Chapter Fourteen

Taxation and Income Distribution

Struggle and contrive as you will, lay your taxes as you please, the traders will shift it off from their own gain.

—JOHN LOCKE

American policy debates about the tax system are dominated by the question of whether its burden is distributed fairly. A sensible discussion of this normative issue requires some understanding of the positive question of how taxes affect the distribution of income. A simple way to determine how taxes change the income distribution would be to conduct a survey in which each person is asked how many dollars he or she pays to the tax collector each year. Simple—but usually wrong. An example demonstrates that assessing correctly the burden of taxation is much more complicated.

Suppose the price of a bottle of wine is $10. The government imposes a tax of $1 per bottle, to be collected in the following way: Every time a bottle is purchased, the tax collector (who is lurking about the store) takes a dollar out of the wine seller's hand before the money is put into the cash register. A casual observer might conclude that the wine seller is paying the tax.

However, suppose that a few weeks after its imposition, the tax induces a price rise to $11 per bottle. Clearly, the proprietor receives the same amount per bottle as he did before the tax. The tax has apparently made him no worse off. Consumers pay the entire tax in the form of higher prices. On the other hand, suppose that after the tax the price increases to only $10.30. In this case, the proprietor keeps only $9.30 for each bottle sold; he is worse off by 70 cents per bottle. Consumers are also worse off, however, because they have to pay 30 cents more per bottle.[1] In this case, producers and consumers share the burden of the tax. Yet another possibility is that after the tax is imposed, the price stays at $10. If this happens, the consumer is no worse off, while the seller bears the full burden of the tax.

statutory incidence

Indicates who is legally responsible for a tax.

economic incidence

The change in the distribution of real income induced by a tax.

The **statutory incidence** of a tax indicates who is legally responsible for the tax. All three cases in the preceding paragraph are identical in the sense that the statutory incidence is on the seller. But the situations differ drastically with respect to who really bears the burden. Because prices may change in response to the tax, knowledge of statutory incidence tells us *essentially nothing* about who really pays the tax. In contrast, the **economic incidence** of a tax is the change in the distribution of private real income induced by a tax. Our focus in this chapter is on the forces that determine the extent to which statutory and economic incidence differ—the amount of **tax shifting.**

[1] Actually, the change in the prices faced by consumers and producers is only part of the story. There is also a burden due to the tax-induced distortion of choice. See Chapter 15.

tax shifting

The difference between statutory incidence and economic incidence.

TAX INCIDENCE: GENERAL REMARKS

Several observations should be kept in mind in any discussion of how taxes affect the distribution of income.

Only People Can Bear Taxes

In a discussion of a tax bill that was once being considered by Congress, a *Wall Street Journal* columnist observed that "the Senate voted to approve a major tax-law revamp that focuses mainly on corporations, but lawmakers also approved important changes that will benefit many people" [Herman, 2004a]. By drawing a sharp distinction between "corporations" and "people," the statement reflects a common fallacy—that businesses have an independent ability to bear a tax. True, the US legal system treats certain institutions such as corporations as if they were people. Although for many purposes this is a convenient fiction, it sometimes creates confusion. From an economist's point of view, people—stockholders, workers, landlords, consumers—bear taxes. A corporation cannot.

Given that only people can bear taxes, how should they be classified for purposes of incidence analysis? Often their role in production—what inputs they supply to the production process—is used. (Inputs are often referred to as *factors of production*.) The focus is on how the tax system changes the distribution of income among capitalists, laborers, and landlords. This is referred to as the **functional distribution of income.**

functional distribution of income

The way income is distributed among people when they are classified according to the inputs they supply to the production process (for example, landlords, capitalists, laborers).

Framing the analysis this way seems a bit old-fashioned. Perhaps in 18th-century England property owners never worked and workers owned no property. But in the contemporary United States, many people who derive most of their income from labor also have savings accounts and/or common stocks. (Often, these assets are held for individuals in pensions.) Similarly, some people own huge amounts of capital and also work full-time. Thus, it seems more relevant to study how taxes affect the way in which total income is distributed among people: the **size distribution of income.** Given information on what proportion of people's income is from capital, land, and labor, changes in the functional distribution can be translated into changes in the size distribution. For example, a tax that lowers the relative return on capital tends to hurt those at the top of the income distribution because a relatively high proportion of the incomes of the rich is from capital.[2]

size distribution of income

The way that total income is distributed across income classes.

Other classification schemes might be interesting for particular problems. When increases in the federal tax on cigarettes are proposed, the incidence by region receives a great deal of attention. (Are people from tobacco-growing states going to

[2] However, some low-income retirees also derive the bulk of their income from capital.

suffer disproportionate harm?) Alternatively, when proposals are made to change the taxation of land in urban areas, analysts often look at incidence by race. It is easy to think of further examples based on sex, age, and so forth.

Both Sources and Uses of Income Should Be Considered

In the previous wine tax example, it is natural to assume that the distributional effects of the tax depend crucially on people's spending patterns. To the extent that the price of wine increases, the people who tend to consume a lot of wine are made worse off. However, if the tax reduces the demand for wine, the factors employed in wine production may suffer income losses. Thus, the tax can also change the income distribution by affecting the sources of income. Suppose that poor people spend a relatively large proportion of their incomes on wine, but that vineyards tend to be owned by the rich. Then on the uses of income side, the tax redistributes income away from the poor, but on the sources side, it redistributes income away from the rich. The overall incidence depends on how both the sources and uses of income are affected. This distinction is important for understanding the debate over former Vice President Gore's proposal to clean up the Florida Everglades. Because the ecology of the Everglades is harmed by the runoff from sugar fields, he argued that sugar products be subjected to a special tax and the proceeds used to finance a cleanup. Opposition came not only from consumer groups who were concerned about the price of products using sugar but also from Florida *workers,* who realized that by reducing the demand for sugar, such a tax would hurt their incomes.

In practice, economists commonly ignore effects on the sources side when considering a tax on a commodity and ignore the uses side when analyzing a tax on an input. This procedure is appropriate if the most *systematic* effects of a commodity tax are on the uses of income and those of a factor tax on the sources of income. The assumption simplifies analyses, but its correctness must be considered for each case.

Incidence Depends on How Prices Are Determined

We have emphasized that the incidence problem is fundamentally one of determining how taxes change prices. Clearly, different models of price determination may give quite different answers to the question of who really bears a tax. This chapter considers several different models and compares the results.

A closely related issue is the time dimension of the analysis. Incidence depends on changes in prices, but change takes time. In most cases, responses are larger in the long run than the short run. Thus, the short- and long-run incidence of a tax may differ, and the time frame that is relevant for a given policy question must be specified.

Incidence Depends on the Disposition of Tax Revenues

Balanced-budget incidence computes the combined effects of levying taxes *and* government spending financed by those taxes. In general, the distributional effect of a tax depends on how the government spends the money. Expenditures on AIDS research have a very different distributional impact than spending on hot lunches for schoolchildren. Some studies assume the government spends the tax revenue exactly as the consumers would if they had received the money. This is equivalent to returning the revenue as a lump sum and letting consumers spend it.

Tax revenues are usually not earmarked for particular expenditures. It is then desirable to be able to abstract from the question of how the government spends the money. The idea is to examine how incidence differs when one tax is replaced with another, holding the government budget constant. This is called *differential tax incidence*. Because differential incidence looks at changes in taxes, a reference point is needed. The hypothetical "other tax" used as the basis of comparison is often assumed to be a **lump sum tax**—a tax for which the individual's liability does not depend upon behavior. (For example, a 10 percent income tax is *not* a lump sum tax because it depends on how much the individual earns. But a head tax of $500 independent of earnings *is* a lump sum tax.)

lump sum tax

A tax whose value is independent of the individual's behavior.

Finally, *absolute tax incidence* examines the effects of a tax when there is no change in either other taxes or government expenditure. Absolute incidence is of most interest for macroeconomic models in which tax levels are changed to achieve some stabilization goal.

Tax Progressiveness Can Be Measured in Several Ways

Suppose that an investigator has managed to calculate every person's real share of a particular tax—the economic incidence as defined previously. The bottom line of such an exercise is often a characterization of the tax as proportional, progressive, or regressive. The definition of **proportional** is straightforward; it describes a situation in which the ratio of taxes paid to income is constant regardless of income level.[3]

proportional

A tax system under which an individual's average tax rate is the same at each level of income.

average tax rate

Ratio of taxes paid to income.

Defining progressive and regressive is not easy and, unfortunately, ambiguities in definition sometimes confuse public debate. A natural way to define these words is in terms of the **average tax rate,** the ratio of taxes paid to income. If the average tax rate increases with income, the system is **progressive;** if it falls, the tax is **regressive.**

progressive

A tax system under which an individual's average tax rate increases with income.

regressive

A tax system under which an individual's average tax rate decreases with income.

Confusion arises because some people think of progressiveness in terms of the **marginal tax rate**—the *change* in taxes paid with respect to a change in income. To illustrate the distinction, consider the following very simple income tax structure. Each individual computes her tax bill by subtracting $3,000 from income and paying an amount equal to 20 percent of the remainder. (If the difference is negative, the individual gets a subsidy equal to 20 percent of the figure.) Table 14.1 shows the amount of tax paid, the average tax rate, and the marginal tax rate for each of several income levels. The average rates increase with income. However, the marginal tax rate is constant at 0.2 because for each additional dollar earned, the individual pays an additional 20 cents, regardless of income level. People could disagree about the progressiveness of this tax system and each be right according to their own definitions. It is therefore very important to make the definition clear when using the terms *regressive* and *progressive*. From here on, we assume they are defined in terms of average tax rates.

marginal tax rate

The proportion of the last dollar of income taxed by the government.

[3] However, the definition of income is not straightforward; see Chapter 17.

Table 14.1 **Tax liabilities under a hypothetical tax system**

Income	Tax Liability	Average Tax Rate	Marginal Tax Rate
$2,000	$−200	−0.10	0.2
3,000	0	0	0.2
5,000	400	0.08	0.2
10,000	1,400	0.14	0.2
30,000	5,400	0.18	0.2

Under this hypothetical tax system, each individual computes her tax bill by subtracting $3,000 from income and paying an amount equal to 20 percent of the remainder. While the marginal tax rate is constant at 20 percent, the average tax rate is increasing as income increases, which means the tax is progressive.

Measuring *how* progressive a tax system is presents an even harder task than defining progressiveness. Many reasonable alternatives have been proposed, and we consider two simple ones. The first says that the greater the increase in average tax rates as income increases, the more progressive the system. Algebraically, let T_0 and T_1 be the true (as opposed to statutory) tax liabilities at income levels I_0 and I_1, respectively (I_1 is greater than I_0). The measurement of progressiveness, v_1, is

$$v_1 = \frac{\frac{T_1}{I_1} - \frac{T_0}{I_0}}{I_1 - I_0} \tag{14.1}$$

Once the analyst computes the values of T_1 and T_0 and substitutes into Equation (14.1), the tax system with the higher value of v_1 is said to be more progressive.

The second possibility is to say that one tax system is more progressive than another if its elasticity of tax revenues with respect to income (i.e., the percentage change in tax revenues divided by percentage change in income) is higher. Here the expression to be evaluated is v_2, defined as

$$v_2 = \frac{(T_1 - T_0)}{T_0} \div \frac{(I_1 - I_0)}{I_0} \tag{14.2}$$

Now consider the following proposal: Everyone's tax liability is to be increased by 20 percent of the amount of tax he or she currently pays. This proposal would increase the tax liability of a person who formerly paid T_0 to $1.2 \times T_0$, and the liability that was formerly T_1 to $1.2 \times T_1$. Member of Congress A says the proposal will make the tax system more progressive, while member of Congress B says it has no effect on progressiveness whatsoever. Who is right? It depends on the progressivity measure. Substituting the expressions $1.2 \times T_0$ and $1.2 \times T_1$ for T_0 and T_1, respectively, in Equation (14.1), v_1 increases by 20 percent. The proposal thus increases progressiveness. On the other hand, if the same substitution is done in Equation (14.2), the value of v_2 is unchanged. (Both the numerator and denominator are multiplied by 1.2, which cancels out the effect.) The lesson here is that even very intuitively appealing measures of progressiveness can give different answers.[4] Again, intelligent public debate requires that people make their definitions clear.

[4] Note also that v_1 and v_2, in general, depend on the level of income. That is, even a single tax system does not usually have a constant v_1 and v_2. This further complicates discussions of the degree of progressiveness.

▸ PARTIAL EQUILIBRIUM MODELS

With preliminaries out of the way, we turn now to the fundamental issue of this chapter: how taxes affect the income distribution. Recall that the essence of the problem is that taxes induce changes in relative prices. Knowing how prices are determined is therefore critical to the analysis. In this section we analyze **partial equilibrium models** of price determination—models that look only at the market in which the tax is imposed and ignore the ramifications in other markets. This kind of analysis is most appropriate when the market for the taxed commodity is relatively small compared to the economy as a whole. The vehicle for our analysis is the supply and demand model of perfect competition.

partial equilibrium models

Models that study only one market and ignore possible spillover effects in other markets.

Unit Taxes on Commodities

We study first the incidence of a **unit tax,** so named because it is levied as a fixed amount per unit of a commodity sold. For example, the federal government imposes a tax on champagne of $3.40 per wine gallon and a tax on cigarettes of 39 cents per pack. Suppose that the price and quantity of champagne are determined competitively by supply (S_c) and demand (D_c) as in Figure 14.1. Before imposition of the tax, the quantity demanded and price are Q_0 and P_0, respectively.

unit tax

A tax levied as a fixed amount per unit of commodity purchased.

Now suppose that a unit tax of $\$u$ per gallon is imposed on each purchase, and the statutory incidence is on buyers. A key step in incidence analysis is to recognize that in the presence of a tax, the price paid by consumers and the price received by suppliers differ. Previously, we could use a supply-demand analysis to determine the *single* market price. Now, this analysis must be modified to accommodate two different prices, one for buyers and one for sellers.

We begin by determining how the tax affects the demand schedule. Consider an arbitrary point a on the demand curve. This point indicates that the *maximum* price

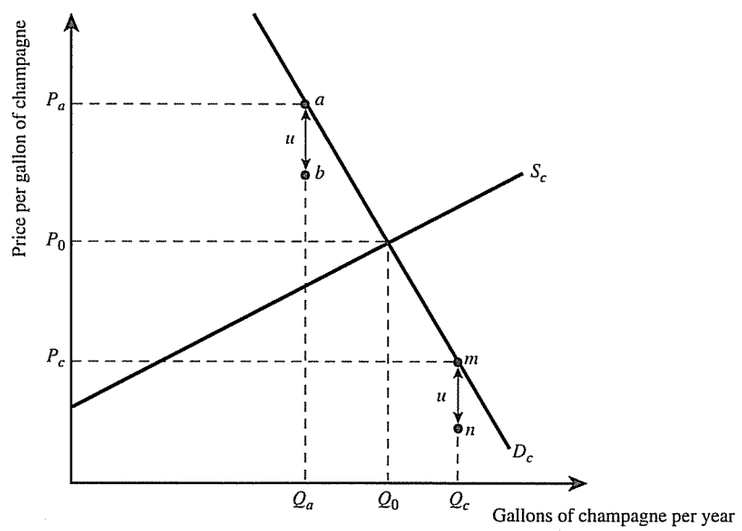

Figure 14.1
Price and quantity before taxation
A unit tax of $\$u$ per gallon changes the demand curve as perceived by suppliers. For example, the maximum price per gallon that people are willing to pay for Q_a is P_a. After the tax, when people pay P_a per gallon, producers only receive $P_a - u$ per gallon (which corresponds to point b). The new demand curve is located exactly u dollars below the old one.

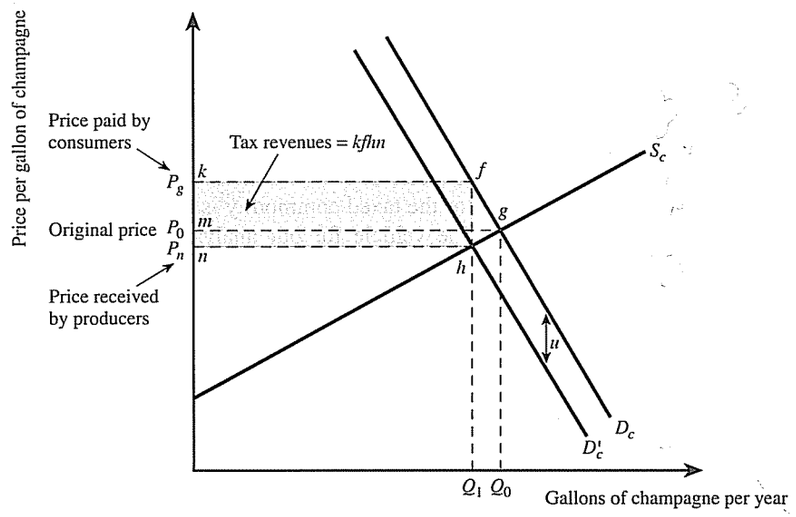

Figure 14.2
Incidence of a unit tax imposed on the demand side
After the imposition of the unit tax on consumers, the new equilibrium quantity is Q_1. The price received by producers is P_n, and the price paid by consumers is P_n plus u, which is P_g.

per gallon that people would be willing to pay for Q_a gallons is P_a. After the unit tax of u is imposed, the most that people would be willing to spend for Q_a is *still* P_a. There is no reason to believe the tax affects the underlying valuation people place on champagne. However, when people pay P_a per gallon, producers no longer receive the whole amount. Instead, they receive only $(P_a - u)$, an amount that is labeled point b in Figure 14.1. In other words, after the unit tax is imposed, a is no longer a point on the demand curve *as perceived by suppliers*. Point b is on the demand curve as perceived by suppliers, because they realize that if Q_a is supplied, they receive only $(P_a - u)$ per gallon. It is irrelevant to the suppliers how much consumers pay per gallon; all that matters to suppliers is the amount they receive per gallon.

Of course, point a was chosen arbitrarily. At any other point on the demand curve, the story is just the same. Thus, for example, after the tax is imposed, the price received by suppliers for output Q_c is at point n, which is found by subtracting the distance u from point m. Repeating this process at every point along the demand curve, we generate a new demand curve located exactly u dollars below the old one. In Figure 14.2, the demand curve so constructed is labeled D_c'. Schedule D_c' is relevant to suppliers because it shows how much they receive for each unit sold.

We are now in a position to find the equilibrium quantity of champagne after the unit tax is imposed. The equilibrium is where the supply equals demand as perceived by suppliers, output Q_1 in Figure 14.2. Thus, the tax lowers the quantity sold from Q_0 to Q_1.

The next step is to find the new equilibrium price. As noted earlier, there are really two prices at the new equilibrium: the price received by producers, and the price paid by consumers. The price received by producers is at the intersection of their effective demand and supply curves, which occurs at P_n. The price paid by consumers is P_n *plus* u, the unit tax. To find this price geometrically, we must go up from P_n a vertical distance exactly equal to u. But by construction, the distance between schedules D_c and D_c' is equal to u. Hence, to find the price paid by

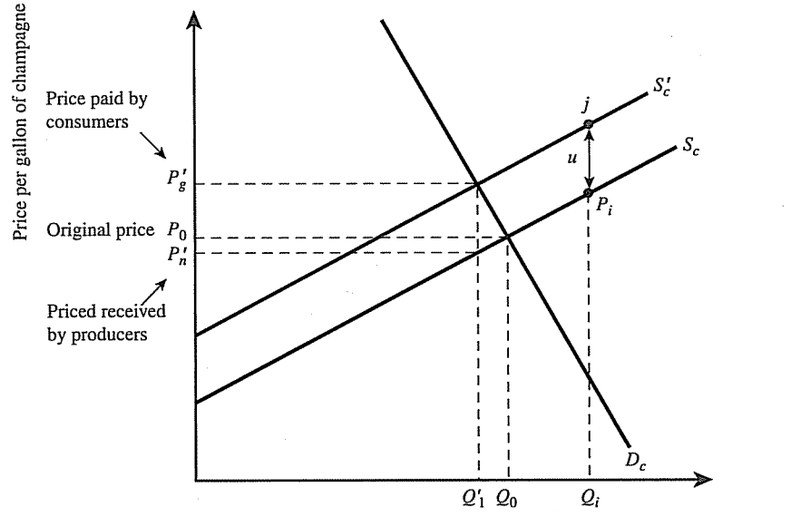

Figure 14.3
Incidence of a unit tax imposed on the supply side
A unit tax imposed on suppliers shifts up the supply curve by the amount of the tax. The posttax equilibrium quantity, price to consumers, and price to suppliers is the same as when the statutory incidence is on consumers.

consumers, we simply go up from the intersection of D_c' and S_c to the original demand curve D_c. The price so determined is P_g. Because P_g includes the tax, it is often referred to as the price *gross* of tax. On the other hand, P_n is the price *net* of tax.

The tax makes consumers worse off because P_g, the new price they face, is higher than the original price P_0. But the consumers' price does not increase by the full amount of the tax—$(P_g - P_0)$ is less than u. Producers also pay part of the tax in the form of a lower price received per gallon. Producers now receive only P_n, while before the tax they received P_0. Thus, the tax makes both producers and consumers worse off.[5] Notice that consumers and producers "split" the tax in the sense that the increase in the consumer price $(P_g - P_0)$ and the decrease in the producer price $(P_0 - P_n)$ just add up to $\$u$.

By definition, revenues collected are the product of the number of units purchased, Q_1, and the tax per unit, u. Geometrically, Q_1 is the width of rectangle *kfhn* and u is its height, so tax revenues are the area of this rectangle.

This analysis has two important implications:

The Incidence of a Unit Tax Is Independent of Whether It Is Levied on Consumers or Producers Suppose the same tax u had been levied on the suppliers of champagne instead of the consumers. Consider an arbitrary price P_i on the original supply curve in Figure 14.3. The supply curve indicates that for suppliers to produce Q_i units, they must receive at least P_i per unit. After the unit tax, suppliers still need to receive P_i per unit. For them to do so, however, consumers must pay price $P_i + u$ per unit, which is shown geometrically as point j. It should

[5] In terms of surplus measures, consumers are worse off by area *mkfg* and producers are worse off by *mghn*. The loss of total surplus exceeds the tax revenues by triangle *fhg;* this is the *excess burden* of the tax, as explained in Chapter 15. For a review of consumer and producer surplus, see the appendix at the end of this book.

now be clear where the argument is heading. To find the supply curve as it is perceived by consumers, S_c must be shifted up by the amount of the unit tax. This new supply curve is labeled S_c'. The posttax equilibrium is at Q_1', where the schedules S_c' and D_c intersect. The price at the intersection, P_g', is the price paid by consumers. To find the price received by producers, we must subtract u from P_g', giving us P_n'. A glance at Figure 14.2 indicates that $Q_1' = Q_1$, $P_g' = P_g$, and $P_n' = P_n$. Thus, the incidence of the unit tax is independent of the side of the market on which it is levied.

This is the same as our statement that the statutory incidence of a tax tells us nothing of the economic incidence of the tax. It is irrelevant whether the tax collector (figuratively) stands next to consumers and takes u dollars every time they pay for a gallon of champagne or stands next to sellers and collects u dollars from them whenever they sell a gallon. Figures 14.2 and 14.3 prove that what matters is the size of the disparity the tax introduces between the price paid by consumers and the price received by producers, and not on which side of the market the disparity is introduced. The tax-induced difference between the price paid by consumers and the price received by producers is referred to as the **tax wedge.**

tax wedge

The tax-induced difference between the price paid by consumers and the price received by producers.

The Incidence of a Unit Tax Depends on the Elasticities of Supply and Demand In Figure 14.2, consumers bear the brunt of the tax—the amount they pay goes up much more than the amount received by producers goes down. This result is strictly determined by the shapes of the demand and supply curves. In general, the more elastic the demand curve, the less the tax borne by consumers, other things being the same. Similarly, the more elastic the supply curve, the less the tax borne by producers, other things being the same. Intuitively, elasticity provides a rough measure of an economic agent's ability to escape the tax. The more elastic the demand, the easier it is for consumers to turn to other products when the price goes up, and therefore more of the tax must be borne by suppliers. Conversely, if consumers purchase the same amount regardless of price, the whole burden can be shifted to them. Similar considerations apply to the supply side.

Illustrations of extreme cases are provided in Figures 14.4 and 14.5. In Figure 14.4, commodity X is supplied perfectly inelastically. When a unit tax is imposed, the effective demand curve becomes D_X'. As before, the price received by producers (P_n) is at the intersection of S_X and D_X'. Note that P_n is exactly u less than P_0. Thus, the price received by producers falls by exactly the amount of the tax. At the same time, the price paid by consumers, P_g ($= P_n + u$), remains at P_0. When supply is perfectly inelastic, producers bear the entire burden. Figure 14.5 represents an opposite extreme. The supply of commodity Z is perfectly elastic. Imposition of a unit tax leads to demand curve D_Z'. At the new equilibrium, quantity demanded is Z_1 and the price received by producers, P_n, is still P_0. The price paid by consumers, P_g, is therefore $P_0 + u$. In this case, consumers bear the entire burden of the tax.[6]

The Cigarette Tax Debate Recently, the United States has been engaging in a major policy debate regarding cigarette taxation. In 2000, the 24-cent-per-pack federal tax was raised to 34 cents, and it is now 39 cents. But certain legislators would like to go further and increase the tax to $1 or more. Some proponents of the higher

[6] Note that as long as input costs are constant, the *long-run* supply curve for a competitive market is horizontal as in Figure 14.5. Hence, under these conditions, in the long run consumers bear the entire burden of the tax.

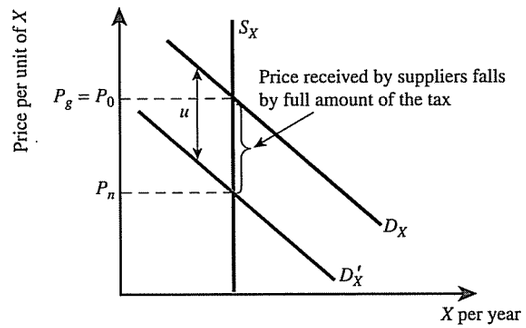

Figure 14.4
Tax incidence when supply is perfectly inelastic
A unit tax on a good that has perfectly inelastic supply causes the price received by producers to fall by exactly the amount of the tax. Producers therefore bear the entire burden of the tax.

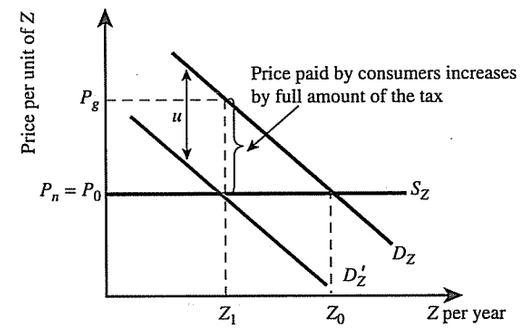

Figure 14.5
Tax incidence when supply is perfectly elastic
A unit tax on a good that has perfectly elastic supply causes the price paid by consumers to increase by exactly the amount of the tax. Consumers therefore bear the entire burden of the tax.

tax seem to be interested primarily in discouraging smoking, and others care more about punishing tobacco producers. Those who want to discourage smoking are implicitly assuming that the tax will drive up the price paid by consumers, and those who want to punish the tobacco producers expect the price they receive to go down. How can one determine which effect would prevail? Our model of tax incidence tells us what we need to find out: the supply and demand elasticities in the cigarette market.

Ad Valorem Taxes

We now turn to the incidence of an **ad valorem tax,** a tax with a rate given as a *proportion* of the price. For example, the state of Tennessee levies a 6 percent tax on purchases of food. Virtually all state and local taxes on restaurant meals and clothing are ad valorem.

ad valorem tax

A tax computed as a percentage of the purchase value.

Luckily, the analysis of ad valorem taxes is very similar to that of unit taxes. The basic strategy is still to find out how the tax changes the effective demand curve and compute the new equilibrium. However, instead of moving the curve down by the same absolute amount for each quantity, the ad valorem tax lowers it by the same *proportion*. To show this, consider the demand (D_f) and supply (S_f) curves for food

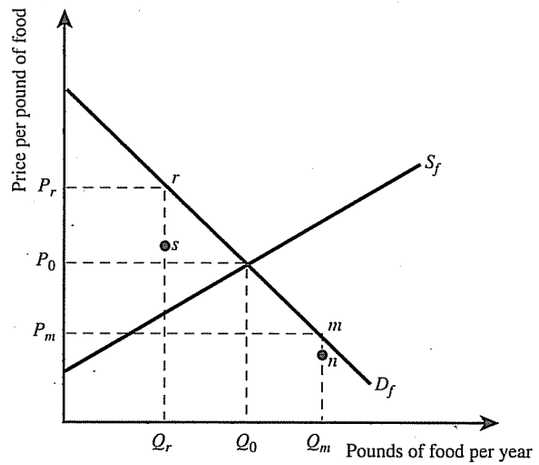

Figure 14.6
Introducing an ad valorem tax
An ad valorem tax on consumers shifts the demand curve down by the same proportion at each level of output.

in Figure 14.6. In the absence of taxation, the equilibrium price and quantity are P_0 and Q_0, respectively. Now suppose that a tax of 25 percent of the gross price is levied on the consumption of food.[7] Consider point m on D_f. After the tax is imposed, P_m is still the most that consumers will pay for Q_m pounds of food; the amount producers will receive is 75 percent of the vertical distance between point m and the horizontal axis, which is labeled point n. Hence, point n is one point on the demand curve perceived by producers. Similarly, the price at point r migrates down one quarter of the way between it and the horizontal axis to point s. Repeating this exercise for every point on D_f, the effective demand curve facing suppliers is determined as D_f' in Figure 14.7. From here, the analysis proceeds exactly as for a unit tax: The equilibrium is where S_f and D_f' intersect, with the quantity exchanged Q_1, the price received by food producers P_n, and the price paid by consumers P_g. As before, the incidence of the tax is determined by the elasticities of supply and demand.

This analysis is applicable to any number of situations. Suppose that Figure 14.7 were relabeled so that it represented the market for rental housing instead of the food market. Then we could show that the burden of the property tax doesn't depend on whether landlords or tenants pay the property tax. This is counter to the usual perception that landlords bear the burden simply because they write the check.

Taxes on Factors

So far we have discussed taxes on goods, but the analysis can also be applied to factors of production.

The Payroll Tax Consider the payroll tax used to finance the Social Security system. As noted in Chapter 11, a tax equal to 7.65 percent of workers' earnings

[7] Measuring ad valorem tax rates involves a fundamental ambiguity. Is the tax measured as a percentage of the net or gross price? In this example, the tax is 25 percent of the gross price, which is equivalent to a rate of 33 percent of net price. If the price paid by the consumer were $1, the tax paid would be 25 cents, and the price received by producers would be 75 cents. Expressing the 25 cent tax bill as a fraction of 75 cents gives us a 33 percent rate as a proportion of the net price.

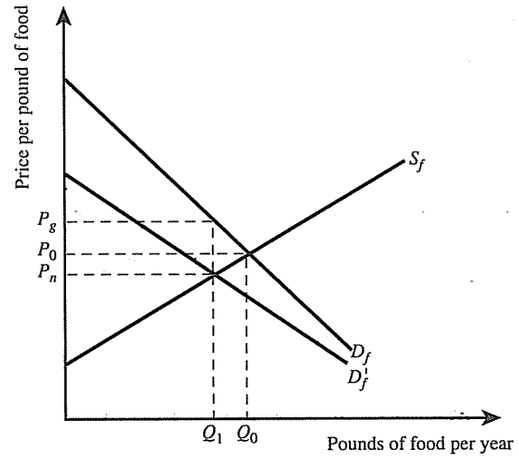

Figure 14.7
Incidence of an ad valorem tax
After the imposition of an ad valorem tax, the new equilibrium quantity is Q_1, the price received by the producers is P_n, and the price paid by consumers is P_g.

must be paid by their employers and a tax at the same rate paid by the workers themselves—a total of 15.3 percent.[8] This division has a long history and is a consequence of our lawmakers' belief that the payroll tax should be shared equally by employers and employees. But the *statutory distinction between workers and bosses is irrelevant*. As suggested earlier, the incidence of this labor tax is determined only by the wedge the tax puts between what employees receive and employers pay.

This point is illustrated in Figure 14.8, where D_L is the demand for labor and S_L is the supply of labor. For purposes of illustration, assume S_L to be perfectly inelastic. Before taxation, the wage is w_0. The ad valorem tax on labor moves the effective demand curve to D_L'. As usual, the distance between D_L' and D_L is the wedge between what is paid for an item and what is received by those who supply it. After the tax is imposed, the wage received by workers falls to w_n. On the other hand, w_g, the price paid by employers, stays at w_0. In this example, despite the statutory division of the tax, the wage rate received by workers falls by exactly the amount of the tax—they bear the entire burden.

Of course, we could have gotten just the opposite result by drawing the supply curve as perfectly elastic. The key point to remember is that nothing about the incidence of a tax can be known without information on the relevant behavioral elasticities. In fact, while estimates of the elasticity of labor supply vary, many economists believe that it is close to zero [Fuchs et al., 1998]. At least in the short run, labor probably bears most of the payroll tax, despite the congressional attempt to split the burden evenly.

Capital Taxation in a Global Economy The strategy for analyzing a tax on capital is essentially the same as that for analyzing a tax on labor—draw the supply and demand curves, shift or pivot the relevant curve by an amount depending on the tax rate, and see how the after-tax equilibrium compares with the original one. In an economy that is closed to trade, it is reasonable to assume that the demand curve slopes down (firms demand less capital when its price goes up), and that the supply

[8] After earnings exceed a certain level, the payroll tax rate falls. See Chapter 11.

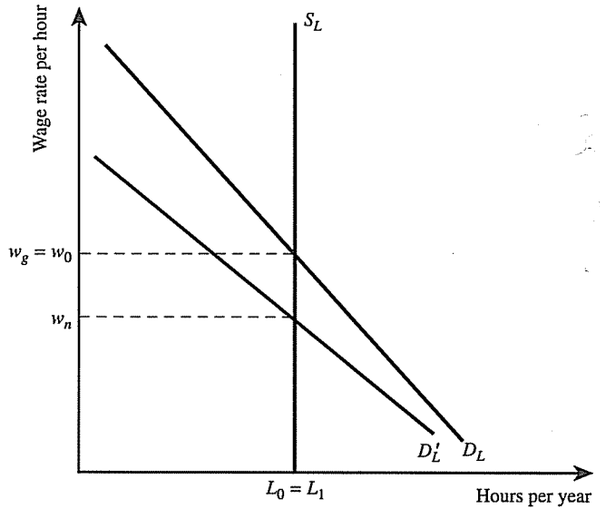

Figure 14.8
Incidence of a payroll tax with an inelastic supply of labor
If labor supply is perfectly inelastic, a payroll tax causes the wage received by workers to fall by the exact amount of the tax. Workers therefore bear the entire burden of the tax.

of capital slopes up (people supply more capital—save more—when the return to saving increases).[9] In this case, the owners of capital bear some of the burden of the tax, the precise amount depending on the supply and demand elasticities.

Suppose now that the economy is open and capital is perfectly mobile across countries. In effect, there is a single global market for capital, and if suppliers of capital cannot earn the going world rate of return in a particular country, they will take it out of that country and put it in another. In terms of a supply and demand diagram, the supply of capital to a particular country is perfectly elastic—its citizens can purchase all the capital they want at the going rate of return, but none whatsoever at a lower rate. The implications for the incidence of a tax on capital are striking. As in Figure 14.5, the before-tax price paid by the users of capital rises by exactly the amount of the tax, and the suppliers of capital bear no burden whatsoever. Intuitively, capital simply moves abroad if it has to bear any of the tax; hence, the before-tax rate of return has to rise.

Now, even in today's highly integrated world economy, capital is not perfectly mobile across countries. Moreover, for a country like the United States whose capital market is large relative to the world market, it is doubtful that the supply curve is perfectly horizontal. Nevertheless, policymakers who ignore globalization will tend to overestimate their ability to place the burden of taxation on owners of capital. To the extent that capital is internationally mobile, taxes on capitalists are shifted to others, and the apparent progressivity of taxes on capital is illusory.

Commodity Taxation without Competition

The assumption of competitive markets has played a major role in our analysis. We now discuss how the results might change under alternative market structures.

[9] However, saving need not increase with the rate of return. See Chapter 18.

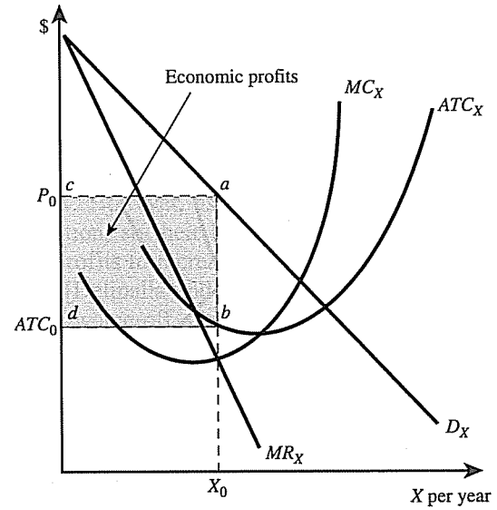

Figure 14.9
Equilibrium of a monopolist
The monopolist produces X_0 per year, charges a price of P_0, and receives profits of area *abdc*.

Monopoly The polar opposite of competition is monopoly—one seller. Figure 14.9 depicts a monopolist that produces commodity X. Before any taxation, the demand curve facing the monopolist is D_X, and the associated marginal revenue curve is MR_X. The marginal cost curve for the production of X is MC_X, and the average total cost curve, ATC_X. As usual, the condition for profit maximization is that production be carried to the point where marginal revenue equals marginal cost, at output X_0 where the price charged is P_0. Economic profit per unit is the difference between average revenue and average total cost, distance *ab*. The number of units sold is *db*. Hence, total profit is *ab* times *db*, which is the area of rectangle *abdc*.

Now suppose that a unit tax of u is levied on X. For exactly the same reasons as before, the effective demand curve facing the producer shifts down by a vertical distance equal to u.[10] In Figure 14.10, this demand curve is labeled D'_X. At the same time, the marginal revenue curve facing the firm also shifts down by distance u because the tax reduces the firm's incremental revenue for each unit sold. The new effective marginal revenue curve is labeled MR'_X.

The profit-maximizing output, X_1, is found at the intersection of MR'_X and MC_X. Using output X_1, we find the price received by the monopolist by going up to D'_X, the demand curve facing him, and locate price P_n. The price paid by consumers is determined by adding u to P_n, which is shown as price P_g on the diagram. After-tax profit per unit is the difference between the price *received by the monopolist* and average total cost, distance *fg*. Number of units sold is *if*. Therefore, monopoly economic profits after tax are measured by area *fghi*.

What are the effects of the tax? Quantity demanded goes down ($X_1 < X_0$), the price paid by consumers goes up ($P_g > P_0$), and the price received by the monopolist goes down ($P_n < P_0$). Note that monopoly profits are lower under

[10] Alternatively, we could shift the marginal cost curve *up* by u. The final outcomes are identical.

Figure 14.10
Imposition of a unit tax on a monopolist
The imposition of a unit tax on a monopolistically produced good shifts the effective demand curve and the marginal revenue curve down by the amount of the tax. The tax reduces the equilibrium quantity from X_0 to X_1, increases the price paid by consumers from P_0 to P_g, decreases the price received by the producer from P_0 to P_n, and decreases the monopolist's profits from area *abcd* to area *fghi*.

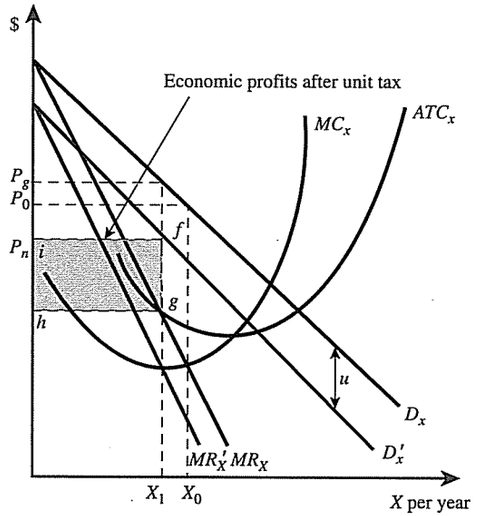

the tax—area *fghi* in Figure 14.10 is smaller than area *abdc* in Figure 14.9. Despite its market power, a monopolist is generally made worse off by a unit tax on the product it sells. Public debates often assume that a firm with market power can simply pass on all taxes to consumers. This analysis shows that even a completely greedy and grasping monopolist must bear some of the burden. As before, the precise share of the burden borne by consumers depends on the elasticity of the demand schedule.

It is straightforward to repeat the exercise for an ad valorem tax on the monopolist (D_X and MR_X pivot instead of moving down in a parallel fashion); this is left as an exercise for the reader.

Oligopoly Between the polar extremes of perfect competition and monopoly is the oligopoly market structure in which there are a "few" sellers. Unfortunately, there is no well-developed theory of tax incidence in oligopoly. The reason for this embarrassing fact is simple: Incidence depends primarily on how relative prices change when taxes are imposed, but there is no generally accepted theory of oligopolistic price determination.

Still, we can get a sense of the issues involved by imagining the problem faced by the firms in an oligopolistic market. From the firms' point of view, the ideal situation would be for them to collude and jointly produce the output that maximizes the profits of the entire industry. This output level is referred to as the *cartel solution*. (A cartel is just a group of producers that act together to maximize profits. The international oil cartel OPEC is the most famous example.) The cartel solution requires each firm to cut its output to force up the market price. The problem for the firms is that the cartel solution is very difficult to obtain. Why? Once an agreement about how much each firm should produce is reached, each firm has an incentive to cheat on that agreement—to take advantage of the higher price and produce more than its quota of output. (Again, think about OPEC, and the problems it has in keeping its members from producing "too much" oil.) Consequently, output in an oligopolistic market is typically higher than the cartel solution. The firms would all be better off if there were some mechanism to force all of them to reduce their output.

What happens when this industry's output is subjected to a tax? As is the case both for competition and monopoly, the firms reduce their output. However, unlike the other market structures, this is not necessarily bad for the oligopolistic firms. To be sure, for any given level of before-tax profits, the firms are worse off, because they have to pay the tax. However, as the firms contract their outputs, they move closer to the cartel solution, so their before-tax profits increase. It is theoretically possible for before-tax profits to increase by so much that even after paying the tax, the firms are better off [Delipalla and O'Donnell, 2001]. Of course, it is also possible for the firms to be worse off. One needs more information on just how much the firms cut back their output to obtain a definitive answer.

As economic behavior under oligopoly becomes better understood, improved models of incidence will be developed. In the meantime, most economists feel fairly comfortable in relying on the predictions produced by competitive models, although they realize these are only approximations.

Profits Taxes

So far we have been discussing taxes based on sales. Firms can also be taxed on their **economic profits,** defined as the return to owners of the firm in excess of the opportunity costs of the factors used in production. (Economic profits are also referred to as *supranormal* or *excess* profits.) We now show that for profit-maximizing firms, a tax on economic profits cannot be shifted—it is borne only by the owners of the firm.

economic profit
The return to owners of a firm above the opportunity costs of all the factors used in production. Also called supranormal or excess profit.

Consider first a perfectly competitive firm in short-run equilibrium. The firm's output is determined by the intersection of its marginal cost and marginal revenue schedules. A proportional tax on economic profits changes neither marginal cost nor marginal revenue. Therefore, no firm has the incentive to change its output decision. Because output does not change, neither does the price paid by consumers, so they are no worse off. The tax is completely absorbed by the firms. Here's another way to get to the same result: If the tax rate on economic profits is t_p, the firm's objective is to maximize after-tax profits, $(1 - t_p)\Pi$, where Π is the pretax level of economic profits. But it is just a matter of arithmetic that whatever strategy maximizes Π is identical to the one that maximizes $(1 - t_p)\Pi$. Hence, output and price faced by consumers stay the same, and the firm bears the whole tax.

In long-run competitive equilibrium, a tax on economic profits has no yield, because economic profits are zero—they are all competed away. For a monopolist, there may be economic profits even in the long run. But for the same reasons given in the preceding paragraph, the tax is borne by the owners of the monopoly. If a firm is maximizing profits before the profits tax is imposed, the tax cannot be shifted.[11]

Because they distort no economic decisions, taxes on economic profits might appear to be very attractive policy alternatives. In 2006, for example, certain

[11] On the other hand, if the firm is following some other goal, it may raise the price in response to a profits tax. One alternative to profit maximization is revenue maximization; firms try to make their sales as large as possible, subject to the constraint that they earn a "reasonable" rate of return.

members of both political parties called for a "profits tax" on oil companies. However, profits taxes receive very little support from public finance specialists. The main reason is the tremendous problems in making the theoretical notion of economic profits operational. Economic profits are often computed by examining the rate of return that a firm makes on its capital stock and comparing it to some "basic" rate of return set by the government. Clearly, how the capital stock is measured is important. Should the original cost be used, or the cost of replacing it? And what if the rate of return is high not because of excess profits, but because the enterprise is very risky and investors have to be compensated for this risk? Considerations like these lead to major difficulties in administration and compliance.

Tax Incidence and Capitalization

Several years ago the coastal city of Port Hueneme, California, levied a special tax on beach properties. The tax was determined in part by how close the properties were to the ocean. For owners close to the water, the extra tax was $192 per year. Owners of beachfront property complained vociferously.

This episode leads us to consider the special issues that arise when land is taxed. For these purposes, the distinctive characteristics of land are that it is fixed in supply and it is durable. Suppose the annual rental rate on land is $\$R_0$ this year. It is known that the rental will be $\$R_1$ next year, $\$R_2$ two years from now, and so on. How much should someone be willing to pay for the land? If the market for land is competitive, its price is just equal to the present discounted value of the stream of the rents. Thus, if the interest rate is r, the price of land (P_R) is

$$P_R = \$R_0 + \frac{\$R_1}{1+r} + \frac{\$R_2}{(1+r)^2} + \cdots + \frac{\$R_T}{(1+r)^T} \tag{14.3}$$

where T is the last year the land yields its services (possibly infinity).

Now it is announced that a tax of $\$u_0$ will be imposed on land now, $\$u_1$ next year, $\$u_2$ two years from now, and so forth. From Figure 14.4 we know that because land is fixed in supply, the annual rental received by the owner falls by the full amount of the tax. Thus, the landlord's return initially falls to $\$(R_0 - u_0)$, in year 1 to $\$(R_1 - u_1)$, in year 2 to $\$(R_2 - u_2)$, and so on. Prospective purchasers of the land take into account the fact that if they purchase the land, they buy a future stream of tax liabilities as well as a future stream of returns. Therefore, the most a purchaser is willing to pay for the land after the tax is announced (P_R') is

$$P_R' = \$(R_0 - u_0) + \frac{\$(R_1 - u_1)}{1+r} + \frac{\$(R_2 - u_2)}{(1+r)^2} + \cdots + \frac{\$(R_T - u_T)}{(1+r)^T} \tag{14.4}$$

Comparing Equations (14.4) and (14.3), we see that as a consequence of the tax, the price of land falls by

$$u_0 + \frac{u_1}{1+r} + \frac{u_2}{(1+r)^2} + \cdots + \frac{u_T}{(1+r)^T}$$

capitalization

The process by which a stream of tax liabilities becomes incorporated into the price of an asset.

Thus, at the time the tax is imposed, the price of the land falls by the present value of *all future tax payments*. This process by which a stream of taxes becomes incorporated into the price of an asset is referred to as **capitalization.**

Because of capitalization, the person who bears the full burden of the tax *forever* is the landlord at the time the tax is levied. To be sure, *future* landlords write checks to the tax authorities, but such payments are not really a "burden" because they just balance the lower price paid at purchase. Capitalization complicates attempts to assess the incidence of a tax on any durable item that is fixed in supply. Knowing the identities of current owners is not sufficient—one must know who the landlords *were* at the time the tax was imposed. It's no wonder the owners of beach property in Port Hueneme were so upset![12]

▸ GENERAL EQUILIBRIUM MODELS

A great attraction of partial equilibrium models is their simplicity—examining only one market at a time is relatively uncomplicated. In some cases, however, ignoring feedback into other markets leads to an incomplete picture of a tax's incidence. Suppose, for example, that the tax rate on cigarettes is increased. To the extent that the demand for cigarettes decreases, so does the demand for tobacco. Farmers who formerly grew tobacco on their land may turn to other crops, perhaps cotton. As the supply of cotton increases, its price falls, harming the individuals who were already producing cotton. Thus, cotton producers end up bearing part of the burden of a cigarette tax.

More generally, when a tax is imposed on a sector that is "large" relative to the economy, looking only at that particular market may not be enough. **General equilibrium analysis** takes into account the ways in which various markets are interrelated.

general equilibrium analysis

The study of how various markets are interrelated.

Another problem with partial equilibrium analysis is that it gives insufficient attention to the question of just who the "producers" of a taxed commodity are. Think again of the cigarette tax and the desire of some policymakers to use it as an instrument to punish "the tobacco industry." Only people can pay taxes, and the producers of tobacco include the shareholders who finance the purchase of machinery, farmers who own the land on which the tobacco is grown, the workers in the factories, and so on. The division of the tax burden among these groups is often important. General equilibrium analysis provides a framework for investigating it.

Before turning to the specifics of general equilibrium analysis, note that the fundamental lesson from partial equilibrium models still holds: Because of relative price adjustments, the statutory incidence of a tax generally tells *nothing* about who really bears its burden.

Tax Equivalence Relations

The idea of dealing with tax incidence in a general equilibrium framework at first appears daunting. After all, thousands of different commodities and inputs are traded in the economy. How can we keep track of all their complicated interrelations? Luckily,

[12] When a land tax is anticipated before it is levied, presumably it is borne at least in part by the owner at the time the anticipation becomes widespread. If so, even finding out the identity of the landowner at the time the tax was imposed may not be enough.