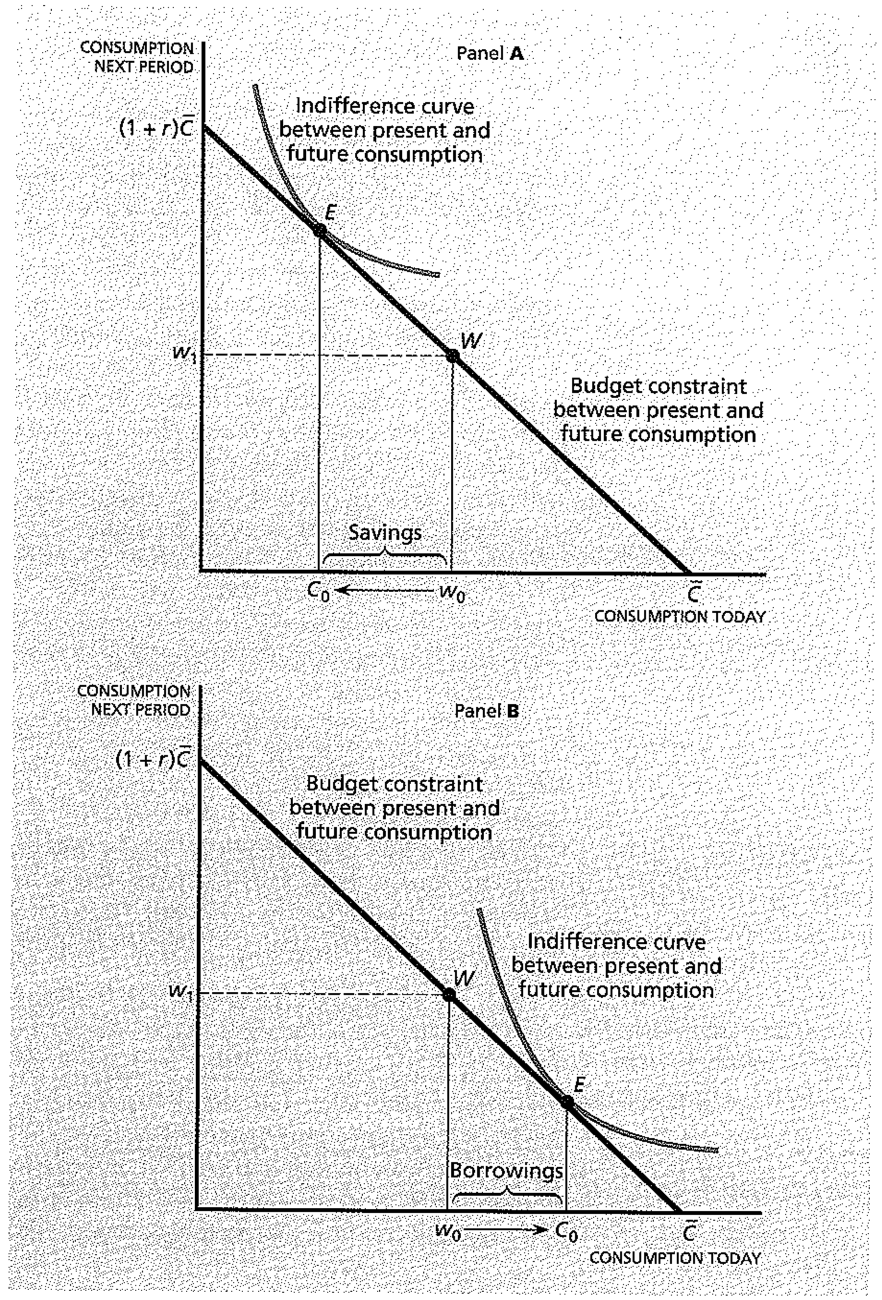

FIGURE 19.10 **Consumption, Savings, and Borrowings** The individual allocates her income between consumption this period and next period. In panel A the individual saves, while in panel B she borrows.

ence curve has the shape depicted. The individual chooses that point, denoted by E, on her budget constraint which is tangent to her indifference curve.

Panel A of Figure 19.10 illustrates a situation where the individual wishes to consume less than her wage income the first period and so saves the rest, while in panel B the individual wishes to consume more than her wage income the first period and so borrows the difference.

Consider now the effect of a tax, at the rate t, on interest income. (We assume that if interest income is negative—that is, the individual is a borrower—there is a negative tax; in other words, the borrower receives money from the government.) For a saver, someone whose first-period consumption is less than her (after-tax) first-period wage income, the tax has both an income effect and a substitution effect. Because the individual is worse off, she normally will reduce her consumption in both periods. Thus, the income effect leads to a lower current consumption. (Remember that savings is just the difference between after-tax wage income the first period and first-period consumption.) But because the individual receives a lower return from postponing consumption, the substitution effect discourages future consumption and encourages current consumption; it leads individuals to reduce their savings. The net effect on current consumption—and hence on savings—is ambiguous. If the substitution effect is large enough, savings are reduced.

If the substitution effect and the income effect were exactly to cancel each other, leaving savings unchanged, would this imply that the tax is nondistortionary? No, for the tax is distortionary so long as it causes the individual to substitute between current and future consumption along her indifference curve.

Figure 19.11 depicts the case where the income effect of the interest tax (the movement from E to $\hat{E}$) is just offset by the substitution effect (the movement from $\hat{E}$ to E^*). Therefore savings, $w_0 - C_0$, are the same before and after the tax. Nonetheless, there is a substantial distortion in second-period consumption.

We could contrast the effect of the interest income tax with a lump-sum tax, a tax that shifted the budget constraint down in a parallel manner. Again, it is straightforward to show that such a tax will, for any given effect on the individual's utility, raise more revenue (the deadweight loss is measured by E^*F), or that for any given level of revenue, individuals will be better off with the lump-sum tax than with the interest income tax. The magnitude of the distortion depends on the magnitude of the substitution effect, which in turn depends on how substitutable current and future consumption are.

QUANTIFYING THE EFFECTS OF AN INTEREST INCOME TAX

Most empirical estimates suggest that the substitution effect slightly outweighs the income effect, so that an interest income tax has a slight negative effect on savings. While from one perspective this is good news—the tax system may not be reducing savings by much—from another perspective it is bad news: government is unlikely to encourage savings by much through tax incentives.

The fact that the net effect is small does not, of course, mean that the distortionary effect is small. That depends on the magnitude of the substitu-

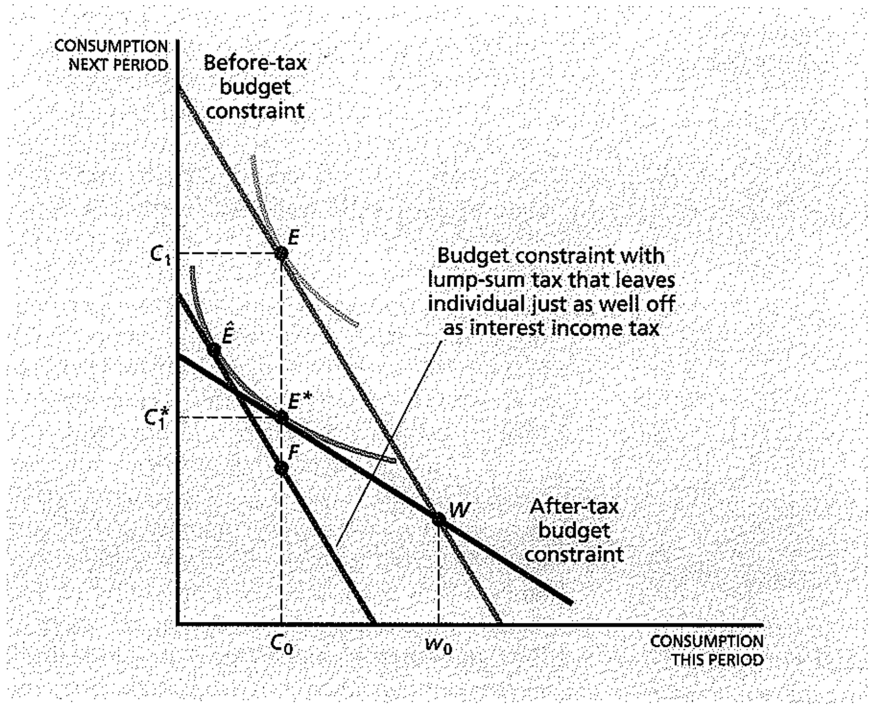

FIGURE 19.11 **Effect of Interest Income Tax** The income effect of the interest tax is just offset by the substitution effect in the first period. But there is still a deadweight loss of E^*F.

tion effect. But since most Americans save a relatively small fraction of their incremental income, the income effect is relatively small, implying that the substitution effect is also relatively small (since the two cancel out).

As we noted, distortionary effects increase with the square of the tax. Since much of saving is done by the rich, who face relatively high taxes, the deadweight loss may still be significant, even if the elasticity of substitution is small.

The precise magnitude of the deadweight loss remains a subject of controversy, with some economists arguing that the deadweight loss is actually quite large.

TAXATION OF LABOR INCOME

Exactly the same kind of analysis can be applied to labor supply decisions. There are, of course, many dimensions to labor supply—number of hours worked, effort exerted on the job, years of education (age of entry into the labor force), and age of retirement. The same principles apply to each. Here, we illustrate the analysis for the decision concerning the number of hours worked. As with the savings decision, we can model the labor supply decision in terms of the choice between two commodities. Here the two commodities are leisure and all consumption of goods and services. Figure 19.12 shows the individual's budget constraint, with hours of leisure on the horizontal axis and consumption (income)[6] on the vertical axis. The wage tax (like the beer tax) rotates the budget constraint. If the individual does not work at all, he faces no tax—he still has sixteen hours of leisure a day (ignoring the eight hours of sleep). At a 50 percent wage tax, his consumption—at any given level of work (leisure) is reduced by half. Again, there is an income effect and a substitution effect. The substitution effect, as before, makes him work less (enjoy more leisure). But the income effect makes him work more: when the individual is poorer, he "consumes" less of all "goods," including leisure. The income and substitution effects work in opposite directions. The figure illustrates a case where the two effects are essentially offsetting; there is no effect on hours worked.

The fact that the labor supply curve is relatively inelastic—that income and substitution effects are offsetting—does not mean that the income tax is not distortionary. It is, so long as there is a substitution effect. Indeed, Figure 19.13 shows a case where the income effect outweighs the substitution effect, so that the labor supply curve is backward-bending (at lower wages, individuals actually supply more labor). A tax in that case actually increases the labor supply. But, nonetheless, since there is a substitution effect, the tax is distortionary—that is, there is a deadweight loss associated with the tax.

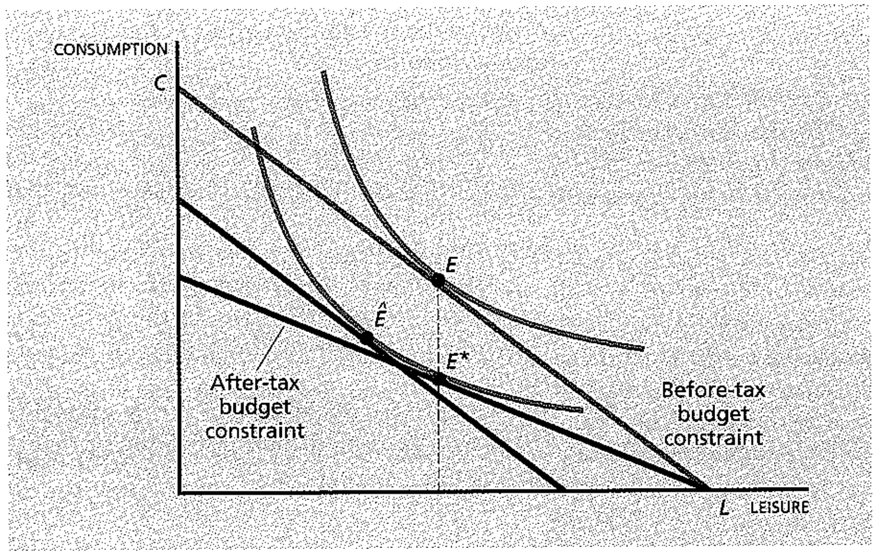

FIGURE 19.12 **Wage Taxation** Taxes on labor reduce the return to working. The substitution effect leads individuals to work less (enjoy more leisure), while the income effect leads individuals to work more. The two effects are offsetting.

[6] For purposes of this section, we assume there is no savings, so that consumption and income are identical.

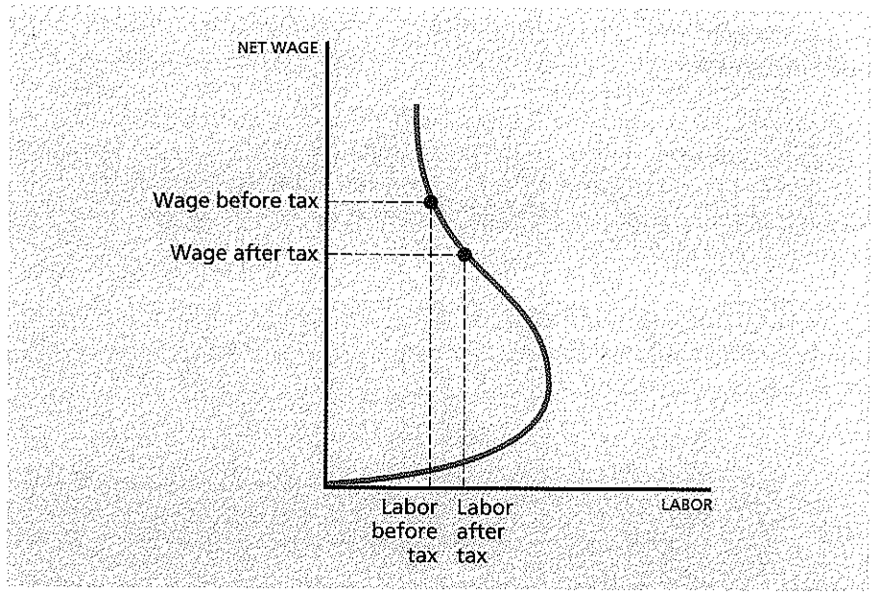

FIGURE 19.13 **Backward-Bending Labor Supply Curve** If the income effect outweighs the substitution effect, the labor supply curve will be backward-bending; increases in wages will lead to less labor supply. Then a tax which reduces wages (net of taxes) may actually increase the labor supply. The tax still has a deadweight loss.

EFFECTS OF PROGRESSIVE TAXATION

So far in this chapter, we have focused on *proportional taxes.* The beer tax was a per unit tax, so that as consumption of beer increased, tax payments increased proportionally. Similarly for the interest income tax. Some wage taxes, such as the social security payroll tax, are proportional (up to some maximum), but other taxes, such as the individual income tax, are not.

EARNED INCOME TAX CREDIT Figure 19.14 shows the effect of taxes on the budget constraint facing a low-wage individual who receives a *wage subsidy* (under the earned income tax credit) up to some level, and faces a tax beyond a certain (higher) level.[7] In the interval *LA*, the new budget constraint is actually steeper than the before-tax budget constraint *LL*; in the interval *AB*, the new and old budget constraints are parallel; and in the interval *BC* the after-tax budget constraint is much flatter. For an individual who works little (chooses a point in the interval *LA*), the income and substitution effects work in opposite directions, with the income effect leading to less work (since the individual is better off) and the substitution effect to more work (since the return to working has increased). For an individual who works a moderate amount and chooses a point in the interval *AB*, there is only an income effect: he unambiguously works less than before the subsidy. For the hardworking individual, in the in-

[7] The analysis simplifies the full complexity of the tax law, by ignoring, for instance, both state and social security taxes.

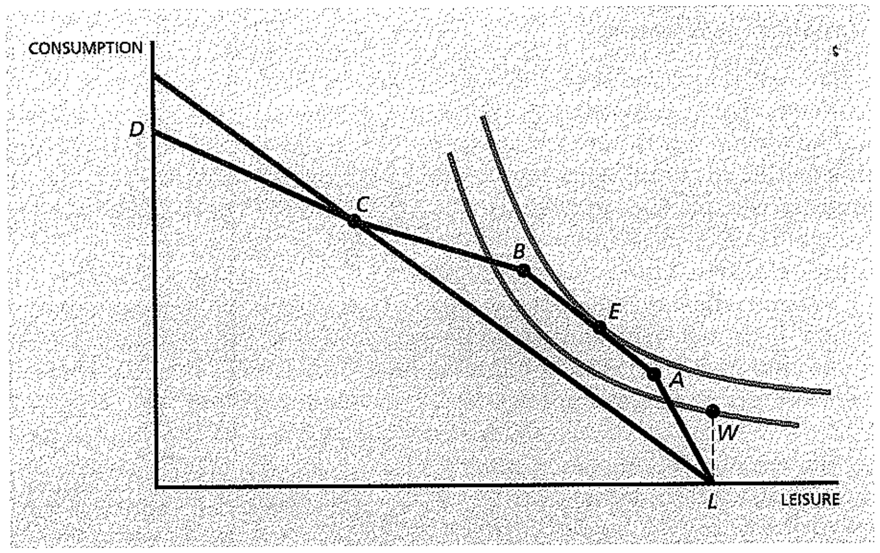

FIGURE 19.14 **Budget Constraint for Low-Income Individuals: The Effect of the Earned Income Tax Credit** Low-wage individuals who work little face positive incentives to work from the earned income tax credit. The tax credit affects both hours worked and labor force participation.

terval *BC*, the income and substitution effects are reinforcing: he is better off, and the return to working is lowered, so the reduction in work effort is even greater. Finally, in the interval *CD*, the individual's EITC benefit is completely exhausted and he now faces an income tax. Thus, he is worse off than he was in the absence of any tax/subsidy program. Now the income and substitution effects work in opposite directions. (Point *C*, where EITC is fully faced out, can occur to the left or the right of the before-tax budget constraint.)

PARTICIPATION VERSUS HOURS DECISIONS Even though for many individuals (those in the interval *BC*), marginal incentives to work are reduced by the EITC, *incentives to participate in the labor force are increased.* Assume, as an alternative to working, individuals can receive a fixed amount of welfare payments, represented in the figure by point *W* (where they enjoy full leisure). We have drawn the indifference curve through *W*.[8] It is clear that, with the earned income tax credit, the individual who would have been content simply to receive welfare now prefers to work.

Most economists believe that the participation decision is far more important than the hours decision. First, attachment to the labor force brings with it education, skills, and a sense of belonging to society which contributes to social stability. Second, for many jobs, discretion over the number of hours is far more limited than the discussion we have presented suggests.[9] There is either a workweek of five eight-hour-days, or, say, a workweek of five seven-hour-days. To be sure, over time, the number of hours worked does adjust in response to economic forces—the average number of hours worked declined markedly during the first half of the twentieth century, while more recently, as wages, particularly at the lower end of the income distribution, have failed to grow or even declined, hours worked have increased—the individual, in deciding to work, is not usually in a position to bargain about whether he should work thirty-five hours or thirty-seven hours.

[8] In practice, under welfare individuals do not lose all benefits if they work, but rather face a high marginal tax rate. There is a flat budget constraint through *W*. We assume it is sufficiently flat that the individual chooses *W*.

HIGH-INCOME INDIVIDUALS Figure 19.15 analyzes the effect of taxation on high-income individuals in the 36 percent tax bracket. Someone in the 36 percent tax bracket faces a budget constraint with four segments, represent-

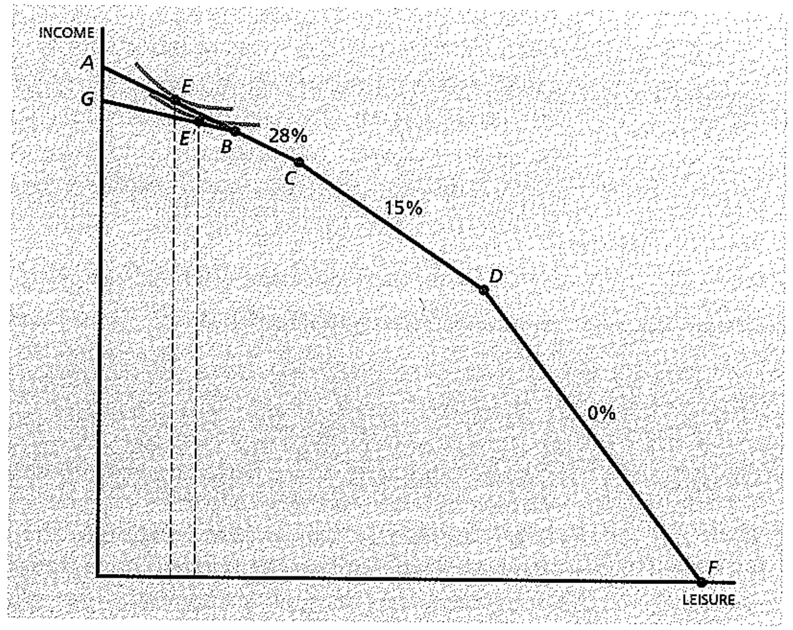

FIGURE 19.15 **Impact of 1993 Tax Change** In 1993, the tax rate on upper-income individuals was increased. The curve *ABCDF* depicts schematically the original budget constraint, with three marginal tax rates, at 0 percent, 15 percent, and 28 percent. After the change, there is a fourth segment, *BG*, at a tax rate of 36 percent. Originally, the high-income individual chooses point *E*. Afterward, he chooses *E'*. The income effect is small, since he only has to pay the higher tax on the excess of income over a very high threshold. The substitution effect is large relative to the income effect, so the individual works less (enjoys more leisure).

[9] This observation is consistent with a recent study of high-income physicians in Britain, which found that while self-employed physicians exhibited considerable sensitivity to the marginal tax rate, those who are employees have no discernible sensitivity. See Mark H. Showalter and Norman K. Thurston, "Taxes and Labor Supply of High-Income Physicians," *Journal of Public Economics* 66 (1997): 73–97.

THE 1993 TAX REFORM

In 1993, as the government faced ever-increasing deficits, some tax increases appeared imperative. The Clinton administration proposed that those who had benefited most from the economic expansion and the tax cuts of the 1980s should bear the brunt of these tax increases. Only the top 1.2 percent of taxpayers experienced rate increases. For example, married couples with incomes in excess of $140,000 had their income tax rates increased, from a marginal tax rate of approximately 28 percent to 36 or 39.6 percent.*

Critics, such as Professor Martin Feldstein of Harvard University, who was chairman of the Council of Economic Advisers under the Reagan administration, predicted that the tax increase would raise less revenue than both Congress and the administration had estimated, because of large responses in labor supply. As was illustrated in Figure 19.15, many of those facing tax increases were clustered near point *B*, the income level at which taxes were increased. For these individuals, as we have noted, there was little income effect from the tax but a large substitution effect, and therefore, Feldstein argued, there would be a large labor supply response. To support his theoretical arguments he cited evidence that after the 1986 tax reform, which had reduced marginal rates for upper-income individuals, there had been a marked increase in income tax collections from upper-income individuals.

But, as is so often the case, reading the evidence is not easy. The 1986 tax reform also closed a number of loopholes which had allowed upper-income individuals to avoid taxation; and in the expansion of the economy which had begun in 1983, as the economy recovered from its worst recession since World War II, there was a strong trend of increased inequality—with earnings at the top growing far more rapidly than earnings in the middle. There was little direct evidence that upper-income individuals had worked either longer hours or harder as a result of the decrease in marginal tax rates.

As it turned out, the tax revenues raised on upper-income individuals in the years following 1993 were far higher than had been anticipated, and indeed these increased revenues were largely responsible for the elimination of the deficit in the late 1990s. And there was little evidence that the upper-income individuals had reduced their labor supply in the way Feldstein had predicted.

* This ignores the Medicare tax, which was also raised.

ing the 0 percent, 15 percent, 28 percent, and 36 percent brackets. Assume initially that the top bracket was 28 percent. An increase in the top bracket is represented by a downward rotation of the budget constraint at point *B*. Because of the kink in the budget constraint, many individuals—with different preferences between leisure and work—may be clustered at the level of income associated with the kink.

For those near the kink, the income effect is small relative to the substitution effect, and hence their labor supply would have been expected to decrease, as depicted. (If the tax rate on all income had increased by the same percentage points, there would have been a large income effect; but for someone with an income just above the level at which the higher rate set in, the income effect was in fact negligible.)

SECONDARY LABOR FORCE PARTICIPANTS

Historically, in two-earner families with primary and secondary earners, the labor supply response of the secondary earner (typically the wife) has been markedly different from that of the primary earner. There are obvious reasons for this, particularly during the years in which there are small children at home. The *net* income—after subtracting out the costs of child care (which would not have to be paid if she did not work), costs of

commuting, and so forth—may be far lower than the gross income. Thus, if in 1997 a woman earned $15,000 a year, and child care expenses amounted to $5000, the net income was just $10,000. The tax rate she faced is determined by *household* income, including that of her husband. If her husband earned more than $56,050 she faced a 28 percent (or higher) income tax on all of her income. (This ignores the payroll tax and state taxes.) But a 28 percent tax on her total income translates into a 42 percent tax on her net income. Her incentives to participate in the labor force are greatly reduced. Secondary workers have, accordingly, shown much more sensitivity in their labor force participation to changes in tax rates.

MEASURING THE EFFECTS OF TAXES ON LABOR SUPPLIED

The fact that theoretically the effect of a tax on wages is indeterminate makes it all the more important to attempt to determine empirically what its effects in fact have been. Research in this area has been extensive and has yielded important (but controversial) results. Two main methods have been employed to study these questions: statistical analyses using market data, and experiments.[10]

STATISTICAL TECHNIQUES USING MARKET DATA

The first method entails using statistical techniques to analyze how individuals in the past have responded to changes in their after-tax wages. In general, we do not have data on how particular individuals responded to changes in wages. Rather, we have data on how many hours individuals who earn different wages work. Those who earn higher wages seem to work more hours. From this we can calculate the "average" effect of wages on hours worked.

Up to now, we have simply described a *correlation*, an observed relationship between two economic variables. We now wish to use this to make an *inference*, a prediction or a statement about the effect of lowering take-home wages resulting from, say, the imposition of a tax. To make such an inference, we must make an assumption: that the reason individuals who receive higher wages are observed to work more is that they choose to work more because of the higher wage; in other words, that an individual who receives a higher wage is essentially like one who receives a lower wage; the only important difference is the difference in pay, and it is this difference that leads to a difference in the number of hours worked. There are, of course, other important differences, and more sophisticated statistical analyses attempt to take as many of these as possible (say, age or occupation or sex) into account; they attempt to see, among individuals of the same age, occupation, or sex (or who have other characteristics in common), whether those who receive higher wages work more.

The vast literature on labor supply suggests that estimated labor supply elasticities may depend on the precise statistical methods used as well as on the data employed.[11] There appears to be widespread agreement (though not unanimity) on the following:

- Labor supply of married males is fairly unresponsive to changes in the wage rate.[12]

[10] Earlier, several studies approached the problem using *qualitative* approaches, simply asking individuals whether taxes led them to work more or less. The different responses reflected the presence of income effects (leading to more work) and substitution effects (leading to less work). See D. M. Holland, "The Effect of Taxation on Effort: Some Results for Business Executives," *National Tax Association Proceedings of the Sixty-Second Annual Conference*, 1969; and George Break, "Income Taxes and Incentives to Work: An Empirical Study," *American Economic Review* 47 (1957): 529–49.

[11] For instance, Jim Heckman of the University of Chicago has argued that reporting errors may have obscured a larger decline in hours worked by men and may also account for much of the observed decline in real wages. This measurement error may also bias labor supply elasticities toward zero. See James J. Heckman, "What Has Been Learned about Labor Supply in the Past Twenty Years?" *American Economic Association Papers and Proceedings*, May 1993, pp. 116–21.

[12] For instance, one study estimated that the tax system (as of 1983) had resulted in only a 2.6 percent reduction in hours worked by married men. See Robert K. Triest, "The Effect of Income Taxation on Labor Supply in the United States," *Journal of Human Resources*, summer 1990, pp. 491–516. An earlier study, using similar data and model, had suggested a somewhat larger effect. See J. Hausman, "Labor Supply," in *How Taxes Affect Economic Behavior*, ed. H. J. Aaron and J. Pechman (Washington, D.C.: Brookings Institution, 1981), pp. 27–72.

- The compensated labor supply elasticity also appears to be small, though there appears to be more disagreement over this finding. (The uncompensated labor supply elasticity can be small either because of small offsetting substitution and income effects or because of large offsetting substitution and income effects.)[13]
- Labor supply elasticity of married females is larger, but more problematic to estimate.[14]
- Labor supply elasticity of female heads of household is somewhere in between that of married men and married women.
- The participation decision is more sensitive to the wage rate than are marginal hours of work.[15]
- Labor supply parameters estimated from market data tend to be larger than those estimated from experimental data (discussed in the next section).

Large changes in tax rates, of the kind that occurred in 1981, 1986, and 1993, provide natural tests of the effect of taxation on labor supply.

The 1993 tax changes have yet to be subjected to detailed analysis. But each of these changes was complex; different individuals faced different changes in tax rates, and a variety of loopholes were closed (or in some cases opened up). While the evidence is that the 1986 tax law led to an increase in taxable income faster than would have been predicted based on the assumption of fixed incomes,[16] there is controversy about the reason.

[13] Jerry Hausman of MIT, using a technique meant to capture the effects of the nonlinearities in the budget constraints described earlier, obtains much larger estimates than do others. See his "Labor Supply," cited above. His techniques have been criticized by Tom MaCurdy of Stanford, who claims that his technique "forces higher estimates of substitution effects or lower estimates of income effects than are obtained from other procedures. This . . . raises serious questions about the reliability of evidence cited by much of the literature to support recent tax reforms aimed at lowering marginal tax rates." See Thomas MaCurdy, "Work Disincentive Effects of Taxes: A Reexamination of Some Evidence," *American Economic Association Papers and Proceedings*, May 1992, pp. 243–49.

[14] See Thomas A. Mroz's widely cited study, "The Sensitivity of an Empirical Model of Married Women's Hours of Work to Economic and Statistical Assumptions" (*Econometrica*, July 1987, pp. 765–99). After noting that the estimated elasticity is sensitive to the specification of the model, Mroz concludes that when the correct specification is used, "factors such as wage rates, taxes and nonlabor incomes have a small impact on the labor supply behavior of working married women" (p. 795). His analysis only focused on working women, ignoring the participation decision.

[15] This could be either because of constraints imposed by the employer or because of the fixed costs of working.

Triest estimates that, despite the small responsiveness of hours worked, participation is so responsive that the U.S. (state and federal) tax system reduced total hours worked by wives by as much as 30 percent, depending on the specification employed. See "The Effect of Income Taxation on Labor Supply in the United States," cited in note 12.

[16] See Martin Feldstein, "The Effect of Marginal Tax Rates on Table Income: A Panel Study of the 1986 Tax Reform Act, *Journal of Political Economy* 1-3, no. 3 (1995): 551–72.

Did the changes, for instance, following the 1986 tax reform, reflect mostly the closing of loopholes? Or did they largely reflect underlying trends? Increasing inequality was placing a larger fraction of the nation's income in the hands of those facing higher tax rates.

Barry Bosworth and Gary Burtless of the Brookings Institution argued that the observed changes had little to do with changes in taxes, but mostly reflected underlying economic trends.[17] While it appears that after the tax cuts, labor supply was higher than it would have been had past trends continued (by 1989, for men, work effort was 6 percent above what it would have been had the 1967–1980 trend continued; and for women it was 5.4 percent higher), the interpretation of what went on is not so clear. If tax reform were the impetus for the growth in labor supply, we would expect higher-income workers (who experienced a greater change in marginal tax rates) to have increased their labor supply by proportionately more than lower-income workers (who were less affected by tax reform). Yet it was the lower-income individuals whose labor supply seems to have increased the most.[18] The fact that income tax revenues among the very rich continued to increase so robustly after 1993, in spite of the large increase in marginal tax rates, is consistent with the hypothesis that it was long-term trends and the closing of loopholes, not incentive effects on labor supply, that accounted for the increased tax revenues after the 1986 reform.

EXPERIMENTS

The second approach to obtaining a quantitative estimate of the magnitude of the labor supply responses to tax changes is an experimental one. We are interested in the question: What would happen to the labor supply if we raised or lowered tax rates (or changed the tax structure in some other way)? One approach is to say, "Let's change the tax structure and see what happens." This could be an expensive approach: the change might have a very negative effect on labor supply, but before the effects were recognized and the tax structure changed again, considerable damage (welfare loss) could have occurred.

But we can learn something by changing the tax structure for just a small portion of the population. Just as opinion polls can give fairly accurate estimates of how voters will vote in an election, simply by asking a small sample of the population (often fewer than 1000 individuals), so too the response of a small sample may give a fairly reliable indication of how other, similar individuals would respond facing the alternative tax structure. Opinion polls are careful to obtain a representative sample of views—they make sure that views of young and old, of the rich and poor, of skilled and unskilled workers, of married and unmarried individuals, and so on, are all

[17] Barry Bosworth and Gary Burtless, "Effects of Tax Reform on Labor Supply, Investment, and Saving," *Journal of Economic Perspectives,* winter 1992, pp. 3–25.

[18] Men in the lowest quintile increased hours worked by 31 percent, whereas those in the highest quintile increased hours worked by 3.2 percent; the corresponding numbers for women were 16.7 percent and 11.8 percent. (Bosworth and Burtless, "Effects of Tax Reform.")

represented; and in forming their estimate of how the population as a whole will vote, they weight the relative importance of the various groups in the population (when they are attempting to predict the outcome of elections, they assign weights corresponding to the known likelihood that members of different groups vote).

Between 1968 and 1982, a series of such experiments attempted to ascertain in particular the effects of changes in the tax structure and welfare system on the labor supplied by poorer individuals. Different individuals were confronted with different levels of guaranteed income and tax structures, making it possible, in principle, not only to estimate the overall effect of tax changes but to separate out the income effects from the substitution effects. The results were consistent with the view that the overall effect of taxes on labor supply is relatively small. The report on the first such experiment, conducted in New Jersey, described it as presenting "a picture of generally small absolute labor supply differentials between" those who were confronted with the alternative tax/welfare structures and those who faced the existing tax/welfare structure. "Only among wives, whose mean labor supply is quite small to begin with, are the differentials large in relative terms."[19] (In subsequent years, labor participation of wives has increased enormously, so that the aggregate effect of such adverse incentives is now far more significant.) The experiments yielded some further results concerning the possible effects of changes in the welfare/tax system. Providing more income to the poor resulted in their searching longer for a job when they became unemployed.

While the early experiments focused on the effect of alternative tax-subsidy schemes on labor supply (and related variables, like job search), later studies attempted to ascertain whether there were other effects as well. For instance, an experiment in Gary, Indiana, found a higher birth weight of babies—an indication of the health of the child—in families whose income had been increased. A large-scale experiment sponsored by the U.S. Department of Health, Education, and Welfare and conducted in Seattle, Washington, and Denver, Colorado, found that providing women with a guaranteed income, as the negative income tax does, might contribute to the breakup of families. However, the most generous negative income tax programs in the Seattle-Denver experiment had the *least* effect on family dissolution rates. It has been argued that income guarantees have two opposing effects on dissolution rates: On the one hand, they stabilize marriages by improving the family's ability to buy essential goods and services; on the other hand, they destabilize marriages by improving the economic viability of alternatives to marriage. Under this theory, the experimental results suggest that for low guaranteed income levels, the second effect (the "independence effect") dominates the first.

The experiments represented an important advance in the tools that are available to social scientists. At the same time, there are some important

[19] U.S. Department of Health, Education, and Welfare, *Summary Report: New Jersey Graduated Work Incentive Experiments* (Washington, D.C.: Government Printing Office, 1973).

limitations to the experimental approach that need to be borne in mind in evaluating the results.

First, there is a well-known phenomenon called the **Hawthorne effect,** which plagues all experimental work with individuals: When an individual is included in an experiment, and she knows her behavior is being examined, her behavior is often altered.

Second, there are problems associated with ensuring that the sample is representative. Since participation in the experiment is voluntary, there may be systematic biases associated with the kinds of individuals who refuse to participate.

Third, the response of individuals to short-run changes may differ from their responses to long-run changes. On the one hand, a temporary change in the tax structure that leads them to be better off has a smaller effect on lifetime income than a permanent change in the tax structure; hence the income effect may be understated. On the other hand, since the experiment discussed above often involved individuals facing a higher or lower marginal tax rate during the course of the experiment, the after-tax wage was temporarily reduced or increased; a temporary reduction in the wage may have different effects than a permanent reduction. In the absence of costs of adjustment there is a presumption that individuals will reduce their work (increase their leisure) more than they would with a permanent wage reduction. Thus an individual who was planning to take some time off from work (say, a woman thinking of having children in the near future) might have taken advantage of the temporary availability of a large subsidy combined with a high marginal tax rate. If this is true, the experiments overstate the effects relative to what they would be with a permanent change. On the other hand, costs of adjustment may be very high; an individual might be reluctant to quit his current job, knowing that he will want it back in three years' time (when the experiment is over), because he believes it will be difficult to get it back then. If these effects are important, the experiment may have understated not only the income effects but the substitution effects as well. Some of the more recent experiments have attempted to ascertain the magnitude of the biases in the estimates resulting from the fact that the change in tax structure/welfare payments was only temporary, by guaranteeing to the individual the same tax structure/welfare structure over a more extended period (up to twenty years).

A final important qualification on interpreting how accurately the experiments describe the extent to which labor supply is affected by changes in tax laws or welfare programs relates to the role of institutions in determining the length of the workweek. We commented earlier that, in the short run, institutional practices play an important role in restricting individuals' choices over the number of hours worked. But in the long run, these institutional practices themselves change, partly in response to changes in the economic environment. Thus many of the individuals in the experiment may have had only limited discretion over the number of hours they worked; but if everyone in society were confronted with the new tax/welfare payments structure, pressures might develop to alter these institutional practices to bring them more into conformity with individuals' preferences.

The high cost and ambiguous results of such experiments have meant that there have been few experiments of the scale and scope of the earlier studies. On the other hand, more care is placed in the design of pilot programs, so that stronger inferences can be made concerning what works and what does not work. There have been, for instance, a large number of studies of training programs and of programs designed to move people from welfare to work. Still, primary reliance has to be placed on "natural experiments," the experiments that occur as a result of, for example, different states' trying different programs. For instance, before the 1996 welfare reform, several states had experimented with time-limited welfare programs and welfare programs with work requirements. Such experiments suggested that the welfare reform would result in significantly reduced welfare dependency, a prediction borne out in the months after passage of the legislation—with the reduction in welfare roles far greater than could be explained by the declining unemployment rates.[20]

REVIEW AND PRACTICE

SUMMARY

1 The imposition of a tax that is not a lump-sum tax introduces inefficiencies. The magnitude of the inefficiencies is measured by the deadweight loss, the difference in revenues that could be obtained from a lump-sum tax as compared to a distortionary tax, with the same effect on the level of welfare of consumers.

2 The effect of any tax can be decomposed into an income effect and a substitution effect. There is an income effect associated with a lump-sum tax, but no substitution effect. The greater the substitution effect, the greater the deadweight loss.

3 There is also a deadweight loss associated with the reduction in the price received by producers as a result of the imposition of a tax. The reduction in their profits exceeds the tax revenues they effectively pay to the government.

4 For a tax on a commodity, both the income effect and the substitution effect usually lead to a reduction in the level of consumption of that

[20] On negative income tax experiments, see Robins, "A Comparison of the Labor Supply Findings from the Four Negative Income Tax Experiments," *Journal of Human Resources* 20, no. 4 (fall 1985): 567–82; David Greenberg and Harlan Halsey, "Systematic Misreporting and Effects of Income Maintenance Experiments on Work Effort: Evidence from the Seattle-Denver Experiments," *Journal of Labor Economics* 1, no. 4 (October 1983): 380–407; and Robert G. Spiegelman and K. E. Yaeger, "The Seattle and Denver Income Maintenance Experiments: Overview," *Journal of Human Resources* 15, no. 4 (fall 1980): 463–79. On natural experiments of tax changes, see two articles in *Empirical Foundations of Household Taxation*, ed. Martin Feldstein and James Porteba (Chicago and London: University of Chicago Press, 1996): Nada Eissa, "Labor Supply and the Economic Recovery Tax Act of 1981" (pp. 5–32); and James J. Heckman, "Comment on Labour Supply and the Economic Recovery Tax Act of 1981" (pp. 32–38).

commodity. For an interest income tax as viewed by a saver, the income effect typically leads to an increase in savings and the substitution effect leads to a decrease in savings: the net effect is ambiguous. But even if the net effect is to leave savings unchanged, there is still a distortion associated with the interest income tax. For workers, the income and substitution effects of an increase in wages have opposite effects; thus higher wages may lead to either an increase or a decrease in labor supply.

5 Empirical evidence suggests that for males, the substitution and income effects of wage taxes virtually cancel, so that the total effect of the tax on the male labor supply is probably not large. For females, there may be a marked effect on labor force participation. On the other hand, even though the total effect may be small for males, the substitution effect, and hence the deadweight loss associated with the tax, may be significant.

KEY CONCEPTS

Income effect
Substitution effect
Producer surplus
Hawthorne effect

QUESTIONS AND PROBLEMS

1 If savings do not respond to changes in the interest rate, does it mean that there is no deadweight loss associated with the taxation of interest?

2 What is the deadweight loss from the mineral tax in problem 1, Chapter 18? What is the relationship between deadweight loss and supply curves? Relate this to the discussion of lump-sum taxes.

3 Taxes and government expenditure programs affect a variety of other aspects of household behavior. Some economists, for instance, argue that they affect birth rates. What provisions of the tax system might affect the decision to have a child? What government expenditure programs?

4 Instead of representing the individual's decisions as a choice between consumption and leisure, they could be represented in terms of a choice between consumption and work. Draw the indifference curves, and identify the income and substitution effects resulting from a change in the tax rate on labor.

5 Compare the effects of a wage tax and a lump-sum tax raising the same revenue. In particular, show that the individual's utility is higher with the lump-sum tax than with the income tax.

6 Compare the effects of a proportional income tax and a progressive flat-rate income tax (i.e., one in which there is a lump-sum grant from the government of, say, $3000, and then a constant marginal tax rate on all income). In particular, show that if the two taxes raise the same revenue,

and all individuals have the same income, utility will be higher with the proportional tax.

7 Prior to 1981, the government imposed only a 67 percent (instead of a 100 percent) marginal tax rate on income earned by a mother receiving AFDC. Draw the budget constraint before 1981 and after 1981. Draw the indifference curve of someone who prefers to remain out of the labor force under both regimes. Draw the indifference curve of someone who worked before 1981 but chose not to work after 1981. Show how, for this person, lowering the tax rate will increase utility, reduce costs to the welfare system, and increase labor supply. Finally, draw the indifference curve of someone who worked both before and after 1981. Show how, for this person, the lower tax rate affects AFDC costs and affects labor supply. What can you say about government policy if there are some individuals of the first type, some of the second type, and some of the third type?

8 What would be the effect of a switch to taxing individuals on the basis of their own income (rather than family income) on labor force participation of wives?

9 Describe the income and substitution effects of an increase in the interest rate for a borrower. What does this imply for the effect of eliminating tax deductibility of interest payments?